


the law of large numbers & the CLT



$$\Pr \left(\lim_{n \rightarrow \infty} \left(\frac{X_1 + \cdots + X_n}{n} = \mu \right) \right) = 1$$


If X, Y are independent, what is the distribution of $Z = X + Y$?

Discrete case:

$$p_Z(z) = \sum_x p_X(x) \cdot p_Y(z-x)$$

Continuous case:

$$f_Z(z) = \int_{-\infty}^{+\infty} f_X(x) \cdot f_Y(z-x) dx$$





E.g. what is the p.d.f. of the sum of 2 normal RV's?

$W = X + Y + Z$? Similar, but double sums/integrals

$V = W + X + Y + Z$? Similar, but triple sums/integrals

If X and Y are *uniform*, then $Z = X + Y$ is *not*; it's *triangular* (like dice):



Intuition: $X + Y \approx 0$ or ≈ 1 is rare, but many ways to get $X + Y \approx 0.5$

moment generating functions

aka transforms; b&t 229

Powerful math tricks for dealing with distributions

We won't do much with it, but mentioned/used in book, so a very brief introduction:

The k^{th} moment of r.v. X is $E[X^k]$; M.G.F. is $M(t) = E[e^{tX}]$

$$e^{tX} = X^0 \frac{t^0}{0!} + X^1 \frac{t^1}{1!} + X^2 \frac{t^2}{2!} + X^3 \frac{t^3}{3!} + \dots$$

$$M(t) = E[e^{tX}] = E[X^0] \frac{t^0}{0!} + E[X^1] \frac{t^1}{1!} + E[X^2] \frac{t^2}{2!} + E[X^3] \frac{t^3}{3!} + \dots$$

$$\frac{d}{dt} M(t) = 0 + E[X^1] + E[X^2] \frac{t^1}{1!} + E[X^3] \frac{t^2}{2!} + \dots$$

$$\frac{d^2}{dt^2} M(t) = 0 + 0 + E[X^2] + E[X^3] \frac{t^1}{1!} + \dots$$

$$\left. \frac{d}{dt} M(t) \right|_{t=0} = E[X]$$

$$\left. \frac{d^2}{dt^2} M(t) \right|_{t=0} = E[X^2]$$

$$\dots \left. \frac{d^k}{dt^k} M(t) \right|_{t=0} = E[X^k] \dots$$

Closely related to Laplace transforms, which you may have seen.

An example:

MGF of normal(μ, σ^2) is $\exp(\mu t + \sigma^2 t^2 / 2)$

Two key properties:

1. MGF of *sum* independent r.v.s is *product* of MGFs:

$$M_{X+Y}(t) = E[e^{t(X+Y)}] = E[e^{tX} e^{tY}] = E[e^{tX}] E[e^{tY}] = M_X(t) M_Y(t)$$

2. Invertibility: MGF uniquely determines the distribution.

e.g.: $M_X(t) = \exp(at + bt^2)$, with $b > 0$, then $X \sim \text{Normal}(a, 2b)$

Important example: *sum of independent normals is normal:*

$$X \sim \text{Normal}(\mu_1, \sigma_1^2) \quad Y \sim \text{Normal}(\mu_2, \sigma_2^2)$$

$$\begin{aligned} M_{X+Y}(t) &= \exp(\mu_1 t + \sigma_1^2 t^2 / 2) \cdot \exp(\mu_2 t + \sigma_2^2 t^2 / 2) \\ &= \exp[(\mu_1 + \mu_2)t + (\sigma_1^2 + \sigma_2^2)t^2 / 2] \end{aligned}$$

So $X+Y$ has mean $(\mu_1 + \mu_2)$, variance $(\sigma_1^2 + \sigma_2^2)$ (duh) *and is normal!*
(way easier than slide 2 way!)

Consider i.i.d. (independent, identically distributed) R.V.s

$$X_1, X_2, X_3, \dots$$

Suppose X_i has $\mu = E[X_i] < \infty$ and $\sigma^2 = \text{Var}[X_i] < \infty$.

What are the mean & variance of their sum?

$$E\left[\sum_{i=1}^n X_i\right] = n\mu \text{ and } \text{Var}\left[\sum_{i=1}^n X_i\right] = n\sigma^2$$

So limit as $n \rightarrow \infty$ *does not exist* (except in the degenerate case where $\mu = 0$; note that if $\mu = 0$, the *center* of the data stays fixed, but if $\sigma^2 > 0$, then the variance is unbounded, i.e., its *spread* grows with n).

Consider i.i.d. (independent, identically distributed) R.V.s

$$X_1, X_2, X_3, \dots$$

Suppose X_i has $\mu = E[X_i] < \infty$ and $\sigma^2 = \text{Var}[X_i] < \infty$

What about the *sample mean* $M_n = \frac{1}{n} \sum_{i=1}^n X_i$, as $n \rightarrow \infty$?

$$E[M_n] = E\left[\frac{X_1 + \dots + X_n}{n}\right] = \mu$$

$$\text{Var}[M_n] = \text{Var}\left[\frac{X_1 + \dots + X_n}{n}\right] = \frac{\sigma^2}{n}$$

So, limits *do* exist; mean is independent of n , variance shrinks.

weak law of large numbers

Continuing: iid RVs X_1, X_2, X_3, \dots ; $\mu = E[X_i]$; $\sigma^2 = \text{Var}[X_i]$; $M_n = \frac{1}{n} \sum_{i=1}^n X_i$

$$E[M_n] = E\left[\frac{X_1 + \dots + X_n}{n}\right] = \mu \quad \text{Var}[M_n] = \text{Var}\left[\frac{X_1 + \dots + X_n}{n}\right] = \frac{\sigma^2}{n}$$

Expectation is an important guarantee.

BUT: observed values may be far from expected values.

E.g., if $X_i \sim \text{Bernoulli}(1/2)$, the $E[X_i] = 1/2$, but X_i is **NEVER** $1/2$.

Is it also possible that sample mean of X_i 's will be far from $1/2$?

Always? Usually? Sometimes? Never?

For any $\varepsilon > 0$, as $n \rightarrow \infty$

$$\Pr (|M_n - \mu| > \varepsilon) \rightarrow 0$$

Proof: (assume $\sigma^2 < \infty$; theorem true without that, but harder proof)

$$\mathbb{E} [M_n] = \mathbb{E} \left[\frac{X_1 + \cdots + X_n}{n} \right] = \mu$$

$$\text{Var} [M_n] = \text{Var} \left[\frac{X_1 + \cdots + X_n}{n} \right] = \frac{\sigma^2}{n}$$

By Chebyshev inequality,

$$\Pr (|M_n - \mu| > \varepsilon) \leq \frac{\sigma^2}{n\varepsilon^2} \xrightarrow{n \rightarrow \infty} 0$$

i.i.d. (independent, identically distributed) random vars

X_1, X_2, X_3, \dots

$$M_n = \frac{1}{n} \sum_{i=1}^n X_i$$

X_i has $\mu = E[X_i] < \infty$

$$\Pr \left(\lim_{n \rightarrow \infty} \left(\frac{X_1 + \dots + X_n}{n} \right) = \mu \right) = 1$$

Strong Law \Rightarrow Weak Law (but not vice versa)

Strong law implies that for any $\epsilon > 0$, there are only a finite number of n satisfying the weak law condition $|M_n - \mu| \geq \epsilon$ (almost surely, i.e., with probability 1)

Supports the intuition of probability as long term frequency

Weak Law:

$$\lim_{n \rightarrow \infty} \Pr \left(\left| \frac{X_1 + \cdots + X_n}{n} - \mu \right| > \epsilon \right) = 0$$

Strong Law:

$$\Pr \left(\lim_{n \rightarrow \infty} \left(\frac{X_1 + \cdots + X_n}{n} \right) = \mu \right) = 1$$

How do they differ? Imagine an infinite 2-D table, whose rows are indep infinite sample sequences X_i . Pick ϵ . Imagine cell m, n lights up if average of 1st n samples in row m is $> \epsilon$ away from μ .

WLLN says fraction of lights in n^{th} column goes to zero as $n \rightarrow \infty$. It does not prohibit every row from having ∞ lights, so long as frequency declines.

SLLN also says only a vanishingly small fraction of rows can have ∞ lights.

weak vs strong laws – supplement

The differences between the WLLN & SLLN are subtle, and not critically important for this course, but for students wanting to know more (e.g., not on exams), here is my summary. Both “laws” rigorously connect long-term averages of repeated, independent observations to mathematical expectation, justifying the intuitive motivation for $E[\cdot]$. Specifically, both say that the sequence of (non-i.i.d.) rvs $M_n = \sum_{i=1}^n X_i/n$ derived from any sequence of i.i.d. rvs X_i converge to $E[X_i]=\mu$. The strong law totally subsumes the weak law, but the latter remains interesting because (a) of its simple proof (Khinchine, early 20th century; using Chebyshev’s inequality (1867)) and (b) historically (WLLN was proved by Bernoulli ~1705, for Bernoulli rvs, >150 years before Chebyshev [Ross, p391]). The technical difference between WLLN and SLLN is in the definition of convergence.

Definition: Let Y_i be any sequence of rvs (i.i.d. not assumed) and c a constant.

Y_i converges to c in probability if

$$\forall \epsilon > 0, \lim_{n \rightarrow \infty} \Pr (|Y_n - c| > \epsilon) = 0$$

Y_i converges to c with probability 1 if

$$\Pr (\lim_{n \rightarrow \infty} Y_n = c) = 1$$

The weak law is the statement that M_n converges *in probability* to μ ; the strong law states it converges *with probability 1* to μ . The strong law subsumes the weak law since convergence with probability 1 implies convergence in probability for any sequence Y_i of rvs (B&T problem 5.5-15). B&T ex 5.15 illustrates the failure of the converse. A second counterexample is given on the following slide.

weak vs strong laws – supplement

Example: Consider the sequence of rvs $Y_n \sim \text{Ber}(1/n)$

Recall the definition of convergence in probability:

$$\forall \epsilon > 0, \lim_{n \rightarrow \infty} \Pr(|Y_n - c| > \epsilon) = 0$$

Then Y_n converges to $c = 0$ in probability since $\Pr(Y_n > \epsilon) = 1/n$ for any $0 < \epsilon < 1$, hence the limit as $n \rightarrow \infty$ is 0, satisfying the definition.

Recall that Y_n converges to c with probability 1 if

$$\Pr(\lim_{n \rightarrow \infty} Y_n = c) = 1$$

However, I claim that $\lim_{n \rightarrow \infty} Y_n$ does not exist, hence doesn't equal zero with probability 1. Why no limit? A 0/1 sequence will have a limit if and only if it is all 0 after some finite point (i.e., contains only a finite number of 1's) or vice versa. But the expected number of 0's & 1's in the sequence are both infinite; e.g.:


$$E[\sum_{i>0} Y_i] = \sum_{i>0} E[Y_i] = \sum_{i>0} \frac{1}{i} = \infty$$

Thus, Y_n converges in probability to zero, but does not converge with probability 1.


Revisiting the “lightbulb model” 2 slides up, w/ “lights on” $\Leftrightarrow 1$, column n has a decreasing fraction ($1/n$) of lit bulbs, while all but a vanishingly small fraction of rows have infinitely many lit bulbs.

(For an interesting contrast, consider the sequence of rvs $Z_n \sim \text{Ber}(1/n^2)$.)


sample mean \rightarrow population mean





sample mean \rightarrow population mean



demo







Weak Law:

$$\lim_{n \rightarrow \infty} \Pr \left(\left| \frac{X_1 + \cdots + X_n}{n} - \mu \right| > \epsilon \right) = 0$$

Strong Law:

$$\Pr \left(\lim_{n \rightarrow \infty} \left(\frac{X_1 + \cdots + X_n}{n} \right) = \mu \right) = 1$$

How do they differ? Imagine an infinite 2-D table, whose rows are indep infinite sample sequences X_i . Pick ϵ . Imagine cell m, n lights up if average of 1st n samples in row m is $> \epsilon$ away from μ .

WLLN says fraction of lights in n^{th} column goes to zero as $n \rightarrow \infty$. It does not prohibit every row from having ∞ lights, so long as frequency declines.

SLLN also says only a vanishingly small fraction of rows can have ∞ lights.

Note: $D_n = E[| \sum_{1 \leq i \leq n} (X_i - \mu) |]$ grows with n , but $D_n/n \rightarrow 0$


Justifies the “frequency” interpretation of probability

“Regression toward the mean”

Gambler’s fallacy: “I’m *due* for a win!”

“Swamps, but does not compensate”

“Result will usually be close to the mean”



Many web demos, e.g.

<http://stat-www.berkeley.edu/~stark/Java/Html/lln.htm>


Recall

normal random variable

X is a normal random variable $X \sim N(\mu, \sigma^2)$

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}$$

$$E[X] = \mu \quad \text{Var}[X] = \sigma^2$$



the central limit theorem (CLT)

i.i.d. (independent, identically distributed) random vars

$$X_1, X_2, X_3, \dots$$

X_i has $\mu = E[X_i] < \infty$ and $\sigma^2 = \text{Var}[X_i] < \infty$

As $n \rightarrow \infty$,





$$M_n = \frac{1}{n} \sum_{i=1}^n X_i \rightarrow N\left(\mu, \frac{\sigma^2}{n}\right)$$

Restated: As $n \rightarrow \infty$,





$$\frac{X_1 + X_2 + \dots + X_n - n\mu}{\sigma\sqrt{n}} \longrightarrow N(0, 1)$$


Note: on slide 5, showed sum of normals is exactly normal. Maybe not a surprise, given that sums of almost *anything* become approximately normal...

demo



CLT applies even to whacky distributions





CLT *also* holds under weaker assumptions than stated above, and is the reason many things appear normally distributed
Many quantities = sums of (roughly) independent random vars

Exam scores: sums of individual problems

People's heights: sum of many genetic & environmental factors

Measurements: sums of various small instrument errors


“Noise” in sci/engr applications: sums of random perturbations

...

Human height is approximately normal.

Why might that be true?

R.A. Fisher (1918) noted it would follow from CLT if height were the sum of many independent random effects, e.g. many genetic factors (plus some environmental ones like diet). I.e., suggested part of *mechanism* by looking at *shape* of the curve.



normal approximation to binomial

Let S_n = number of successes in n trials (with prob. p).

$$S_n \sim \text{Bin}(n,p) \quad E[S_n] = np \quad \text{Var}[S_n] = np(1-p)$$

Poisson approx: good for n large, p small (np constant)

Normal approx: For large n , (p stays fixed):

$$S_n \approx Y \sim N(E[S_n], \text{Var}[S_n]) = N(np, np(1-p))$$

Rule of thumb: Normal approx “good” if $np(1-p) \geq 10$

DeMoivre-Laplace Theorem:


As $n \rightarrow \infty$:

$$Pr \left(a \leq \frac{S_n - np}{\sqrt{np(1-p)}} \leq b \right) \longrightarrow \Phi(b) - \Phi(a)$$

Equivalently:

$$Pr(a \leq S_n \leq b) \longrightarrow \Phi \left(\frac{b - np}{\sqrt{np(1-p)}} \right) - \Phi \left(\frac{a - np}{\sqrt{np(1-p)}} \right)$$

normal approximation to binomial



normal approximation to binomial

Ex: Fair coin flipped 40 times. Probability of 15 to 25 heads?

Exact (binomial) answer:

$$P_{\text{bin}}(15 \leq X \leq 25) = \sum_{i=15}^{25} \binom{40}{i} \left(\frac{1}{2}\right)^{40} \approx \boxed{0.9193}$$

Normal approximation:

$$\begin{aligned} P_{\text{norm}}(15 \leq X \leq 25) &= P\left(\frac{15 - 20}{\sqrt{10}} \leq \frac{X - 20}{\sqrt{10}} \leq \frac{25 - 20}{\sqrt{10}}\right) \\ &\approx P\left(-1.58 \leq \frac{X - 20}{\sqrt{10}} \leq 1.58\right) \\ &= \Phi(1.58) - \Phi(-1.58) \approx \boxed{0.8862} \end{aligned}$$

R Sidebar

```
> pbinom(25, 40, .5) - pbinom(14, 40, .5)
[1] 0.9193095
```

```
> pnorm(5/sqrt(10)) - pnorm(-5/sqrt(10))
[1] 0.8861537
```

```
> 5/sqrt(10)
[1] 1.581139
```

```
> pnorm(1.58) - pnorm(-1.58)
[1] 0.8858931
```

SIDEBARS

I've included a few sidebar slides like this one (a) to show you how to do various calculations in R, (b) to check my own math, and (c) occasionally to show the (usually small) effect of some approximations. Feel free to ignore them unless you want to pick up some R tips.

normal approximation to binomial

Ex: Fair coin flipped 40 times. Probability of 20 or 21 heads?

Exact (binomial) answer:

$$P_{\text{bin}}(X = 20 \vee X = 21) = \left[\binom{40}{20} + \binom{40}{21} \right] \left(\frac{1}{2} \right)^{40} \approx \boxed{0.2448}$$

Normal approximation:

$$\begin{aligned} P_{\text{norm}}(20 \leq X \leq 21) &= P \left(\frac{20 - 20}{\sqrt{10}} \leq \frac{X - 20}{\sqrt{10}} \leq \frac{21 - 20}{\sqrt{10}} \right) \\ &\approx P \left(0 \leq \frac{X - 20}{\sqrt{10}} \leq 0.32 \right) \\ &\approx \Phi(0.32) - \Phi(0.00) \approx \boxed{0.1241} \end{aligned}$$

Hmmm... A bit disappointing.

R Sidebar

```
> sum(dbinom(20:21, 40, .5))
```

```
[1] 0.2447713
```

```
> pnorm(0) - pnorm(-1/sqrt(10))
```

```
[1] 0.1240852
```

```
> 1/sqrt(10)
```

```
[1] 0.3162278
```

```
> pnorm(.32) - pnorm(0)
```

```
[1] 0.1255158
```

normal approximation to binomial

Ex: Fair coin flipped 40 times. Probability of 20 heads?

Exact (binomial) answer:

$$P_{\text{bin}}(X = 20) = \binom{40}{20} \left(\frac{1}{2}\right)^{40} \approx \boxed{0.1254}$$

Normal approximation:


$$\begin{aligned} P_{\text{norm}}(20 \leq X \leq 20) &= P\left(\frac{20 - 20}{\sqrt{10}} \leq \frac{X - 20}{\sqrt{10}} \leq \frac{20 - 20}{\sqrt{10}}\right) \\ &\approx P\left(0 \leq \frac{X - 20}{\sqrt{10}} \leq 0\right) \\ &= \Phi(0.00) - \Phi(0.00) = \boxed{0.0000} \end{aligned}$$

Whoa! ... Even more disappointing.




DeMoivre-Laplace and the “continuity correction”

The “continuity correction”:

Imagine *discretizing* the normal density by shifting probability mass at non-integer x to the nearest integer (i.e., “rounding” x). Then, probability of binom r.v. falling in the (*integer*) interval $[a, \dots, b]$, inclusive, is \approx the probability of a normal r.v. with the same μ, σ^2 falling in the (*real*) interval $[a - \frac{1}{2}, b + \frac{1}{2}]$, even when $a = b$.



Bin($n=10, p=.5$) vs Norm($\mu=np, \sigma^2 = np(1-p)$)



normal approx to binomial, revisited

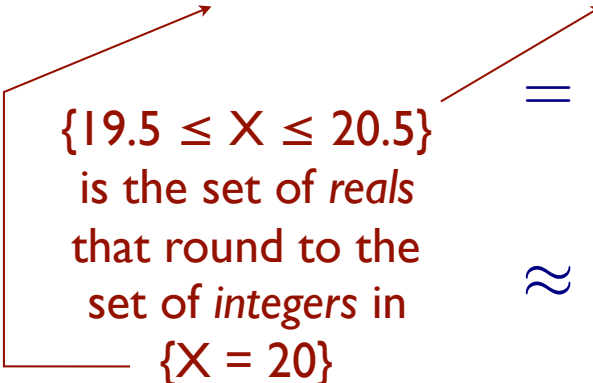
Ex: Fair coin flipped 40 times. Probability of 20 heads?

Exact (binomial) answer:

$$P_{\text{bin}}(X = 20) = \binom{40}{20} \left(\frac{1}{2}\right)^{40} \approx \boxed{0.1254}$$

Normal approximation:

$$P_{\text{bin}}(X = 20) \approx P_{\text{norm}}(19.5 \leq X \leq 20.5)$$



$\{19.5 \leq X \leq 20.5\}$
is the set of *reals*
that round to the
set of *integers* in
 $\{X = 20\}$

$$= P_{\text{norm}}\left(\frac{19.5 - 20}{\sqrt{10}} \leq \frac{X - 20}{\sqrt{10}} < \frac{20.5 - 20}{\sqrt{10}}\right)$$

$$\approx P_{\text{norm}}\left(-0.16 \leq \frac{X - 20}{\sqrt{10}} < 0.16\right)$$

$$= \Phi(0.16) - \Phi(-0.16) \approx \boxed{0.1272}$$

normal approx to binomial, revisited

Ex: Fair coin flipped 40 times. Probability of 20 or 21 heads?

Exact (binomial) answer:

$$P_{\text{bin}}(X = 20 \vee X = 21) = \left[\binom{40}{20} + \binom{40}{21} \right] \left(\frac{1}{2} \right)^{40} \approx \boxed{0.2448}$$

Normal approximation:

$$P_{\text{bin}}(20 \leq X < 22) = P_{\text{norm}}(19.5 \leq X \leq 21.5)$$

$$= P_{\text{norm}} \left(\frac{19.5 - 20}{\sqrt{10}} \leq \frac{X - 20}{\sqrt{10}} \leq \frac{21.5 - 20}{\sqrt{10}} \right)$$

$$\approx P_{\text{norm}} \left(-0.16 \leq \frac{X - 20}{\sqrt{10}} \leq 0.47 \right)$$

$$\approx \Phi(0.47) - \Phi(-0.16) \approx \boxed{0.2452}$$

$\{19.5 \leq X < 21.5\}$
is the set of *reals*
that round to the
set of *integers* in
 $\{20 \leq X < 22\}$

One more note on continuity correction: Never wrong to use it, but it has the largest effect when the set of integers is small. Conversely, it's often omitted when the set is large.

R Sidebar

```
> pnorm(1.5/sqrt(10)) - pnorm(-.5/sqrt(10))
```

```
[1] 0.2451883
```

```
> c(0.5, 1.5)/sqrt(10)
```

```
[1] 0.1581139 0.4743416
```

```
> pnorm(0.47) - pnorm(-0.16)
```


```
[1] 0.244382
```

continuity correction is applicable beyond binomials

Roll 10 6-sided dice

X = total value of all 10 dice

Win if: $X \leq 25$ or $X \geq 45$



$$E[X] = E\left[\sum_{i=1}^{10} X_i\right] = 10E[X_1] = 10(7/2) = 35$$

$$\text{Var}[X] = \text{Var}\left[\sum_{i=1}^{10} X_i\right] = 10\text{Var}[X_1] = 10(35/12) = 350/12$$

$$P(\text{win}) = 1 - P(25.5 \leq X \leq 44.5) =$$

$$1 - P\left(\frac{25.5 - 35}{\sqrt{350/12}} \leq \frac{X - 35}{\sqrt{350/12}} \leq \frac{44.5 - 35}{\sqrt{350/12}}\right)$$

$$\approx 2(1 - \Phi(1.76)) \approx 0.079$$

example: polling

Poll of 100 randomly chosen voters finds that K of them favor proposition 666.

So: the *estimated proportion* in favor is $K/100 = q$

Suppose: the *true proportion* in favor is p .

Q. Give an upper bound on the probability that your estimate is off by > 10 percentage points, i.e., the probability of $|q - p| > 0.1$

A. $K = X_1 + \dots + X_{100}$, where X_i are Bernoulli(p), so by CLT:

$K \approx$ normal with mean $100p$ and variance $100p(1-p)$; or:

$q \approx$ normal with mean p and variance $\sigma^2 = p(1-p)/100$

Letting $Z = (q-p)/\sigma$ (a standardized r.v.), then $|q - p| > 0.1 \Leftrightarrow |Z| > 0.1/\sigma$

By symmetry of the normal

$$P_{\text{Ber}}(|q - p| > 0.1) \approx 2 P_{\text{norm}}(Z > 0.1/\sigma) = 2 (1 - \Phi(0.1/\sigma))$$

Unfortunately, p & σ are unknown, but $\sigma^2 = p(1-p)/100$ is maximized when $p = 1/2$, so $\sigma^2 \leq 1/400$, i.e. $\sigma \leq 1/20$, hence

$$2 (1 - \Phi(0.1/\sigma)) \leq 2(1 - \Phi(2)) \approx \boxed{0.046}$$

Exercise: How much smaller can σ be if $p \neq 1/2$?

i.e., less than a 5% chance of an error as large as 10 percentage points.

Distribution of $X + Y$: summations, integrals (or MGF)

Distribution of $X + Y \neq$ distribution X or Y in general

Distribution of $X + Y$ is normal if X and Y are normal (*)
(ditto for a few other special distributions)

Sums generally don't "converge," but averages do:

Weak Law of Large Numbers

Strong Law of Large Numbers

Most surprisingly, averages often converge to the *same* distribution:
the Central Limit Theorem says sample mean \rightarrow normal

[Note that (*) essentially a prerequisite, and that (*) is exact, whereas CLT is approximate]