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 37: Error Correction Models

Hello all, welcome to this course on time series modeling and forecasting using R. Now, again, a very quick refresher as always: where are we standing this week? So, we started the last session, which was the first session this week, with a broad idea called co-integration, right? There, we specifically mentioned that the idea of co-integration between two vectors—let's say Xt and Yt—firstly, what exactly do you mean by co-integration? But even before that, we started the session, if you remember, by developing an understanding of the term called spurious regression, right. And there, we took up multiple examples of what would happen if Xt and Yt, even though they are not related, the correlation structure between Xt and Yt or the regression model between Xt and Yt shows otherwise, right.

So, such relationships are called spurious, or the underlying regression between Yt and Xt is called spurious regression, right. And then we sort of broadened the idea by developing what you mean by co-integration. So, again, just a very quick revision: even though Xt and Yt are integrated. So, the idea of integration is what? So, Xt and Yt are I(1).

So, I(1) means what? So, there has to be some trend in both series. So, the underlying Xt and Yt are not stationary, of course. So, both Xt and Yt are non-stationary since they contain some trend, but if there exists some linear combination or a co-integrated vector between Xt and Yt. So, again, if you remember, we brought in the idea of a beta vector, right, such that beta prime multiplied by Yt, Xt—so this entire collection becomes stationary or rather I(0)—then we can say that Xt and Yt are co-integrated, right.

$$\nabla Y_t = \alpha(\beta X_{t-1} - Y_{t-1}) + \gamma \nabla X_t + \epsilon_t,$$

And of course, towards the end, we studied why exactly we care about co-integration and a couple of examples, right. So, one simple point that we will start today's session with is to mention that if XT and YT are co-integrated, a couple of things happen. So, in the short run, they could deviate from each other, but there has to be some long-term equilibrium between XT and YT. So, in the long run, XT and YT should converge. So, this is the idea of co-integration, right.

So, now in the same spirit, we will start the second session by mentioning what you mean by error correction models or, in short, ECM, right. So, in short, this is called ECM or error correction models. So, again, this is a kind of extension or the very next step to developing the idea of co-integration. So, once we know that XT and YT are co-integrated, how do you deal with that, okay? So, here ECM comes into play.

So, firstly, we will try to understand what exactly this error correction model is. So, an error correction model or, in short, ECM is a type of time series model used to analyze the relationship between two or more non-stationary variables which share a stable long-term equilibrium relationship. So, again, if you repeat this definition or if you repeat these lines a number of times in your mind or in your head, you can understand that the idea of co-integration is the underlying idea here, right. Because once you study or once you understand that XT and YT are co-integrated, then how you deal with such co-integrations, and this is exactly where the ECM model comes into play, right. So, again, just to repeat, an error correction model or, in short, ECM is a type of time series model used to analyze the relationship between two or more non-stationary variables.

So, again, when we say non-stationary, we are kind of hinting at that I(1) kind of process, right. So, I(1) again—if you borrow some ideas from the last session—there we saw that both X(t) and Y(t) were I(1). So, since they were integrated with an order of 1, hence both were non-stationary. So, two or more non-stationary variables which share a stable long-term equilibrium relationship, okay. Similarly, if two variables move together over time but can temporarily diverge in shorter spans, ECM helps us to understand and model both their long-term and short-term interactions.

So, again, whatever is written here or whatever you see here kind of points exactly to the idea of cointegration, right? Because when cointegration exists, there also exists a long-term equilibrium, and in shorter terms, both X(t) and Y(t) can deviate, of course, okay. So, similarly, if two variables move together over time but can temporarily diverge, ECM helps us to understand and model both their long-term as well as short-term interactions or short-term dependencies. So, of course, now we will try to understand what exactly this ECM model is like and how the ECM model is actually helpful in correcting cointegration or trying to model the underlying cointegration. So, firstly, why ECM? So, when two variables are cointegrated—now again, this is just a quick revision—they move together in the longer run.

So, when two variables are cointegrated, any deviation from their shared equilibrium is temporary, right? I will give you a very simple example. So, let us say if you have two assets: the first asset is moving like that, right, and similarly, the second asset is probably moving like this, right. So, again, if you see that in the shorter run, they can actually deviate. So, let us say the second asset goes like that, while the first asset temporarily might go like that.

So, can you see that both the assets have created a small gap here. So, even though they have a common mean, but then the first asset in a shorter term has risen while the lower one has dropped right and hence you can see this gap, but eventually if X t and Y t or if these two assets are co-integrated what would happen in the longer run is that again they would converge to the mean. So, the first asset has to drop back and then the second asset has to go. So, in a shorter run one can actually capture some deviations, but both X t and Y t or in this example both these assets or both these examples share a long term equilibrium all right. So, any deviation from their shared equilibrium is only temporary.

So, or rather in a short term. So, an ECM sort of corrects these deviations by adjusting the model to pull the variables back towards the equilibrium over a time. So, if you apply this ECM on this collection of both the assets, the error correction model, then the ECM has a tendency of pulling back the risen asset back to the mean and rising the dropped asset again back to the mean, so that both the assets again converge to the same mean. So, finally ECM is essential in areas like finance and economics because it provides a way to capture both the long-term stability and short-term adjustments between related time series making predictions more accurate and theoretically sound. So, in general, the ECM has the capability or ECM has the power of modelling or taking care of both the long-term equilibrium as well as short-term deviations right.

And of course, by applying this ECM on this collection. So, predictions would be superior, forecasts would be superior right and estimations would be more accurate and theoretically sound. So, this is exactly why we require ECM right? So, since XT and YT are co-integrated in a shorter time span they have a tendency of deviating from the mean.

So, if you apply an ECM or a rather proper ECM, then ECM has a tendency of pulling back both the series towards its mean such that they sort of retain their long-term equilibrium ok.

So, firstly, how ECM works, right? So, an ECM separates changes in a variable into two things. So, short-term effects and long-term correction effects, right? So, since ECM has the capability of capturing both the short-term deviations and the long-term equilibrium or deviations, the first thing that ECM does is separate changes in a variable into two aspects. So, the first one is short-term effects.

So, what exactly do we mean by short-term effects? So, short-term effects mean that it captures the immediate impact of changes in related variables in a shorter time span, and on the other hand, long-term correction. So, long-term correction adjusts for deviations from the equilibrium relationship, pulling the variables back in line over a longer time. So, again, as seen before, let us say the first asset is moving downward, and suddenly it deviates from the mean, while the second asset is moving parallel to the first asset, but then just in the short term, it can deviate. So, let us say it drops.

So, a gap is created, which is this one. So, if you sort of model this collection using an ECM, right? So, ECM, again as discussed earlier, has a tendency of moving. Sort of pulling both the series back toward the mean and then correcting those short-term gaps. So, now we will understand, or we will try to delve deeper into the ECM characteristics and the structure of ECM.

So, what exactly are the key components of ECM? So, the first term is called the error correction term or ECT. So, the first component or the first key component under ECM is ECT, which stands for error correction term. And what is the role of ECT? So, ECT simply measures the gap between the current value and the long-term equilibrium value.

So, ECT is nothing, but—so let me go back a slide since we have already drawn the plot there. So, what ECT does is measure the gap between short-term deviations and the long-term equilibrium value. So, here the gap between the long-term mean and the short-term deviations is nothing but the gap between these two peaks and the long-term mean, right? And this is exactly what the ECT does—okay, this is exactly what the ECT does. So, ECT measures the gap between the current value and the long-term equilibrium value. So, in other words, if the ECT is large, the model will make a larger adjustment in the next period, right? And similarly, if the ECT is small or not that large, the model will make smaller adjustments in the next period.

And at the same time, you have short-term coefficients. So, what do short-term coefficients tell you? So, short-term coefficients show the immediate effect of changes in one variable on the others without needing to consider the equilibrium relationship. So, on one hand, you have the ECT term or the error correction term, which measures the gap between the current value and the long-term equilibrium value. And on the other hand, you have some short-term coefficients, which show the immediate effect of changes in one variable on the others without needing to consider the long-term equilibrium relationship, right? So, we will start with a small example.

So, let us say suppose we are modeling the relationship between inflation and interest rates ok. So, let us say we want to model collectively the relationship or the interdependency between inflation and interest rates which tend to move together in the longer run. So, in this regard an ECM would allow us to model the short-term changes in interest rates based on recent changes in inflation. So, the first point is the ECM would allow us to model any short-term deviations or short-term changes in let us say interest rates based on recent changes in the other variable which is inflation right. And second point is the ECM would also allow us to capture the error correction process where a high interest rate could slowly adjust downward if inflation is low.

So, again just to summarize this slide takes up a particular example a concrete example from economics let us say modeling the relationship between inflation and interest rates and particularly tells us that how the ECM comes into play here right. So, ECM can sort of control two things here. So, ECM would allow one to model any short-term deviations between the two variables. So, interest rates and inflation. And also, to capture the error correction process wherein a high interest rate could slowly adjust downwards if inflation is low.

So, again similar to the last plot. So, let us say this is the interest rate right? So, let us say I will draw a hypothetical plot. So, interest rates are rising and again they are kind of stagnating here and let us say inflation is behaving like that in a horizontal manner right? So, here since the interest rates and inflation are co-integrated assuming.

So, what the ECM would do is correct all the short-term deviations and ensure that the interest rates drop back to their means, which would be closer to the inflation values, right? So, ECM would cover both the short-term deviations as well as try to control the long-term equilibrium, okay? So, now pay attention to what exactly the ECM equation structure looks like, right? So, let us say again, if you have two variables for a two-variable ECM involving two variables—let us say x_t and y_t—the model, or the ECM model, can be expressed as this. So, on the left-hand side, we have ∇y_t.

On the right-hand side, we have a couple of coefficients. So, alpha, beta, and gamma operate on different variations of X_t and their lags, and Y_t and their lags. For example, alpha multiplied by beta multiplied by X_t minus 1 minus Y_t minus 1. So, these two are lag variables of X_t and Y_t, plus another constant gamma applied on ∇X_t, plus a usual random error. Okay. And here, ∇ is the usual differencing operator; this term that you see inside the bracket.

So, beta multiplied by X_t minus 1 minus Y_t minus 1 is the error correction term. And why exactly? So, again, can you visualize that this gives you the difference between the two series at one prior lag? So, X_t minus 1 minus Y_t minus 1. So, this sort of captures the short-term deviations between X_t and Y_t.

And of course, you have this extra coefficient for some shifting here and there, right. So, this extra beta sort of gives you some room to play around with, right. But this entire term that you see inside this bracket here is called the error correction term, which we discussed a couple of slides back. So, this is nothing but the ECT, which represents the deviation from the long-term equilibrium at time t minus 1. Similarly, what is alpha?

So, alpha is the speed of adjustment coefficient, which adjusts y t in response to deviations from the equilibrium. So, alpha takes values and tells you how swiftly or how slowly the two series should come back in the long run. So, if x t and y t deviate drastically in the short term, one can actually put a higher or lower value of alpha to bring both series, x t and y t, together in the long run, wherein gamma is the short-term coefficient. So, the gamma coefficient that you see here captures any short-term deviation. So, in general, you have three coefficients: alpha, beta, and gamma.

So, this beta is attached to the ECT, this alpha is the speed of adjustment coefficient, and gamma is nothing but the short-term coefficient, and this is the ECM equation. So, a bit more on the variable. So, the first idea, or the first aspect, or the first component in the equation we saw earlier, is the error correction term, or the ECT. So, what exactly is the role of the ECT? So, this term represents the degree of equilibrium at time t minus 1.

So, the value is high, which means the series is far from equilibrium, requiring larger adjustments to move back to equilibrium in the next period, okay. So, similarly, again, if you go back a slide very quickly. So, this term gives you the ECT. So, if the value of this

ECT is large or high, it sorts of points to or indicates that the series is far from equilibrium, requiring larger adjustments to move back to its mean or to its equilibrium in the next period. And then the next thing is alpha.

So, alpha is the speed of adjustment coefficient. So, this alpha determines how quickly the system corrects back to equilibrium. So, it sort of captures the speed of how slowly or quickly the two series come back to equilibrium. So, clearly, a higher value of alpha indicates faster adjustment, while a lower value of alpha suggests a slow correction process. So, by the way, the first point is pointing toward this beta coefficient, right?

So, because beta is attached to that ECT, right? And then the second point points to alpha. So, alpha captures the speed of adjustment. So, a higher value of alpha means faster adjustment, and a lower value of alpha means slower adjustment. And the last one is gamma. So, gamma points to short-term dynamics.

So, the term gamma nabla xt models the short-term impact of changes in x and y, independent of the long-term relationship. So, what happens in shorter durations or what happens in shorter time spans is captured by this term, which is gamma nabla xt, OK? And by the way, this gamma nabla xt is independent of what goes on in the long term because, for controlling the long-term behavior, we already have beta and alpha with us, right? So, there should not be any mixing between gamma and either alpha or beta. So, gamma is focused on the short term, while beta is focused on the long term, while alpha is focused on how slow or how quick the adjustment is taking place.

Make sense? So, a bit more understanding about all the variables—alpha, beta, and gamma. So, now the next thing is interpretation. So, the first component is ECT, or the error correction term. So, if ECT has a significant coefficient—now again, what exactly is the coefficient here?

So, the coefficient is nothing but beta. So, if the ECT coefficient beta is significant, it means yt adjusts to bring the system back to equilibrium whenever it deviates, OK? So, the first component is ECT. So, the coefficient attached to ECT, which is beta, is significant. It sort of tells you that yt adjusts accordingly to bring the system back to equilibrium whenever it deviates away from xt.

And the secondly short-term coefficient. So, if the short-term coefficients example gamma because gamma are nothing but the short-term coefficients. So, if all the short-term coefficients in the ECM are significant, it suggests that changes in X have immediate effects on Y in the short term. So, hopefully all these individual ideas would be clear right. I mean beta is to capture the long run equilibrium right and gamma is to capture the short-term deviations and alpha is to capture the speed of the adjustment ok.

So, this is this is a short interpretation of what do you mean if the ECT has a significant coefficient or what do you mean if gamma is significant and so on and so forth basically. Okay. So, now, the next thing is we will outline a few steps to fit an ECM. So, how exactly do we go ahead and then fit an ECM? So, the very first thing that one has to ensure is one has to perform some tests for non-stationarity.

So, one can either use an augmented Dickey-Fuller or in short ADF or similarly other tests to ensure that each of the variables is I1, right? Because if XT and YT to begin with are stationery, there is no argument of co-integration right. So, if both XT and YT are stationary individually then there exists no non-stationarity and we do not require an ECM there right. I mean one can actually model them using individual white noise processes or if you want to put forward a vector model then let us say VMA or something like that.

So, the very first test is to ensure that both the individual series is indeed integrated. So, I1 or I2 something like that. So, for that we require some tests for non-stationarity. So, you know ADF test or some other test etcetera. Now, the second test is test for cointegration.

So, one can actually conduct a Johansson cointegration test. So, this was a test given by Johansson. So, one can actually conduct a Johansson cointegration test to determine if a cointegrated relationship exists indeed or not right. So, once you identify non-stationarity, but even though Xt and Yt are integrated, that would not mean that XT and YT are co-integrated, is not it? So, individually XT and YT could be integrated, right?

But does there exist a co-integration between XT and YT? that has to be checked, right? So, the second step is once you ensure that XT and YT are non-stationary to begin with, then on top of that is there some co-integration existing between XT and YT? And in this period one can actually apply this Johansson co-integration test to determine if a co-integrated relationship exists, right? And if cointegration exists or if cointegration is found proceed with an ECM ok. So, so specifically ECM is predominant with cointegrated variables only in the in the not just non stationary.

There has to be some cointegration, or there has to be some underlying cointegration between XT and YT for a successful ECM to run through, all right. And now, of course, once you proceed with ECM, you have to estimate all the coefficients. So, estimate the cointegration equation to find the equilibrium relationship, usually using OLS or ordinary least squares. And lastly, estimate the ECM with the error correction term from the cointegration equation. So, now that you have ensured that the series is non-stationary and there exists some cointegration between Xt and Yt, then how do you develop or propose the ECM equation is of importance, right?

And the ECM equation we saw in the earlier slide, right. I mean, you have to estimate alpha, you have to estimate beta, and you have to estimate gamma, OK? And once alpha, beta, and gamma are estimated, you actually get the complete estimated ECM equation, which can be used to model any cointegration between Xt and Yt, all right. So, now the next thing is a couple of advantages and disadvantages of using an ECM. So, a couple of advantages here.

So, the first one is ECMs combine short-term and long-term dynamics, providing more accurate models for cointegrated time series. So, ECM has the capability, as discussed earlier, to capture any short-term deviations or long-term dynamics or long-term equilibriums—or rather disequilibriums—providing a more accurate model for cointegrated time series. The second advantage is that they are effective in systems where relationships drift but maintain an equilibrium over time. So, even though one can actually see or sense some short-term deviations or short-term drifts, overall, in the longer run, there has to be some equilibrium. And such ideas or such tendencies are very common in, let us say, economic or financial data, etc.

But what about limitations? So, ECM is again not completely foolproof. So, there have to be some limitations as well. So, the first limitation is that ECMs require the variables to be co-integrated, of course. As discussed, multiple times throughout this lecture, one has to ensure non-stationarity and co-integration as well.

So, ECMs require that variables should be co-integrated. So, they are unsuitable for time series without a long-term relationship. So, again, just to summarize or stress this fact: even though XT and YT are individually non-stationary, cointegration ensures that there is some long-term equilibrium between XT and YT. And the whole idea of deploying an ECM is to capture that long-term equilibrium, right? So, if cointegration does not exist, then ECM is not suitable for time series without a long-term relationship.

And then the second one is if the model is misspecified. So, what do you mean by misspecification? So, for example, incorrect cointegration rank or including too many coefficients. So, let us say multiple gamma coefficients, multiple beta coefficients, etc. Results can be misleading, so the ECM equation as such is not that easy to control, right?

So, there could be some misspecified models which could arrive at some errors in their results. So, a couple of limitations of using ECM exist, but at the same time, we have some advantages where cointegration exists, okay? So, now, lastly, as always, we will discuss some practical examples, right, and then close this session. So, the first one is stock prices and dividends. So, this is one area where ECM could be correctly deployed.

So, ECMs are useful in analyzing stock price behavior relative to their dividends, where prices might diverge in the short term from dividend-based valuations, yet a cointegrated long-term relationship persists. So, the first example is stock prices and dividends. So, ECMs are useful in analyzing stock price behavior relative to dividends, where prices might diverge in the short term from dividend-based valuations. Yet, if such a thing exists, then a cointegrated long-term relationship persists. And as you all know by this time, whenever a cointegrated or long-term relationship persists, then ECM is kind of very helpful.

As a second example, we can talk about interest rates and inflation. So, such an effect is, by the way, called the Fisher effect. So, the relationship between interest rates and inflation. So, in this period, ECMs capture the long-term relationship between nominal interest rates and inflation, allowing short-term deviations to correct over time. And such a model helps central banks and investors understand inflation expectations in relation to interest rates.

So, here the whole idea is how do you capture the dependencies between interest rates and inflation, and even further, how do you capture the short-term deviations between the two as well as the long-term equilibrium between interest rates and inflation. So, in this regard, one can actually make use of ECMs very predominantly, right? Thirdly, oil prices and exchange rates. So, ECMs model the relationship between oil prices and exchange rates, especially for oil-exporting countries. And since oil prices can temporarily deviate from currency values, an ECM can show the adjustment process very nicely, right?

So, the third example is slightly different from the energy market. So, oil prices and exchange rates, and lastly, electricity demand and temperature. So, this is again one more area where one can actually try to model using an ECM because surely there would be some cointegration existing. So, energy providers use ECMs to model how electricity demand relates to temperature changes. So, for example, short-term spikes in demand occur during extreme weather, but demand typically returns to average levels following long-term equilibrium trends.

So, again, just to summarize this entire idea. So, what could happen here is that in the short run, let us say you have two series again plotted side by side. So, and then let us say the top one is electricity demand and then the bottom one is temperature. So, for a rising electricity demand, right, there could be a rise in temperatures also. So, for example, in summers, one can have more demand for electricity; in winters, where temperatures are low, one does not have that much demand for, let us say, ACs or fans, etcetera.

So, as and when the electricity demand rises, the temperature also rises, but there could be some short-term deviations, right? I mean, let us say even if the temperature drops for a short-term duration, even though the temperature is dropping, the energy demand can still go on rising. So, maybe because people are kind of using both ACs as well as geysers, right? So, there could be these short-term deviations between energy demand and temperatures. So, in this regard, if one applies an appropriate ECM. So, the ECM would ensure that the energy demand or the electricity demand kind of converges back to the mean as and when the temperature also rises.

So, in the longer run, there will be some equilibrium between electricity demand and temperature. So, hence, just to conclude this session, broadly speaking, the idea of cointegration we saw in the earlier lecture, and then the idea of modeling such cointegration using ECM or error correction models, we saw in this particular lecture. Now, in the subsequent lecture this week, we will delve deeper into the ideas of, let us say, estimation, and then a few more examples of cointegration, a few more tests of cointegration, let us say. And down the line, we will close this session or close this week with a practical session in R.

Thank you.