

So now we can come to the central topic of our lecture, which describes the conditions under which a Markov chain reaches steady state.

The question that we are asking, and which we motivated in the previous lecture by looking at an example with a simple, two-state Markov chain is the following-- we are asking whether the probability of being in state j at time n , given that you started at time 0 in state i , converges to some constant, π_j .

In fact, that question consists of two parts.

Do we have convergence?

And is it independent of i ?

We have seen an example where it is not always the case.

For example, in this Markov chain, you have two recurrent classes.

This is one recurrent class here.

And then there's a second recurrent class here.

And we know that if we are interested in the long-time probability of being in that state, assuming that you started in one of these states here, the probability will be 0 to be here.

But if you started in one of these two states, the probability would be positive.

So clearly here, the initial conditions will matter whenever you have two or more recurrent classes.

So what would happen if you have only one recurrent class?

So let's remove this one and consider this situation here, where you have only one recurrent class.

In that case, what we have seen-- this is still not sufficient.

Indeed, if you look at this recurrent class and you are interested in 9 and assume that you started at 9, then at time 1, you will be either here or here.

And at time 2, you will be back at 9 for sure.

And in general, for time n even, the probability will be one, and for time n odd, it will be zero.

So that specific n -step transition probability in that situation here will never converge.

It will keep oscillating between 0 and 1.

So the issue here is that we had a periodic recurrent class, and the period in that case was 2.

So let us consider now the final case where you have only one recurrent class.

And that recurrent class is not periodic.

And how do we realize that this is not periodic here?

Well, we have a self transition here.

So now that we have one recurrent class, and this recurrent class is aperiodic, the question is-- do you have this kind of convergence here?

And it turns out-- and this is the big theory of Markov chains under the name of the steady-state convergence theorem-- that indeed, yes.

The r_{ij} 's do converge to a steady-state limit, which we call a steady-state probability as long as these two conditions are satisfied.

So in summary, not only these two conditions are necessary, like we had seen with our counter example, but they are sufficient.

We're not going to prove this theorem here.

It's a little bit complicated.

But what is the intuitive idea behind this theorem?

Well, let us think intuitively as to why the initial state does not matter, when the chain has a single recurrent class and no periodic states.

The technique is pretty classical.

The idea is the following-- think about two independent copies of that Markov chain, starting at two different initial conditions.

So for example, think about a red copy.

And the red copy would initially start at state 2.

And then at each unit of time will jump to the next state, according to the transition probabilities of this Markov chain.

So for example, so this is at time 0, which was here.

Time 1 might come here.

Time 2, 3, 4, 5 and so forth.

So this is one copy of the Markov chain.

So think about another copy, the blue copy.

And assume that the blue copy started here at time 0.

And again, independently of the other, but during the same unit of time, it will jump from state to state according to the transition probabilities again.

So think that maybe this one will go here.

Then will go here.

Times 2, 3 will be here, 4 here, maybe 5 here.

And so forth.

Now look at these two independent copies.

There will be a time and in that case, for our little example here, is the time 4, where for the first time they collide, in the sense that they jump to the same state at the same time.

So at time 4, both of them are here.

Now, think a little bit about the future evolution of these two independent copies, given that they are in state 4 now.

And here we are using the Markov property to say that the future evolution of the blue copy is independent of the previous path.

Given that you are in state 4, the fact that you started in 1 does not matter for the future evolution of that blue copy.

And for the red copy, given that you are in 4, the fact that you started in 2 does not matter to characterize the future evolutions of that red copy.

So in some sense, probabilistically speaking, these two copies cannot be distinguished for their future evolutions, given that they both are at state 4.

So this means that the initial conditions for these two copies, given that these two copies met at a given state, at a given time-- probabilistically speaking, nothing can differentiate them in the future because of the Markov property.

That's essentially the high-level idea of this proof.

Now, the key thing here mathematically is to prove that whenever you have a Markov chain that has a single recurrent class and this single recurrent class is not periodic, and you start from any initial conditions, the two copies will eventually meet in a given state at a given time with probability 1.

OK.

So now let's assume that the theorem holds.

That means that yes, indeed, we have proved the existence of these steady-state probabilities.

The question is now how to calculate them.

Well, the way to do it is to start from our key recursion that we had for the m-state transition probabilities.

So where we assume here that we have m states, and we are going to take the limits on both sides of this equality.

So when n goes to infinity, we know that r_{ij} of n will go to π_j .

And here, when n goes to infinity, in some sense n minus 1 also goes to infinity.

And so r_{ik} of n minus 1 should go to π_k .

And so we are using that property.

And again, we take the limit as n goes to infinity.

And we say that r_{ij} of n converges to here.

Now, the limit on this side-- you have a limit of a finite sum.

You can exchange the summation and the limit.

And so you take the limit inside.

The limit of $\sum_{k=1}^n r_{ik}$ of n minus 1, when n goes to infinity, goes to π_i of k .

And then you have the resulting term.

And so from that, taking the limit again as n goes to infinity on both sides, you end up with this equation here for j .

Now, you can do that for any of the j of your Markov chain.

So you have m states, so you end up having m equations.

And you have m unknowns, the m π_j 's.

So this is a system of m equations with m unknowns.

Unfortunately, this system is singular and it has multiple solutions.

And one way to see that is the solution $\pi_j = 0$ for all j is a valid solution to the system.

Zero equals zero.

So clearly this is not very informative.

So maybe we need one more condition to get a uniquely solvable system of linear equations.

It turns out that the system of equations has a unique solution if you impose an additional condition, which is pretty natural, which means that the π_j 's are actually probabilities.

They should all sum to 1.

In other words, in the future, if you ask yourself what is the probability of being in state j , and you get π_j , the summation of all of the possible states have to be 1.

If you consider that additional equation, plus that system here, so if you consider this extended system, then you can show that this has a unique solution.

And this unique solution cannot be this one.

And so in conclusion, we can find the steady-state probabilities of the Markov chain by just solving these linear equations, which should be numerically a straightforward procedure.