UNDERSTANDING ERGODICITY

Michael Keane
Mathematics and Computer Science Department, Wesleyan University, Middletown, Connecticut
mkeane@wesleyan.edu

Received: 10/14/10, Revised: 3/16/11, Accepted: 7/15/11, Published: 12/2/11

Abstract

In this article I try to explain some aspects of the nature of ergodic theory, and in particular why results concerning dynamical averages are interesting, and what we feel to be the natural proofs of convergence of such averages.

1. Introduction

This article is based on a part of several lectures given in the instructional week of the conference *NUMERATION* at the University of Leiden in June 2010. In it we try to explain some aspects of the nature of ergodic theory, and in particular why results concerning dynamical averages are interesting, and what we feel to be the natural proofs of convergence of such averages. We have attempted to make the article accessible to a broad audience.

2. Averages of 0-1 Sequences

Let

$$x := (x_0, x_1, \cdots)$$

be a sequence such that for each nonnegative integer n, x_n is either 0 or 1. For each $n \geq 1$ set

$$a_n = \frac{x_0 + \cdots + x_{n-1}}{n}.$$

Then a_n is the average of the sequence up to time (or position) n. Clearly

$$a := (a_1, a_2, \dots)$$

is also a sequence, the sequence of *averages*, and for each $n \geq 1$,

$$0 \leq a_n \leq 1.$$

In particular,

$$0 \leq \underline{a} = \liminf a_n \leq \limsup a_n = \overline{a} \leq 1,$$

and the limit of this sequence exists if and only if $\underline{a} = \overline{a}$.

3. The Quasi-Cauchy Property of Averages

Averages of 0-1 sequences (and many real sequences otherwise defined) "seem to converge". That is, for each $n \geq 1$,

$$(n+1)a_{n+1} = x_0 + \cdots x_n = na_n + x_n,$$

so

$$a_{n+1} - a_n = \frac{1}{n}(x_n - a_{n+1})$$

and

$$|a_{n+1} - a_n| = \frac{1}{n}|x_n - a_{n+1}| \leq \frac{1}{n}.$$

In particular,

$$\lim(a_{n+1} - a_n) = 0.$$

That is, the sequence of averages is a so-called *quasi-Cauchy sequence* (see [1]).

4. Convergence of Quasi-Cauchy Sequences

Quasi-Cauchy sequences arise in diverse situations, and it is often difficult to determine whether or not they converge, and if so, to which limit. It is easy to construct a zero-one sequence such that the quasi-Cauchy average sequence does not converge. The usual constructions have a somewhat artificial feeling. Nevertheless, there are sequences which seem natural, have the quasi-Cauchy property, and do not converge. We know of two interesting examples which are not sequences of averages.

Example 1. Let n be a positive integer. In a group of n people, each person selects at random and simultaneously another person of the group. All of the selected persons are then removed from the group, leaving a (random) number $n_1 < n$ of people which form a new group. The new group then repeats independently the selection and removal thus described, leaving $n_2 < n_1$ persons, and so forth until either one person remains, or no persons remain. Denote by p_n the probability that, at the end of this iteration initiated with a group of n persons, one person remains. Then the sequence

$$p = (p_1, p_2, \cdots)$$

is a quasi-Cauchy sequence, and $\lim p_n$ does not exist (see [8]).

Example 2. Again let n be a positive integer. In a group of n people, each person selects independently and at random one of three subgroups to which to belong, resulting in three groups with (random) numbers n_1, n_2, n_3 of members; $n_1 + n_2 + n_3 = n$. Each of the subgroups is then partitioned (independently) in the same manner to form three subsubgroups, and so forth. Subgroups having no members or having only one member are removed from the process. Denote by t_n the expected value of the number of iterations up to complete removal, starting initially with a group of n people. Then the sequence

$$(t_1, t_2/2, t_3/3, \cdots)$$

is a (bounded) nonconvergent quasi-Cauchy sequence (see [J]).

To our knowledge, it is not easy to prove either of the above claims, nor to determine the inferior and superior limits of these and similar sequences. At first sight, one would think that both sequences converge.

5. Dynamical Averages

We next consider a general way in which averages of 0-1 sequences arise.

Let X be any set, and let

$$S : X \to X$$

be a mapping from the set into itself. Let A be any subset of X. Using this data, define a zero-one sequence for each $x \in X$ by setting $x_n = 1$ if $S^n(x) \in A$ and $x_n = 0$ if not. That is

$$x_n = 1_A(S^n x).$$

The corresponding averages are then functions of the point x given by

$$a_n(x) = \frac{1}{n} \sum_{k=0}^{n-1} 1_A(S^k x) = \frac{x_0 + \cdots + x_{n-1}}{n},$$

for $n \geq 1$. We call such averages *dynamical averages*. Let us note in passing that as before, for each $x \in X$,

$$\underline{a}(x) = \liminf a_n(x)$$

$$\overline{a}(x) = \limsup a_n(x)$$

then $0 \le \underline{a}(x) \le \overline{a}(x) \le 1$.

Interesting here is that the asymptotic behaviour for the averages arising from a point x and its image Sx are the same. That is, for $n \ge 1$,

$$na_n(x) = \sum_{k=0}^{n-1} 1_A(S^k x)$$

$$na_n(Sx) = \sum_{k=1}^{n} 1_A(S^k x),$$

so that we have

$$|a_n(x) - a_n(Sx)| \le \frac{1}{n}.$$

Hence the values of the functions $\underline{a}$ and $\overline{a}$ are the same on the *forward orbit* $x, Sx, \cdots$ of x. In particular, convergence of averages is a property of orbits of S.

6. Do Dynamical Averages Converge?

Up to now the discussion has centered on a broad general situation. To motivate further the interest in this situation, we briefly describe a classical unsolved problem which falls in this area.

Example 3. Let $X := [0, 1)$. For $x \in X$ define $Sx := 10x \bmod 1$ and set $A := [3/10, 4, 10)$. Consider now the point $x := \sqrt{2} - 1 = 0.41421359\cdots$. Then it should be clear that x_n has the value 1 if the n-th digit in the decimal expansion of x is 3, and otherwise has the value 0. The dynamical average $a_n(x)$ is then the proportion of threes in the first n digits of the decimal expansion of x. Computer calculation rapidly leads us to believe that the limit $\lim a_n(x)$ exists and is equal to 1/10. However, up to now no one has been able to prove this, or any similar statement. Even worse, we can't even show yet that $\underline{a} > 0$. The choices of 10 and 3 have no meaning here - replace 10 by any integer greater than one and 3 with any nonnegative integer less than it, and the difficulty persists. Numbers which possess the correct frequencies of digits in a given base are said to be *simply normal*, and it is unknown whether any irrational algebraic number is simply normal (or not) in any base.

One could think (and perhaps correctly) that the basic difficulty in the preceding problem arises from number theoretic complications - we know from many sources that questions in this area can be very intricate. So here is another example in which number theory seems to play less of a role.

Example 4. Instead of defining the sequence we wish to consider in a dynamical manner, we shall give it directly as a sequence of the symbols 1 and 2 as follows:

$$x := 22\ 11\ 2\ 1\ 22\ 1\ 22\ 11\ \cdots$$

In this sequence, ones and twos alternate in groups of lengths either one or two, and the sequence itself gives the length numbers. So the first 2 indicates that the first group 22 has length two, the second 2 that the second group 11 has length two, and so on, which defines the sequence uniquely given that we start with a group of twos. Again extensive calculation yields the frequency of 1's (and of 2's) to have the limit one-half, but no proof has yet been found. This sequence is called the Kolakowski sequence; for an interesting discussion see [2].

7. Goals of Ergodic Theory

Dynamical averages have many types of behaviour, and currently problems such as the ones briefly discussed in the previous section are not within reach. Ergodic theory is an attempt to prove convergence of dynamical averages not for all points x in X, but for a "large number" of such points, without explicitly determining for exactly which points the convergence holds. At the same time, an attempt is made to give a rule for calculation of this limit, and a large part of the theory also deals with classification of those dynamical averages which have the same "behaviour", coming from different setups. In our discussion we shall first limit our attention to the simplest basic ergodic theorem, which provides existence and limit identification for a wide class of dynamical averages. In further sections we try to indicate the extension of these results to more intricate situations. The field is closely linked to probability theory and number theory through independence and combinatorics respectively. In this view, basic ergodic theory is:

- An attempt to prove existence of $\lim a_n(x)$ for a large class of sequences obtained by the above procedure.
- An attempt to calculate the values of such limits, if they exist.
- An attempt to determine, given two different mappings, whether they produce dynamical sequences which exhibit the same average behaviour.

More generally and less explicitly, ergodic theory is often seen as an attempt to explain and classify the long-term behaviour of iteratively defined (i.e., dynamical) systems.

8. A Case of Success - The Weyl Equidistribution Theorem

If $X = \boldsymbol{R}/\boldsymbol{Z}$, and if $Sx = x + \sigma$, where σ is irrational, then for every interval A in X and every point x of X, the sequence of averages described above converges to the length of the interval A. This is the content of the Weyl equidistribution theorem, which certainly many in this audience have studied. Dynamical systems which behave in this manner are called Kronecker systems. The generalization to rotations of compact Abelian groups is quite straightforward. These are the first examples usually studied in ergodic theory. The reason that we understand convergence here is the prevalence of functions which transform simply under the mapping S, namely, the so-called characters of the group. Each such character is a continuous function f from the group to the complex numbers with modulus one, which is an eigenfunction under the mapping S, the eigenvalue being the value of the character on the element which defines the rotation. Thus the sequence to be averaged consists of successive multiples of a single complex number (the value of the character at x) by successive powers of a complex number of modulus one (the value of the character at σ), and the limit of the average is seen to be either zero or the single complex number. The theorem then follows easily from the uniform density of linear combinations of characters in the continuous functions on the group. Interesting here is also that some functions, namely the characters, have simple behaviour under the mapping S, and that any integrable function can be closely approximated by such simple functions. This yields also easily almost everywhere convergence when the set A is any measurable set, and in the following we shall show that a similar strategy can be carried out for any measure preserving transformation of a probability space. A further development is provided by so-called nilmanifolds, to which currently active research is being devoted. If you have not been exposed to rotations, then it is a worthwhile exercise to verify the above ideas in detail.

9. Interval Exchange Transformations

In this section we meet the first example of a seemingly natural sequence of dynamical averages which does not converge. Let n be an integer greater than one. Suppose that α is a probability vector of length n and that π is a finite permutation of $\{1, \cdots, n\}$. Cut the unit interval X into successive subintervals of lengths determined by α and rearrange these intervals in the order given by π. This defines (almost - up to finitely many cut points) a mapping S which is called an interval exchange transformation. Note that if $n = 2$ and if π is not the identity, then we obtain (essentially) the transformation at the beginning of the previous section, with σ being the second coordinate of α. Now previous to the Weyl theorem a simple argument due to Kronecker based on the pigeonhole principle was known (and you probably have encountered this) showing that if σ is irrational, then the orbit under S of any point in X is dense in X. When this occurs, we say that the transformation S is minimal, because then the only closed invariant sets are the empty set and the entire space X. It is a remarkable fact that if the only rational relation among the coordinates of α is that their sum is one, and if the permutation π fixes no initial segment $\{1, \cdots, k\}$ (except when $k = n$), then the corresponding interval exchange transformation is minimal - every orbit is dense (see [5]). At first sight, one might then also expect a proof of the Weyl theorem, and this holds indeed for the case in which $n = 3$. However, for $n = 4$ an example has been constructed of an "irrational" α (with the permutation π which reverses the order of the four intervals) for which there are points of the space for which the dynamical averages, say, with A being any of the four intervals, do not converge! (See [6].) In this example, the vector α can be made explicit, but we don't seem to be able to determine a point of nonconvergence, although we know that many such exist.

10. Hypotheses of the Basic Ergodic Theorem

The basic ergodic theorem provides a large class of dynamical sequences for which the limiting averages exist. Unfortunately, it gives less than satisfactory answers, as we shall discuss afterwards. But first, we formulate the statement which we can prove.

Suppose that the set X is provided with a measure-theoretic structure which makes it into a probability space. That is, there is a σ-algebra $\mathcal{A}$ of subsets of X, called the measurable subsets, and a measure μ assigning a number $\mu(A)$ between 0 and 1 to each measurable subset which has the properties of a probability measure. Suppose also that the mapping S *preserves* the measure μ, in the sense that for any measurable set A, we have $\mu(A) = \mu(S^{-1}A)$.

Basic Ergodic Theorem. For any fixed measurable set A, the averages $a_n(x)$ indeed converge, except for an exceptional set of points x of X having μ-measure zero. That is, for μ-almost every point $x \in X$,

$$\lim \frac{1}{n} \sum_{k=0}^{n-1} 1_A(S^k x)$$

exists.

Some Comments

- It can (and often does) occur that although the ergodic theorem is valid, we cannot decide for any particular point of interest whether $\lim a_n(x)$ exists or not.
- It can also occur that even when we can prove that the limit exists, we cannot determine its value.
- If A is an *invariant* set (i.e., $S^{-1}A = A$), then clearly the limit is one for points belonging to A and zero for points not belonging to A; however, often we have trouble determining the σ-algebra of invariant sets.
- Generalizations do not usually cause trouble. For instance, once the basic ergodic theorem is properly understood, averages of the form

$$\frac{1}{n}\sum_{k=0}^{n-1} f(S^k x),$$

where the function f is μ-integrable, can be shown to exist.

A mathematically fully equivalent formulation merits attention:

Probabilistic Formulation of The Basic Ergodic Theorem: If x_n is a stationary sequence of 0-1 valued random variables, then the average sequence a_n converges almost surely.

In many situations it is more convenient to use probabilistic language. Here we choose to remain in the dynamical setting.

11. The Proof

As in our explanation of the Weyl theorem, we shall begin by slightly modifying the measurable set A to obtain a more "well-behaved" set A^*. Let $\epsilon > 0$. For each $x \in X$ there is a positive integer N such that

$$a_N(x) \geq \overline{a}(x) - \epsilon.$$

To be specific (although any choice would work), we take $N(x)$ to be the least such integer. Note that if $x \in A$, then $N(x) = 1$. Choose now a positive integer M large enough so that the μ-measure of the set of $x \notin A$ with $N(x) > M$ is less than ϵ, and define

$$A^* := A \cup \{x \notin A : N(x) > M\}.$$

Then, using the notation a^* for averages of visit numbers to A^*, we see that $\mu(A^*) \leq \mu(A) + \epsilon$ and that for every $x \in X$, there is a positive integer $M(x)$ which is less than or equal to M such that $a^*_{M(x)} \geq \overline{a}(x) - \epsilon$.

It should now be clear that the averages $a^*_n(x)$ are well–behaved, in the sense that for all sufficiently large positive integers n (not depending on the point x), $a^*_n(x) \geq \overline{a} - 2\epsilon$. That is, the amount of time needed to reach the lim sup within ϵ using $a*$ is bounded by M for any point of the space, so that if n is large with respect to M, we can break the sequence into pieces all of length at most M and each of which is close to the lim sup.

Choosing any such integer n, we have therefore for every $x \in X$ that

$$\frac{1}{n}\sum_{k=0}^{n-1} 1_{A^*}(S^k x) = a^*_n(x) \geq \overline{a}(x) - 2\epsilon.$$

Integration by μ now yields

$$\int \overline{a}d\mu - 2\epsilon \leq \mu(A^*) \leq \mu(A) + \epsilon,$$

and, letting ϵ tend to zero we obtain

$$\int \overline{a}d\mu \leq \mu(A).$$

If the set A is now replaced by its complement and we figure out what the inequality

$$\int \overline{a}d\mu \leq \mu(A)$$

means for the complement, then we see that $\overline{a}$ for the complement becomes $1 - \underline{a}$ and the measure of the complement becomes $1 - \mu(A)$. Hence we obtain

$$\mu(A) \leq \int \underline{a}d\mu.$$

So the integrals of $\underline{a}$ and $\overline{a}$ must be equal, implying equality of the lim sup and lim inf μ-almost everywhere.

12. Comments

It is fortunate that we have been able to prove convergence in a simple manner for a large set of starting points and every measurable set. However, the situation is far from being satisfactory:

- It can (and often does) occur that we cannot decide for a particular point of interest and a specific set of interest whether the limiting average exists (see Examples 1 and 2 above).
- Even when we know that the limit exists, we may not be able to determine its value.
- If A is an invariant set (i.e., $s^{-1}A = A$), then clearly for each n and$x$,

$$a_n(x) = 1_A(x).$$

  But it is often unknown which sets are invariant. In particular, S is said to be *ergodic* if the only invariant sets are those of measures 0 or 1. The situation sketched above for interval exchange transformations with four intervals arises because of non-ergodicity, and it is in many situations not clear whether an S is ergodic or not.

On the other hand, there is little trouble in extending the basic ergodic theorem to iteration of functions instead of sets. For instance, if f is a μ-integrable function on X, the same proof essentially shows that

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

converges for μ-almost every $x \in X$. (Hint: if we denote the lim sup by $\overline{f}$, then the basic inequality works using $\min \overline{f}, K$ for any nonnegative integrable f and nonnegative constant K.) Another case which works well is the subadditive ergodic theorem (see [4] for the first proof) which follows almost trivially using the argument given here.

The original idea of simplifying the set appeared in [3]; a similar, perhaps more convoluted description was given earlier by [7]. Previous proofs are elegant and sometimes very short, but somehow do not convey the essence of why such a theorem is true, in our opinion.

13. More General Actions

Finally, we'd like to briefly discuss the situation in which we have *two* mappings S and T, both taking X to itself and preserving the measure μ; let us suppose in addition that S and T commute. So for each pair (m, n) of nonnegative integers and each $x \in X$, we have a point $S^m T^n x \in X$, and we'd like to investigate the convergence of sums of the form

$$\frac{1}{MN}\sum_{k=0}^{m-1}\sum_{l=0}^{n-1} 1_A(S^k T^l x)$$

as M and N tend to infinity. It seems that the same procedure should work, but there is one difficulty, namely, the point in the proof where, after having found a "well–behaved" set close to A, we need to actually show that the sums for this set reach the lim sup minus ϵ for all large MN. To accomplish this, we must produce a collection of rectangles over each of which the average is close to the lim sup, and which are pairwise disjoint and almost cover sufficiently large rectangles. If one uses a so-called greedy algorithm to choose the largest possible rectangle sequentially, then it is easy to see that roughly one-quarter of the large rectangle will be covered. In order to get more, one repeats this procedure on different scales, using the property of the lim sup that it is attained within ϵ infinitely often. A detailed description involves a number of pages of definitions and calculations, which are omitted. This can be carried out for actions of discrete (countable) amenable groups, with some effort.

References

[1] D. Burton and J. Coleman, *Quasi-Cauchy sequences*, Amer. Math. Monthly 117, 328-333. (April 2010)

[2] F.M. Dekking, *What is the long range order in the Kolakoski sequence?*, in: The mathematics of long-range aperiodic order (Waterloo, ON, 1995), 115-125, NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., 489, (Eds.: R.V.Moody and J.Patera), Kluwer Acad. Publ., Dordrecht, 1997.

[3] T. Kamae, *A simple proof of the ergodic theorem using nonstandard analysis*, Israel J. Math. 42-4 (1982) pp.284-290.

[4] Y.Katznelson and B.Weiss, *A Simple Proof of Some Ergodic Theorems* , Israel J. Math., Vol. 42, 1982, pp. 291-296.

[5] M. Keane, *Interval Exchange Transformations*, Math. Z. 14, 25–31 (1975).

[6] M. Keane, *Non–ergodic Interval Exchange Transformations*, Israel J. of Math. 26, 188–196 (1977).

[7] Paul Shields, personal communication.

[8] P. Winkler, *Mathematical Puzzles: A Connoisseur's Collection*, A.K.Peters LTD, 2004, ISBN 1-56881-201-9.