

Chapter 23

Ergodic Properties and Ergodic Limits

Section 23.1 gives a general orientation to ergodic theory, which we will study in discrete time.

Section 23.2 introduces dynamical systems and their invariants, the setting in which we will prove our ergodic theorems.

Section 23.3 considers time averages, defines what we mean for a function to have an ergodic property (its time average converges), and derives some consequences.

Section 23.4 defines asymptotic mean stationarity, and shows that, with AMS dynamics, the limiting time average is equivalent to conditioning on the invariant sets.

23.1 General Remarks

To begin our study of ergodic theory, let us consider a famous¹ line from Gnedenko and Kolmogorov (1954, p. 1):

In fact, all epistemological value of the theory of probability is based on this: that large-scale random phenomena in their collective action create strict, nonrandom regularity.

Now, this is how Gnedenko and Kolmogorov introduced their classic study of the limit laws for *independent* random variables, but most of the random phenomena we encounter around us are not independent. Ergodic theory is a study of how large-scale *dependent* random phenomena nonetheless create nonrandom regularity. The classical limit laws for IID variables X_1, X_2, \dots assert that,

¹Among mathematical scientists, anyway.

under the right conditions, sample averages converge on expectations,

$$\frac{1}{n} \sum_{i=1}^n X_i \rightarrow \mathbf{E}[X_i]$$

where the sense of convergence can be “almost sure” (strong law of large numbers), “ L_p ” (p^{th} mean), “in probability” (weak law), etc., depending on the hypotheses we put on the X_i . One meaning of this convergence is that sufficiently large random samples are representative of the entire population — that the sample mean makes a good estimate of $\mathbf{E}[X]$.

The ergodic theorems, likewise, assert that for *dependent* sequences X_1, X_2, \dots , time averages converge on expectations

$$\frac{1}{t} \sum_{i=1}^t X_i \rightarrow \mathbf{E}[X_\infty | \mathcal{I}]$$

where X_∞ is some limiting random variable, or in the most useful cases a *non*-random variable, and \mathcal{I} is a σ -field representing some sort of asymptotic information. Once again, the mode of convergence will depend on the kind of hypotheses we make about the random sequence X . Once again, the interpretation is that a *single* sample path is representative of an entire distribution over sample paths, *if* it goes on long enough. The IID laws of large numbers are, in fact, special cases of the corresponding ergodic theorems.

Section 22.3 proved a mean-square (L_2) ergodic theorem for weakly stationary continuous-parameter processes.² The next few chapters, by contrast, will develop ergodic theorems for non-stationary discrete-parameter processes.³ This is a little unusual, compared to most probability books, so let me say a word or two about why. (1) Results we get will include stationary processes as special cases, but stationarity fails for many applications where ergodicity (in a suitable sense) holds. So this is more general and more broadly applicable. (2) Our results will all have continuous-time analogs, but the algebra is a lot cleaner in discrete time. (3) Some of the most important applications (for people like you!) are to statistical inference and learning with dependent samples, and to Markov chain Monte Carlo, and both of those are naturally discrete-parameter processes. We will, however, stick to continuous state spaces.

23.2 Dynamical Systems and Their Invariants

It is a very remarkable fact — but one with deep historical roots (von Plato, 1994, ch. 3) — that the way to get regular limits for stochastic processes is to first turn them into irregular deterministic dynamical systems, and then let averaging smooth away the irregularity. This section will begin by laying out dynamical systems, and their invariant sets and functions, which will be the foundation for what follows.

²Can you identify X_∞ and \mathcal{I} for this case?

³In doing so, I’m ripping off Gray (1988), especially chapters 6 and 7.

Definition 299 (Dynamical System) A dynamical system consists of a measurable space Ξ , a σ -field \mathcal{X} on Ξ , a probability measure μ defined on \mathcal{X} , and a mapping $T : \Xi \mapsto \Xi$ which is \mathcal{X}/\mathcal{X} -measurable.

Remark: Measure-preserving transformations (Definition 53) are special cases of dynamical systems. Since (Theorem 52) every strongly stationary process can be represented by a measure-preserving transformation, namely the shift (Definition 48), the theory of ergodicity for dynamical systems which we'll develop is easily seen to include the usual ergodic theory of strictly-stationary processes as a special case. Thus, at the cost of going to the infinite-dimensional space of sample paths, we can always make it the case that the time-evolution is completely deterministic, and the only stochastic component to the process is its initial condition.

Lemma 300 (Dynamical Systems are Markov Processes) Let Ξ, \mathcal{X}, μ, T be a dynamical system. Let $\mathcal{L}(X_1) = \mu$, and define $X_t = T^{t-1}X_1$. Then the X_t form a Markov process, with evolution operator K defined through $Kf(x) = f(Tx)$.

PROOF: For every $x \in \Xi$ and $B \in \mathcal{X}$, define $\kappa(x, B) \equiv \mathbf{1}_B(Tx)$. For fixed x , this is clearly a probability measure (specifically, the δ measure at Tx). For fixed B , this is a measurable function of x , because T is a measurable mapping. Hence, $\kappa(x, B)$ is a probability kernel. So, by Theorem 106, the X_t form a Markov process. By definition, $\mathbf{E}[f(X_1)|X_0 = x] = Kf(x)$. But the expectation is in this case just $f(Tx)$. \square

Notice that, as a consequence, there is a corresponding operator, call it U , which takes signed measures (defined over \mathcal{X}) to signed measures, and specifically takes probability measures to probability measures.

Definition 301 (Observable) A function $f : \Xi \mapsto \mathbb{R}$ which is \mathbb{B}/\mathcal{X} measurable is an observable of the dynamical system Ξ, \mathcal{X}, μ, T .

Pretty much all of what follows would work if the observables took values in any real or complex vector space, but that situation can be built up from this one.

Definition 302 (Invariant Functions, Sets and Measures) A function is invariant, under the action of a dynamical system, if $f(Tx) = f(x)$ for all $x \in \Xi$, or equivalently if $Kf = f$ everywhere. An event $B \in \mathcal{X}$ is invariant if its indicator function is an invariant function. A measure ν is invariant if it is preserved by T , i.e. if $\nu(C) = \nu(T^{-1}C)$ for all $C \in \mathcal{X}$, equivalently if $U\nu = \nu$.

Lemma 303 (Invariant Sets are a σ -Algebra) The class \mathcal{I} of all measurable invariant sets in Ξ forms a σ -algebra.

PROOF: Clearly, Ξ is invariant. The other properties of a σ -algebra follow because set-theoretic operations (union, complementation, etc.) commute with taking inverse images. \square

Lemma 304 (Invariant Sets and Observables) *An observable is invariant if and only if it is \mathcal{I} -measurable. Consequently, \mathcal{I} is the σ -field generated by the invariant observables.*

PROOF: “If”: Pick any Borel set B . Since $f = f \circ T$, $f^{-1}(B) = (f \circ T)^{-1}(B) = T^{-1}f^{-1}B$. Hence $f^{-1}(B) \in \mathcal{I}$. Since the inverse image of every Borel set is in \mathcal{I} , f is \mathcal{I} -measurable. “Only if”: Again, pick any Borel set B . By assumption, $f^{-1}(B) \in \mathcal{I}$, so $f^{-1}(B) = T^{-1}f^{-1}(B) = (f \circ T)^{-1}(B)$, so the inverse image of under Tf of any Borel set is an invariant set, implying that $f \circ T$ is \mathcal{I} -measurable. Since, for every B , $f^{-1}(B) = (f \circ T)^{-1}(B)$, we must have $f \circ T = f$. The consequence follows. \square

Definition 305 (Infinitely Often, i.o.) *For any set $C \in \mathcal{X}$, the set C infinitely often, $C_{i.o.}$, consists of all those points in Ξ whose trajectories visit C infinitely often, $C_{i.o.} \equiv \limsup_t T^{-t}C$.*

Lemma 306 (“Infinitely often” implies invariance) *For every $C \in \mathcal{X}$, $C_{i.o.}$ is invariant.*

PROOF: Exercise 23.1. \square

Definition 307 (Invariance Almost Everywhere) *A measurable function is invariant μ -a.e., or almost invariant, when*

$$\mu \{x \in \Xi \mid \forall n, f(x) = K^n f(x)\} = 1 \tag{23.1}$$

A measurable set is invariant μ -a.e., when its indicator function is almost invariant.

Remark 1: Some of the older literature annoyingly calls these objects *totally invariant*.

Remark 2: Invariance implies invariance μ -almost everywhere, for any μ .

Lemma 308 (Almost-invariant sets form a σ -algebra) *The almost-invariant sets form a σ -field, \mathcal{I}' , and an observable is almost invariant if and only if it is measurable with respect to this field.*

PROOF: Entirely parallel to that for the strict invariants. \square

Let’s close this section with a simple lemma, which will however be useful in approximation-by-simple-function arguments in building up expectations.

Lemma 309 (Invariance for simple functions) *A simple function, $f(x) = \sum_{k=1}^m a_k \mathbf{1}_{C_k}(x)$, is invariant if and only if all the sets $C_k \in \mathcal{I}$. Similarly, a simple function is almost invariant iff all the defining sets are almost invariant.*

PROOF: Exercise 23.2. \square

23.3 Time Averages and Ergodic Properties

For convenience, let's re-iterate the definition of a time average. The notation differs here a little from that given earlier, in a way which helps with discrete time.

Definition 310 (Time Averages) *The time-average of an observable f is the real-valued function*

$$A_t f(x) \equiv \frac{1}{t} \sum_{i=0}^{t-1} f(T^i x) \tag{23.2}$$

where A_t is the operator taking functions to their time-averages.

Lemma 311 (Time averages are observables) *For every t , the time-average of an observable is an observable.*

PROOF: The class of measurable functions is closed under finite iterations of arithmetic operations. \square

Definition 312 (Ergodic Property) *An observable f has the ergodic property when $A_t f(x)$ converges as $t \rightarrow \infty$ for μ -almost-all x . An observable has the mean ergodic property when $A_t f(x)$ converges in $L_1(\mu)$, and similarly for the other L_p ergodic properties. If for some class of functions \mathcal{D} , every $f \in \mathcal{D}$ has an ergodic property, then the class \mathcal{D} has that ergodic property.*

Remark. Notice that what is required for f to have the ergodic property is that almost every initial point has *some* limit for its time average,

$$\mu \left\{ x \in \Xi \mid \exists r \in \mathbb{R} : \lim_{t \rightarrow \infty} A_t f(x) = r \right\} = 1 \tag{23.3}$$

This does *not* mean that there is some *common* limit for almost every initial point,

$$\exists r \in \mathbb{R} : \mu \left\{ x \in \Xi \mid \lim_{t \rightarrow \infty} A_t f(x) = r \right\} = 1 \tag{23.4}$$

Similarly, a class of functions has the ergodic property if all of their time averages converge; they do not have to converge uniformly.

Definition 313 (Ergodic Limit) *If an observable f has the ergodic property, define $Af(x)$ to be the limit of $A_t f(x)$ where that exists, and 0 elsewhere. The domain of A consists of all and only the functions with ergodic properties.*

Observe that

$$Af(x) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=0}^t K^n f(x) \tag{23.5}$$

That is, A is the limit of an arithmetic mean of conditional expectations. This suggests that it should itself have many of the properties of conditional expectations. In fact, under a reasonable condition, we will see that $Af = \mathbf{E}[f|\mathcal{I}]$, expectation conditional on the σ -algebra of invariant sets. We'll check first that A has the properties we'd want from a conditional expectation.

Lemma 314 (Linearity of Ergodic Limits) *A is a linear operator, and its domain is a linear space.*

PROOF: If c is any real number, then $A_t c f(x) = c A_t f(x)$, and so clearly, if the limit exists, $A c f(x) = c A f(x)$. Similarly, $A_t(f + g)(x) = A_t f(x) + A_t g(x)$, so if f and g both have ergodic properties, then so does $f + g$, and $A(f + g)(x) = A f(x) + A g(x)$. \square

Lemma 315 (Non-negativity of ergodic limits) *If $f \in \text{Dom} A$, and, for all $n \geq 0$, $f T^n \geq 0$ a.e., then $A f(x) \geq 0$ a.e.*

PROOF: The event $A f(x) < 0$ is a sub-event of $\bigcup_n \{f(T^n(x)) < 0\}$. Since the union of a countable collection of measure zero events has measure zero, $A f(x) \geq 0$ almost everywhere. \square

Notice that our hypothesis is that $f T^n \geq 0$ a.e. for all n , not just that $f \geq 0$. The reason we need the stronger assumption is that the transformation T might map every point to the bad set of f ; the lemma guards against that. Of course, if $f(x) \geq 0$ for all, and not just almost all, x , then the bad set is non-existent, and $A f \geq 0$ follows automatically.

Lemma 316 (Constant functions have the ergodic property) *The constant function 1 has the ergodic property. Consequently, so does every other constant function.*

PROOF: For every n , $1(T^n x) = 1$. Hence $A_t 1(x) = 1$ for all t , and so $A 1(x) = 1$. Extension to other constants follows by linearity (Lemma 314). \square

Remember that for any time-evolution operator K , $K 1 = 1$.

Lemma 317 (Ergodic limits are invariantifying) *If $f \in \text{Dom}(A)$, then, for all n , $f \circ T^n$ is too, and $A f(x) = A f \circ T^n(x)$. Or, $A K^n f(x) = A f(x)$.*

PROOF: Start with $n = 1$, and show that the discrepancy goes to zero.

$$A K f(x) - A f(x) = \lim_t \frac{1}{t} \sum_{i=0}^t (K^{i+1} f(x) - K^i f(x)) \quad (23.6)$$

$$= \lim_t \frac{1}{t} (K^t f(x) - f(x)) \quad (23.7)$$

Since $A f(x)$ exists a.e., we know that the series $t^{-1} \sum_{i=0}^{t-1} K^i f(x)$ converges a.e., implying that $(t + 1)^{-1} K^t f(x) \rightarrow 0$ a.e.. But $t^{-1} = \frac{t+1}{t} (t + 1)^{-1}$, and for large t , $t + 1/t < 2$. Hence $(t + 1)^{-1} K^t f(x) \leq t^{-1} K^t f(x) \leq 2(t + 1)^{-1} K^t f(x)$, implying that $t^{-1} K^t f(x)$ itself goes to zero (a.e.). Similarly, $t^{-1} f(x)$ must go to zero. Thus, overall, we have $A K f(x) = A f(x)$ a.e., and $K f(x) \in \text{Dom}(A)$. \square

Lemma 318 (Ergodic limits are invariant functions) *If $f \in \text{Dom}(A)$, then $A f$ is an invariant, and \mathcal{I} -measurable.*

PROOF: Af exists, so (previous lemma) AKf exists and is equal to Af (almost everywhere). But $AKf(x) = Af(Tx)$, by definition, hence Af is invariant, i.e., $KAf = AKf = Af$. Measurability follows from Lemma 304. \square

Lemma 319 (Ergodic limits and invariant indicator functions) *If $f \in \text{Dom}(A)$, and B is any set in \mathcal{I} , then $A(\mathbf{1}_B(x)f(x)) = \mathbf{1}_B(x)Af(x)$.*

PROOF: For every n , $\mathbf{1}_B(T^n x)f(T^n x) = \mathbf{1}_B(x)f(T^n x)$, since $x \in B$ iff $T^n x \in B$. So, for all finite t , $A_t(\mathbf{1}_B(x)f(x)) = \mathbf{1}_B(x)A_t f(x)$, and the lemma follows by taking the limit. \square

Lemma 320 (Ergodic properties of sets and observables) *All indicator functions of measurable sets have ergodic properties if and only if all bounded observables have ergodic properties.*

PROOF: A standard approximation-by-simple-functions argument, as in the construction of Lebesgue integrals. \square

Lemma 321 (Expectations of ergodic limits) *Let f be bounded and have the ergodic property. Then Af is μ -integrable, and $\mathbf{E}[Af(X)] = \mathbf{E}[f(X)]$, where $\mathcal{L}(X) = \mu$.*

PROOF: Since f is bounded, it is integrable. Hence $A_t f$ is bounded, too, for all t , and $A_t f(X)$ is an integrable random variable. A sequence of bounded, integrable random variables is uniformly integrable. Uniform integrability, plus the convergence $A_t f(x) \rightarrow Af(x)$ for μ -almost-all x , gives us that $\mathbf{E}[Af(X)]$ exists and is equal to $\lim \mathbf{E}[A_t f(X)]$ via Fatou's lemma. (See e.g., Theorem 117 in the notes to 36-752.)

Now use the invariance of Af , i.e., the fact that $Af(X) = Af(TX)$ μ -a.s.

$$0 = \mathbf{E}[Af(TX)] - \mathbf{E}[Af(X)] \quad (23.8)$$

$$= \lim \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[K^n f(TX)] - \lim \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[K^n f(X)] \quad (23.9)$$

$$= \lim \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[K^n f(TX)] - \mathbf{E}[K^n f(X)] \quad (23.10)$$

$$= \lim \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[K^{n+1} f(X)] - \mathbf{E}[K^n f(X)] \quad (23.11)$$

$$= \lim \frac{1}{t} (\mathbf{E}[K^t f(X)] - \mathbf{E}[f(X)]) \quad (23.12)$$

Hence

$$\mathbf{E}[Af] = \lim \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[K^n f(X)] = \mathbf{E}[f(X)] \quad (23.13)$$

as was to be shown. \square

Lemma 322 (Convergence of Ergodic Limits) *If f is as in Lemma 321, then $A_t f \rightarrow f$ in $L_1(\mu)$.*

PROOF: From Lemma 321, $\lim \mathbf{E}[A_t f(X)] = \mathbf{E}[f(X)]$. Since the variables $A_t f(X)$ are uniformly integrable (as we saw in the proof of that lemma), it follows (Proposition 4.12 in Kallenberg, p. 68) that they also converge in $L_1(\mu)$. \square

Lemma 323 (Cesàro Mean of Expectations) *Let f be as in Lemmas 321 and 322, and $B \in \mathcal{X}$ be an arbitrary measurable set. Then*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=0}^{t-1} \mathbf{E}[\mathbf{1}_B(X) K^n f(X)] = \mathbf{E}[\mathbf{1}_B(X) f(X)] \quad (23.14)$$

where $\mathcal{L}(X) = \mu$.

PROOF: Let's write out the expectations explicitly as integrals.

$$\left| \int_B f(x) d\mu - \frac{1}{t} \sum_{n=0}^{t-1} \int_B K^n f(x) d\mu \right| \quad (23.15)$$

$$= \left| \int_B f(x) - \frac{1}{t} \sum_{n=0}^{t-1} K^n f(x) d\mu \right|$$

$$= \left| \int_B f(x) - A_t f(x) d\mu \right| \quad (23.16)$$

$$\leq \int_B |f(x) - A_t f(x)| d\mu \quad (23.17)$$

$$\leq \int |f(x) - A_t f(x)| d\mu \quad (23.18)$$

$$= \|f - A_t f\|_{L_1(\mu)} \quad (23.19)$$

But (previous lemma) these functions converge in $L_1(\mu)$, so the limit of the norm of their difference is zero. \square

Boundedness is not essential.

Corollary 324 (Replacing Boundedness with Uniform Integrability) *Lemmas 321, 322 and 323 hold for any integrable observable $f \in \text{Dom}(A)$, bounded or not, provided that $A_t f$ is a uniformly integrable sequence.*

PROOF: Examining the proofs shows that the boundedness of f was important only to establish the uniform integrability of $A_t f$. \square

23.4 Asymptotic Mean Stationarity

Next, we come to an important concept which will prove to be necessary and sufficient for the most important ergodic properties to hold.

Definition 325 (Asymptotically Mean Stationary) *A dynamical system is asymptotically mean stationary when, for every $C \in \mathcal{X}$, the limit*

$$m(C) \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=0}^{t-1} \mu(T^{-n}C) \tag{23.20}$$

exists, and the set function m is its stationary mean.

Remark 1: It might've been more logical to call this "asymptotically measure stationary", or something, but I didn't make up the names...

Remark 2: Symbolically, we can write

$$m = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=0}^{t-1} U^n \mu$$

where U is the operator taking measures to measures. This leads us to the next proposition.

Lemma 326 (Stationary Implies Asymptotically Mean Stationary) *If a dynamical system is stationary, i.e., T preserves the measure μ , then it is asymptotically mean stationary, with stationary mean μ .*

PROOF: If T preserves μ , then for every measurable set, $\mu(C) = \mu(T^{-1}C)$. Hence every term in the sum in Eq. 23.20 is $\mu(C)$, and consequently the limit exists and is equal to $\mu(C)$. \square

Proposition 327 (Vitali-Hahn Theorem) *If m_t are a sequence of probability measures on a common σ -algebra \mathcal{X} , and $m(C)$ is a set function such that $\lim_t m_t(C) = m(C)$ for all $C \in \mathcal{X}$, then m is a probability measure on \mathcal{X} .*

PROOF: This is a standard result from measure theory. \square

Theorem 328 (Stationary Means are Invariant Measures) *If a dynamical system is asymptotically mean stationary, then its stationary mean is an invariant probability measure.*

PROOF: For every t , let $m_t(C) = \frac{1}{t} \sum_{n=0}^{t-1} \mu(T^{-n}(C))$. Then m_t is a linear combination of probability measures, hence a probability measure itself. Since, for every $C \in \mathcal{X}$, $\lim m_t(C) = m(C)$, by Definition 325, Proposition 327 says that $m(C)$ is also a probability measure. It remains to check invariance.

$$m(C) - m(T^{-1}C) \tag{23.21}$$

$$\begin{aligned} &= \lim \frac{1}{t} \sum_{n=0}^{t-1} \mu(T^{-n}(C)) - \lim \frac{1}{t} \sum_{n=0}^{t-1} \mu(T^{-n}(T^{-1}C)) \\ &= \lim \frac{1}{t} \sum_{n=0}^{t-1} \mu(T^{-n-1}C) - \mu(T^{-n}C) \end{aligned} \tag{23.22}$$

$$= \lim \frac{1}{t} (\mu(T^{-t}C) - \mu(C)) \tag{23.23}$$

Since the probability measure of any set is at most 1, the difference between two probabilities is at most 1, and so $m(C) = m(T^{-1}C)$, for all $C \in \mathcal{X}$. But this means that m is invariant under T (Definition 53). \square

Remark: Returning to the symbolic manipulations, if μ is AMS with stationary mean m , then $Um = m$ (because m is invariant), and so we can write $\mu = m + (\mu - m)$, knowing that $\mu - m$ goes to zero under averaging. Speaking loosely (this can be made precise, at the cost of a fairly long excursion) m is an eigenvector of U (with eigenvalue 1), and $\mu - m$ lies in an orthogonal direction, along which U is contracting, so that, under averaging, it goes away, leaving only m , which is like the projection of the original measure μ on to the invariant manifold of U .

The relationship between an AMS measure μ and its stationary mean m is particularly simple on invariant sets: they are equal there. A slightly more general theorem is actually just as easy to prove, however, so we'll do that.

Lemma 329 (Expectations of Almost-Invariant Functions) *If μ is AMS with limit m , and f is an observable which is invariant μ -a.e., then $\mathbf{E}_\mu[f] = \mathbf{E}_m[f]$.*

PROOF: Let C be any almost invariant set. Then, for any t , C and $T^{-t}C$ differ by, at most, a set of μ -measure 0, so that $\mu(C) = \mu(T^{-t}C)$. The definition of the stationary mean (Equation 23.20) then gives $\mu(C) = m(C)$, or $\mathbf{E}_\mu[\mathbf{1}_C] = \mathbf{E}_m[\mathbf{1}_C]$, i.e., the result holds for indicator functions. By Lemma 309, this then extends to simple functions. The usual arguments then take us to all functions which are measurable with respect to \mathcal{I}' , the σ -field of almost-invariant sets, but this (Lemma 308) is the class of all almost-invariant functions. \square

Lemma 330 (Limit of Cesàro Means of Expectations) *If μ is AMS with stationary mean m , and f is a bounded observable,*

$$\lim_{t \rightarrow \infty} \mathbf{E}_\mu[A_t f] = \mathbf{E}_m[f] \tag{23.24}$$

PROOF: By Eq. 23.20, this must hold when f is an indicator function. By the linearity of A_t and of expectations, it thus holds for simple functions, and so for general measurable functions, using boundedness to exchange limits and expectations where necessary. \square

Lemma 331 (Expectation of the Ergodic Limit is the AMS Expectation) *If f is a bounded observable in $\text{Dom}(A)$, and μ is AMS with stationary mean m , then $\mathbf{E}_\mu[Af] = \mathbf{E}_m[f]$.*

PROOF: From Lemma 323, $\mathbf{E}_\mu[Af] = \lim_{t \rightarrow \infty} \mathbf{E}_\mu[A_t f]$. From Lemma 330, the latter is $\mathbf{E}_m[f]$. \square

Remark: Since Af is invariant, we've got $\mathbf{E}_\mu[Af] = \mathbf{E}_m[Af]$, from Lemma 329, but that's not the same.

Corollary 332 (Replacing Boundedness with Uniform Integrability) *Lemmas 330 and 331 continue to hold if f is not bounded, but $A_t f$ is uniformly integrable (μ).*

PROOF: As in Corollary 324. \square

Theorem 333 (Ergodic Limits of Bounded Observables are Conditional Expectations) *If μ is AMS, with stationary mean m , and the dynamics have ergodic properties for all the indicator functions, then, for any bounded observable f ,*

$$Af = \mathbf{E}_m [f|\mathcal{I}] \tag{23.25}$$

with probability 1 under both μ and m .

PROOF: Begin by showing this holds for the indicator functions of measurable sets, i.e., when $f = \mathbf{1}_C$ for arbitrary measurable C . By assumption, this function has the ergodic property.

By Lemma 318, $A\mathbf{1}_C$ is an invariant function. Pick any set $B \in \mathcal{I}$, so that $\mathbf{1}_B$ is also invariant. By Lemma 319, $A(\mathbf{1}_B\mathbf{1}_C) = \mathbf{1}_BA\mathbf{1}_C$, which is invariant (as a product of invariant functions). So Lemma 329 gives

$$\mathbf{E}_\mu [\mathbf{1}_BA\mathbf{1}_C] = \mathbf{E}_m [\mathbf{1}_BA\mathbf{1}_C] \tag{23.26}$$

while Lemma 331 says

$$\mathbf{E}_\mu [A(\mathbf{1}_B\mathbf{1}_C)] = \mathbf{E}_m [\mathbf{1}_B\mathbf{1}_C] \tag{23.27}$$

Since the left-hand sides are equal, the right-hand sides must be equal as well, so

$$m(B \cap C) = \mathbf{E}_m [\mathbf{1}_B\mathbf{1}_C] \tag{23.28}$$

$$= \mathbf{E}_m [\mathbf{1}_BA\mathbf{1}_C] \tag{23.29}$$

Since this holds for all invariant sets $B \in \mathcal{I}$, we conclude that $A\mathbf{1}_C$ must be a version of the conditional probability $m(C|\mathcal{I})$.

From Lemma 320, every bounded observable has the ergodic property. The proof above then extends naturally to arbitrary bounded f . \square

Corollary 334 (Ergodic Limits of Integrable Observables are Conditional Expectations) *Equation 23.25 continues to hold if $A_t f$ are uniformly μ -integrable, or f is m -integrable.*

PROOF: Exercise 23.3. \square

23.5 Exercises

Exercise 23.1 (“Infinitely often” implies invariant) *Prove Lemma 306.*

Exercise 23.2 (Invariant simple functions) *Prove Lemma 309.*

Exercise 23.3 (Ergodic limits of integrable observables) *Prove Corollary 334.*