We now follow a program that parallels our development for the case of the Bernoulli process.

We will study the time until the first arrival, a random variable that we denote by T1.

We're interested in finding the probability distribution of this random variable.

And later on, we will continue and try to study the time until the kth arrival.

Now T1 is a continuous random variable, because the Poisson process runs in continuous time.

And therefore, it has a PDF.

But instead of finding the PDF directly, we will first find the CDF of this random variable.

So we fix a certain time, T. And we're asking for the probability that the first arrival happens during this interval.

Now this is 1 minus the probability that the first arrival happens outside this interval.

So we can write this probability as 1 minus the probability that T1 is bigger than t.

But what is this event?

The first arrival occurring after time, little t, is the same as saying that there were no arrivals in the time interval from 0 to little t.

And this probability of 0 arrivals in a time interval of length t is something for which we already have a formula.

Take this formula and replace k by 0, tau by t.

When k is equal to 0, this term is something to the 0-th power equal to 1.

Using our convention, that 0 factorial is equal to 1, we're left just with e to the minus lambda t.

And this is the answer for the CDF of the time until the first arrival.

We then take the derivative.

And we find that the PDF of the time until the first arrival has this form, which is the PDF of an exponential random variable.

Of course, this calculation here is only valid for t's that are non-negative.

For negative t's, the PDF of T1 is, of course, 0.

For the exponential random variable, we have seen that it has certain memorylessness properties.

Namely, if I condition on an event that nothing has occurred until a certain time, t, and I am interested in the time from now until the first arrival occurs, this remaining until the first arrival is again an exponential distribution.

That is, looking ahead from this time, I will still wait an exponentially distributed amount of time until I see the first arrival.

Whatever happened in the past and how long I have been waiting doesn't matter.

Starting from this time, I will still wait an exponentially distributed amount of time.

This is essentially an expression of a fresh start property of the Poisson process, which is analogous to the fresh start properties for the Bernoulli process.

And we will be discussing this fresh start property a lot more.

Having figured out the distribution of the time of the first arrival, let us now study the time of the k-th arrival, a random variable that we denote by Y sub k, similar to the case of the Bernoulli process.

This random variable is a continuous one, because arrivals happen in continuous time, so it takes continuous values.

And therefore, it will be described by a PDF.

And this is what we want to find.

In order to find it, we will make use of the Poisson PMF that we have already derived for the number of arrivals during an interval of a fixed length.

One approach to finding the PDF of Yk is the usual program, similar again to what we did for the case of the first arrival time.

We can first find CDF, and then differentiate to find the PDF.

So what is the CDF?

We want to calculate the probability that Yk is less than or equal to some number, little y.

Now what is this event?

The k-th arrival occurs by time y.

This means that by time y, we've had at least k arrivals.

We've had k arrivals, or maybe k plus 1, or maybe k plus 2.

We've had some number of arrivals, n, in an interval of length, y.

And this is an event that happens with this probability.

But we need to take into account all of the possible values of n that are at least as large as k.

Now we have a formula for this probability, the probability of n arrivals in an interval of given length.

This is the Poisson PMF with appropriate changes of symbols.

So we can take this expression, substitute it here, and then differentiate to do some algebra and find the answer.

Instead of carrying out this algebra, however, we will proceed in a more intuitive way that will get us there perhaps faster.

And the derivation that we would follow actually parallels the one that we went through in the case of the Bernoulli process.

The intuitive argument that we will use will rest on the interpretation of a PDF in terms of probabilities of small intervals.

So the PDF evaluated at some particular point, y, times delta, is approximately the probability that our random variable falls within a delta interval from this number, little y, that we're considering.

So here's time 0, here's time y, and here's time y plus delta.

We want to find or to say something about the probability of falling inside this small interval.

Now what does it mean for the k-th arrival to fall inside this interval?

This is an event that can happen as follows.

The k-th arrival falls in this interval, and we've had k minus 1 arrivals during the previous interval.

What is the probability of this event?

A basic assumption about the Poisson process is the independence assumption.

Therefore, having k minus 1 arrivals in this interval and having one arrival in this interval are independent events.

Therefore, the probability of this scenario is the product of the probabilities that we've had k minus 1 arrivals in an interval of length, y, times the probability that we've had one arrival in an interval of length delta.

And that latter probability is approximately equal to lambda times delta.

So I should write here an approximate equality instead of an exact equality, to indicate that there are other terms, order of delta squared, for example, but which are much smaller compared to the delta.

However, this is not the only way that we can get the k-th arrival in this interval.

There's an alternative scenario.

We might have had k minus 2 arrivals during this interval, and then two arrivals during that little interval.

In this case, the k-th arrival again occurs within that little interval.

So we need to also calculate the probability of this scenario.

The probability of that scenario is the probability of k minus 2 arrivals in an interval of length, y, times the probability of two arrivals.

But the probability of two arrivals is something that's order of delta squared.

And order of delta squared is much smaller than this term, which is linear in delta.

And so this term can be ignored as long as we're just keeping track of the dominant terms, those are linear in delta.

And then, they would be similar expressions.

For example, the scenario that we have three arrivals up to time y, and then three more arrivals during that little interval, which is again an event of probability, order of delta squared, that we get three arrivals there.

And all of those terms are insignificant, and we can ignore them.

And we end up with an approximate equality between this term and this expression here.

Delta shows up on both sides, so we can cancel delta.

And therefore, we have ended up with a formula for the PDF.

In particular, the PDF is equal to this probability times lambda.

What is this probability?

We have a formula for it.

But we just need to substitute.

Put k minus 1 in the place of k, and put y in the place of tau.

This gives us lambda y to the power k minus 1, e to the minus lambda y, divided by k minus 1 factorial.

And then we have the extra factor of lambda, which can be put together with this lambda to the k minus 1 here.

And we end up with this final formula for the PDF of Yk.

The distribution that we have here is called an Erlang distribution.

But actually, it's not just one distribution.

We have different distributions depending on what k we're considering.

The distribution of the time of the third arrival is different from the distribution of the 10th arrival.

So if we fix a particular k, then we say that we have an Erlang distribution of order k.

For the case where k is equal to 1, this term here disappears, k minus 1 is equal to 0.

And the denominator term disappears, and we end up with lambda times e to the minus lambda y.

But this is the exponential distribution that we had already derived with a different method earlier.

As you increase k, of course, you get different distributions.

And these tend to shift towards the right.

This makes sense.

The time of the second arrival is likely to take certain values.

The time of the third arrival is likely to take values that are higher.

And the more you increase k, the more the distribution will be shifting towards the right.