

Large Sample Theory

In statistics, we are interested in the properties of particular random variables (or “estimators”), which are functions of our data. In asymptotic analysis, we focus on describing the properties of estimators when the sample size becomes arbitrarily large. The idea is that given a reasonably large dataset, the properties of an estimator even when the sample size is finite are similar to the properties of an estimator when the sample size is arbitrarily large.

In these notes we focus on the large sample properties of *sample averages* formed from i.i.d. data. That is, assume that $X_i \sim i.i.d.F$, for $i = 1, \dots, n, \dots$. Assume $EX_i = \mu$, for all i . The sample average after n draws is $\bar{X}_n \equiv \frac{1}{n} \sum_i X_i$.

We focus on two important sets of large sample results:

(1) **Law of large numbers:** $\bar{X}_n \rightarrow EX$ as $n \rightarrow \infty$.

(2) **Central limit theorem:** $\sqrt{n}(\bar{X}_n - EX) \rightarrow N(0, \cdot)$. That is, \sqrt{n} times a sample average looks like (in a precise sense to be defined later) a normal random variable as n gets large.

An important endeavor of asymptotic statistics is to show (1) and (2) under various assumptions on the data sampling process.



Consider a sequence of random variables Z_1, Z_2, \dots, Z_n .

Convergence in probability:

$Z_n \xrightarrow{p} Z \iff$ for all $\epsilon > 0$, $\lim_{n \rightarrow \infty} Prob(|Z_n - Z| < \epsilon) = 1$.

More formally: for all $\epsilon > 0$ and $\delta > 0$, there exists $n_0(\epsilon, \delta)$ such that for all $n > n_0(\epsilon, \delta)$:

$$Prob(|Z_n - Z| < \epsilon) > 1 - \delta.$$

Note that the limiting variable Z can also be a random variable. Furthermore, for convergence in probability, the random variables Z_1, Z_2, \dots and Z should be defined on the same probability space: if this common sample space is Ω with elements ω , then the statement of convergence in probability is that:

$$Prob(\omega \in \Omega : |Z_n(\omega) - Z(\omega)| < \epsilon) > 1 - \delta.$$

Corresponding to this convergence concept, we have the *Weak Law of Large Numbers* (WLLN), which is the result that $\bar{X}_n \xrightarrow{p} \mu$.

Earlier, we had used Chebyshev's inequality to prove a version of the WLLN, under the assumption that X_i are i.i.d. with mean μ and variance σ^2 .

There is a variant by Khinchine which only requires that the mean exists (i.e., is finite).

■■■

More general definition: the above definition of convergence in probability applies to random variables taking values on the real line. In general, for random variables mapping $\omega \in \Omega$ to some metric space \mathbb{M} with distance function $\rho(\cdot, \cdot)$, convergence in probability says that

$$Z_n \xrightarrow{p} Z \Leftrightarrow \forall \epsilon > 0, \lim_{n \rightarrow \infty} P(\omega : \rho(Z_n(\omega), Z(\omega)) < \epsilon) = 1.$$

■■■

Some definitions and results:

- Define: the *probability limit* of a random sequence Z_n , denoted $\text{plim}_{n \rightarrow \infty} Z_n$, is a non-random quantity α such that $Z_n \xrightarrow{p} \alpha$.
- Stochastic orders of magnitude:
 - “big-O.p.”: $Z_n = O_p(n^\lambda) \Leftrightarrow$ there exists an $O(1)$ nonstochastic sequence a_n such that $\frac{Z_n}{n^\lambda} - a_n \xrightarrow{p} 0$.
 $Z_n = O_p(1) \Leftrightarrow$ for every δ , there exists a finite $\Delta(\delta)$ and $n^*(\delta)$ such that $\text{Prob}(|Z_n| > \Delta) < \delta$ for all $n \geq n^*(\delta)$.
 - “little-o.p.”: $Z_n = o_p(n^\lambda) \Leftrightarrow \frac{Z_n}{n^\lambda} \xrightarrow{p} 0$.
 $Z_n = o_p(1) \Leftrightarrow Z_n \xrightarrow{p} 0$.
- **Plim operator theorem:** Let Z_n be a k -dimensional random vector, and $g(\cdot)$ be a function which is continuous at a constant k -vector point α . Then

$$Z_n \xrightarrow{p} \alpha \implies g(Z_n) \xrightarrow{p} g(\alpha).$$

In other words: $g(\text{plim } Z_n) = \text{plim } g(Z_n)$.

Don't confuse this: $\text{plim } \frac{1}{n} \sum_i g(Z_i) = E g(Z_i)$, from the LLN, but this is distinct from $\text{plim } g(\frac{1}{n} \sum_i Z_i) = g(E Z_i)$, using plim operator result. The two are generally not the same; if $g(\cdot)$ is convex, then $\text{plim } \frac{1}{n} \sum_i g(Z_i) \geq \text{plim } g(\frac{1}{n} \sum_i Z_i)$.

- A useful intermediate result (“squeezing”): if $Z_n \xrightarrow{p} \alpha$ (a constant) and Z_n^* lies between Z_n and α with probability 1, then $Z_n^* \xrightarrow{p} \alpha$.

(Quick proof: for all $\epsilon > 0$, $|X_n - \alpha| < \epsilon \Rightarrow |X_n^* - \alpha| < \epsilon$. Hence $Prob(|X_n^* - \alpha| < \epsilon) \geq Prob(|X_n - \alpha| < \epsilon) \rightarrow 1$.)



Convergence almost surely

$$Z_n \xrightarrow{as} Z \iff Prob(\omega : \lim_{n \rightarrow \infty} Z_n(\omega) = Z(\omega)) = 1.$$

As with convergence in probability, almost sure convergence makes sense when the random variables Z_1, \dots, Z_n, \dots and Z are all defined on the same sample space. Hence, for each element in the sample space $\omega \in \Omega$, we can associate the sequence $Z_1(\omega), \dots, Z_n(\omega), \dots$ as well as the limit point $Z(\omega)$.

To understand the probability statement more clearly, assume that all these RVs are defined on the sample space $(\Omega, \mathcal{B}(\Omega), P)$.

Define some set-theoretic concepts. Consider a sequence of sets S_1, \dots, S_n , all $\in \mathcal{B}(\Omega)$. Unless this sequence is monotonic, it’s difficult to talk about a “limit” of this sequence. So we define the “liminf” and “limsup” of this set sequence.

-

$$\begin{aligned} \liminf_{n \rightarrow \infty} S_n &\equiv \bigcup_{n=1}^{\infty} \bigcap_{m=n}^{\infty} S_m \\ &= \{\omega \in \Omega : \omega \in S_n, \forall n \geq n_0(\omega)\}. \end{aligned}$$

Note that the sequence of sets $\bigcap_{m=n}^{\infty} S_m$, for $n = 1, 2, 3, \dots$ is an increasing sequence of sets. By taking the union, the liminf is this the limit of this increasing sequence of sets: this we can also write

$$\liminf S_n = \lim_{n \rightarrow \infty} \bigcap_{m=n}^{\infty} S_m.$$

That is, for all $\omega \in \liminf S_n$, there exists some number $n_0(\omega)$ such that $\omega \in S_n$ for all $n \geq n_0(\omega)$. This is the of ω ’s which belong in all the S_n for n large enough: hence we can say that $\liminf S_n$ is the set of outcomes which occur “eventually”.

•

$$\begin{aligned} \limsup_{n \rightarrow \infty} S_n &\equiv \bigcap_{n=1}^{\infty} \bigcup_{m=n}^{\infty} S_m = \lim_{n \rightarrow \infty} \bigcup_{m=n}^{\infty} S_m \\ &= \{ \omega \in \Omega : \text{for every } m, \omega \in S_{n_0(\omega, m)} \text{ for some } n_0(\omega, m) \geq m \}. \end{aligned}$$

Note that the sequence of sets $\bigcup_{m=n}^{\infty} S_m$, for $n = 1, 2, 3, \dots$ is a decreasing sequence of sets. Hence, limsup is limit of this decreasing sequence of sets.

$\limsup S_n$ is the set of outcomes ω which occurs within every tails of the set sequence S_n . Hence we say that an outcome $\omega \in \limsup S_n$ occurs “infinitely often”.¹

Note that

$$\liminf S_n \subset \limsup S_n.$$

Hence

$$P(\limsup S_n) \geq P(\liminf S_n)$$

and

$$P(\limsup S_n) = 0 \rightarrow P(\liminf S_n) = 0.$$

- **Borel-Cantelli Lemma:** if $\sum_{i=1}^{\infty} P(S_i) < \infty$ then $P(\limsup_{n \rightarrow \infty} S_n) = 0$.

Quick proof: for each m , we have

$$\begin{aligned} P(\limsup_{n \rightarrow \infty} S_n) &= \lim_{n \rightarrow \infty} P\left(\bigcup_{m=n}^{\infty} S_m\right) \\ &\leq \lim_{n \rightarrow \infty} \sum_{m \geq n} P(S_m) \end{aligned}$$

which equals zero (by the assumption that $\sum_{i=1}^{\infty} P(S_i) < \infty$).

The first equality (interchange of “limit” and “Prob” operations) is an application of the Monotone Convergence Theorem, which holds only because the sequence $\bigcup_{m=n}^{\infty} S_m$ is monotone.

To apply this to almost-sure convergence, consider the sequence of sets is

$$S_n \equiv \{ \omega : |Z_n(\omega) - Z(\omega)| > \epsilon \},$$

¹However, note that all outcomes $\omega \in \liminf S_n$ also occur an infinite number of times along the infinite sequence S_n .

for $\epsilon > 0$. Almost sure convergence is the statement that $P(\limsup S_n) = 0$.

Then $\bigcup_{m=n}^{\infty} S_m$ denotes the ω 's such that the sequence $|Z_n(\omega) - Z(\omega)|$ exceeds ϵ in the tail beyond n .

The $\limsup \lim_{n \rightarrow \infty} \bigcup_{m=n}^{\infty} S_m$ denotes, then, all ω 's such that the sequence $|Z_n(\omega) - Z(\omega)|$ exceeds ϵ in every tail. For these ω 's, $\lim_{n \rightarrow \infty} Z_n(\omega) \neq Z(\omega)$. For almost-sure convergence, you require the probability of this set to equal zero, i.e. $\Pr(\lim_{n \rightarrow \infty} \bigcup_{m=n}^{\infty} S_m) = 0$.

■■■

A usual way to show almost-sure convergence would be to show that $\sum_{n=1}^{\infty} P(S_n) < \infty$, and then use the Borel-Cantelli Lemma.

Example: $X_i \sim i.i.d. U[0, 1]$. Show that $X_{(1:n)} \equiv \min_{i=1, \dots, n} X_i \xrightarrow{as} 0$.

Take $S_n \equiv \{|X_{(1:n)}| > \epsilon\}$. For all $\epsilon > 0$, $P(|X_{(1:n)}| > \epsilon) = P(X_{(1:n)} > \epsilon) = P(X_i > \epsilon, i = 1, \dots, n) = (1 - \epsilon)^n$.

Hence, $\sum_{n=1}^{\infty} (1 - \epsilon)^n = 1/\epsilon < \infty$. So the conclusion follows by the BC Lemma.

■■■

Corresponding to this convergence concept, we have the *Strong Law of Large Numbers* (SLLN), which is the result that $\bar{X}_n \xrightarrow{as} \mu$. Theorem 5.5.9 in CB has a version which requires a finite mean and variance.

To relate the SLLN to the previous section, consider that the sample averages are random variables $\bar{X}_1(\omega), \bar{X}_2(\omega), \dots$ defined on the same probability space, say (Ω, \mathcal{B}, P) . Then each ω indexes a sequence a sequence $\bar{X}_1(\omega), \bar{X}_2(\omega), \bar{X}_3(\omega), \dots$

Then consider the set sequence

$$S_n = \{\omega : |\bar{X}_n(\omega) - \mu| > \epsilon\}.$$

SLLN is the statement that $P(\limsup S_n) = 0$.

Proof (sketch; Davidson, pg. 296): From the above discussion, we see that the two statements $\bar{X}_n \xrightarrow{as} \mu$ and $P(|\bar{X}_n| > \epsilon, i.o.) = 0$ for any $\epsilon > 0$ are equivalent. Therefore, one way of proving a SLLN is to verify that $\sum_{i=1}^{\infty} P(|\bar{X}_n| > \epsilon) < \infty$ for all $\epsilon > 0$, and then apply the Borel-Cantelli Lemma (cf. Section 5.8.4 in CB). Note that Chebyshev's inequality not good enough by itself; it tells us

$$P(|\bar{X}_n| > \epsilon) \leq \frac{\sigma^2}{n\epsilon^2}; \quad \text{but} \quad \sum_{n=1}^{\infty} \frac{1}{n\epsilon^2} = \infty.$$

However, consider the subsequence with indices n^2 : $\{2, 4, 9, 16, \dots\}$. Again using Chebyshev's inequality, we have

$$P(|\bar{X}_{n^2}| > \epsilon) \leq \frac{\sigma^2}{n^2 \epsilon^2}; \quad \text{and} \quad \sum_{n=1}^{\infty} \frac{1}{n^2 \epsilon^2} \approx \frac{1.64}{\epsilon^2} < \infty.$$

Hence, by BC lemma, we have that $\bar{X}_{n^2} \xrightarrow{a.s.} 0$. Next, examine the "omitted" terms from the sum $\sum_{i=1}^{\infty} P(|\bar{X}_n| > \epsilon)$. Define $D_{n^2} = \max_{n^2 \leq k < (n+1)^2} |\bar{X}_k - \bar{X}_{n^2}|$. Note that:

$$\bar{X}_k - \bar{X}_{n^2} = \left(\frac{n^2}{k} - 1\right) \bar{X}_{n^2} = \frac{1}{k} \sum_{t=n^2+1}^k X_t$$

and, by independence of X_i , the two terms on RHS are uncorrelated. Hence

$$\begin{aligned} \text{Var}(\bar{X}_k - \bar{X}_{n^2}) &= \left(1 - \frac{n^2}{k}\right)^2 \frac{\sigma^2}{n^2} + \frac{k - n^2}{k^2} \sigma^2 \\ &= \sigma^2 \left(\frac{1}{n^2} - \frac{1}{k}\right) \leq \sigma^2 \left(\frac{1}{n^2} - \frac{1}{(n+1)^2}\right). \end{aligned}$$

By Chebyshev's inequality, then, we have

$$P(D_{n^2} > \epsilon) \leq \frac{\sigma^2}{\epsilon} \left(\frac{1}{n^2} - \frac{1}{(n+1)^2}\right); \quad \text{so}$$

$$\sum_{n^2} P(D_{n^2} > \epsilon) \leq \frac{\sigma^2}{\epsilon^2} \sum_{n^2} \left(\frac{1}{n^2} - \frac{1}{(n+1)^2}\right) < \frac{\sigma^2}{\epsilon^2} \sum_n \left(\frac{1}{n^2} - \frac{1}{(n+1)^2}\right) = \frac{\sigma^2}{\epsilon^2}$$

implying, using BC Lemma, that $D_{n^2} \xrightarrow{a.s.} 0$.

Now, we have, for $n^2 \leq k < (n+1)^2$,

$$\begin{aligned} |\bar{X}_k| &= |\bar{X}_k - \bar{X}_{n^2} + \bar{X}_{n^2}| \\ &\leq |\bar{X}_k - \bar{X}_{n^2}| + |\bar{X}_{n^2}| \end{aligned}$$

By the above discussion, the RHS converges a.s. to 0; that is, almost surely, for every $\epsilon > 0$, there exists $(n^2)^*$ such that $RHS_{n^2} < \epsilon$ for $n^2 > (n^2)^*$. Consequently, for $k > (n^2)^*$, $|\bar{X}_k| < \epsilon$; that is, $\bar{X}_k \xrightarrow{a.s.} 0$. ■

■■■

Example: $X \sim U[0, 1]$, and define $Z_n \equiv \frac{X}{n}$, for $n = 1, 2, 3, \dots$ ²

²For instance, if $x = 1/2$, then the Z sequence is $1/2, 1/4, 1/6, \dots$

Show that $Z_n \xrightarrow{p} 0$.

For all $\epsilon > 0$:

$$\begin{aligned} \text{Prob}(|Z_n| > \epsilon) &= \text{Prob}(|X/n| > \epsilon) \\ &= \text{Prob}(X/n > \epsilon) \text{ because } X \text{ non-negative} \\ &= 1 - \min(1, n\epsilon) \rightarrow 1 - 1 = 0 \text{ (as } n \rightarrow \infty). \end{aligned}$$

What about $Z_n \xrightarrow{as} 0$?

For all $x \in [0, 1]$, $\lim_{n \rightarrow \infty} x/n = 0$. Hence $\text{Prob}(X : \lim_{n \rightarrow \infty} X/n = 0) = 1$.

Does $Z_n \xrightarrow{as} 0$ if

$$Z_n = \begin{cases} 100 & \text{if } X = \frac{1}{2} \\ X/n & \text{otherwise} \end{cases} ?$$

Theorem: $Z_n \xrightarrow{as} Z \implies Z_n \xrightarrow{p} Z$.

■■■

More general definition of a.s. convergence, for random variables taking values on metric space:

$$Z_n \xrightarrow{as} Z \Leftrightarrow P(\omega : \rho(Z_n(\omega), Z(\omega)) \rightarrow 0) = 1.$$

■■■

Convergence in Distribution: $Z_n \xrightarrow{d} Z$.

A sequence of real-valued random variables Z_1, Z_2, Z_3, \dots converges in distribution to a random variable Z if

$$\lim_{n \rightarrow \infty} F_{Z_n}(z) = F_Z(z) \quad \forall z \text{ s.t. } F_Z(z) \text{ is continuous.} \quad (1)$$

This is a statement about the CDFs of the random variables Z , and Z_1, Z_2, \dots . These random variables do not need to be defined on the same probability space. F_Z is also called the “limiting distribution” of the random sequence Z_n .

■■■

Alternative definitions of convergence in distribution:

- **Corollary:** Letting $F([a, b]) \equiv F(b) - F(a)$ for $a < b$, convergence (1) is equivalent to

$$F_{Z_n}([a, b]) \rightarrow F_Z([a, b]) \quad \forall [a, b] \text{ s.t. } F_Z \text{ continuous at } a, b.$$

- Another equivalent definition of weak convergence is the following: for all bounded, continuous, and real-valued functions $g(\cdot)$,

$$\int g(z)dF_{Z_n}(z) \rightarrow \int g(z)dF_Z(z). \quad (2)$$

This definition of weak convergence is more useful in advanced settings, because it is extendable to setting where Z_n and Z are general random elements taking values in metric space.

- **Levy’s continuity theorem:** A sequence $\{X_n\}$ of random variables converges in distribution to random variable X if and only if the sequence of characteristic function $\{\phi_{X_j}(t)\}$ converges pointwise (in t) to a function $\phi_X(t)$ which is continuous at the origin. Then ϕ_X is the characteristic function of X .

Hence, this ties together convergence of characteristic functions, and convergence in distribution. This theorem will be used to prove the CLT later.



Distributional convergence, as defined above, is usually called “weak convergence”. This is because there are stronger notions of distributional convergence. One such notion is:

$$\sup_{A \in \mathbb{B}(\mathbb{R})} |P_{Z_n}(A) - P_Z(A)| \rightarrow 0$$

which is convergence in *total variation norm*.

Example: Multinomial distribution on $[0, 1]$. $X_n \in \{\frac{1}{n}, \frac{2}{n}, \frac{3}{n}, \dots, \frac{n}{n} = 1\}$ each with probability $\frac{1}{n}$. Consider $A = \mathbb{Q}$ (set of rational numbers in $[0, 1]$).



Some definitions and results:

- **Slutsky Theorem:** If $Z_n \xrightarrow{d} Z$, and $Y_n \xrightarrow{p} \alpha$ (a constant), then
 - $Y_n Z_n \xrightarrow{d} \alpha Z$
 - $Z_n + Y + n \xrightarrow{d} Z + \alpha$.
- **Theorem *:**

- $Z_n \xrightarrow{p} Z \implies Z_n \xrightarrow{d} Z$
- $Z_n \xrightarrow{d} Z \implies Z_n \xrightarrow{p} Z$ if Z is a constant.

Note: convergence in probability implies that (roughly speaking), the random variables Z_n (for n large enough) and Z frequently have the same numerical value. Convergence in distribution need not imply this, only that the CDF's of Z_n and Z are similar.

Example: Let Y_1, Y_2, \dots denote independent tosses of a fair coin, where $Y_i = 1$ if the outcome of the toss is heads, and $Y_i = 0$ otherwise. Let $S \sim U[0, 1]$.

Consider the sequence of random variables $Z_n = \mathbf{1}(\frac{S}{n} < \epsilon) * (1 - Y_n) + \mathbf{1}(\frac{S}{n} \geq \epsilon) * Y_n$, for some fixed $\epsilon > 0$.

In this case, Z_1, \dots, Z_n, \dots and Y_1, \dots, Y_n, \dots are all defined on the same probability space. Obviously, $Z_n \xrightarrow{d} Y_1$, but $Z_n \not\xrightarrow{p} Y_1$ (because Z_n and Y_1 yield numerically different values for the two possible outcomes H and T). Note that $Z_n \xrightarrow{p, a.s.} 1 - Y_1$.



[can skip] Combining two results above, we have

$$Z_n \xrightarrow{a.s.} Z \implies Z_n \xrightarrow{d} Z.$$

The converse is not generally true. (Indeed, $\{Z_1, Z_2, Z_3, \dots\}$ need not be defined on the same probability space, in which case s.d. convergence makes no sense.)

However, consider

$$Z_n^* = F_{Z_n}^{-1}(U), \quad Z^* = F_Z^{-1}(U); \quad U \sim U[0, 1].$$

Here $F_Z^{-1}(U)$ denotes the quantile function corresponding to the CDF $F(z)$:

$$F^{-1}(\tau) = \inf \{z : F(z) > \tau\} \quad \tau \in [0, 1].$$

We have $F(F^{-1}(\tau)) = \tau$. (Note that quantile function is also right-continuous; discontinuity points of the quantile function arise where the CDF function is “flat”.)

Then

$$P(Z_n^* \leq z) = P(F_{Z_n}^{-1}(U) \leq z) = P(U \leq F_{Z_n}(z)) = F_{Z_n}(z)$$

so that $Z_n^* \stackrel{d}{=} Z_n$. (Even though their domains are different!) Similarly, $Z^* \stackrel{d}{=} Z$. The notation “ $\stackrel{d}{=}$ ” means “identically distributed”.

Moreover, it turns out (quite intuitive) that the convergence $F_{Z_n}^{-1}(U) \rightarrow F_Z^{-1}(U)$ fails only at points where $F_Z^{-1}(U)$ is discontinuous (corresponding to flat portions

of $F_Z(z)$). Since these points of discontinuity are a countable set, their probability (under $U[0, 1]$) is equal to zero, so that $F_{Z_n}^{-1}(U) \rightarrow F_Z^{-1}(U)$ for almost-all U .

So what we have here, is a result that for real-valued random variables $Z_n \xrightarrow{d} Z$, we can construct identically distributed variables such that both $Z_n^* \xrightarrow{d} Z^*$ and $Z_n^* \xrightarrow{as} Z^*$.

This construction is called the “Skorokhod representation theorem”. More generally, it holds for random elements defined on arbitrary metric spaces.

■■■

Note that the LLN tells us that $\bar{X}_n \xrightarrow{p,as} \mu$, which implies (trivially) that $\bar{X}_n \xrightarrow{d} \mu$, a *degenerate* limiting distribution.

This is not very useful for our purposes, because we are interested in knowing (say) how far \bar{X}_n is from μ , which is unknown.

How do we “fix” \bar{X}_n , so that it has a non-degenerate limiting distribution?

Central Limit Theorem: statement that, under various assumptions,

$$\sqrt{n}(\bar{X}_n - \mu) \xrightarrow{d} N(0, \cdot).$$

Lindeberg-Levy CLT: Let X_i be i.i.d. with mean μ and variance σ^2 . Then

$$\frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \xrightarrow{d} N(0, 1).$$

\sqrt{n} is also called the “rate of convergence” of the sequence \bar{X}_n . By definition, a rate of convergence is the lowest polynomial in n for which $n^p \bar{X}_n$ converges in distribution to a nondegenerate distribution.

The “blowing up” factor of \sqrt{n} make sense: if you blow up \bar{X}_n by a constant (no matter how big), you still get a degenerate limiting distribution. If you blow up by n , then the sequence $S_n \equiv n\bar{X}_n = \sum_{i=1}^n X_n$ will diverge.

Proof via characteristic functions.

■■■

Asymptotic approximations for \bar{X}_n

CLT tells us that $\sqrt{n}(\bar{X}_n - \mu)/\sigma$ has a limiting standard normal distribution. We can use this “true” result to say something about the distribution of \bar{X}_n , even when n is finite. That is, we use the CLT to derive an *asymptotic approximation* for the finite-sample distribution of \bar{X}_n .

The approximation is as follows: we “flip over” the result of the CLT:

$$\bar{X}_n \stackrel{a}{\sim} \frac{\sigma}{\sqrt{n}}N(0, 1) + \mu \Rightarrow \bar{X}_n \stackrel{a}{\sim} N\left(\mu, \frac{1}{n}\sigma^2\right).$$

The notation “ $\stackrel{a}{\sim}$ ” makes explicit that what is on the RHS is an approximation. Note that $\bar{X}_n \xrightarrow{d} N\left(\mu, \frac{1}{n}\sigma^2\right)$ is definitely not true!

This approximation intuitively makes sense: under the assumptions of the LLCLT, we know that $E\bar{X}_n = \mu$ and $Var(\bar{X}_n) = \sigma^2/n$. What the asymptotic approximation tells us is that the distribution of \bar{X}_n is approximately normal.³

Example: Bernoulli random variables. $X_i = B(1, p)$, $i = 1, 2, 3, \dots$. Under the LLCLT, we know that

$$\frac{\sqrt{n}(\bar{X}_n - p)}{\sqrt{p(1-p)}} \xrightarrow{d} N(0, 1)$$

so approximately

$$\bar{X}_n \stackrel{a}{\sim} N\left(p, \frac{1}{n}p(1-p)\right).$$

For a fixed n , we can describe the exact distribution of \bar{X}_n . Note that \bar{X}_n is just n times a $B(n, p)$ binomial random variable, so that \bar{X}_n is a discrete random variable with support $\{0, \frac{1}{n}, \frac{2}{n}, \dots, 1\}$ with probabilities

$$P(\bar{X}_n = k/n) = \binom{n}{k} p^k (1-p)^{n-k}.$$

This is much more complicated than $N\left(p, \frac{1}{n}p(1-p)\right)$!



Asymptotic approximations for functions of \bar{X}_n

Oftentimes, we are not interested *per se* in approximating the finite-sample distribution of the sample mean \bar{X}_n , but rather functions of a sample mean. (Later, you will see that the asymptotic approximations for many statistics and estimators that you run across are derived by expressing them as sample averages.)

Continuous mapping theorem: Let $g(\cdot)$ be a continuous function. Then

$$Z_n \xrightarrow{d} Z \implies g(Z_n) \xrightarrow{d} g(Z).$$

³Note that the approximation is exactly right if we assumed that $X_i \sim N(\mu, \sigma^2)$, for all i .

(Note: you still have to figure out what the limiting distribution of $g(Z)$ is. But if you know F_X , then you can get $F_{g(X)}$ by the change of variables formulas.)

Note that for any linear function $g(\bar{X}_n) = a\bar{X}_n + b$, deriving the limiting distribution of $g(\bar{X}_n)$ is no problem (just use Slutsky's Theorem to get $a\bar{X}_n + b \stackrel{a}{\sim} N(a\mu + b, \frac{1}{n}a^2\sigma^2)$).

The “problem” in deriving the distribution of $g(\bar{X}_n)$ is that the \bar{X}_n is “inside” the g function:

1. Use the **Mean Value Theorem**: Let g be a continuous function on $[a, b]$ that is differentiable on (a, b) . Then, there exists (at least one) $\lambda \in (a, b)$ such that

$$g'(\lambda) = \frac{g(b) - g(a)}{b - a} \Leftrightarrow g(b) - g(a) = g'(\lambda)(b - a).$$

Using the MVT, we can write

$$g(\bar{X}_n) - g(\mu) = g'(X_n^*)(\bar{X}_n - \mu) \quad (3)$$

where X_n^* is an RV strictly between \bar{X}_n and μ .

2. On the RHS of Eq. (3):

- (a) $g'(X_n^*) \xrightarrow{p} g'(\mu)$ by the “squeezing” result, and the plim operator theorem.
- (b) If we multiply by \sqrt{n} and divide by σ , we can apply the CLT to get $\sqrt{n}(\bar{X}_n - \mu)/\sigma \xrightarrow{d} N(0, 1)$.
- (c) Hence,

$$\sqrt{n} * \text{RHS} = \sqrt{n}[g(\bar{X}_n) - g(\mu)] \xrightarrow{d} N(0, g'(\mu)^2\sigma^2), \quad (4)$$

using Slutsky's theorem.

3. Now, in order to get the asymptotic approximation for the distribution of $g(\bar{X}_n)$, we “flip over” to get

$$g(\bar{X}_n) \stackrel{a}{\sim} N(g(\mu), \frac{1}{n}g'(\mu)^2\sigma^2).$$

4. Check that g satisfies assumptions for these results to go through: continuity and differentiability for MVT, and continuous at μ for plim operator theorem.

Examples: $1/\bar{X}_n$, $\exp(\bar{X}_n)$, $(\bar{X}_n)^2$, etc.

Of course, Eq. (4) is a general result, known as the **Delta method**. For the purposes of this class, I want you to derive the approximate distributions for $g(\bar{X}_n)$ from first principles (as we did above).



Functional Central limit theorems

[**can skip**] There are a set of distributional convergence results, which are known as functional CLT's (or Donsker theorems), because they deal with convergence of random *functions* (or, interchangeably, *processes*). These are indispensable tools in finance.

One of the simplest random functions is the **Wiener process** $\mathcal{W}(t)$, which is viewed as a random function on the unit interval $t \in [0, 1]$. This is also known as a Brownian motion process.

Features of the Wiener process:

1. $\mathcal{W}(0) = 0$
2. Gaussian marginals: $\mathcal{W}(t) \stackrel{d}{=} N(0, t)$; that is,

$$P(\mathcal{W}(t) \leq a) = \frac{1}{\sqrt{2\pi t}} \int_{-\infty}^a \exp\left(\frac{-u^2}{2t}\right) du.$$

3. Independent increments: Define $x_t \equiv \mathcal{W}(t)$. For any set $0 \leq t_0 \leq t_1 \leq \dots \leq 1$, the differences

$$x_{t_1} - x_{t_0}, x_{t_2} - x_{t_1}, \dots$$

are independent.

4. Given the two above features, we have that the increments are themselves normally distributed:

$$x_{t_i} - x_{t_{i-1}} \stackrel{d}{=} N(0, t_i - t_{i-1}).$$

Moreover, from

$$\begin{aligned} t_1 - t_0 = V(x_1 - x_0) &= E[(x_1 - x_0)^2] \\ &= E x_1^2 - 2E x_1 x_0 + E x_0^2 \\ &= t_1 + t_0 - 2E x_1 x_0 \end{aligned}$$

implying

$$E x_1 x_0 = Cov(x_1, x_0) = t_0.$$

5. Furthermore, we know that any finite collection $(x_{t_1}, x_{t_2}, \dots)$ is jointly distributed multivariate normal, with mean 0 and variance matrix Σ as described above.

Take the conditions of the LLCLT: X_1, X_2, \dots are iid with mean zero and finite variance σ^2 . For a given n , define the “partial sum” $S_k = X_1 + X_2 + \dots + X_k$ for $k \leq n$. Now, for $t \in [0, 1]$, define the normalized “partial sum process”

$$\tilde{S}_n(t) = \frac{S_{[tn]}}{\sigma\sqrt{n}} + (tn - [tn]) * \frac{1}{\sigma\sqrt{n}} X_{[tn]+1}$$

where $[a]$ denotes the largest integer $\leq a$. $\tilde{S}_n(t)$ is a random function on $[0, 1]$. (Graph)

Then the functional CLT is the following result:

Functional CLT: $\tilde{S}_n \xrightarrow{d} \mathcal{W}$.

Note that \tilde{S}_n is a random element taking values in $C[0, 1]$, the space of continuous functions on $[0, 1]$. Proving this result is beyond our focus in this class. \blacksquare

Note that the functional CLT implies more than the LLCLT. For example, take $t = 1$. Then $\tilde{S}_n(1) = \frac{\sum_{i=1}^n X_i}{\sigma\sqrt{n}} = \frac{\sqrt{n}\bar{X}_n}{\sigma}$. By the FCLT, we have

$$\begin{aligned} P(\tilde{S}_n(1) \leq a) &\rightarrow P(\mathcal{W}(1) \leq a) \\ &= P(N(0, 1) \leq a) \end{aligned}$$

which is the same result as the LLCLT.

Similarly, for $k_n < n$ but $\frac{k_n}{n} \rightarrow \tau \in [0, 1]$, we have

$$\tilde{S}_n\left(\frac{k}{n}\right) = \frac{\sum_{i=1}^k X_i}{\sigma\sqrt{n}} = \frac{k \cdot \bar{X}_k}{\sigma\sqrt{n}} \xrightarrow{d} N(0, \tau).$$

Also, it turns out, by a functional version of the continuous mapping theorem, we can get convergence results for functionals of the partial-sum process. For instance, $\sup_t \tilde{S}_n(t) \xrightarrow{d} \sup_t \mathcal{W}(t)$, and it turns out

$$P(\sup_t \mathcal{W}(t) \leq a) = \frac{2}{\sqrt{2\pi}} \int_0^a \exp(-u^2/2) du; \quad a \geq 0.$$

(Recall that $\mathcal{W}_0 = 0$, so $a < 0$ makes no sense.)

⁴See, for instance, Billingsley, *Convergence of Probability Measures*, ch. 2, section 10.