We now discuss how to come up with confidence intervals when we try to estimate the unknown mean of some random variable, or of some distribution, using the sample mean as our estimator.

So here X1 up to Xn are independent, identically distributed random variables that are drawn from a distribution that has a certain mean theta, the quantity that we want to estimate, and some variance sigma squared.

Let us say that we want to construct a 95% confidence interval.

Our starting point will be the fact that the sample mean, according to the central limit theorem, can be described approximately using normal distributions.

And we look up at the normal table, and we observe the following-- that if we take a standard normal random variable, then there is probability, 97.5% of falling below this number, 1.96, which means that there is probability 2 1/2% of falling above that number.

And by symmetry, the probability of falling below minus 1.96 is also 2 1/2%.

This means that this middle interval here has probability 95%, and we exploit this fact as follows.

If we take the sample mean, subtract the true mean, and then divide by the standard deviation of the sample mean, then we obtain a random variable, which is approximately a standard normal.

Therefore, what we have here is the probability of an approximately standard normal random variable.

Or actually, its absolute value falling below 1.96.

This is just the event that our standard normal falls inside this middle interval here, according to this entry from the normal tables and the previous discussion, this probability is going to be approximately 95%.

And now we take this statement, send this term to the other side of the inequality, and then interpret what it means for an absolute value to be less than something.

And we obtain an equivalent statement.

This event here is algebraically identical to the event that we have up there, and this provides us with the desired confidence interval.

We think of this quantity here as the lower end of the confidence interval.

This quantity here is the upper end of the confidence interval.

And this statement tells us that there is probability approximately equal to 95% that the confidence interval constructed this way contains the true value of the unknown parameter.

So this is how we obtain a 95% confidence interval.

If instead we wanted a 90% confidence interval, we would proceed in more or less in the same way.

Here, we would want to have the number 0.95.

Why is that?

We want this middle interval to have probability 90%, which means that we want to have probability 5% at each one of the tails.

And then we look up at the normal tables, and we find that the entry that gives us probability 95% of being below that value is 1.645.

So if we use 1.645 in place of 1.96, we obtain a 90% confidence interval, and similarly for other choices.

For example, if we want a 99% confidence interval.

There's only one issue that's left to discuss, and this is the following.

In order to obtain numerical values for the endpoints of the confidence interval, we need to know sigma, the standard deviation of the random variables that we are observing.

But if we do not know the value of sigma, then we may have to do some additional work.