

This is a draft; please send me corrections and/or suggestions.

Convergence of the sample mean

Suppose that $\langle X_i \rangle_{i=1}^N$ are a sequence of i.i.d. random variables¹ representing the outcomes from some experiment repeated N times. We have discussed properties of the sample mean,

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X_i.$$

In particular, $\mathbb{E}[\bar{X}] = \mathbb{E}[X]$, and $\text{Var}(\bar{X}) = \frac{1}{N} \text{Var}(X)$. A related result — the weak law of large numbers (WLOLN) — says that as N grows large, $\mathbb{P}(|\bar{X} - \mathbb{E}[X]| > \varepsilon) \rightarrow 0$ for all $\varepsilon > 0$. Taken together, these results say that the more observations we make, the closer the sample mean should be to the true mean. We will extend this logic slightly for an example.

Corollary: *let \bar{X}_N and \bar{X}_{N+1} be the sample means of $\langle X_i \rangle_{i=1}^N$ and $\langle X_i \rangle_{i=1}^{N+1}$, respectively. Then as $N \rightarrow +\infty$,*

$$\text{Var}(\bar{X}_{N+1} - \bar{X}_N) \rightarrow 0.$$

Proof: this corollary says that not only are sample means coming arbitrarily close to the population mean, they are also tending arbitrarily close to one another (or rather, arbitrarily close to the one following). This is fairly intuitive: if things are converging to a particular point, then the distance between two consecutive observations should be shrinking. A formal proof is a matter of applying what we know about independent random variables.

$$\begin{aligned} \text{Var}(\bar{X}_{N+1} - \bar{X}_N) &= \text{Var}\left(\frac{1}{N+1} \sum_{i=1}^{N+1} X_i - \frac{1}{N} \sum_{i=1}^N X_i\right) \\ &= \text{Var}\left(\frac{N}{N(N+1)} \sum_{i=1}^N X_i - \frac{N+1}{N(N+1)} \sum_{i=1}^N X_i + \frac{1}{N+1} X_{N+1}\right) \\ &= \text{Var}\left(\frac{1}{N+1} X_{N+1} - \frac{1}{N(N+1)} \sum_{i=1}^N X_i\right) \\ &= \text{Var}\left(\frac{1}{N+1} X_{N+1}\right) + \sum_{i=1}^N \text{Var}\left(-\frac{1}{N(N+1)} X_i\right) \\ &= \left(\frac{1}{N+1}\right)^2 \text{Var}(X_{N+1}) + \sum_{i=1}^N \left(\frac{1}{N(N+1)}\right)^2 \text{Var}(X_i) \\ &= \frac{\text{Var}(X)}{(N+1)^2} + \frac{N \text{Var}(X)}{N^2(N+1)^2} \\ &= \left(\frac{1}{N(N+1)}\right) \text{Var}(X). \end{aligned}$$

Since $\text{Var}(X)$ is fixed, as $N \rightarrow +\infty$ the right-hand multiple will go to 0, so the variance of the difference will go to 0. □

¹When the sequence $\langle X_i \rangle_{i=1}^N$ is i.i.d., we generally drop the subscript and think of each being distributed identically to some other random variable X . This allows us to avoid having to justify “which X_i we are talking about.”

An interesting consequence here is that not only does the variance of the difference go to 0, but it does so *quickly*, on order N^2 . That means that with 10 observations, the variance of the difference of sample means will be only 1% of the variance of the underlying random variable! This is because not only is the variance of an individual sample mean decreasing, but also because an extra observation will have so little effect on the sample mean, since it is downweighted by $\frac{1}{N+1}$.

In section, I sampled the class to determine how long people sleep on a normal school night². We collected the following data:

	X_1	X_2	X_3	X_4	X_5	X_6	X_7	X_8	X_9	X_{10}
X_i	7	7	4	6	5	6	8	7	5	5
\bar{X}_N	7	7	6	6	$\frac{29}{5}$	$\frac{35}{6}$	$\frac{43}{7}$	$\frac{50}{8}$	$\frac{55}{9}$	6
	7	7	6	6	5.8	5.833	6.143	6.25	6.111	6
$\bar{X}_N - \bar{X}_{N-1}$	-	0	-1	0	-0.2	0.033	0.310	0.107	-0.139	-0.111

We can read what we want into this data — this speaks to the question of when, precisely, is N large? — but we can see that with a low number of observations the difference in sample means is swinging wildly about, while with a larger number the difference in sample means is coming closer and closer to 0.

If people send me the data points from the 1:00pm section, I will fill in a section on Chebyshev's inequality here.

Estimation potpourri

You have two [possibly-biased] coins. The first lands heads-up with probability p_1 , while the second lands heads-up with probability p_2 . You run the following experiment: flip the first coin; if it lands heads-up, flip the second coin. Let H_1 and H_2 be two random variables; $H_1 = 1$ when the first coin lands heads-up and $H_1 = 0$ otherwise; $H_2 = 1$ when the second coin lands heads-up and $H_2 = 0$ otherwise (including if the second coin is not flipped).

- What is the joint PMF of H_1 and H_2 ?
- What are the marginal PMFs of H_1 and H_2 ?
- What is the correlation between H_1 and H_2 , $\rho_{H_1 H_2}$?
- Let X represent the number of heads obtained in the experiment. What is the PMF of X ?
- You compute X_i over 10 experiments; you find

X_1	X_2	X_3	X_4	X_5	X_6	X_7	X_8	X_9	X_{10}
2	0	0	1	0	1	2	1	0	1

What is the maximum-likelihood estimate for (p_1, p_2) ?

- Suppose $p_1 = \frac{3}{5}$ and $p_2 = \frac{1}{3}$. Use facts about linear combinations of random variables (H_1 and H_2) to compute $\mathbb{E}[X]$ and $\text{Var}(X)$.
- You run this experiment 7 times. Use Chebyshev's inequality to place an upper bound on $P(\bar{X} \leq \frac{1}{5})$.


²Several people cheated and tried to give ranges; tough cookies.

Continuous random variables


Up until this point we have been dealing with discrete random variables; these are good for addressing whether or not something happens, what kind of thing happens, how many times it happens, etc. This concept will only capture so many features of the world. Consider, for example, the number of gallons of water you use daily, as an Angeleno: it could be 100 gallons or it could be 110 gallons. It could also be 105.1 gallons 105.11 gallons, 105.111 gallons, etc. The crux is this: the amount of water you use could be any among a continuous set of values; a discrete random variable cannot capture this³!

There is a slight hiccup, though. With discrete random variables, any outcome may happen with a strictly positive probability (hence, *probability mass*). But with the infinite outcomes a continuous random variable may take⁴, if each occurred with positive probability we would certainly have an overall probability greater than 1; an [uncountably] infinite quantity of positive numbers must sum to infinity. So instead of speaking of probability mass, we speak of *probability density*.

The analogy that I like to use is this: imagine a brick. This brick has a definite mass. If we take a small



portion of this brick, it also has mass. However, once we take an *infinitesimally small* portion of the brick,



it has no mass. That is, since it has no size, it contains nothing and therefore has no mass. However, it still

³Although it could say you use between 108 and 109 gallons of water, between 109 and 110 gallons of water, etc.

⁴The math is slightly more nuanced than this, but this is good enough for government work.

has a density! In fact, to determine the overall mass of the brick, we can integrate over the densities of each infinitesimal slice. That is, since density is mass per unit volume, if we integrate over the entire volume we will get the mass back.

A continuous random variable functions in a similar way: each individual outcome (analogous to the infinitesimal slice of the brick) has zero probability mass, but it still has a probability density. When we integrate over the probability density, we get back the probability mass that we are used to dealing with. This motivates the fact that rather than probability mass functions (PMFs), we now deal with *probability density functions*, or PDFs.

If the PDF of a random variable X is f_X , we have the following:

$$P(a \leq X \leq b) = \int_a^b f_X(x) dx.$$

PDFs have properties that more or less align with the properties we learned about PMFs:

- $f_X(x) \geq 0$;
- $\int_{\text{Support}(X)} f_X(x) dx = 1$.

In particular, many of the things we did before with PMFs may be done with PDFs, substituting the summation (\sum) with integration (\int).

Question: the PDF of X is $f_X(x) = ax$, and X has support $[0, 1]$. What is a ?

Question: the PDF of X is $f_X(x) = a(k+x)(k-x)$, ($a > 0$) and X has support $[-B, B]$.

- (a) What values can B take?
- (b) What is a as a function of B and k ?

Henceforth, assume $B = k$.

- (c) What is a ?
- (d) What is $P(X \leq \frac{k}{2})$?
- (e) What is $P(-\frac{k}{2} \leq X \leq \frac{k}{2})$?