Time Series Modelling and Forecasting with Applications in R

Prof. Sudeep Bapat

Shailesh J. Mehta School of Management

Indian Institute of Technology Bombay

Week 08

Lecture 38: Tests for Cointegration

Hello all, welcome to this course on time series modeling and forecasting using R. Now, again just to give you a very quick overview of where we are standing in this week and again just to revise a few ideas about the main topic of cointegration. So, on this week we have been discussing about what exactly does one mean by cointegration and when can we say that two series are cointegrated. So, the idea is that let us say if you have two series x and y then even if x and y are integrated with certain order let us say 1 or something like that. So, in short if you remember we had introduced this notation right.

So, xt and yt would be something like I1. So, this I1 stands for both xt and yt are integrated with an order of 1. Right? But then cointegrated means that if there exists a certain constant or a certain combination of X t and Y t. So, let us say in some sense you are able to combine X t and Y t through some cointegrating vector which is beta such that any linear combination of X t and Y t happens to be Now, again I 0 stands for that the linear combination is stationary right and obviously, I 1 stands for that both X t and Y t individually are integrated with an order of 1 ok. And by the way just again to emphasize on the point that integration means non-stationarity.

So, the moment any time series is integrated with a certain order let us say 1, 2, 3 etcetera. So, it is sort of means that the underlying time series is non-stationary. And again, at the same time, the moment the series converts itself to I0, so I0 stands for stationarity. So, here initially we had taken even the last lecture or the one before that, that Xt and Yt are individually I1, but certain linear combination containing that co-integration vector beta happens to be I0. And in this context, we say that Xt and Yt are co-integrated.

Now again, why exactly do we study co-integration? So, all these ideas we have studied so far. Let us say in the last lecture or the one before that. But again, a very quick explanation as to why this entire idea of co-integration is important is that if Xt and Yt are co-integrated, then there should exist some long-term equilibrium between the two series. So, Xt and Yt.

And whatever deviations happen in the short term—so, let us say, if you remember, we have drawn multiple graphs where, let us say, this is the overall mean, and the series Xt sort of converges or diverges like that, and Yt sort of converges or diverges like that. So, one can actually see that in the long term, they have to both revert to some possible mean structure, right? So, in the longer run, they come to some equilibrium, but of course, in the shorter duration, there might be some deviations from the actual mean, right? Now, in this lecture, as the title suggests, we will go forward by discussing a few tests for cointegration and, of course, a few applications toward the end. Now, firstly, what do you mean by tests for cointegration?

So, all these tests would basically tell us whether Xt and Yt are indeed cointegrated or not. Alright, so before we discuss all these tests that you see in front of you on this slide, we will give a very brief introduction to why testing is required. So, we will introduce co-integration tests now. So, co-integration tests are used to determine whether two or more non-stationary time series have a stable long-term relationship. So, this is the essence here.

So, if XT and YT are co-integrated, they should actually have a long-term equilibrium or a long-term stable relationship. On the other hand, if time series are co-integrated, it means that although they may drift apart in the short term, they tend to move together in the long term. So, this is exactly what we discussed a short while back in the introduction part of today's lecture, and we thoroughly discussed both these points in the last lecture as well. So, in the short term, the two series might deviate from each other, but in the long term, they possess some stable long-term relationship. So, in this context, the following are a few tests in this regard.

So, how does one test if two series, XT and YT, are indeed co-integrated or not? So, again, just to give a few names here. So, the first one is the Engle-Granger two-step test. So, of course, this was given by two people: Engle and Granger. The second one is the Johansen test, right?

Then, the third one is the Phillips-Ouliaris cointegration test. Now again, I understand that all these names are hard to pronounce, but one cannot help it, right? So, the third one is the Phillips-Ouliaris cointegration test. Then, the next one is the Durbin-Watson cointegration test.

And the last one we discuss in today's lecture is called the autoregressive distributed lag, or in short, ARDL bounds test, ok. So, again, just to give you a quick overview once more, all these tests' key idea is to test for the underlying cointegration between x t and y t, ok. So, again, if x t and y t are sort of integrated with some order. So, now, let me in general put I d. So, the order of integration is I d, or rather the order of integration is d, but if there exists a certain cointegration vector beta such that the linear combination of x t and y t becomes So, once you go through this entire process of converting into a linear combination, then if both the series get converted to I 0, not individually.

So, this is one important part: individually, they would be I d or rather I 1, let us say, right, but any linear combination containing this beta vector, which is called a cointegration vector, happens to be I 0. So, then x t and y t would be cointegrated, right now. Now, again, all these tests are meant to sort of test this underlying characteristic: whether x t and y t are indeed cointegrated or not, right? Okay. So, we will start with the first one, which is one of the famous ones.

It is called the Engle-Granger two-step test. Now, what exactly does one mean by Engle-Granger two-step and so on? So, the Engle-Granger test is one of the first methods developed to test for cointegration, specifically between any two variables. And it involves two main steps. So, the first one is the first step: estimating the long-run relationship.

Now, as we have discussed so far, if XT and YT are cointegrated, then there should exist some long-term equilibrium or long-run relationship. And on the other hand, there should be some short-term deviations as well. So, the first step under this Engel-Granger two-step test is to estimate the long-run relationship—but how exactly? So, through a simple regression. So, the equation that you see here in front of you is nothing but a simple regression equation, where alpha is the intercept, beta is the slope, xt is the first series, yt is the second series, and epsilon t is the random error.

$$Y_t = \alpha + \beta X_t + \epsilon_t$$

Okay. All right. So, how do you progress with such a test? You regress one variable on the other (or others, depending on how many series you have) to obtain the residuals, of course, right? And what exactly do you mean by residuals? So, residuals are nothing but the estimated deviations from the long-run equilibrium.

So, here, as discussed already, Xt is one of the series, and Yt is the other series, and we will be regressing Yt on Xt, right? So, beta is the slope, alpha is the intercept, and epsilon t is the random error. So, once we run this regression, we will look at the residuals obtained from this regression, okay? So, this is the first step. And now, the second step is to test for stationarity of these residuals.

So, once you obtain the residuals from the earlier regression equation, then we actually have to test for the stationarity of those residuals. And one can apply the famous augmented Dickey-Fuller test. So, which is also called the ADF test in short, to test if the residuals are indeed stationary or not. So, I can actually apply this ADF test on the residuals epsilon t. And if it turns out that the residuals are stationary, The other term for stationarity, as we have seen throughout this week, is I(0).

So, I(0) stands for a stationary series, right? So, if the residuals are indeed stationary or integrated in order of 0, then one can actually safely say that the variables are co-integrated. So, again, just to summarize very quickly. So, the first step is to generate that regression equation we saw earlier, where you simply form a regression between one of the series, let us say X t, and the other series Y t, using a simple linear regression kind of structure. So, Y t equals alpha plus beta X t plus random error, obviously. Then run the regression, obtain the residuals of the regression, and test if the residuals are indeed stationary or not stationary, okay?

So, here, what could be the null hypothesis and the alternative hypothesis? So, for this Engle-Granger test, the null hypothesis or H0 is that the residuals are not stationary, and if the residuals are not stationary, the series is not co-integrated. Of course, right? Because if X t and Y t or the residuals are not stationary from that regression, then one can actually say that X t and Y t are individually, let us say, I(1) or I(d) in general, okay? So, the null hypothesis is that the residuals are not stationary, which implies that the series is not co-integrated or the two series are not co-integrated. And similarly, the alternative hypothesis or H1 is that the residuals are stationary, so that the series are indeed co-integrated.

So, I guess this is slightly confusing, but one has to keep in mind that a null hypothesis points to a situation where the series is not cointegrated, while the alternative points to a situation where the series is indeed cointegrated. Now, how about a limitation of this test though? So, are there any limitations of this Engle-Granger test? And the answer is yes.

So, this test is only suitable for two variables and may lack power if more than two variables are involved, okay.

So, for example, even the example we discussed in the last slide, we took the example of two variables. So, X_t and Y_t. So, the moment you go beyond two series, let us say X_t and Y_t, one might encounter some problems or some limitations, okay. So, specifically, this Engle-Granger test is kind of superior in a situation where one only has two variables. So, or rather two series, let us say X_t and Y_t, and this may lack power if more than two variables are involved, okay. Okay, so now the second test in the array of tests is called the Johansen test, and this was, of course, given by this person named Johansen.

So, the Johansen test is a more comprehensive test for cointegration, which can handle multiple variables, okay. So, one drawback that we saw for the Engle-Granger test in the last slide was that it cannot perform really well if we extend to more than two variables or more than two series. And this is exactly where the Johansen test comes to the rescue. So, again, just to summarize, the Johansen test is a much more comprehensive test for cointegration, which can handle multiple variables. So, let us say if you have more than two series for some reason.

So, Xt, Yt, Zt, etcetera. So, even in such a situation, the Johansson test proves to be superior. Now, what exactly is the underlying mechanism of this test? So, this Johansson test is based on a vector autoregression, or in short, a VAR model structure, and examines the rank of the cointegration matrix to determine the number of cointegrated relationships. So, in short, in the last week, if you remember, we discussed thoroughly about VAR models, right, or rather vector autoregressive models.

So, this is one of the underlying ingredients of this Johansson test. So, it sort of makes use of a VAR model and examines the rank of the cointegration matrix. So, assume that you have multiple variables, right? For example, Xt, Yt, Zt, Wt, Ut, Vt. So, if you have more than two series, you can actually get a cointegration matrix. So, what you can do is pairwise look at cointegrating factors.

So, whether Xt and Yt are cointegrated, or whether Xt and Zt are cointegrated, or whether Yt and Zt are cointegrated. So, in short, you can get a cointegration matrix which is pairwise. And such a cointegration matrix helps to determine the number of cointegrating relationships. So, out of all these different series, you can actually get hold of all such pairs which might be cointegrated in a way. So, the Johansson test is sort of powerful in that sense that by analyzing the cointegration matrix in a multivariate kind of structure, or rather a VAR kind of structure, you can sort of determine the number of cointegrating relationships among all these different series.

So, hopefully, the idea behind both these tests is clear now. All right. So now, coming back to the actual Johansson test. So, how do you actually take the approach of testing using the Johansson test? So, under this Johansson test, you have two more steps to follow.

So, the first one is the trace test. So, what do you mean by the trace test? So, this trace test tests the hypothesis that the number of cointegrating vectors is less than or equal to some constant R against the alternative of more than R. So, just in the previous slide, we have seen that the Johansson test identifies the cointegration matrix, or rather the cointegrating matrix, and tries to identify how many pairs or how many such series are indeed cointegrated. So, in this part, the first step is to actually test the hypothesis that the number of cointegrating vectors is less than or equal to some constant R against the alternative of more than R.

And in the same spirit, I can actually test using a second approach, which is the maximum eigenvalue test. So, what exactly does one mean by this? So, such a test tests the hypothesis that the number of cointegrating vectors is R against the alternative of R plus 1. So, in a way, both the trace test and the maximum eigenvalue test sort of put forward a constant R, and then the hypothesis is whether the number of cointegrating vectors is less than that threshold, which is R, or more than R. Now, again pointing to what exactly is the null hypothesis for the Johansson test.

So, again, H naught is that there are R co-integrating relationships, right? And on the other hand, the alternative hypothesis or H1 is that there are more than R co-integrating relationships. So, of course, if one relies on H1. So, once you test using the Johansson test and the conclusion is to reject the null, if you are rejecting the null, you have to go with the alternative. And if you are going with the alternative, it sort of means that you have more than R co-integrating relationships among all the C's, which also means that there is some underlying strong cointegration which exists. Alright, so now, similarly, the third one.

So, the third one is the Phillips-Oliaris cointegration test or Phillips-Oliaris cointegration test, right? Now, what exactly is this test like? So, the Phillips-Oliaris test is kind of similar to the Engel-Granger test. It is a residual-based test for cointegration, but uses different test statistics which can be more robust in certain cases. Now, again, just for the second, if you go back to the Engel-Granger test, right?

So, what exactly happened there is that we created that regression, if you remember. So, y_t equals alpha plus beta x_t plus some random error epsilon_t, and through this regression and the behavior of, let us say, the slope coefficient and the intercept coefficients, we kind of argued whether x_t and y_t are co-integrated or not, okay? So, in the same spirit, this Phillips-Odiaris test is slightly different from this Engel-Granger structure. In a sense that this is again a residual-based test for cointegration, but uses different test statistics other than those used by the Engel-Granger test, which can be more robust in certain cases. So, again, just a couple of points to mention here: if you have already done a course on regression, then whatever I will speak now would be kind of more in relationship to all of you—that if you are setting up a regression model like this, let us say between X_t and Y_t.

So, any regression has some assumptions, right? So, what are the assumptions in regression? So, it has to be a linear model, and the errors have to be independent, right? Then, one of the other underlying assumptions is that the errors have to follow a normal distribution. And lastly, the errors should have equal variance.

So, if any of these assumptions is violated, then we say that a simple linear regression or multiple linear regression does not make sense. So, in the same spirit, if you come back to the Philips-O'Leary test, what it does is create a similar regression model but using a different test statistic, which is slightly different from the model written here, and it turns out that it can be robust in certain cases. So, even if some of the assumptions of the underlying regression structure are not met, in such a situation, the Philips-O'Leary test does not fail, right? So, this is the underlying idea of why you should use the Philips-O'Leary test compared to the Engle-Granger test. Of course, not in all situations, but at least in some where you require that robustness.

Make sense? Now, again as before, what is the null hypothesis or H0? So, H0 stands for no cointegration existing between the series, and similarly, the alternative hypothesis or H1 is that cointegration exists between the series, right? So, in all the tests we have seen so far, the null hypothesis is kind of constant, which points to the fact that there is no cointegration between the series, right? And similarly, the alternative points to the fact that there is some cointegration existing between the series.

Now, by the way, the name of this test statistic is nothing but the Phillips-Ouliaris test statistic. So, this underlying test uses the Phillips-Ouliaris test statistic to evaluate the stationarity of the residuals from a long-run regression. So, let us say you create this regression with slight modification, you obtain the residuals, but then the test statistic under the Phillips-Ouliaris is slightly different than the Engle-Granger. Make sense? All right. So, the next one is called the Durbin-Watson cointegration test.

So, what exactly goes in this is that the Durbin-Watson test is not a direct cointegration test as written down here. But in other words, it can help in testing for spurious relationships. Now again, I am pretty sure that in the very first lecture of this week, we talked about spurious regressions. Now again, just to give you a quick overview of what you mean by spurious regressions or spurious correlations or relationships is that even though two variables, let us say x t and y t, are not at all related, but once you create a normal regression model between XT and YT or you try to find out the correlation between XT and YT, it turns out that the correlation value is high or the regression is also seeing a very good fit. So, such relationships are called spurious because even though XT and YT have no relationship between themselves, the correlation or the underlying regression is pretty strong.

So, in such cases where a regression's Durbin-Watson statistic is close to 0, it suggests the possibility of a spurious regression. So, in situations where you have spurious regressions, I can actually find out the value of the Durbin-Watson statistic. And if the value of the Durbin-Watson statistic is close to 0, it suggests the possibility of a spurious regression, indicating that the variables may not be cointegrated. So, again, you have to go back to all the examples you have discussed in the prior lectures this week that the moment a spurious regression exists between Xt and Yt, it sort of indicates that the variables may not be cointegrated. Secondly, if the Durbin-Watson statistic is above some threshold, so let us say generally above 1.5 or 2, it suggests that the cointegration relationship might exist between Xt and Yt, okay.

So, a couple of things here—the underlying idea is about spurious regression. So, if spurious regression exists between Xt and Yt, one can actually proceed by using a test such as the Durbin-Watson test, right? And one should determine the value of this Durbin-Watson test statistic. If the value of the test statistic is close to 0, it suggests the possibility of spurious regression, indicating that the variables may not be cointegrated. On the other hand, if the Durbin-Watson statistic is above some threshold, let's say 1.5 or 2 generally, then this suggests that a cointegration relationship might exist, okay? So, let me write down Ts. So, Ts stands for the test statistic value. So, if the test statistic value is close to 0, then cointegration does not exist, while if the test statistic value is greater than 0—and generally within the threshold of 1.5 to 2—then one might expect that some cointegration exists, okay? All right. Now, the last one is called the autoregressive distributed lag, or in short, the ARDL bounds test.

So, again, just to quickly remind you of how many tests we have seen so far. So, again, let me quickly sift through all the previous slides. So, the very recent one was Durbin-Watson. Just before that, we discussed Phillips-Ouliaris. Just before that, we discussed Johansen, and the very first one was Engle-Granger.

So, you have seen four tests so far. So, Engle-Granger, then Johansen, then Philips-Ouliaris, and then lastly, Durbin-Watson. And now, the very last one we will discuss is the ARDL bounds test, or autoregressive distributed lag bounds test. Now, what exactly goes inside this ARDL bounds test is that the ARDL bounds test is suitable when the underlying series are of mixed integration orders. So, firstly, one should understand what you mean by mixed integration orders.

So, some are I(0), some are I(1). So, I will give you an example. So, let us say you are taking multiple series again. So, let us say X_t, Y_t, Z_t, and let us say W_t. OK.

And then, let us say you see combinations between all these series. So, X_t, Y_t, Z_t, W_t, and it so turns out that let us say X_t and Y_t are individually I(1), and probably let us say Z_t might be I(0), W_t again might be I(1), right? So, you sort of have a mixed integration order. So, some of the series are I(0), and some of the other series are I(1) in this entire set, OK. So, let me write down that some are I(0), some are I(1).

So, there might be one problem when you handle such a set right. I mean let us say if X t and Y t are I 1. So, there should exist some linear combination of X t and Y t such that it is I 0 which means that X t and Y t are co-integrated right. But in the beginning itself if let us say Z t is I 0 ok. So, let us say Z t is assumed to be I 0 while Y t is assumed to be I 1 as an example right.

So, Z t is I 0 Y t is I 1. If at all there exists any linear combination between yt and zt, which is I0, can we still say that they are co-integrated? Because zt itself is I0, isn't it? So, all such intricacies come into picture when some of them are I1, some of them are I0, some might be I2, etc. So, in all such situations this autoregressive distributed lag bounce test is kind of suitable.

So, again just to summarize this ARDL bounce test is suitable when the underlying series are of mixed integration order. So, some are I 0 some are I 1 right. The ARDL bounce test provides critical values to determine if a long run relationship exists between the variables or not. So, as seen before for all the other tests also. So, again what one gets out of this ARDL bounce test is that it sort of produces some critical values to again determine if a long run relationship exists between the variables or not.

So, again, what exactly is the null hypothesis? So, H naught is no cointegration. So, there does not exist any long-run relationship between the series, and the alternative is that cointegration exists, as before. So, the ARDL approach is widely used in fields like economics and finance, especially when stationarity requirements are mixed, and it is flexible with various sample sizes. So, these are some advantages of using the ARDL bounds test: such an approach can be widely used in areas like economics and finance, especially when stationarity requirements are mixed, and it is flexible with various sample sizes.

Okay, so right toward the end, we will discuss a couple of practical examples quickly and then close today's session. So, the first combination is house prices and rent prices. So, the Engel-Granger test helps determine if house prices and rental prices share a long-term relationship. So, cointegration between these two suggests that, despite any short-term imbalances, house prices should align with rental values over time, providing insights for pricing in real estate markets. Then, the second example could be something like commercial real estate prices and economic indicators.

So, what exactly goes inside this one is that cointegration between commercial real estate prices and broader economic indicators, such as, let's say, GDP or employment, etc., suggests that real estate prices move with the economy, and such a relationship helps investors and policymakers assess the sensitivity of real estate prices to economic cycles. Then, the third one could be, let's say, healthcare spending and GDP. So, in healthcare, the test helps determine if healthcare spending is aligned with economic growth. If healthcare expenditures and GDP are cointegrated, this would imply that spending grows with the economy in the long term, which can inform sustainable budgeting and health policy decisions. And the last one could be drug prices and R&D costs.

So, pharmaceutical companies may analyze whether drug prices and research and development costs, or in short, R&D costs, are co-integrated or not. If they are co-integrated—or rather, if confirmed—it suggests that R&D costs will reflect long-term drug prices, validating the cost recovery model for pharmaceutical pricing strategies. So, in short, just to summarize, all these examples point to the fact that one should actually test for co-integration in the first place. So, if a couple of series are co-integrated, one can actually go forward and analyze multiple other things and then provide some indicators as to why that might happen or some solutions to handle such co-integrations in practice.

Now, in the next lecture, we will discuss a few more ideas about co-integration. And again, the lecture after that will be completely practical.

Thank you.