

8 The Laws of Large Numbers

The traditional interpretation of the *probability* of an event E is its *asymptotic frequency*: the limit as $n \rightarrow \infty$ of the fraction of n repeated, similar, and independent trials in which E occurs. Similarly the “expectation” of a random variable X is taken to be its *asymptotic average*, the limit as $n \rightarrow \infty$ of the average of n repeated, similar, and independent replications of X . For statisticians trying to make inference about the underlying probability distribution $f(x | \theta)$ governing observed random variables X_i , this suggests that we should be interested in the probability distribution for large n of quantities like the sample average of the RVs, $\bar{X}_n := \frac{1}{n} \sum_{i=1}^n X_i$ or the partial sum $S_n := \sum_{i=1}^n X_i$. Three of the most celebrated theorems of probability theory concern this quantity. For independent random variables X_i , all with the same probability distribution satisfying $E|X_i|^3 < \infty$, set $\mu := EX_i$, $\sigma^2 := E|X_i - \mu|^2$, and $S_n := \sum_{i=1}^n X_i$. The three main results are:

Laws of Large Numbers (LLN):

$$\frac{S_n - n\mu}{\sigma n} \rightarrow 0 \quad (\text{pr. and a.s.})$$

Central Limit Theorem (CLT):

$$\frac{S_n - n\mu}{\sigma\sqrt{n}} \Rightarrow \text{No}(0, 1) \quad (\text{in dist.})$$

Law of the Iterated Logarithm (LIL):

$$\limsup_{n \rightarrow \infty} \pm \frac{S_n - n\mu}{\sigma\sqrt{2n \log \log n}} = 1.0 \quad (\text{a.s.})$$

Together these three give a clear picture of how quickly and in what sense $\frac{1}{n}S_n$ tends to μ . We begin with the Law of Large Numbers (LLN), first in its “weak” form (asserting convergence *pr.*) and then in its “strong” form (convergence *a.s.*). There are several versions of both theorems. The simplest requires the X_i to be IID and L_2 ; stronger results allow us to weaken (but not eliminate) the independence requirement, permit non-identical distributions, and consider what happens if we relax the L_2 requirement and allow the RVs to be only L_1 (or worse!).

The text covers these things well; to complement it I am going to: (1) Prove the simplest version, and with it the Borel-Cantelli theorems; and (2) Show what happens with Cauchy random variables, which don’t satisfy the requirements (the LLN fails).

8.1 Proofs of the Weak and Strong Laws

Here are two simple versions (one Weak, one Strong) of the Law of Large Numbers; first we prove an elementary but very useful result:

Proposition 1 (Markov’s Inequality) *Let $\phi(x) \geq 0$ be non-decreasing on \mathbb{R}_+ . For any random variable $X \geq 0$ and constant $a \in \mathbb{R}_+$,*

$$P[X \geq a] \leq P[\phi(X) \geq \phi(a)] \leq E[\phi(X)]/\phi(a)$$

To see this, set $Y := \phi(a)\mathbf{1}_A$ for the event $A := \{\phi(X) \geq \phi(a)\}$ and note $Y \leq \phi(X)$ so $EY \leq E\phi(X)$.

Theorem 1 (L_2 WLLN) Let $\{X_n\}$ be independent random variables with the same mean $\mu = E[X_n]$ and uniformly bounded variance $E(X_n - \mu)^2 \leq B$ for some fixed bound $B < \infty$. Set $S_n := \sum_{j \leq n} X_j$ and $\bar{X}_n := S_n/n = \frac{1}{n} \sum_{j \leq n} X_j$. Then, as $n \rightarrow \infty$,

$$(\forall \epsilon > 0) \quad P[|\bar{X}_n - \mu| > \epsilon] \rightarrow 0. \quad (1)$$

Proof.

$$E(S_n - n\mu)^2 = \sum_{i=1}^n E(X_i - \mu)^2 \leq nB$$

so for $\epsilon > 0$

$$\begin{aligned} P[|\bar{X}_n - \mu| > \epsilon] &= P[(S_n - n\mu)^2 > (n\epsilon)^2] \\ &\leq E[(S_n - n\mu)^2]/n^2\epsilon^2 \\ &\leq B/n\epsilon^2 \rightarrow 0 \quad \text{as } n \rightarrow \infty. \end{aligned}$$

□

This Law of Large Numbers is called *weak* because its conclusion is only that \bar{X}_n converges to zero *in probability* (Eqn (1)); the *strong* Law of Large Numbers asserts convergence of a stronger sort, called *almost sure* convergence (Eqn (2) below). If $P[|\bar{X}_n - \mu| > \epsilon]$ were *summable* then by B-C we could conclude almost-sure convergence; unfortunately we have only the bound $P[|\bar{X}_n - \mu| > \epsilon] < c/n$ which tends to zero but isn't summable. It is summable along the *subsequence* n^2 , however; our approach to proving a strong LLN is to show that $|S_k - S_{n^2}|$ isn't too big for any $n^2 \leq k < (n+1)^2$.

Theorem 2 (L_2 SLLN) Under the same conditions,

$$P[\bar{X}_n \rightarrow \mu] = 1. \quad (2)$$

Proof. Without loss of generality take $\mu = 0$ (otherwise subtract μ from each X_n), and fix $\epsilon > 0$. Set $S_n := \sum_{j \leq n} X_j$. Then

$$\begin{aligned} P[|S_{n^2}| > n^2\epsilon] &\leq E|S_{n^2}|^2/(n^2\epsilon)^2 \leq n^2B/n^4\epsilon^2 = B/n^2\epsilon^2 \\ P[|S_{n^2}| > n^2\epsilon \text{ i.o.}] &= 0 \text{ by B-C} \Rightarrow S_{n^2}/n^2 \rightarrow 0 \text{ a.s. Set} \\ D_n &:= \max_{n^2 \leq k < (n+1)^2} |S_k - S_{n^2}| \\ ED_n^2 &= E \left[\max_{n^2 \leq k < (n+1)^2} |S_k - S_{n^2}|^2 \right] \\ &\leq E \sum_{n^2 \leq k < (n+1)^2} |S_k - S_{n^2}|^2 = \sum_{n^2 \leq k < (n+1)^2} E|S_k - S_{n^2}|^2 \\ &\leq \sum_{n^2 \leq k < (n+1)^2} (k - n^2)B \leq 4n^2B, \text{ so} \\ P[D_n > n^2\epsilon] &\leq 4n^2B/n^4\epsilon^2 = 4B/n^2\epsilon^2 \Rightarrow D_n/n^2 \rightarrow 0 \text{ a.s.} \\ \left| \frac{S_k}{k} \right| &\leq \frac{|S_{n^2}| + D_n}{n^2} \rightarrow 0 \text{ a.s. as } k \rightarrow \infty, \text{ where } n = \lfloor \sqrt{k} \rfloor. \end{aligned}$$

□

Each of these LLNs required only that $\text{Cov}(X_n, X_m) \leq 0$, not pairwise (let alone full) independence. To see that, suppose $\mathbb{E}X_n = \mu$ for all n , $\text{Cov}(X_n, X_m) \leq 0$ for $n \neq m$, and $\forall X_n \leq B$. Then

$$\begin{aligned} \mathbb{E}(S_n - n\mu)^2 &= \mathbb{E} \sum_{i,j=1,1}^{n,n} (X_i - \mu)(X_j - \mu) \\ &= \sum_{i=1}^n \mathbb{E}(X_i - \mu)^2 + 2 \sum_{1 \leq i < j \leq n} \mathbb{E}(X_i - \mu)(X_j - \mu) \\ &\leq nB. \end{aligned}$$

We'll see below in Section (8.4) that even positive correlations are okay if they fall off fast enough— *e.g.*, if $|\text{Cov}(X_n, X_m)| \leq ar^{|n-m|}$ for some $a > 0$, $0 < r < 1$ — with a similar proof. The uniform L_2 bound isn't necessary either. There are a variety of LLNs with either or both of the L_2 bound and independence weakened in some way, but they can't be dispensed with altogether— consider iid Cauchy random variables, for example, to show L_2 isn't entirely superfluous, or $X_n \equiv X_1$ with any nontrivial distribution to show the need for at least a modicum of independence.

8.2 Other Strong Laws

Let's first state two lemmas:

Lemma 1 (Lévy) *If $\{X_n\}$ is an independent sequence then $\sum_{n=1}^{\infty} X_n$ converges pr. if and only if it converges a.s.*

Lemma 2 (Kronecker) *Suppose $\{x_n\} \subset \mathbb{R}$ and $0 < a_n \nearrow \infty$. Then*

$$\sum_{k=1}^{\infty} \frac{x_k}{a_k} \text{ converges} \Rightarrow \frac{1}{a_n} \sum_{k=1}^n x_k \rightarrow 0.$$

and, from these, prove two useful theorems. First,

Theorem 3 (Kolmogorov Convergence Criterion) *Let $\{X_n\} \subset L_2$ be an independent sequence with means $\mu_n := \mathbb{E}X_n$ and variances $\sigma_n^2 := \mathbb{V}(X_n)$. Then*

$$\sum_{n=1}^{\infty} \sigma_n^2 < \infty \Rightarrow \sum_{n=1}^{\infty} (X_n - \mu_n) \text{ converges a.s.}$$

Proof. Without loss of generality take $\mu_n \equiv 0$. Fix $\epsilon > 0$. For $1 \leq n \leq N < \infty$,

$$\begin{aligned} \mathbb{P}[|S_N - S_n| > \epsilon] &= \mathbb{P}\left[\left|\sum_{n < k \leq N} X_k\right| > \epsilon\right] \\ &\leq \mathbb{E}\left(\sum_{n < k \leq N} X_k\right)^2 / \epsilon^2 && \text{by Markov's inequality} \\ &\leq \sum_{n < k \leq \infty} \sigma_k^2 / \epsilon^2 && \text{by independence} \\ &\rightarrow 0 \text{ as } n \rightarrow \infty && \text{since } \sum \sigma_n^2 < \infty, \end{aligned}$$

so S_n is a convergent Cauchy sequence in probability and, by Lemma 1, also converges *a.s.* \square

Theorem 4 (Kronecker L_2 SLLN) *Let $\{X_n\} \subset L_2$ be an independent sequence, and let $\{b_n\} \subset \mathbb{R}_+$ be a monotone sequence such that*

$$\sum_n \mathbb{V}\left(\frac{X_n}{b_n}\right) < \infty.$$

Then

$$\frac{S_n - \mathbb{E}S_n}{b_n} \rightarrow 0 \text{ a.s.}$$

In particular, for iid $\{X_n\}$ we may take $b_n = n$ to see $\bar{X}_n \rightarrow \mu$ *a.s.*, a SLLN for independent but non-identically-distributed L_2 random variables.

Proof. Again take $\mathbb{E}[X_n] \equiv 0$, and set $Y_n := X_n/b_n$, with variance σ_n^2 . By hypothesis $\sum \sigma_n^2 < \infty$, so $\sum Y_n$ converges *a.s.* by Theorem 3. Thus, for almost every ω , $\sum X_n(\omega)/b_n$ converges and, by Lemma 2, also $S_n/b_n \rightarrow 0$. \square

Kolmogorov's L_1 SLLN

The most-cited and most-used version of the Strong Law for iid sequences is that due to Kolmogorov, with no moment assumptions:

Theorem 5 (Kolmogorov's L_1 SLLN) *Let $\{X_n\}$ be iid and set $S_n := \sum_{i \leq n} X_i$. There exists $c \in \mathbb{R}$ such that*

$$\bar{X}_n = S_n/n \rightarrow c \text{ a.s.}$$

if and only if $\mathbb{E}|X_1| < \infty$, in which case $c = \mathbb{E}X_1$.

This has the pleasant consequence that the usual estimators for the mean and variance of iid sequences are consistent:

Corollary 1

$$\begin{aligned} \{X_n\} \subset L_1 &\Rightarrow \bar{X}_n \rightarrow \mu \text{ a.s.} \\ \{X_n\} \subset L_2 &\Rightarrow \frac{1}{n} \sum_{i \leq n} (X_i - \bar{X}_n)^2 \rightarrow \sigma^2 \text{ a.s.} \end{aligned}$$

Here's a quick summary of some LLN facts:

I. Weak version, non-*iid*, L_2 : $\mu_i = \mathbf{E}X_i$, $\sigma_{ij} = \mathbf{E}[X_i - \mu_i][X_j - \mu_j]$

A. $Y_n = (S_n - \sum \mu_i)/n$ satisfies $\mathbf{E}Y_n = 0$, $\mathbf{E}Y_n^2 = \frac{1}{n^2} \sum_{i \leq n} \sigma_{ii} + \frac{2}{n^2} \sum_{i < j \leq n} \sigma_{ij}$;

1. If $\sigma_{ii} \leq M$ and $\sigma_{ij} \leq 0$ or $|\sigma_{ij}| < Mr^{|i-j|}$, $r < 1$,
Chebychev $\implies Y_n \rightarrow 0$, *pr.*

2. (pairwise) IID L_2 is OK

II. Strong version, non-*iid*, L_2 : $\mathbf{E}X_i = 0$, $\mathbf{E}X_i^2 \leq M$, $\mathbf{E}X_i X_j \leq 0$.

A. $\mathbf{P}[|S_n| > n\epsilon] < \frac{Mn}{n^2\epsilon^2} = \frac{M}{n\epsilon^2}$

1. $\mathbf{P}[|S_{n^2}| > n^2\epsilon] < \frac{M}{n^2\epsilon^2}$, $\sum_n \mathbf{P}[|S_{n^2}| > n^2\epsilon] < \frac{M\pi^2}{6\epsilon^2} < \infty$

2. Borel-Cantelli: $\mathbf{P}[|S_{n^2}| > n^2\epsilon \text{ i.o.}] = 0$, $\therefore \frac{1}{n^2} S_{n^2} \rightarrow 0$ *a.s.*

3. $D_n := \max_{n^2 \leq k < (n+1)^2} |S_k - S_{n^2}|$, so
 $\mathbf{E}D_n^2 \leq 2n\mathbf{E}|S_{(n+1)^2-1} - S_{n^2}| \leq 4n^2M$

4. Chebychev: $\mathbf{P}[D_n > n^2\epsilon] < \frac{4n^2M}{n^4\epsilon^2}$, $\therefore D_n/n^2 \rightarrow 0$ *a.s.*

B. $|S_k/k| \leq \frac{|S_{n^2}| + D_n}{n^2} \rightarrow 0$ *a.s.* as $k \rightarrow \infty$, QED

1. Bernoulli RVs, normal number theorem, Monte Carlo integration.

III. Weak version, pairwise-*iid*, L_1

A. Equivalent sequences: If $\sum_n \mathbf{P}[X_n \neq Y_n] < \infty$, then:

1. $\sum_n |X_n - Y_n| < \infty$ *a.s.*

2. $\sum_{i=1}^n X_i$ converges iff $\sum_{i=1}^n Y_i$ converges

3. If $(\exists a_n \nearrow)(\exists X) \ni \frac{1}{a_n} \sum_{i \leq n} X_i \rightarrow X$ then $\frac{1}{a_n} \sum_{i \leq n} Y_i \rightarrow X$ too.

B. If $\{X_n\}$ are iid L_1 then $X_n, Y_n := X_n 1_{|X_n| \leq n}$ are equivalent

IV. Strong version, *iid*, L_1

A. Kolmogorov: For $\{X_n\}$ IID, $(\exists c \ni \bar{X}_n \rightarrow c \text{ a.s.}) \Leftrightarrow (X_n \in L_1)$.

V. Counterexamples: Cauchy,

A. $X_i \sim \frac{dx}{\pi[1+x^2]} \implies \mathbf{P}[|S_n|/n \leq \epsilon] \equiv \frac{2}{\pi} \tan^{-1}(\epsilon) \not\rightarrow 1$, WLLN fails.

B. $\mathbf{P}[X_i = \pm n] = \frac{c}{n^2}$, $n \geq 1$; $X_i \notin L_1$, and $S_n/n \not\rightarrow 0$ *pr.* or *a.s.*

C. $\mathbf{P}[X_i = \pm n] = \frac{c}{n^2 \log n}$, $n > 1$; $X_i \notin L_1$, but $S_n/n \rightarrow 0$ *pr.* and not *a.s.*

D. Medians: for ANY RVs $X_n \rightarrow X_\infty$ *pr.*, then $m_n \rightarrow m_\infty$ if m_∞ is unique.

8.3 An example where the LLN fails: iid Cauchy RVs

Let X_i be iid standard Cauchy RVs, with

$$\mathbb{P}[X_1 \leq t] = \int_{-\infty}^t \frac{dx}{\pi[1+x^2]} = \frac{1}{2} + \frac{1}{\pi} \arctan(t)$$

and characteristic function (we'll learn more about these next week)

$$\mathbb{E} e^{i\omega X_1} = \int_{-\infty}^{\infty} e^{i\omega x} \frac{dx}{\pi[1+x^2]} = e^{-|\omega|}.$$

The sample mean $\bar{X}_n := S_n/n$ has characteristic function

$$\mathbb{E} e^{i\omega S_n/n} = \mathbb{E} e^{i\frac{\omega}{n}[X_1+\dots+X_n]} = \left(\mathbb{E} e^{i\frac{\omega}{n} X_1} \right)^n = (e^{-|\frac{\omega}{n}|})^n = e^{-|\omega|}.$$

Thus S_n/n also has the standard Cauchy distribution with $\mathbb{P}[S_n/n \leq t] = \frac{1}{2} + \frac{1}{\pi} \arctan(t)$. In particular, S_n/n does not converge almost surely, or even in probability.

8.4 A LLN for Correlated Sequences

In many applications we would like a Law of Large Numbers for sequences of random variables that are *not* independent. For example, in Markov Chain Monte Carlo integration, we have a stationary Markov chain $\{X_t\}$ (this means that the distribution of X_t is the same for all t and that the conditional distribution of a future value X_u for $u > t$, given the past $\{X_s \mid s \leq t\}$, depends only on the present X_t) and want to estimate the population mean $\mathbb{E}[\phi(X_t)]$ for some function $\phi(\cdot)$ by the sample mean

$$\mathbb{E}[\phi(X_t)] \approx \frac{1}{T} \sum_{t=0}^{T-1} \phi(X_t).$$

Even though they are identically distributed, the random variables $Y_t := \phi(X_t)$ won't be independent if the X_t aren't independent, so the LLN we already have doesn't quite apply.

Theorem 6 *If an L_2 sequence $\{Y_t\}$ has a constant mean $\mu = \mathbb{E}Y_t$ and a summable autocovariance $\gamma_{st} := \mathbb{E}(Y_s - \mu)(Y_t - \mu)$ that satisfies $\sum_{t=-\infty}^{\infty} |\gamma_{st}| \leq c < \infty$ uniformly in s , then the random variables $\{Y_t\}$ obey a LLN:*

$$\mathbb{E}[Y_t] = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} Y_t.$$

Proof. Let S_T be the sum of the first T Y_t s and (as usual) set $\bar{Y}_T := S_T/T$. The variance of S_T is

$$\begin{aligned} \mathbb{E}[(S_T - T\mu)^2] &= \sum_{s=0}^{T-1} \sum_{t=0}^{T-1} \mathbb{E}[(Y_s - \mu)(Y_t - \mu)] \\ &\leq \sum_{s=0}^{T-1} \sum_{t=-\infty}^{\infty} |\gamma_{st}| \leq Tc, \end{aligned}$$

so \bar{Y}_T has variance $V[\bar{Y}_T] \leq c/T$ and by Chebychev's inequality

$$\begin{aligned} P[|\bar{Y}_T - \mu| > \epsilon] &\leq \frac{E[(\bar{Y}_T - \mu)^2]}{\epsilon^2} = \frac{E[(S_T - T\mu)^2]}{T^2\epsilon^2} \\ &\leq \frac{Tc}{T^2\epsilon^2} = \frac{c}{T\epsilon^2} \rightarrow 0 \quad \text{as } T \rightarrow \infty. \end{aligned}$$

As in the iid case, this WLLN can be extended to a SLLN by first applying it along the sequence $\{T^2 : T \in \mathbb{N}\}$, applying B/C, then filling in the gaps.

A sequence of random variables Y_t is called *stationary* if each Y_t has the same probability distribution and, moreover, each finite set $(Y_{t_1+h}, Y_{t_2+h}, \dots, Y_{t_k+h})$ has a joint distribution that doesn't depend on h . The sequence is called " L_2 " if each Y_t has a finite variance σ^2 (and hence also a well-defined mean μ); by stationarity it also follows that the *covariance*

$$\gamma_{st} = E[(Y_s - \mu)(Y_t - \mu)]$$

is finite and depends only on the absolute difference $|t - s|$ (write: $\gamma_{st} = \gamma(s - t) = \gamma(t - s)$).

Corollary 2 *If a stationary L_2 sequence $\{Y_t\}$ has a summable covariance function, i.e., satisfies $\sum_{t=-\infty}^{\infty} |\gamma(t)| \leq c < \infty$, then*

$$E[Y_t] = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} Y_t.$$

Note we didn't need full stationarity. It would be good enough for $\{Y_t\}$ to be " 2^{nd} -order stationary," i.e., to have a common mean $\mu = EY_t$ and a covariance function $\gamma_{st} = E(Y_s - \mu)(Y_t - \mu) = \gamma(s - t)$ that depends only on the absolute difference $|s - t|$.

8.4.1 Examples

1. **IID:** If X_t are independent and identically distributed, and if $Y_t = \phi(X_t)$ has finite variance σ^2 , then Y_t has a well-defined finite mean μ and $\bar{Y}_T \rightarrow \mu$.

$$\text{Here } \gamma_{st} = \begin{cases} \sigma^2 & \text{if } s = t \\ 0 & \text{if } s \neq t \end{cases}, \text{ so } c = \sigma^2 < \infty.$$

2. **AR₁:** If Z_t are iid $\text{No}(0, 1)$ for $-\infty < t < \infty$, $\mu \in \mathbb{R}$, $\sigma > 0$, $-1 < \rho < 1$, and

$$\begin{aligned} X_t &:= \mu + \sigma \sum_{s=0}^{\infty} \rho^s Z_{t-s} \\ &= \rho X_{t-1} + \alpha + \sigma Z_t, \end{aligned} \tag{*}$$

where $\alpha = (1 - \rho)\mu$, then the X_t are identically distributed (all with the $\text{No}(\mu, \frac{\sigma^2}{1-\rho^2})$ distribution) but not independent (since $\gamma_{st} = \frac{\sigma^2}{1-\rho^2} \rho^{|s-t|} \neq 0$ for $s \neq t$). This is called an "autoregressive process" (because of equation (*), expressing X_t as a regression of previous

X_s s) of order one (because only one earlier X_s appears in (*)), and is about the simplest non-*iid* sequence occurring in applications. Since the covariance is summable,

$$\sum_{t=-\infty}^{\infty} |\gamma(t)| = \frac{\sigma^2}{1-\rho^2} \frac{1+|\rho|}{1-|\rho|} = \frac{\sigma^2}{(1-|\rho|)^2} < \infty,$$

we again have $\bar{X}_T \rightarrow \mu$ *a.s.* as $T \rightarrow \infty$.

3. **Geometric Ergodicity:** If for some $0 < \rho < 1$ and $c > 0$ we have $\gamma_{st} \leq c\rho^{|s-t|}$ for a Markov chain Y_t the chain is called *Geometrically Ergodic* (because $c\rho^t$ is a geometric sequence), and the same argument as for AR_1 shows that \bar{Y}_t converges. Meyn & Tweedie (1993), Tierney (1994), and others have given conditions for MCMC chains to be Geometric Ergodic, and hence for the almost-sure convergence of sample averages to population means. The mean squared error $MSE := E|\bar{Y}_t - \mu|^2$ for a geometrically ergodic sequence is bounded by $MSE \leq \frac{1+\rho}{1-\rho}(\sigma^2/t) \asymp 1/t$, but the constant grows without bound as $\rho \rightarrow 1$. Irreducible aperiodic finite-state Markov chains are geometrically ergodic, with ρ the second-largest eigenvalue of the one-step transition matrix.
4. **General Ergodicity:** Consider the three sequences of random variables on (Ω, \mathcal{F}, P) with $\Omega = (0, 1]$ and $\mathcal{F} = \mathcal{B}(\Omega)$, each with $X_0(\omega) = \omega$:
 - (a) $X_{n+1} := 2X_n \pmod{1}$;
 - (b) $X_{n+1} := X_n + \alpha \pmod{1}$ (Does it matter if α is rational?);
 - (c) $X_{n+1} := 4X_n(1 - X_n)$.

For each there exists a probability measure P (a distribution for X_0) such that the X_n are identically distributed—the $Un(0,1)$ distribution for (a) and (b), and the “arcsin” distribution $Be(\frac{1}{2}, \frac{1}{2})$ for (c), and each is of the form $X_n = X_0(T^n(\omega))$ for a measure-preserving transformation $T : \Omega \rightarrow \Omega$, a measurable mapping $T : \Omega \rightarrow \Omega$ for which $P(E) = P(T^{-1}(E))$ for each $E \in \mathcal{F}$. Such a transformation is called *ergodic* if the only T -invariant events (those that satisfy $E = T^{-1}(E)$) are “almost trivial” in the sense that $P[E] = 0$ or $P[E] = 1$. A sequence $X_n := X_0 \circ T^n$ is called ergodic if T is. Each of the sequences (a)–(c) is ergodic (can you show that?).

Birkhoff’s *Ergodic Theorem* asserts that, if $X_n \in L_1(\Omega, \mathcal{F}, P)$ is an integrable ergodic sequence, then \bar{X}_n converges almost-surely to a T -invariant limit X_∞ as $n \rightarrow \infty$. Since only constants are T -invariant for ergodic sequences, it follows that $\bar{X}_n \rightarrow \mu := EX_n$ *a.s.* for ergodic sequences. The conditions here are weaker than those for the usual LLN; in all three cases above, for example, all ergodic, each X_n is completely determined by X_0 so there is complete dependence, with $\sigma(X_n) \subset \sigma(X_m) \subset \sigma(X_0)$ for all $0 \leq m \leq n$!

For any L_1 distribution $\mu(dx)$ on $(\mathbb{R}, \mathcal{B})$, we can construct iid random variables $\{X_n\} \stackrel{iid}{\sim} \mu(dx)$ on the product probability space $(\Omega = \mathbb{R}^\infty, \mathcal{F} = \mathcal{B}^\infty, P = \otimes \mu)$ and a measure-preserving transformation $T : \Omega \rightarrow \Omega$ called the *left-shift* by

$$X_n(\omega_1, \omega_2, \omega_3, \dots) := \omega_n \qquad T(\omega_1, \omega_2, \omega_3, \dots) := (\omega_2, \omega_3, \omega_4, \dots).$$

The σ -algebra $\mathcal{T} = \{A : A = T^{-1}(A)\}$ of T -invariant sets is just the tail σ -algebra for the independent random variables $\{X_n\}$, so by Kolmogorov’s zero-one law \mathcal{T} is almost-trivial

and so T is ergodic. It follows from Birkhoff's Ergodic Theorem that the sample mean $\bar{X}_n := \frac{1}{n} \sum_{j=1}^n X_j$ converges almost-surely to a T -invariant and hence almost-surely constant random variable whose value must be μ , proving a strong LLN for iid random variables that assumes only L_1 :

Theorem 7 (L_1 iid SLLN) *Let $\{X_n\}$ be iid $L_1(\Omega, \mathcal{F}, \mathbb{P})$ random variables with mean $\mu = \mathbb{E}[X_n]$. Set $S_n := \sum_{j \leq n} X_j$ and $\bar{X}_n := S_n/n = \frac{1}{n} \sum_{j \leq n} X_j$. Then:*

$$\mathbb{P}[\bar{X}_n \rightarrow \mu] = 1.$$

Thus the “space average” $\int_{\mathbb{R}} X_n d\mu$ and the limiting “time average” $\lim_{n \rightarrow \infty} \bar{X}_n$ coincide for ergodic sequences.