

In this tutorial, you're going to learn about what I believe is the most important theorem in all of statistics. It's called the Central Limit Theorem. It's a big deal. So it starts with some basic understandings of the characteristics of distributions-- center, spread, and shape. And those are characteristics that any distribution will have. In this particular example, we're going to be talking about a specific distribution called a sampling distribution.

The Central Limit Theorem deals with the shape of a sampling distribution. Center and spread are talked about more in another tutorial. But the short version is this-- is that in a sampling distribution, the center is the same as the center of the original distribution. That is to say, the mean of all the  $\bar{x}$  averages is the same as the mean of the original distribution.

Now, how about spread? The spread, or standard deviation, is the same as the original standard deviation divided by the square root of sample size. So that is to say, the standard deviation of all the  $\bar{x}$  bars is equal to the original standard deviation divided by square root of  $n$ ,  $n$  being the sample size. But we don't have anything to say yet about shape, and that's where the Central Limit Theorem comes in.

So consider this spinner. Consider the sampling distributions caused by averaging different numbers of spins. Well, it's pretty obvious that if you spun it once, you would get a one about three out of eight times, a two about one out of eight times, a three about two out of eight times, and a four about two out of eight times, making the distribution look something like this.

One being the most common, three and four being equally common and next most common, and two being the least common. But what about the sampling distributions if you averaged four spins? Well, you wouldn't just have options of one, two, three, and four anymore. You'd have options of  $1/4$ ,  $1/2$ , one and  $3/4$ , two, two and  $1/4$ , two and  $1/2$ , et cetera.

You'd have more options, and that would necessarily decrease the likelihood of getting all ones or all fours. The distribution would look something like this. Getting all four ones would be extremely unlikely, and getting all four four's would be extremely unlikely. It seems like the most likely scenario is you getting two and  $1/4$ , which makes sense. There are a little bit more ones than there are anything else.

And so it pulls the mean down a little bit from maybe you thought it would be two and  $1/2$ , right in the middle of two and three, but it's pulled down a little bit. This is a slightly skewed to the right distribution. What about if you sampled nine and averaged nine spins? Well, the probability, for instance, that you get all ones therefore averaging a one goes down even further.

And the probability that you get all fours goes down, way, way, way down to almost zero. Let's take a look at that sampling distribution. Well, it looks like it's possible to get all ones, though not very likely. It looks like it's going to be a lot more common to get something between two and three. And in fact, what you might be seeing is that the spread of the sampling distribution is decreasing as  $n$  gets bigger.

But what's happening to the shape? Let's look at averaging 20 spins. Now it's almost impossible to average a one or average of four, or even average something close to three. You're almost guaranteed to average something between two and three. The spread, again, is decreasing. But what happened to the shape? Well, look at the evolution here. What would you say is happening to the shape?

The shape as  $n$  increases is becoming more normal. And this is what the Central Limit Theorem says is that when the sample size is large, the shape of the sampling distribution of means becomes very nearly normal. Now, when it's stated like this, there's an obvious question, because there's one word here that's not very well defined. And it's the word large.

How large is considered a large enough for the distribution to be considered approximately normal? For instance, if you look back to this page, is it approximately normal here when  $n$  is nine? Is it approximately normal here when  $n$  is 20? What constitutes a large enough sample so that when you show the distribution of all the averages, you get a normal distribution?

Well, it's different depending on what the original distribution looked like. Our original distribution was almost uniform, so it didn't take very many trials. If the distribution had been heavily skewed, it would have taken more trials to average out some of those high numbers with some of those low numbers. So to be on the safe side, we're going to say 30 is going to be a good sample size such that when we average the 30 observations, we're going to get something close to what we expect.

And the ones that are off from what we expect will tail off in a normal shape. So for almost all distributions, a sample size of 30 is exactly what we want. As it turns out, it's the Central Limit Theorem that explains why so many real world processes are normally distributed. So to recap, a sampling distribution is the distribution of all possible means of a given size.

And the Central Limit Theorem outlines that when the sample size is large, for most distributions, that means 30 or larger, the distribution of sample means will be approximately normal. And occasionally, you need to make it even bigger still than 30. If the parent distribution that we started with is very, very skewed or has outliers. And there's your Central Limit Theorem, probably the most important idea in statistics. So good luck, and we'll see you next time.