

# Week 3

## 27 May 2021

### 11 The binomial distribution

We're now going to look at a generalization of flipping a coin (aka, doing an action multiple times where the actions are independent of one another). We saw this partially when we talked about the Bernoulli  $p$  distribution.

Suppose that I have a fair coin and I flip it multiple times. I don't necessarily care about the order that the flips come in, but I do care about the end result. For example, if I flip the coin four times and I get  $(H, T, T, H)$ , I just care that I got two heads and two tails. Let's see how many possibilities we have when keep flipping the coin.

The first time I flip it I can either get a heads or a tails:

The second time I flip, I again can either get a heads or tails. But since we don't care about order  $(H, T)$  and  $(T, H)$  are the same. So we have

Let's flip a third time. We end up getting:

Notice that this just gives us the binomial distribution which we talked about in the counting section in the first week. We can draw these numbers using *Pascal's triangle*:

Wikipedia: [Pascal's Triangle](#)




Figure 3.1: Pascal's Triangle

We can now use this to figure out probabilities given an arbitrary (not necessarily fair) coin. If the coin has probability  $p$  of landing on heads, then we can label each edge with  $p$  for the probability we flip a heads and a  $q = 1 - p$  to represent a tails flip. So for example, if I flip 3 times, the probability that I will get 2 heads and 1 tails is  $3p^2q$ .

The name “binomial” comes from the expansion of the binomial  $(x + y)^n$ . For example:

We can keep extending Pascal's triangle infinitely in order to calculate any  $n$  that we want. And so we might wonder, how does this help with flipping coins? This is where it helps to recall that the numbers in Pascal's triangle are given by the choice function. In particular if we go to the  $n$ th row and look at the  $k$ th number, that number is equal to  $\binom{n}{k} = \frac{n!}{(n-k)!k!}$ . In other words, if I flip my coin  $n$  times and  $k$  of them turn out to be heads then we know the probability is equal to:

This is known as the binomial distribution.

To be more precise, if we have  $n$  independent trials with probability of  $p$  that a particular trial is successful, then the *binomial probability formula* is given by:

For a fixed  $n$  and  $p$ , the binomial probabilities for all  $k$  give a probability distribution over the numbers  $\{0, 1, \dots, n\}$  called the *binomial distribution*.

Notice how we have 0 in our distribution! This is why we let  $0! = 1$ ; or else we couldn't work with factorial properly.

Wikipedia: [Binomial Distribution](#)  
 Note that the book uses the term “Binomial  $(n, p)$  distribution”.


**Example 11.1** Let's look at an example of something other than flipping a coin. Say we have a six sided dice and we want to figure out what the probability of getting 3 sixes out of 7 rolls. First, we know that each roll is independent and so we know we can use the binomial probability function. We know that  $p = \frac{1}{6}$  and since we are doing 7 rolls we have  $n = 7$ . As we want 3 sixes we also know that  $k = 3$ . Therefore we have:

**Example 11.2** Now we get a little more complicated. Say we're in a coin flipping tournament and there are five contestants. Each contestant must flip six coins and are allowed into the next round of the competition if they get more than three heads. What are the chances that at least three contestants proceed to the next round?

## 12 Consecutive Odds Ratios

It turns out that for the binomial distribution, odds are actually a nicer way to look at things. Remember that odds compare two things and say what the ratio between two things occurring is. (So “Horse  $A$  has 2 : 3 odds in favour of winning against horse  $B$ ” means out of ever 5 races, horse  $A$  will win 2 and horse  $B$  will win 3 on average) We’re going to look at consecutive odds in the binomial distribution and see what they give.

**Example 12.1** Let’s look at the distribution when  $p = \frac{1}{2}$  and  $n = 10$ . We end up with the following:



The distribution is given by:

Let’s compare them consecutively:

It’s easy to see that this generalizes generally:

The above only gives us consecutive odds if the probability of success is equal to  $\frac{1}{2}$ , but it’s not hard to see how this expands to arbitrary probability.

**Theorem 12.2** (Consecutive odds for the binomial distribution) *For independent events with a probability of success  $p$ . Then the odds of  $k$  successes relative to  $k - 1$  success are given by:*

**Example 12.3** Let's see how this helps. Say that I have hat with five numbers in it and every week I pull out a number for seven weeks. What is the probability distribution of pulling out the number 4.

### 13 Most probable outcome

Notice how in the previous example  $P(1)$  had the highest value. This implies that the number 4 has the highest probability of appearing exactly 1 time. So it makes sense to ask if we can generalize this. If I have  $n$  independent trials and each trial has probability  $p$  of being successful then we might expect to have roughly \_\_\_\_\_ successes. (For example, if I flip a coin 4 times then I expect \_\_\_\_\_ heads.) If we let  $m$  be the trial which has the highest probability, *i.e.*,  $P(m) > P(k)$  for all  $k \neq m$ , we call  $m$  the *mode* of the probability. We end up with the following:

**Theorem 13.1 (Mode of the binomial distribution)** *If  $p \in (0, 1)$  where  $p$  is the probability of success of an independent trial and we suppose that we perform  $n$  trials. Then the mode of the binomial distribution is given by  $m$  where*

A related idea in probability is the expected value. The *expected value* is the expected number of successes. . This is different than the mode which is the most likely number of successes. The expected value is defined to be  $np$  and it is usually denoted by  $\mu$  for the binomial distribution.

Wikipedia: [Expected value](#)

The book uses the term *mean* as the expected value of a binomial distribution.

The difference between the two is best seen through examples.

**Example 13.2** First let's see when they are the same. If  $n = 4$  and  $p = \frac{1}{2}$  (aka we flip a coin 4 times) we expect the number of successes to be:

This is the expected value.

The most likely number of successes (aka the mode) is given by:

Notice that they are the same!

**Example 13.3** So let's look at when they aren't the same. let's flip the coin a fifth time. So  $n = 5$  and  $p = \frac{1}{2}$  again.

In this case, the expected number of successes is:

So we expect roughly two and a half heads to happen. That is the expected value.

Obviously this doesn't make sense since coin flips are static! We should have whole numbers in this case. That is where the mode kicks

in and gives us an actual number. The most likely number of successes is given by:


So we expect roughly 3 heads to occur.

This will be more discussed in more detail when we hit chapter three in the text.

## 14 Normal Distribution

Recall when we were looking at the binomial distribution we had the following example.


**Example 14.1** The binomial distribution when  $p = \frac{1}{2}$  and  $n = 10$  is given by:



If you notice, we can kind of create a curve from these bars. This gives us what's called the normal curve.

The *normal curve* is the curve with equation:

The normal curve is often called the *bell curve* in the non-math world.



Five of these symbols you should know:

- $x$  is a variable
- $y$  is a variable (sometimes represented as  $f(x)$ ).
- $\pi$  is a constant roughly equal to 3.14159265358979...
- $e$  is a constant roughly equal to 2.7182818285...
- $\mu$  is the *expected value*.

Wikipedia: [Standard deviation](#)

The only variable we haven't discussed so far is  $\sigma$  which is known as the *standard deviation*. Note that  $\mu$  can be *any* real number while  $\sigma$  must be a strictly positive real number.

“Infinitesimally small width” should make you think of integration/calculus.

The best way to think about the normal curve is to think about it like a continuous histogram. In essence, it's when we start taking our boxes and making their width infinitesimally small. The  $\mu$  tells us where the peak of the distribution is and the  $\sigma$  tells us how it's spread out.

Note that the constant  $\frac{1}{\sigma\sqrt{2\pi}}$  is only there to make sure that the area under the curve is equal to 1.

Wikipedia: [Normal distribution](#)

The normal distribution is sometimes called the “Gaussian distribution”.


Although we have a curve, that doesn't actually define a distribution. Recall that for a distribution we need to define a probability map. The *normal distribution with expected value  $\mu$  and standard deviation  $\sigma$*  is the distribution over the  $x$ -axis defined by the areas under the normal curve with  $\mu$  and  $\sigma$ . If \_\_\_\_\_ then we call this distribution the *standard normal distribution*.

Another way to look at the normal distribution is through a cumulative distribution function. For this we let  $z = \frac{(x-\mu)}{\sigma}$  and we say that  $z$  is  *$x$  in standard units*. This terminology comes from the fact that in the standard normal distribution  $z = x$ . We let  $\Phi(x)$  denote the probability of everything to the left of  $z$ . The function  $\Phi$  is called the *standard normal cumulative distribution function*.

Wikipedia: [Cumulative distribution function](#)

We'll normally just refer to CDF to mean “cumulative distribution function”.





Since we're looking at area under the curve, we know that the standard normal CDF is given by the formula:

$$\Phi(z) =$$

The function \_\_\_\_\_ is known as the *standard normal distribution*. For a normal distribution the probability between  $a$  and  $b$  is given by:

If you recall from the end of your calculus class,  $e^{-\frac{1}{2}x^2}$  doesn't have a simple indefinite integral, so  $\Phi(z)$  has no simple exact formula unfortunately. But there are methods to calculate it.

According to the book, there's a nice approximation function for  $\Phi(z)$  given by:


$$\Phi(z) \approx 1 - \frac{1}{2} \left( 1 + 0.196854z + 0.115194z^2 + 0.000344z^3 + 0.019527z^4 \right)^{-4}$$

Checkout Appendix 5 if you want some values just for fun.

This approximation will give an approximation up to three significant figures for every  $z \geq 0$ .

Even without exact formulas there are some things we can deduce directly from the normal curve. For example, we know that

$$\Phi(-z) = 1 - \Phi(z) \text{ and } \Phi(0) = \frac{1}{2}$$



If we let  $\Phi(a, b)$  denote the probability on the interval  $(a, b)$  then by the difference rule of probabilities we have:

We can also combine some formulas to get:

It's best not to memorize all of these formulas. It's better to try and know how we are deriving them. You can always look them up if you need to double check.

Plugging in some numbers we see:

$$\Phi(-1, 1) \approx \text{___}\% \quad \Phi(-2, 2) \approx \text{___}\% \quad \Phi(-3, 3) \approx \text{_____}\%$$


From these probabilities you can deduce almost any other probability.

**Example 14.2** What is  $\Phi(0, 1)$ ?

What is  $\Phi(2, \infty)$ ?

## 15 The normal approximation to the binomial distribution

When we defined the normal curve, we used  $\mu$  and  $\sigma$ , but we also mentioned that the idea came from looking at the curve of the binomial distribution.



Since the binomial distribution is dependent on  $n$  and  $p$ , we might ask how are the two pairs of numbers related to one another? One of them,  $\mu$ , we already know. The expected value  $\mu$  is just equal to \_\_\_\_\_. But what about  $\sigma$ ? We won't go into details now, but it turns out that for sufficiently large  $n$  then  $\sigma =$  \_\_\_\_\_. Remember that we called this  $\sigma$  the *standard deviation*.

Recall that in the normal curve, it doesn't make sense to ask about the probability at a certain point. So instead, let  $P(a \text{ to } b)$  be the probability of getting  $i$  successes where  $i \in [a, b]$  in  $n$  independent trials with probability  $p$  of success. Then we have that

$$P(a \text{ to } b) \text{ is}$$

$$\approx$$

$$=$$

In other words, the *normal approximation to the binomial distribution* for  $n$  independent trials with probability  $p$  of success is given by:

Let's see how this helps.

**Example 15.1** Suppose that I randomly toss a fair coin  $n = 100$  times. We want to know the probability of getting 50 heads.

By the binomial distribution we have:

By the numerical approximation we have that  $a = b = 50$  and  $\mu = np = 100 \cdot \frac{1}{2} = 50$ .

**Fluctuations** If we make  $n$  grow more and more, we notice that our box sizes get smaller and smaller. This means that our approximations get better and better as  $n$  gets bigger. The size of the random fluctuations is of the order of  $\sigma = \sqrt{np(1-p)}$  as it tells us how far from  $\mu$  we are. So as  $n \rightarrow \infty$  our  $\sigma$  gets larger and more of the distribution is covered.

Likewise, the proportion of successes fluctuates over time too. The typical size of random fluctuations in the relative frequency of successes is of order  $\frac{\sqrt{p(1-p)n}}{n} = \sqrt{p(1-p)/n}$ . Since  $p(1-p)$  is maximized when  $p = (1-p) = \frac{1}{2}$ , we know  $\sqrt{p(1-p)/n} \leq \frac{1}{\sqrt{4n}} = \frac{1}{2\sqrt{n}}$ . And as  $n \rightarrow \infty$  we get that  $\frac{1}{\sqrt{n}} \rightarrow 0$  making it so that this order goes to 0 as  $n$  grows.

**Theorem 15.2 (Square root law)** *In  $n$  independent trials with  $n$  sufficiently large and probability  $p$  of success then*

- *The number of success will (with high probability) lie in a relatively small interval of numbers, centred on  $\mu = np$ , with width a moderate multiple of  $\sqrt{n}$ .*
- *The proportion of successes will (with high probability) lie in a small interval centred on  $p$ , with width a moderate multiple of  $\frac{1}{\sqrt{n}}$ .*

This implies one of my favourite theorems in probability:

**Theorem 15.3 (Law of large numbers)** *Suppose that  $n$  is the number of independent trials and that  $n$  is very large. Let  $p$  be the probability of success on each trial. Then for every  $\varepsilon > 0$*

$P(\text{proportion of successes in } n \text{ trials differs from } p \text{ by less than } \varepsilon) \rightarrow 1$

Wikipedia: [Law of large numbers](#)

as  $n \rightarrow \infty$ .

What this is saying is that the more trials we conduct, the closer to  $p$  our relative frequency will become.

Note that this is sometimes known as the “strong law of large numbers”.