

MEASURE THEORY THROUGH DYNAMICAL EYES

ANATOLE KATOK AND VAUGHN CLIMENHAGA

These notes are a somewhat embellished version of two rather informal evening review sessions given by the first author on July 14 and 15 2008 at the Będlewo summer school (see http://www.math.psu.edu/katok_a/Bedlewo/school.html), which provide a brief overview of some of the basics of measure theory and its applications to dynamics which are foundational to the various courses at this school.

Most results are quoted without proof, or with at most a bare sketch of a proof; references are given where full proofs may be found. Similarly, most basic definitions are assumed to be known, and we defer their reiteration to the references.

In light of the above, we emphasise that this presentation is *not* meant to be either comprehensive or self-contained; the reader is assumed to have some knowledge of the basic concepts of measure theory, ergodic theory, and hyperbolic dynamics, which will appear without any formal introduction. The tone is meant to be conversational rather than authoritative, and the goal is to make accessible various concepts which should eventually be examined thoroughly in the appropriate references.

For a full presentation of the concepts in §1, which concerns abstract measure theory, we refer the reader to Halmos' book [Ha] (for more basic facts) and to Rokhlin's article [Ro1]. The topics in measurable dynamics mentioned in §2 receive a more complete treatment in a later article by Rokhlin [Ro2], and the account in §3 of the relationship between foliations and measures in smooth dynamics draws on Barreira and Pesin's book [BP], along with two articles by Ledrappier and Young [LY1, LY2].

The second author would like to thank Andrey Gogolev and Misha Guysinsky for providing useful references and comments.

1. ABSTRACT MEASURE THEORY

1.1. Points, sets, and functions. There are three “lenses” through which we can view measure theory; we may think of it in terms of *points*, in terms of *sets*, or in terms of *functions*. To put that a little more concretely, suppose we have a triple (X, \mathcal{B}, μ) comprising a measurable space, a σ -algebra, and a measure. Then we may focus our attention either on the space X (and concern ourselves with points), or on the σ -algebra \mathcal{B} (and concern ourselves with sets), or on the space $L^2(X, \mathcal{B}, \mu)$ (and concern ourselves with functions).

All three points of view play an important role in dynamics, but for the time being we will focus on the first two, and on the correspondence between partitions (of the space X) and σ -algebras (we will consider sub- σ -algebras $\mathcal{A} \subset \mathcal{B}$).

First let us consider the set of all partitions of X . This is a partially ordered set, with ordering given by refinement; given two partitions ξ, η , we say that ξ is a *refinement* of η , written $\xi \geq \eta$, iff every $C \in \xi$ is contained in some $D \in \eta$. In this

case, we also say that η is a *coarsening* of ξ . The finest partition (which in this notation may be thought of as the “largest”) is the partition into points, denoted ε , while the coarsest (the “smallest”) is the trivial partition $\{X\}$, denoted ν .

As on any partially ordered set, we have a notion of *join* and *meet*, corresponding to least upper bound and greatest lower bound, respectively. Following [Ro2], we shall refer to these as the *product* and *intersection*, and we briefly recall their definitions. Given two partitions ξ and η , their product (join) is

$$(1.1) \quad \xi \vee \eta := \{C \cap D \mid C \in \xi, D \in \eta\}$$

This is the coarsest partition which refines both ξ and η , and is also sometimes referred to as the *joint partition*. The intersection (meet) of ξ and η is the finest partition which coarsens both ξ and η , and is denoted $\xi \wedge \eta$; in general, there is no analogue of (1.1) for $\xi \wedge \eta$.

So much for partitions; what do these have to do with σ -algebras?¹ Suppose we partition our space X into some finite number of subsets C_1, \dots, C_n —denote this partition by ξ_1 . Then we may consider the σ -algebra which consists of all measurable sets (elements of \mathcal{B}) which are unions of none, some, or all of the C_i —denote this σ -algebra by \mathcal{B}_1 . If all the sets C_i are measurable, then \mathcal{B}_1 will contain 2^n sets.

Carry this a step further, and partition each C_i into $C_{i,1}, \dots, C_{i,k_i}$. Then we obtain another, larger σ -algebra \mathcal{B}_2 whose elements are unions of none, some, or all of the $C_{i,j}$. Iterating this procedure, we have a sequence of partitions

$$(1.2) \quad \xi_1 < \xi_2 < \dots$$

each of which is a refinement of the previous partitions, and a sequence of σ -algebras

$$(1.3) \quad \mathcal{B}_1 \subset \mathcal{B}_2 \subset \dots$$

The idea here is the same as that involved in constructing the σ -algebra of Lebesgue sets on $[0, 1]$ by using intervals, but without the extraneous geometric information which enters in that case.

In this picture, what do individual points of X correspond to? Provided we do our refining intelligently, each point will be represented by a “funnel”

$$(1.4) \quad C_{i_1} \supset C_{i_1, i_2} \supset C_{i_1, i_2, i_3} \supset \dots$$

of decreasing subsets within the sequence of partitions. A useful image to keep in mind here is the standard picture of the construction of a Cantor set, in which the unit interval is first divided into two pieces, then four, then eight, and so on—these “cylinders” (to use the terminology arising from symbolic dynamics) are the various sets denoted by $C_i, C_{i,j}$, etc. in (1.4).

1.2. Lebesgue spaces. In a state of affairs which Tolkein would surely render quite poetically,² there is in some sense one master σ -algebra, to which all others (or at least the “nice” ones) are isomorphic.³

¹Note that the set of σ -algebras $\mathcal{A} \subset \mathcal{B}$ carries a partial ordering as well, coming from the partial ordering by containment on the power set of \mathcal{B} .

²Particularly if we were to adopt a slightly different line of exposition and consider σ -rings instead of σ -algebras.

³We should really be speaking of *measured* σ -algebras, since we consider not just the σ -algebra, but the measure it carries. This will be implicit in our discussion throughout this section.

What do we mean by this? Perhaps the simplest imaginable class of measure spaces is the collection of *atomic* spaces, for which X is a finite or countable set, \mathcal{B} is the entire power set of X , and μ is defined by the sequence of numbers $\mu(x_i), x_i \in X$.

The canonical example of a measure space which is not quite so trivial is the interval $[0, 1]$ with Lebesgue measure; the σ -algebra of Lebesgue sets is the “master” example referred to above.

Under what conditions is a σ -algebra isomorphic to the Lebesgue sets on the interval? Put another way, what are the invariants of σ -algebras? Of course the number of atoms⁴ and the weight they carry is one invariant; are there any others?

There turns out to be just one. We may introduce a pseudo-metric on the σ -algebra \mathcal{B} by the formula

$$d_\mu(A, B) = \mu(A \triangle B)$$

where $A \triangle B$ denotes the symmetric difference $(A \cup B) \setminus (A \cap B)$. If we pass to the quotient space of \mathcal{B} by the equivalence relation of having measure zero symmetric difference, this is a true metric space; we say that the σ -algebra is *separable* if this metric space is separable. That is, \mathcal{B} is separable iff there exists a countable set $\{A_n\}_{n \in \mathbb{N}}$ which is dense in \mathcal{B} with the d_μ pseudo-metric. This allows us to state the main classification result for σ -algebras:

Theorem 1.1. *A separable σ -algebra with no atoms is isomorphic to the σ -algebra of Lebesgue sets on the unit interval.*

Proof. See [Ha, Section 41, Theorem C] □

As it stands, this is a classification result for the σ -algebra \mathcal{B} , not the space X .⁵ What condition on X will guarantee that it is isomorphic to $[0, 1]$ with Lebesgue measure?

We need the notion of a *basis*, which is to generate both the σ -algebra \mathcal{B} and the space X . To this end, consider an increasing sequence $\xi_1 < \xi_2 < \dots$ of finite partitions, as in (1.2), which is generating in the following ways:

- (1) It generates the σ -algebra \mathcal{B} ; that is, \mathcal{B} is the union of the corresponding σ -algebras $\mathcal{B}_n := \mathcal{B}(\xi_n)$ from (1.3).
- (2) It generates the space X ; that is, every “funnel” $C_{i_1} \supset C_{i_1, i_2} \supset \dots$ as in (1.4) has intersection containing at most one point. Equivalently, any two points x and y are separated by some partition ξ_n , and so $\bigvee_{n=1}^{\infty} \xi_n = \varepsilon$, the partition into points.

Note that the existence of an increasing sequence of partitions satisfying (1) is equivalent to separability of the σ -algebra.

Such a generating sequence of partitions is known as a *basis*—it is convenient to choose a sequence ξ_n such that at each stage, each cylinder set C is partitioned into exactly two smaller sets. This gives a one-to-one correspondence between sequences in $\Sigma_2^+ := \{0, 1\}^{\mathbb{N}}$ and “funnels” as in (1.4).

Exercise 1.1. Determine the correspondence between the above definition of a basis and the definition given in §1.2 of [Ro2].

⁴An *atom* is a measurable set A of positive measure such that every subset B has either $\mu(B) = 0$ or $\mu(B) = \mu(A)$.

⁵While we will not discuss the third way of thinking about measures, the L^2 spaces, it is worth comparing Theorem 1.1 to the classification theorem for separable Hilbert spaces. See [Ro2] for further continuations of this train of thought.

Since each “funnel” corresponds to some element of \mathcal{B} which is either a singleton or empty, we have associated to each subset of Σ_2^+ an element of \mathcal{B} , and so μ yields a measure on Σ_2^+ . Thus we have a notion of “almost all funnels”—we say that the basis is *complete* if almost every funnel contains exactly one point.⁶ That is, the set of funnels whose intersection is empty should be measurable, and should have measure zero. Equivalently, a basis defines a map from X to Σ_2^+ which takes each point to the “funnel” containing it; the basis is complete if the image of this map has full measure.

The existence of a complete basis is the final invariant needed to classify “nice” measure spaces.

Theorem 1.2. *If (X, \mathcal{B}, μ) is separable, non-atomic, and possesses a complete basis, then it is isomorphic to Lebesgue measure on the unit interval.*

Proof. See [Ro1]. □

Definition 1.3. A separable measure space (X, \mathcal{B}, μ) with a complete basis is called a *Lebesgue space*.

By Theorem 1.2, every Lebesgue space is isomorphic to the union of unit interval with at most countably many atoms.

It is also worth noting that any separable measure space admits a completion, just as is the case for metric spaces. The procedure is quite simple; take a basis for X which is not complete, and add to X one point corresponding to each empty “funnel”. Thus we need not concern ourselves with non-complete spaces.

Exercise 1.2. Show that every separable measure space which is *not* complete is isomorphic to a set of outer measure one in a Lebesgue space.

Since the measure spaces which arise in conjunction with dynamics are all separable, we will from now on restrict our attention to Lebesgue spaces.

1.3. Partitions and σ -algebras. We have seen how to go from a sequence of partitions to a sequence of σ -algebras, and in general, one can always pass from a partition ξ to a σ -algebra $\mathcal{A} \subset \mathcal{B}$ by taking for the elements of \mathcal{A} precisely those elements of \mathcal{B} which are unions of elements of ξ . We will usually denote this σ -algebra by $\mathcal{B}(\xi)$.

For example, consider the unit square with Lebesgue measure, and take the partition into vertical lines. To this partition we associate the σ -algebra \mathcal{A} whose elements are direct products of Lebesgue sets on the (horizontal) interval with the (vertical) interval $[0, 1]$.

The procedure in the reverse direction is somewhat more delicate. How do we go from a σ -algebra to a partition? Given \mathcal{A} , can we reconstruct ξ , at least in some situations? It turns out that we can; Theorem 1.5 below gives what turns out to be a sort of one-sided inverse to $\mathcal{B}(\cdot)$.

First we need an appropriate notion of equivalence for partitions and σ -algebras.

Definition 1.4. Two sets $E, F \in \mathcal{B}$ are *equivalent mod zero* if $d_\mu(E, F) = \mu(E \Delta F) = 0$, and we write $E \doteq F$.

Two partitions ξ, η of X are *equivalent mod zero* if there exists a set $E \subset X$ of full measure such that

$$\eta = \{C \cap E \mid C \in \xi\},$$

⁶Recall that there is also a notion of completeness for σ -algebras, distinct from this one.

and we write $\xi \doteq \eta$.

Two σ -algebras $\mathcal{A}, \mathcal{A}' \subset \mathcal{B}$ are equivalent mod zero if they have the same completion—that is, if given any set $E \in \mathcal{B}$, we have $E \doteq A$ for some $A \in \mathcal{A}$ iff we also have $E \doteq A'$ for some $A' \in \mathcal{A}'$. In this case we also write $\mathcal{A} \doteq \mathcal{A}'$.

For simplicity, from now on we will consider sets, partitions, and σ -algebras only up to equivalence mod zero (in particular, we will not distinguish between a σ -algebra and its completion), and we will write $=$ in place of \doteq .

Theorem 1.5. *Given a Lebesgue space (X, \mathcal{B}, μ) and a sub- σ -algebra $\mathcal{A} \subset \mathcal{B}$, there exists a partition ξ of X into measurable sets such that \mathcal{A} and $\mathcal{B}(\xi)$ are equivalent mod zero. ξ is unique up to equivalence mod zero, and we will denote it by $\Xi(\mathcal{A})$.*

Proof. Since the measure space is separable, we may take a generating system of partitions for the σ -algebra \mathcal{A} ; that is, $\xi_1 < \xi_2 < \dots$ such that the associated σ -algebras $\mathcal{B}(\xi_n)$ exhaust \mathcal{A} (i.e. $\mathcal{A} = \cup_n \mathcal{B}(\xi_n)$), and consider the partition of X into “funnels”. Thus our desired partition is obtained as $\xi = \vee_{n=1}^{\infty} \xi_n$.

The proof of uniqueness is left as an exercise. \square

Note that the construction in the proof mirrors the definition of a basis, but requires only that the system of partitions generates the σ -algebra \mathcal{A} , and not necessarily the space X , so some funnels may contain more than one point—in fact, some must contain more than one point unless $\Xi(\mathcal{A}) = \varepsilon$.

1.4. Measurable partitions. We now have a natural way to go from a partition ξ to a σ -algebra $\mathcal{B}(\xi) \subset \mathcal{B}$, and from a σ -algebra $\mathcal{A} \subset \mathcal{B}$ to a partition $\Xi(\mathcal{A})$.

The definition of $\Xi(\cdot)$ in Theorem 1.5 guarantees that it is a one-sided inverse to $\mathcal{B}(\cdot)$, in the sense that $\mathcal{B}(\Xi(\mathcal{A})) = \mathcal{A}$ for any σ -algebra \mathcal{A} (up to equivalence mod zero). So we may ask if the same holds for partitions; is it true that ξ and $\Xi(\mathcal{B}(\xi))$ are equivalent in some sense?

We see that since each set in $\Xi(\mathcal{B}(\xi))$ is measurable, we should at least demand that ξ not contain any non-measurable sets. For example, consider the partition $\xi = \{A, B\}$, where $A \cap B = \emptyset$, $A \cup B = X$, then if A is measurable (and hence B as well), we have

$$\mathcal{B}(\xi) = \{\emptyset, A, B, X\}$$

and $\Xi(\mathcal{B}(\xi)) = \{A, B\}$, while if A is non-measurable, we have

$$\mathcal{B}(\xi) = \{\emptyset, X\}$$

and so $\Xi(\mathcal{B}(\xi)) = \nu$. Thus a “good” partition should only contain measurable sets; it turns out, however, that this is not sufficient, and that there are examples where $\Xi(\mathcal{B}(\xi))$ is not equivalent mod zero to ξ , even though every set in ξ is measurable.

Example 1.6. Consider the torus \mathbb{T}^2 with Lebesgue measure λ , and let ξ be the partition into orbits of a linear flow ϕ_t with irrational slope α ; that is, $\phi_t(x, y) = (x + t, y + t\alpha)$. In order to determine $\mathcal{B}(\xi)$, we must determine which measurable sets are unions of orbits of ϕ_t ; that is, which measurable sets are invariant. Because this flow is ergodic with respect to λ , any such set must have measure 0 or 1, and so up to sets of measure zero, $\mathcal{B}(\xi)$ is the trivial σ -algebra! It follows that $\Xi(\mathcal{B}(\xi))$ is the trivial partition $\nu = \{\mathbb{T}^2\}$.

Definition 1.7. The partition $\Xi(\mathcal{B}(\xi))$ is known as the *measurable hull* of ξ , and will be denoted by $\mathcal{H}(\xi)$. If ξ is equivalent mod zero to its measurable hull, we say that it is a *measurable partition*.

In particular (foreshadowing the next section), let \mathcal{O} be the partition into orbits of some dynamical system, then $\mathcal{H}(\mathcal{O})$ is also known as the *ergodic decomposition* of that system, and is denoted by \mathcal{E} .⁷

It is obvious that in general, the measurable hull of ξ is a coarsening of ξ , and the definition says that if ξ is non-measurable, this is a proper coarsening.⁸

Exercise 1.3. Show that the measurable hull $\mathcal{H}(\xi)$ is the finest measurable partition which coarsens ξ —in particular, if η is any partition with

$$\xi \leq \eta < \mathcal{H}(\xi),$$

then $\mathcal{H}(\eta) = \mathcal{H}(\xi)$, and hence η is non-measurable.

\mathcal{B} gives a map from the class of all partitions to the class of all σ -algebras, and Ξ gives a map in the opposite direction, which is the one-sided inverse of \mathcal{B} . We see that the set of *measurable* partitions is just the image of the map Ξ , on which \mathcal{H} acts as the identity, and \mathcal{B} and Ξ are two-sided inverses.

Thus we have a correspondence between measurable partitions and σ -algebras—one may easily verify that the operations \vee and \wedge on measurable partitions correspond directly to the operations \cup and \cap on σ -algebras, and that the relations \leq and \geq correspond directly to the relations \subset and \supset .

Example 1.6 shows that the orbit partition for an irrational toral flow is non-measurable. This phenomenon is widespread in dynamical systems—for example, the orbit partition \mathcal{O} for *any* ergodic system with more than one orbit will exhibit the same behaviour; we will see in §3 that the partition into unstable foliations is non-measurable whenever entropy is positive.

Aside from finite or countable partitions into measurable sets (which are obviously measurable), a good example of a measurable partition is given by the square $[0, 1] \times [0, 1]$ with Lebesgue measure λ , and the partition into vertical lines:

$$(1.5) \quad \xi_0 = \{ \{x\} \times [0, 1] \mid x \in [0, 1] \}$$

In fact, this is in some sense the *only* measurable partition, just as $[0, 1]$ is, up to isomorphism, the only Lebesgue space.

Theorem 1.8. *Given a measurable partition ξ of a Lebesgue space (X, \mathcal{B}, μ) , there exists a set $E \subset X$ such that*

- (1) *Each element of $\xi|_E$ has positive measure (and hence there are at most countably many such elements).*
- (2) *$\xi|_{X \setminus E}$ is isomorphic to the partition of the unit square with Lebesgue measure into vertical lines given in (1.5).*

Proof. See [Ro1]. □

⁷It should be noted that because we have not yet talked about conditional measures, one may rightly ask just what about this decomposition is ergodic.

⁸Compare this with the action of the Legendre transform on functions—taking the double Legendre transform of any function returns its *convex hull*, which lies on or below the original function, with equality iff the original function was convex.

Theorem 1.8 proves to be quite useful in dealing with measurable partitions; however, it is quite unwieldy to have to resort to the definition every time we want to know if a particular partition is measurable. We would like some alternate characterisations, which may be easier to check in various situations.

One such characterisation may be motivated by recalling that in the “toy” example of a partition into two subsets, the corresponding σ -algebra had four elements in the measurable case, and only two in the non-measurable case. In some sense, measurability of the partition corresponds to increased “richness” in the σ -algebra. This is made precise as follows:

Theorem 1.9. *Let ξ a partition of a Lebesgue space (X, \mathcal{B}, μ) . ξ is measurable iff there exists a countable set $\{A_n\}_{n \in \mathbb{N}} \subset \mathcal{B}(\xi)$ such that for almost every pair $C_1, C_2 \in \xi$, we can find some A_n which separates them in the sense that $C_1 \subset A_n$, $C_2 \subset X \setminus A_n$.*

Sketch of proof. The key observation is the fact that such a set $\{A_n\}_{n \in \mathbb{N}}$ corresponds to a refining sequence of partitions (1.2) defined by

$$\begin{aligned} \eta_k &= \{A_k, X \setminus A_k\} \\ \xi_n &= \bigvee_{k=1}^n \eta_k \end{aligned} \quad \square$$

Exercise 1.4. Complete the proof of Theorem 1.9.

It may not necessarily be immediately clear just what is meant by “almost every pair C_1, C_2 ”; before addressing this point, we state another sufficient criterion for measurability of a partition which follows from this one.

Theorem 1.10. *Let X be a complete metric space, μ a Borel measure on X , Y a second countable topological space, and $\phi: X \rightarrow Y$ a Borel map (that is, preimages of Borel sets⁹ are Borel). Then the partition*

$$\phi^{-1}(\varepsilon_Y) = \{ \phi^{-1}(y) \mid y \in Y \}$$

called the partition into preimages, is measurable.

Proof. . □

Returning to the statement of Theorem 1.9, recall that the natural projection $\pi: X \rightarrow X/\xi$ takes $x \in X$ to the unique partition element $C \in \xi$ containing x . Thus the space of equivalence classes X/ξ carries a measure which is the pushforward of μ under π —given a measurable set $E \subset X/\xi$, we see that

$$\mu_*(E) = \mu(\pi^{-1}(E))$$

This gives a meaning to the notion of “almost every” partition element, and hence to “almost every pair” of partition elements. Another way to parse the statement is to see that we may remove some set E of zero measure from X and pass to the “trimmed-down” partition $\xi|_{X \setminus E}$, for which the statement holds for every C_1, C_2 .

Example 1.11. Let $C \subset [0, 1]$ be the usual ternary Cantor set, which has Lebesgue measure 0 but contains uncountably many points. Then there is a bijection from C to $[0, 1] \setminus C$, and so we may take a partition ξ of $[0, 1]$, all of whose elements

⁹Recall that the Borel σ -algebra is the minimal σ -algebra containing the open sets.

contain exactly two points, one in C and one not in C . Using the characterisation in Theorem 1.9, we see that ξ is measurable, since we may take for our countable collection the set of intervals with rational endpoints. Further, this partition is equivalent mod zero to the partition into points.

The situation described in Example 1.11, where a partition is in some sense finer than it appears to be, happens all the time in ergodic theory. A fundamental example is the so-called *Fubini's nightmare*, in which a partition which seems to divide the space into curves in fact admits a set of full measure intersecting each partition element exactly once, and hence is equivalent mod zero to the partition into points.

This sort of behaviour stands in stark contrast to *absolute continuity*—but in order to make any sense of that notion, we must first discuss conditional measures.

1.5. Conditional measures on measurable partitions. If a partition element C carries positive measure (which can only be true of countably many elements), then we can define a conditional measure on C by the obvious method; given $E \subset C$, the conditional measure of E is

$$(1.6) \quad \mu_C(E) := \frac{\mu(E)}{\mu(C)}$$

However, for many partitions arising in the study of dynamical systems, such as the partitions into stable and unstable manifolds which will be introduced later, we would also like to be able to define a conditional measure on partition elements of zero measure, and to do so in a way which allows us to reconstruct the original measure.

The model to keep in mind is the canonical example of a measurable partition, the square partitioned into vertical lines. Then denoting by λ , λ_1 , and λ_2 the Lebesgue measures on the square, the horizontal unit interval, and vertical intervals, respectively, Fubini's theorem says that for any integrable $f: [0, 1]^2 \rightarrow \mathbb{R}$ we have

$$(1.7) \quad \int_{[0,1]^2} f(x, y) d\lambda(x, y) = \int_{[0,1]} \int_{[0,1]} f(x, y) d\lambda_2(y) d\lambda_1(x)$$

By Theorem 1.8, any measurable partition of a Lebesgue space is isomorphic to the standard example—perhaps with a few elements of positive measure hanging about, but these will not cause any trouble, as we already know how to define conditional measures on them. Taking the pullback of the Lebesgue measures λ_1 and λ_2 under this isomorphism, we obtain a *factor measure* μ_* on X/ξ , which corresponds to the horizontal unit interval (the set of partition elements), and a family of *conditional measures* $\{\mu_C\}_{C \in \xi}$, which correspond to the vertical unit intervals.

Note that the factor measure is exactly the measure on the space of partition elements which was described in the last section. Note also that although the measure λ_2 was the same for each vertical line (up to a horizontal translation), we can make no such statement about the measures μ_C . The key property of these measures is that for any integrable function $f: X \rightarrow \mathbb{R}$, the function

$$(1.8) \quad \xi \mapsto \int_{\xi} f d\mu_C$$

is a measurable function on X/ξ , and we have

$$(1.9) \quad \int_X f d\mu = \int_{X/\xi} \int_\xi f d\mu_C d\mu_*$$

Each μ_C is “supported” on C in the sense that $\mu_C(C) = 1$, but the reader is cautioned that the measure theoretic support of a measure is a different beast than the topological support of a measure, as the following example illustrates.

Example 1.12. Let $A \subset [0, 1]$ be such that both it and its complement intersect every interval in a set of positive measure.¹⁰ Let λ_1 be one-dimensional Lebesgue measure, and define a measure μ on the unit square by

$$\mu(E) = \lambda_1(E \cap A \times \{0\}) + \lambda_1(E \cap ([0, 1] \setminus A) \times \{1\})$$

Then the topological support of μ is the union of the two lines, $[0, 1] \times \{0, 1\}$, and intersects each partition element in two points, but the conditional measures are δ -measure supported on a single point.

We cannot in general write a simple formula for the conditional measures, as we could in the case where partition elements carried positive weight, so on what grounds do we say that these conditional measures exist? We have already alluded to one proof, which relies on the characterisation of measurable partitions given by Theorem 1.8—this is presented in Viana’s notes [Vi], which draw on Rokhlin’s paper [Ro1]. There are other proofs available, such as Furstenberg’s [Fu], which uses methods from functional analysis made available by defining a topology on X in which all partition elements are closed, which has the effect of enlarging the class of continuous functions.

1.6. Measure classes and absolute continuity. Given a measurable space (X, \mathcal{B}) , consider the set \mathcal{M} of all measures on X . This set has various internal structures which may be of importance to us; for the time being, we focus our attention on the fact, guaranteed by the Radon-Nikodym Theorem, that measures come in classes. This theorem addresses the relationship between two measures ν and μ , and allows us to pass from a qualitative statement to a quantitative one; namely, if ν is absolutely continuous with respect to μ ,¹¹ then there exists a measurable function $d\nu/d\mu$, known as the *Radon-Nikodym derivative*, which has the property that

$$\nu(E) = \int_E \frac{d\nu}{d\mu}(x) d\mu(x)$$

for any $E \in \mathcal{B}$.¹²

Given a reference measure μ and any other measure ν , we also have the *Radon-Nikodym decomposition* of ν ; that is, we may write $\nu = \nu_1 + \nu_2$, where $\nu_1 \ll \mu$ and $\nu_2 \perp \mu$ (the latter means that there exists $A \in \mathcal{B}$ such that $\nu_2(A) = 1$ and $\mu(A) = 0$).

The notion of absolute continuity plays an important role in smooth dynamics, where we have a reference measure class given by the smooth structure of the

¹⁰Such a set can be constructed, for instance, by repeatedly removing and replacing appropriate Cantor sets of positive measure.

¹¹This means that if $\mu(E) = 0$, then $\nu(E) = 0$ as well, a state of affairs which is denoted $\nu \ll \mu$.

¹²As an aside, note that if we change the σ -algebra \mathcal{B} , then we also change the Radon-Nikodym derivative, a fact which is crucial to the proof of the Birkhoff Ergodic Theorem in [KH].

manifold in question, and are often particularly interested in measures which are absolutely continuous with respect to this measure class.

Given a partition ξ , we may also speak of ν as being absolutely continuous with respect to μ on the elements of ξ by passing to the conditional measures ν_C and μ_C and applying the above definitions. This is also an important notion in smooth dynamics, where it allows us to ask not just if a measure is absolutely continuous on the manifold as a whole, but if it is absolutely continuous in certain directions, which correspond to the various rates of expansion and contraction given by the Lyapunov exponents. In particular, we are often interested in measures which are absolutely continuous on unstable leaves, so-called *SRB measures*.

2. MEASURABLE DYNAMICS

2.1. Partitions of times past and future. One of the key notions in dynamics is that of *invariance*—given a dynamical system $T: X \rightarrow X$, we are interested in properties and characterisations of sets, functions, measures, etc. which are invariant under the action of f .

We may consider the property of invariance for partitions as well; we say that a partition ξ is invariant if $T(C) \in \xi$ for every $C \in \xi$, that is, if T maps partition elements to partition elements. This is written as $T\xi = \xi$, where

$$T\xi = \{T(C) \mid C \in \xi\}.$$

Given an invariant partition ξ , let π denote the canonical projection $X \rightarrow X/\xi$, as before. Then T induces an action $\pi \circ T \circ \pi^{-1}$ on the space of partition elements X/ξ , and the dynamics of T may be viewed as a sort of skew product over this action.

In light of the correspondence between measurable partitions and σ -algebras discussed in the previous section, we may also consider invariant σ -algebras. It is then natural to ask if there is a natural way to associate to an arbitrary partition or σ -algebra one which is invariant. One obvious way is to take a σ -algebra \mathcal{A} , and consider the σ -subalgebra $\mathcal{A}' \subset \mathcal{A}$ which contains all the T -invariant sets in \mathcal{A} . However, there is another important construction, which we now examine.

Let ξ be a finite partition of X into measurable sets, and define

$$\xi_T := \bigvee_{n \in \mathbb{Z}} T^n \xi = \lim_{n \rightarrow \infty} \bigvee_{j=-n}^n T^j \xi$$

The elements of this partition are given by $\bigcap_{n \in \mathbb{Z}} T^n C_n$, where $C_n \in \xi$. Because each of the partitions $\xi^{(n)} := \bigvee_{j=-n}^n T^j \xi$ have finitely many elements, all measurable, these partitions themselves are measurable, and we have $\xi^{(n)} = \Xi(\mathcal{B}(\xi^{(n)}))$; passing to the limit, we see that $\xi_T = \Xi(\mathcal{B}(\xi_T))$, so ξ_T is measurable as well.

Exercise 2.1. Show that a partition ξ is measurable iff it is the limit ($\xi = \bigvee_{n=1}^{\infty} \xi_n$) of an increasing sequence $\{\xi_n\}_{n \in \mathbb{N}}$ of finite partitions into measurable sets. Further, show that ξ is measurable is the ξ_n are *any* measurable partitions, and also that the limit ($\xi = \bigwedge_{n=1}^{\infty} \xi_n$) of a decreasing sequence of measurable partitions is measurable.

It follows immediately from the construction of ξ_T that it is an invariant partition, whose σ -algebra is very different from the invariant σ -algebra described above.

Example 2.1. Let A be the 0-1 matrix

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}$$

and consider the space

$$X = \Sigma_A := \{ (x_n)_{n \in \mathbb{Z}} \in \{0, 1, 2, 3\}^{\mathbb{Z}} \mid A_{x_n x_{n+1}} = 1 \ \forall n \in \mathbb{Z} \}$$

Let T be the shift $\sigma: (x_n)_{n \in \mathbb{Z}} \mapsto (x_{n+1})_{n \in \mathbb{Z}}$; this is a simple example of a non-transitive subshift of finite type. Equip X with the Bernoulli measure μ which gives each n -cylinder weight $(1/4)^n$.

Geometrically, X may be thought of as the disjoint union of two copies of Σ_2 , each of which is the direct product of two Cantor sets C (representing the one-sided shift space Σ_2^+), and is invariant under f . In this picture, T acts on each copy of $C \times C \subset [0, 1] \times [0, 1]$ as follows; draw a vertical line down the middle of the unit square, take each of the resulting rectangles, contract it in the vertical direction by a factor of two, expand it in the horizontal direction by the same factor, and then stack the resulting rectangles one on top of the other.

Now let ξ be the partition of X into one-cylinders; that is, $\xi = \{C_0, C_1, C_2, C_3\}$, where

$$C_i := \{ (x_n)_{n \in \mathbb{Z}} \in \Sigma_A \mid x_0 = i \}$$

Each one-cylinder C_i corresponds to one of the four vertical rectangles in the above description, and the reader may verify that in this case, ξ_T is the partition into points.

Another important partition, which is *not* necessarily invariant, is

$$\xi^- := \bigvee_{n=-1}^{-\infty} T^n \xi$$

It is useful to think of the various partitions related to ξ as encoding certain types of information, and to describe what we know about the point x and its trajectory under the action of T if we know which element of these partitions it lies in.

If we know which element of ξ_T a point x lies in, when we know which element of ξ all its iterates, both forward and backward, lie in. In this sense, ξ_T corresponds to both the infinite past and the infinite future; ξ^- , by contrast, corresponds to just the infinite past, since points whose forward iterates lie in different elements of ξ may still lie in the same element of ξ^- .

This last statement is just another way of saying that ξ^- is not necessarily invariant under the action of T ; indeed, we have

$$T\xi^- = \xi \vee \xi^- \geq \xi^-$$

that is, ξ^- is an *increasing partition*.¹³

Example 2.2. Let X , T , and ξ be as in Example 2.1; then the elements of ξ^- are the sets

$$C(x) = \{ y \in \Sigma_A \mid y_n = x_n \ \forall n \leq 0 \}$$

¹³Note that ξ^- is increasing in the sense that it is refined by its image; for this to be the case, each individual element must *decrease* in size. Thus one could conceivably define increasing partitions as those for which $T\xi \leq \xi$, which is the convention followed in [LY1, LY2].

each of which is a copy of Σ_2^+ , and corresponds in the geometric picture to a horizontal Cantor set

$$C \times \{t\} \subset C \times C \subset [0, 1] \times [0, 1],$$

where $t \in C$ and C is the Cantor set mentioned previously. Note that passing to the partition $T\xi^-$ corresponds to dividing each of these Cantor sets into two identical pieces.

There are two possibilities for ξ^- ; either $T\xi^- = \xi^-$, and ξ^- is in fact invariant, or $T\xi^-$ is a proper refinement of ξ^- , which is thus not invariant, as in Example 2.2. In the former case, knowledge of the infinite past gives us knowledge of the infinite future; in the latter case, this is not so.

2.2. Entropy. What is the difference between the two cases just discussed, between the case $T\xi^- = \xi^-$ and the case $T\xi^- > \xi^-$? The key word here is *entropy*; recall that the entropy of a transformation T with respect to a partition ξ is defined as

$$(2.1) \quad h_\mu(T, \xi) := \lim_{n \rightarrow \infty} \frac{1}{n} H\left(\bigvee_{k=0}^{n-1} T^{-k}\xi\right).$$

Exercise 2.2. Show that if ξ^- is invariant, then ξ has zero entropy, $h_\mu(T, \xi) = 0$, whereas if $T\xi^-$ is a proper refinement of ξ^- , then the partition carries positive entropy, $h_\mu(T, \xi) > 0$.

The notion of entropy is intimately connected with one more partition canonically associated with ξ , defined as

$$(2.2) \quad \Pi(\xi) := \bigwedge_{n=1}^{\infty} T^{-n}\xi^-$$

Recall that the intersection $\xi \wedge \eta$ of two partitions is the finest partition which coarsens both ξ and η ; if the partitions are measurable, then this corresponds to taking the intersection $\mathcal{B}(\xi) \cap \mathcal{B}(\eta)$ of the σ -algebras.

For the partitions in (2.2), $T^{-(n+1)}\xi^- \leq T^{-n}\xi^-$ for every n , and so

$$(2.3) \quad \bigwedge_{n=1}^N T^{-n}\xi^- = T^{-N}\xi^-$$

The partition in (2.3) corresponds to knowing the infinite past but having forgotten what happened in the most recent N steps. Thinking of (2.2) in the same way, $\Pi(\xi)$ may be thought of as having had an infinite amount of time to forget an infinite history, a characterisation which is somewhat less transparent than the analogous descriptions of ξ_T and ξ^- .

Exercise 2.3. Let X , T , and ξ be as in Examples 2.1 and 2.2. Verify that $\Pi(\xi)$ is the partition $\{C_0 \cup C_1, C_2 \cup C_3\}$, which separates X into two copies of Σ_2 , each homeomorphic to a product $C \times C$ of two Cantor sets.

Exercise 2.4. Let $T: X \rightarrow X$ be a measure-preserving transformation of a Lebesgue space, and let ξ be a measurable partition. Using the results of Exercises 1.3 and 2.1, show that $\Pi(\xi)$ is a coarsening of both the partition into orbits $\mathcal{O}(T)$ and the ergodic decomposition $\mathcal{E}(T)$.

- (1) $\Pi(\xi) \leq \mathcal{O}(T)$;
- (2) $\Pi(\xi) \leq \mathcal{E}(T)$.

The three partitions we have constructed from ξ are related as follows:

$$(2.4) \quad \xi_T \geq \xi^- \geq \Pi(\xi)$$

The reader may also find it useful to consider the partitions

$$\xi_{(N)}^- := \bigvee_{n=N}^{-\infty} T^n \xi,$$

which represent knowledge of everything that happened up to time N . Then $\xi^- = \xi_{(0)}^-$, while ξ_T and $\Pi(\xi)$ may be thought of as the limits of $\xi_{(N)}^-$ as N goes to $+\infty$ and $-\infty$, respectively. In light of Exercise 2.1, all three partitions in (2.4) are measurable; this is important for part (2) of Exercise 2.4.

If $h(T, \xi) = 0$, then $\xi^- = \xi_{(N)}^-$ for all N , and all three partitions ξ_T , ξ^- , and $\Pi(\xi)$ are equal; there is nothing new under the sun, as it were. In the positive entropy case, each is a proper refinement of the next, as in Examples 2.1 and 2.2, and Exercise 2.3. Although all three are measurable, only ξ_T and $\Pi(\xi)$ are always invariant; ξ^- is not invariant except in the zero entropy case.

Exercise 2.5. Show that the partition $\Pi(\xi)$ derived in Exercise 2.3 has zero entropy.

The result of Exercise 2.5 is actually quite general, and we would like to somehow think of $\Pi(\xi)$ as the “zero entropy” coarsening of ξ^- . Notice that that $\Pi(\xi)$ may be a continuous partition, whose elements all have zero measure as in the case of $h(T, \xi) = 0$ discussed above. The proper meaning of “zero entropy” here is that for any finite partition $\eta \leq \Pi(\xi)$ we have $h_\mu(T, \eta) = 0$.

Theorem 2.3. *Let $\eta \leq \Pi(\xi)$ be a finite or countable partition with finite entropy. Then $h_\mu(T, \eta) = 0$.*

Proof. See [Ro2]. □

2.3. The Pinsker partition. We may consider the set of all partitions with the property exhibited by $\Pi(\xi)$ in Theorem 2.3. This set has an infimum in the partially ordered set of all partitions; that is, there exists a partition $\pi(T)$ which is the finest (biggest) partition such that every finite partition coarser (smaller) than it has zero entropy. This is the *Pinsker partition*, and we may rephrase the above statement as the fact that a finite partition η has $h_\mu(T, \eta) = 0$ iff $\eta \leq \pi(T)$.

Equivalently, $\pi(T)$ may be defined through its σ -algebra; consider all finite or countable measurable partitions with zero entropy, and take the union of their associated σ -algebras. This union is the *Pinsker σ -algebra*, whose associated measurable partition is $\pi(T)$.

Even more concretely, we have the following criterion: a set $E \in \mathcal{B}$ is contained in the Pinsker σ -algebra iff the partition $\xi = \{E, X \setminus E\}$ has $h_\mu(T, \xi) = 0$.

The Pinsker partition may be thought of as the canonically defined zero entropy part of a measure preserving transformation; there are two extreme cases. On the one hand, we may have $\pi(T) = \varepsilon$, the partition into points, in which case every finite partition is a coarsening of $\pi(T)$, and hence has zero entropy. Thus T is a zero entropy transformation, $h_\mu(T) = 0$. At the other extreme, we may have $\pi(T) = \nu$, the trivial partition $\{X\}$, in which case every finite partition has positive entropy, and we say that T is a *K-system*.

2.4. Conditional entropy. At this point we must grapple with the difficulty hinted at before Example ???. That is, we would like to make sense of the notion of entropy of T relative to a partition for as broad a class of partitions as possible. The definition (2.1) relies on the following notion of the entropy of a partition (which is measure-theoretic rather than dynamical):

$$H_\mu(\xi) = - \sum_{C \in \xi} \mu(C) \log \mu(C)$$

This only makes sense when ξ is a finite or countable partition whose elements carry positive measure; for a continuous partition such as the partitions into stable and unstable manifolds which will appear in the next section, or the partition of the unit square into vertical lines which we have already seen, this definition is useless, since $\mu(C) = 0$ for each individual partition element C .

The way around this impasse is to recall the definition of conditional entropy, and adapt it to our present situation by making use of a system of conditional measures, which as we have seen may be defined for a measurable partition even when individual elements have measure zero.

To this end, recall that given two finite or countable partitions ξ and η , one definition of the conditional entropy $H_\mu(\xi|\eta)$ is as the expected value of the conditional information function

$$(2.5) \quad I_\mu^{\xi, \eta}: x \mapsto - \log \mu_{\pi_\eta(x)}(\pi_\xi(x))$$

where π_η and π_ξ denote the canonical projections taking x to the elements of η and ξ , respectively, in which it is contained.

The useful feature of (2.5) is that it works for *any* measurable partitions ξ and η , including continuous ones—all we need is a system of conditional measures.

We could also avoid the explicit use of the information function and consider the usual entropy $H_{\mu_C}(\xi|_C)$ on each partition element $C \in \eta$, then integrate using the factor measure to obtain $H_\mu(\xi|\eta)$. Provided $\xi|_C$ has elements of positive conditional measure μ_C , the usual entropy will be well defined, and we are in business.

Exercise 2.6. Let ξ be a finite or countable partition, so that we may apply the usual definition of entropy, and show that

$$h_\mu(T, \xi) = H_\mu(T\xi^-|\xi^-)$$

Further, show that if ξ is increasing ($T\xi \geq \xi$), we have

$$(2.6) \quad h_\mu(T, \xi) = H_\mu(T\xi|\xi)$$

Since the right hand side of (2.6) is defined for any measurable increasing partition, and is shown by Exercise 2.6 to agree with the usual definition of entropy for finite and countable partitions, we may take it as a definition of entropy for an arbitrary measurable increasing partition.

The key fact connecting these considerations to smooth dynamics is the observation that if $h_\mu(T) = 0$, then the conditional entropy on each partition element is 0, which in the context of the next section will imply that conditional measures on stable and unstable leaves must be atomic.

3. FOLIATIONS AND MEASURES

3.1. Uniform hyperbolicity—stable and unstable foliations. Consider now a diffeomorphism $f: M \rightarrow M$, where M is a Riemannian manifold. If $\Lambda \subset M$ is a

hyperbolic set for f , then we are guaranteed the existence of local and global stable and unstable manifolds at each point $x \in \Lambda$.

The local manifolds are characterised as containing all points whose orbit converges to that of x under forward or backward iteration, without ever being too far away:

$$W_{x,\varepsilon}^s = \left\{ y \in M \mid \lim_{n \rightarrow +\infty} d(f^n y, f^n x) = 0 \text{ and } d(f^n y, f^n x) < \varepsilon \forall n \geq 0 \right\}$$

and similarly for $W_{x,\varepsilon}^u$, with $n \rightarrow -\infty$ and $n \leq 0$.

The global manifolds are characterised similarly, without the requirement that the orbits always be close:

$$W_x^s = \left\{ y \in M \mid \lim_{n \rightarrow +\infty} d(f^n y, f^n x) = 0 \right\}$$

For W_x^u , the limit is again taken as $n \rightarrow -\infty$.

The local manifolds are embedded images of \mathbb{R}^k ; the global manifolds, however, are usually only immersed, and have a somewhat strange global topology.¹⁴ For example, they are dense in \mathbb{T}^2 for the Anosov diffeomorphism given by the action of $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, and hence cannot be embedded images.

The connection with the previous two sections comes when we observe that given two points $x, y \in M$, either $W_x^s \cap W_y^s = \emptyset$ or $W_x^s = W_y^s$, and similarly for the unstable manifolds. It follows that the global stable manifolds form a partition of some invariant set $X \supset \Lambda$; we denote this partition into global stable manifolds by Π^- , and its counterpart, the partition into global unstable manifolds, by Π^+ .

For the linear toral automorphism mentioned above, these partitions are exactly the same as the partition into orbits of the irrational linear flow in Example 1.6, and we saw there that such partitions are non-measurable. In fact, such behaviour is quite common.

Theorem 3.1. *Given a C^2 diffeomorphism $f: M \rightarrow M$ and a hyperbolic set $\Lambda \subset M$, the following are equivalent:*

- (1) $h_\mu(f) = 0$;
- (2) Π^- is measurable;
- (3) Π^+ is measurable.

Sketch of proof. We outline a proof which is due to Sinai in the case of absolutely continuous μ . and in the general case¹⁵ can be found in [LY2].

Without loss of generality, assume μ is ergodic; we will sketch the construction of a *leaf-subordinated partition*.

Definition 3.2 ([BP], Theorem 9.4.1). A *leaf-subordinated partition* associated with the global stable manifolds is a measurable partition ξ such that

- (1) For μ -a.e. x , the element of ξ containing x is an open subset of $W^s(x)$ (hence in particular, $\xi \geq \Pi^-$);
- (2) $f\xi \geq \xi$ (ξ is increasing);
- (3) $\xi_T = \varepsilon$;

¹⁴Hence the terminology “strange attractor” which we see in conjunction with various dissipative systems such as the Hénon map.

¹⁵In fact, the argument was known as a “folklore” since the sixties but probably had not appeared in print before the Ledrappier-Young paper

$$(4) \quad \Pi(\xi) = \mathcal{H}(\Pi^-).$$

The existence of such partitions is Lemma 3.1.1 in [LY1], and Theorem 9.4.1 in [BP] (although the latter deals only with the case where μ is absolutely continuous).

One then proves the following lemma:

Lemma 3.3. *For any leaf-subordinated partition ξ associated with the global stable manifold, we have*

$$h_\mu(f) = H_\mu(f\xi|\xi).$$

Proof. Corollary 5.3 in [LY1]. □

Finally, one must show that $H_\mu(f\xi|\xi) = 0$ iff Π^- is measurable; the result for Π^+ follows upon considering f^{-1} .

To construct ξ , divide the manifold into *rectangles*—that is, domains which exhibit the local product structure of the manifold. More precisely, a rectangle is a domain $X \subset M$ which admits a diffeomorphism $\phi: X \rightarrow [0, 1]^N$ such that the connected component of $\phi(W^s(x) \cap X)$ containing x is given by the set of points in $[0, 1]^N$ whose first $N - k$ coordinates match those of $\phi(x)$, and similarly for $W^u(x)$, with the last k coordinates matching; here N is the dimension of M and k the dimension of the stable manifolds.¹⁶

Such a partition into rectangles may be constructed in a variety of ways—for example, by using a triangulation of the manifold M . Further, a standard argument allows us to assume that the boundary of each rectangle has measure zero.

Exercise 3.1. Let $\{S_r\}_{r>0}$ denote the family of concentric spheres around the origin in \mathbb{R}^n , and show that for any measure ν , at most countably many of the S_r have positive measure.

Now consider the partition ξ_0 whose elements are connected components of the stable manifolds W^s intersected with a rectangle. This guarantees part of the first property, that our partition is a refinement of Π^- ; to obtain an expanding partition, pass to the further refinement

$$\xi := \bigvee_{n=0}^{\infty} f^{-n}\xi_0$$

which may be denoted $\xi = (\xi_0)^-$ using our earlier notation. Thus ξ satisfies property (2).

To see that almost every element of ξ contains a ball in W^s , we must be slightly more careful in our construction of the rectangles, choosing them so that the measure of an δ -neighbourhood of the boundary decreases exponentially with δ . Using this fact, and the fact that ξ_0 refines Π^- so that the size of elements in $f^{-n}\xi_0$ grows exponentially, it is possible to show that typical elements of ξ_0 are only cut finitely many times during the refinement into ξ , which establishes property (1).

Because ξ is a refinement of the partition into stable manifolds, we may bound the diameter of elements of $f^n\xi$ from above, and the bound is exponentially decreasing in n . Thus $\xi_T = \varepsilon$, the partition into points, so (3) holds, and we obtain (4) similarly, using the fact that f^{-1} expands elements of ξ exponentially along the

¹⁶Such rectangles are of critical importance in the construction of *Markov partitions*, a key tool in relating smooth dynamics to symbolic dynamics.

leaves W^s , and so $\Pi(\xi) := \bigwedge_{n=0}^{\infty} f^{-n}\xi = \mathcal{H}(\Pi^-)$. Thus ξ is the leaf-subordinated partition we were after.

Now we want to describe the entropy of f in terms of the entropy of ξ ; this is accomplished by Lemma 3.3.

Regarding the proof of this lemma, recall from basic entropy theory that if η is a finite or countable partition with $\eta_T = \varepsilon$, then we say that η is a *generating partition*, and we have

$$h_\mu(T) = h_\mu(T, \eta) = H_\mu(T\eta|\eta^-) = H_\mu(T\eta^-|\eta^-)$$

and so the result would follow if ξ_0 was finite or countable, since $\xi = (\xi_0)^-$. However, ξ_0 is continuous, so its elements have zero measure, and we cannot use this argument directly. In the uniformly hyperbolic case, we can simply use the finite partition η into rectangles, which refines to ξ_0 under iterations of f^{-1} . In the general setting (for in fact versions of this theorem are true beyond the uniformly hyperbolic case), one needs a more subtle argument, as given in [LY1].

For a finite generating partition η , a basic result from entropy theory says that $\Pi(\eta) = \pi_\mu(f)$, the Pinsker partition, and so if $\eta^- = \xi$, property (4) of a leaf-subordinated partition guarantees that

$$\mathcal{H}(\Pi^-) = \pi(f),$$

that is, that the Pinsker partition is the measurable hull of the partition into global unstable manifolds. The general result is Theorem B in [LY1] (stated there in terms of the associated σ -algebras). It follows that Π^- is measurable iff it is equivalent mod zero to the Pinsker partition.

With the lemma in hand, note that $h_\mu(f) = H_\mu(f\xi|\xi) = 0$ iff $H_\mu(f^n\xi|\xi) = 0$ for any (all) $n \geq 0$, and recall that if any element of ξ is split into two elements of positive measure in $f^n\xi$, then information is gained and the conditional entropy is positive. Since we have an exponential upper bound on the size of elements in $f^n\xi$, we see that if μ is not atomic, then there exists n such that the refinement $f^n\xi \vee \xi$ splits some partition element of positive measure into two (or more) elements of positive measure, which guarantees $H_\mu(f^n\xi|\xi) > 0$, and hence $h_\mu(f) > 0$.

Thus if $h_\mu(f) = 0$, then μ is atomic, with at most one atom in each element of ξ . In this case, the set of atoms on each leaf W^s is discrete, and since a discrete set gets denser (rarer) under forward (backward) iteration, and the measure μ is invariant, one can see that each leaf has at most one atom, and so the conditional measures are in fact δ -measures.¹⁷ In particular, taking the union of the supports of these δ -measures, we have a set of full measure which intersects each leaf exactly once (the so-called *Fubini's nightmare*), and hence Π^- is equivalent mod zero to the point partition ε , which in the zero entropy case is also the Pinsker partition $\pi(f)$. Hence $h_\mu(f) = 0$ implies that Π^- is measurable.

Finally, we must prove the implication in the other direction, that positive entropy implies non-measurability of Π^- . Suppose Π^- is measurable; then we have a system of conditional measures on global stable leaves, each of which is finite. The main idea is to use the assumption of positive entropy to obtain arbitrarily small bounds on the conditional measure of any element of ξ , which will then show that all such elements have conditional measure zero, a contradiction since countably many of them cover the leaf, which has positive measure.

¹⁷One must work slightly harder to show that the conditional measure cannot be atomic with dense support—in this case the idea is to focus on the *big* atoms.

Let us make this more explicit: for a given x , let $C_n(x) \in f^{-n}\xi$ denote the element of $f^{-n}\xi$ containing x , and define functions I_n by

$$I_n(x) = -\log \mu_{C_{n+1}(x)}(C_n(x))$$

These are conditional information functions, as in 2.5, for which

$$H_\mu(f^{-(n+1)}\xi|f^{-n}\xi) = \int_X I_n(x) d\mu(x)$$

For any n , the left hand side is equal to $h_\mu(f, \xi) = h_\mu(f)$, and so we see that $h_\mu(f) = \int_X I_n(x) d\mu(x)$. Further, it is apparent that

$$\sum_{k=0}^{n-1} I_k(x) = -\log \mu_{C_n(x)}(C_0(x))$$

and that $\mu_{C_n(x)}$ converges weakly to $\mu_{W^s(x)}$, where the latter comes from the system of conditional probability measures which exists by the assumption that the partition into global stable leaves is measurable. So to obtain our contradiction, we need only show that $\sum_{k=0}^{\infty} I_k(x)$ diverges unless I_k vanishes almost everywhere.

How are the I_k related to each other? Note that given a system μ_C of conditional measures on elements of $f^{-n}\xi$, the pullback $f^*\mu_C$ is a system of conditional measures on elements of $f^{-(n+1)}\xi$, with respect to which we have $I_{k+1} = I_k$. However, since conditional measures are unique up to a constant, we do in fact have $I_{k+1} = I_k$, and the result follows. \square

As an aside, note that the only property of W^s that was used in the proof of Lemma 3.3 was uniform contraction along its leaves. In general, we could take W to be *any* uniformly contracting partition, and we would have a version of the lemma with equality replaced by the inequality

$$H_\mu(f\xi|\xi) \leq h_\mu(f)$$

Taking W to be the foliation in a single stable direction (say a subspace corresponding to a negative Lyapunov exponent), this allows us to speak of the contribution made by certain directions (or certain Lyapunov exponents) to the entropy.

3.2. Conditional measures on global leaves. The theorem on existence of conditional measures only applies to *measurable* partitions, and as we have seen, the partitions into global stable or unstable manifolds are only measurable in the zero entropy case. Thus we cannot apply the theorem directly; however, by restricting our attention to a small section of the manifold, a rectangle, we may consider conditional measures on W^s and W^u within that domain.

Of course, we could choose another rectangle, which may overlap the first, and obtain conditional measures there as well; how will these two sets of conditional measures relate on the intersection?

The answer is as simple as we could hope for, and is best visualised by considering two subsets $A, B \subset X$ of positive measure with nontrivial intersection. Conditional measures μ_A and μ_B are defined in the obvious way, as the normalised restriction of μ to the appropriate domain, and it is easy to see that given $E \subset A \cap B$, we have

$$\mu_A(E) = \frac{\mu(E)}{\mu(A)} = \frac{\mu(B)}{\mu(A)} \frac{\mu(E)}{\mu(B)} = \frac{\mu(B)}{\mu(A)} \mu_B(E)$$

That is, μ_A and μ_B are proportional to each other; a similar result holds for conditional measures on stable and unstable manifolds. If $\mu_{W^s(x)}^{(1)}$ and $\mu_{W^s(x)}^{(2)}$ are two families of conditional measures on stable manifolds coming from different rectangles, then they are proportional on the intersection of the two rectangles. However, because the conditional measure on each leaf is normalised, the constant of proportionality may vary from leaf to leaf.

In this way we may define a σ -finite measure on each leaf, by gluing together conditional measures on rectangles, a procedure that is important for certain constructions in rigidity theory.

In principle, we may think of the conditional measure on a leaf $W^s(x)$ as being the result of a limiting process. Having fixed a rectangle, we have a product structure, and may consider small cylinders C around the leaf, whose cross-sections are transversal to the leaf. Given a set $E \subset W^s(x)$, then, we may consider its product with this transversal cross-section, and approximate $\mu_{W^s(x)}(E)$ by $\mu(E)/\mu(C)$. We would like to say that in the limit as the size of the cross-section goes to zero, this quantity converges to the conditional measure.

The difficulty with this interpretation is that the limit is only guaranteed to exist on almost every leaf,¹⁸ and so it may fail for the particular leaf we are interested in at a given time. This is a manifestation of the fact that even if μ itself is rather “nice”, the conditional measures may have very irregular dependence on the transversal direction.

3.3. Non-uniform hyperbolicity, the Pesin Entropy Formula, and the Ledrappier-Young Theorem. In the non-uniformly hyperbolic setting, all of the above results go through more or less unchanged, with the caveat that now the structure of the foliations is intimately dependent on the measure. We are only guaranteed existence of W^s and W^u at μ -a.e. point, and there are some extra technical difficulties in construction ξ^- , which we shall not get into here.

The word “foliation” must be used guardedly in this setting; here it refers to a family of immersed manifolds which have continuous transversal dependence on particular compact subsets (the Pesin sets), on which all estimates are uniform, but which are not themselves invariant.

Given a diffeomorphism $f: M \rightarrow M$ and an f -invariant measure μ , the Multiplicative Ergodic Theorem of Oseledets guarantees the existence of the Lyapunov exponents $\chi_1(x) < \dots < \chi_k(x)$ at almost every point, along with the corresponding subspaces $E_1(x) \subset \dots \subset E_k(x) = T_x M$. These geometric quantities give infinitesimal rates of expansion and contraction, and are independent of the measure; however, if μ is ergodic then they are constant a.e., and so we occasionally speak of the Lyapunov exponents of an ergodic measure.

The following fundamental inequality is due to Margulis and Ruelle:

Theorem 3.4. *If $f: M \rightarrow M$ is a C^1 diffeomorphism of a smooth compact Riemannian manifold preserving a Borel probability measure μ , and $d_i(x) = \dim(E_i(x))$ is the multiplicity of the i^{th} Lyapunov exponent at x , then*

$$(3.1) \quad h_\mu(f) \leq \int_X \sum_{\chi_i(x) > 0} d_i(x) \chi_i(x) d\mu(x)$$

¹⁸Compare this with the statement of the Lebesgue density theorem, that almost every point is a density point for a given measure.

Proof. Theorem 10.2.1 in [BP]. \square

Pesin gave conditions under which equality holds.

Theorem 3.5 (Pesin Entropy Formula). *If in addition to the above hypotheses we have that f is $C^{1+\alpha}$ and μ is absolutely continuous, then*

$$(3.2) \quad h_\mu(f) = \int_X \sum_{\chi_i(x) > 0} d_i(x) \chi_i(x) d\mu(x)$$

Proof. Theorem 10.4.1 in [BP]. \square

Note that the integral in (3.1) and (3.2) is the exponential rate of volume expansion in the unstable direction, and may also be written as

$$\int_X \log |J_x^u f| d\mu(x)$$

where $J_x^u := J_x|_{W_x^u}$ is the Jacobian on the unstable manifold.

The proof of Pesin's entropy formula relies on the construction of leaf-subordinated partitions outlined in the proof of Theorem 3.1. The key step is to show that the conditional measures on W^u are absolutely continuous, which allows one to establish bounds on the rate at which the volume of elements in the refined partitions decreases.

In fact, it turns out that no particular regularity of μ in the stable direction is required for Pesin's entropy formula to hold, which led Ledrappier and Young to prove the following:

Theorem 3.6 (Ledrappier-Young). *Let $f: M \rightarrow M$ be a C^2 diffeomorphism of a compact Riemannian manifold M preserving a Borel probability measure μ . Then μ has absolutely continuous conditional measures on unstable manifolds iff (3.2) holds.*

Proof. Theorem A in [LY1]. \square

A measure μ satisfying the conditions of the theorem is called an *SRB measure*, after Sinai, Ruelle, and Bowen; despite having absolutely continuous conditional measures on unstable manifolds, such measures are generally singular on M .

SRB measures may or may not exist for a particular system; however, in the Anosov case, they always exist, and in fact one obtains two SRB measures, one corresponding to forward iterations (which is a.c. in the unstable direction), and one corresponding to backward iterations (which is a.c. in the stable direction). The two coincide iff they are absolutely continuous on M .

In fact, Ledrappier and Young proved a more general theorem than Theorem 3.6; in [LY2], they show that (3.2) holds for arbitrary measures μ , when the multiplicities $d_i(x)$ are replaced with coefficients δ_i^μ , which depend on the geometry of μ along the various foliations corresponding to different Lyapunov exponents, but which have no explicit dependence on the dynamics.

For the largest Lyapunov exponent, the coefficient δ_n^μ represents the Hausdorff dimension of the conditional measures on the corresponding foliation. However, this does not extend to intermediate exponents, as shown by a counterexample due to Ruelle and Wilkinson, [RW] for which (3.2) holds, but the conditional measure in the *slow* unstable direction is atomic, and so the foliation is singular. A formula for these coefficients may be found in Theorem 14.1.18 of [BP].

REFERENCES

- [BP] L. Barreira and Y. Pesin, *Nonuniform Hyperbolicity: Dynamics of Systems with Nonzero Lyapunov Exponents*. Cambridge University Press, 2007.
- [Fu] H. Furstenberg, *Recurrence in Ergodic Theory and Combinatorial Number Theory*, Princeton University Press, 1981.
- [Ha] P.R. Halmos, *Measure Theory*. New York: Springer-Verlag, 1974.
- [KH] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge, 1995.
- [LY1] F. Ledrappier and L.-S. Young, *The Metric Entropy of Diffeomorphisms: Part I: Characterization of Measures Satisfying Pesin's Entropy Formula*. Ann. Math. Vol. 122, No. 3 (Nov. 1985), pp. 509–539.
- [LY2] F. Ledrappier and L.-S. Young, *The Metric Entropy of Diffeomorphisms: Part II: Relations between Entropy, Exponents and Dimension*. Ann. Math. Vol. 122, No. 3 (Nov. 1985), pp. 540–574.
- [Ro1] V.A. Rokhlin, *On the Fundamental Ideas of Measure Theory*. Mat. Sb. (N.S.), 1949, 25(67):1, 107–150 (Russian). Trans. Amer. Math. Soc., Series 1, 10 (1962), 1–52 (English).
- [Ro2] V.A. Rokhlin, *Lectures on the Entropy Theory of Measure-Preserving Transformations*. Russ. Math. Surv. Vol. 22 (1967), No. 5, pp. 1–52.
- [RW] D. Ruelle and A. Wilkinson, *Absolutely Singular Dynamical Foliations*, Comm. Math. Phys. 219 (2001), 481–487.
- [Vi] M. Viana, *Disintegration into conditional measures: Rokhlin's theorem*. <http://w3.impa.br/~viana/out/rokhlin.pdf>
- [Wa] P. Walters, *Ergodic Theory*. Springer-Verlag, Berlin, 1975.