Time Series Modelling and Forecasting with Applications in R

Prof. Sudeep Bapat

Shailesh J. Mehta School of Management

Indian Institute of Technology Bombay

Department of Applied Statistics/Finance

Week 11

Lecture 52: Regimes and Nonlinear Models

Hello all, welcome to this course on time series modeling and forecasting using R. So, this week the focus area that we are talking about broadly speaking is non-linear time series models. So, all the models we have studied so far. So, be it autoregressive, moving average or ARMA or rather ARIMA, SARIMA. So, all these were examples of linear structures or linear time series models wherein one can actually encounter.

a practical application and one can put forward a linear sort of a structure to model that ok. But of course, as probably all of you might know that let us say if you bring any practical area of concern where time series modeling and forecasting would be required. So, nonlinearity is kind of predominant right. So, one need not expect every time that the underlying practical data set has to be linear ok. So, we require some sort of a shift from a linear structure to capture the non-linearity from the time series process as well.

And the very first session, in fact, the last session or rather the first session of this week was focusing more on the introduction part or rather the motivation part as to why should one make the transition from linear time series models to non-linear time series models. And even the last session, we talked about several examples of couple of very important non-linear time series models, they say threshold models, etc. And there, if you sort of remember, we talked about an idea called as regimes. So, again, just a very short kind of an overview as to what do you mean by regimes first. And then, of course, as you see, the title of this session is a few more aspects on regimes and a few non-linear models to specify.

Okay. But before we start with anything new again as discussed, a short while back that what exactly does one mean by regimes. So, the word regimes mean nothing but states ok. So, imagine that you have a practical time series split up into different states or different phases. So, all these states are called as different regimes of the time series and one very differentiating point is that any two states or any two regimes are not exactly the same ok.

And all these regimes are kind of divided by threshold variables or rather threshold values or threshold levels. So, as whenever a time series transitions from one state to another or rather one regime to another you see a dynamic shift ok. So, any two regimes or any two states or any two phases need not be exactly equal and hence one requires slightly differentiated sort of an approach to model inside each of the regimes. So, now in today's lecture we will talk more about regimes and a few more non-linear models to start with and then and towards the end of this week I am pretty sure that we will discuss lot of examples and then again we will sort of link back all the examples to the first initial motivation as to why should one in fact transition from this entire area of linear time series processes to non-linear time series processes. So, the first slide is modeling regimes.

So, again the initial part will discuss a few kinds of regimes and then in front of you you have different kinds of regimes starting with the very first one which is called as threshold-based regimes. So, these are kind of very easy to understand and then down the line we will see that you may encounter some slightly advanced kind of regimes which we will discuss very soon, but then threshold-based regimes fall But then threshold based regimes form the base structure and then some of the advanced regimes are kind of extensions of these threshold based regimes. So, before we start with anything very deep into technicalities etcetera, we will define what exactly does one mean by regime. So, regime definition from a mathematical perspective or from a technical perspective.

So, let us say that you have a very small illustration, very easy also, and then here we are considering two regimes. So, regime 1 and regime 2. And let us say gamma is the threshold value. So, this gamma that you see here is the threshold value or the threshold level, and y t minus d is the actual time series progression. So, inside regime 1, what is happening?

$$\text{Regime 1: } y_{t-d} < \gamma$$

$$\text{Regime 2: } y_{t-d} > \gamma$$

So, inside regime 1, we will say that the time series y t and obviously all of its past lags are inside regime 1. As always, as this y t minus d is less than gamma, ok. So, wherever this time series value in terms of t minus d, which is nothing but the lagged value of the time series itself, happens to be strictly less than gamma, we will say that the underlying time series has not shifted yet or, in other words, the underlying time series is still existing in regime 1, ok. But then, since this gamma is the threshold value, as in when the time series transitions from, let us say, being less than gamma to more than gamma, we have a regime shift here. So, regime 2 is nothing but all those values of the time series where y t minus d happens to be bigger than gamma.

So, I guess it is not really difficult to understand the idea of a threshold-based regime. So, again, just to summarize, there should be one threshold value, which is equivalent to gamma here, and then we will sort of compare all the lag values of the time series with that threshold value, which is gamma. And whenever all the lagged values happen to be less than the threshold value, we will say that the time series is sort of performing inside regime 1. And as in when the time series transitions from one regime to another, or rather the lagged values become more than gamma, we will say that the regime has shifted. So, again, just to summarize what exactly the notations mean.

So, as always, y t minus d is nothing but the lagged value of the series or another variable, and gamma in this context is the threshold value, ok. And these are some of the examples where a threshold-based regime can be seen. So, let us say TAR or SETAR sort of models, ok. By the way, the TAR model was kind of discussed very briefly even in the last session. So, the full form is threshold autoregressive structure, ok.

And on top of that, this SETAR model is nothing but an extension of a TAR model. Now, coming to the second kind of regime. So, the second kind of regime is called a smooth transition regime. By the way, we have discussed these three classes in the last session. So, the very last slide of the last session, where if you remember very vaguely, we had listed down all the three

kinds of regimes there, right? So, starting with threshold-based regimes, and the second one was obviously smooth transition regimes, and then lastly, we also discussed probability-based regimes where we talked about Markov switching models and all those, ok. Now, this is the second one. So, smooth transition regimes, and then Under the smooth transition regimes, there has to be some transition function, ok.

$$G(z_{t-d}; \gamma, c) = \frac{1}{1 + e^{-\gamma(z_{t-d} - c)}}$$

Where z_{t-d} is the threshold variable, γ is the slope parameter, and c is the threshold value.

So, this function G that you see, given by capital G, depends on multiple parameters. We will discuss each and every one in detail—do not worry—but the predominant parameter is nothing but the threshold value. So, Z t minus d is nothing but the threshold variable; gamma, in this case, is the slope parameter, okay? And then, C in this case is the threshold value. So, again, just to summarize: Z t minus d is nothing but the threshold variable or the lagged values of the time series. This gamma is nothing but the slope parameter, and then C is nothing but the threshold value. And here, you can see that the transition function is not very difficult to understand; it is rather 1 divided by 1 plus e to the power of minus gamma multiplied by Z t minus d minus C.

So, again, probably one strong suggestion could be to pause the video and then subsume the information given in the transition function. So, try to understand all the components that go inside the transition function, okay? Now, again, as discussed here. So, the transition function majorly depends on two parameters. So, gamma and C, where gamma is nothing but the slope parameter—similar to any regression equation you might have seen—and C, on the other hand, is called the threshold value.

So, in this context here, if you focus on this piece of the transition function, then you can basically argue that if you compare this difference between Z t minus d and C to, let us say, 0—whether the difference is less than 0 or more than 0—we could have a transition. So, in that case, we will say that C is nothing but the threshold value. And some of the examples in this context, or in this spirit, are STAR models. So, we will talk about STAR models in detail, probably in upcoming sessions.

So, we started with one kind of regime, which was threshold-based regimes, which is kind of on the easier side to understand. But what would happen if we could include the regimes inside a function, sort of a thing? So, this gives rise to smooth transition regimes. And now the next kind, or the third kind, is called latent state-based regimes. So, what do you mean by that?

So, latent state-based regimes. So, here the main ingredient has to be some sort of a Markov transition matrix. So, this capital P structure, or this capital P matrix that you see in front of you, comprises four cells. So, P11, P12, P21, and P22 are called nothing but the Markov transition matrix. And again, where each of the inside notations, let us say p i j to be specific, is nothing but the probability of transitioning from regime i to j. So, all these individual entries, for example, p 1 2.

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$$

So, what do you mean by p 1 2? So, again, you can pause the video and try to think in your mind as to what would be the definition of p 1 2 or how would we define, let us say, p 2 2, for example. So, when it comes to p 1 2, p 1 2 is nothing but the probability of transitioning from regime 1 to regime 2. So, just to mention very quickly here that if somebody is transitioning from regime 1 to regime 2, then what exactly is the probability of transitioning from regime 1 to regime 2? So, in a way, all these four entries in this matrix are nothing but some sort of probability.

And one major example in this spirit is nothing but called a Markov switching model. So, in this context, so far we have studied three different sorts of regimes. So, threshold-based regimes, then latent state-based regimes, which is the ingredient of the current slide where we throw in some probabilities, and just in the last slide we talked about smooth transition matrices or smooth transition regimes where you require some transition function to sort of decide on the transitioning between one regime to another. So, now the next thing we will talk about very briefly is what exactly are the benefits or what do we get from doing all this, or rather modeling the regimes and so on and so forth. So, the very first advantage is improved forecasting.

So, what do you mean by that? So, by modeling distinct regimes, these models can in fact improve accuracy, especially for systems with abrupt changes. Now again, just to sort of give you the same visual example we talked about in the last session, so let us say if you have a time series which is transitioning like that, and let us say if you have a threshold value which is like that, and the moment the time series sort of crosses the threshold value, you can see that the time series behaves as if it is highly volatile here. Okay, which means that this could be a threshold value-based regime. So, where this is the threshold value, let us say gamma or something like that.

So, as in where the time series crosses that threshold value, you can see very high volatility in the underlying time series. And if you are able to capture the entire system very nicely using some of the regime modeling ideas, then it can sort of lead to some improved forecasting. And why exactly? Because what we will do here is that we will try to model these separately. So, this would be regime 1 in our case, and then this might be regime 2. So, if you are able to propose sort of differentiated models for each of the regimes and then overall we try to connect them, it can give us improved forecasting rather than putting forward one single model to capture the entire structure.

So, hopefully this small point is kind of beneficial or hopefully all of you understood this that if you divide the time series based on regimes or different characteristics of practicalities, then it might lead us to improve forecasting down the line. Now, the second advantage could be interpretability. So, what do you mean by that? So, clearly defined regimes provide insight into the underlying system. So, if you divide any time series into different slots or different regimes, then it can actually tell us predominantly what are the characteristics inside each of the regimes.

So, clearly defined regimes provide insight into the underlying system as well. And thirdly, flexibility. So, it can capture a wide variety of non-linear behaviors. So, let us say again, just to give you an example, let us say if you have a time series and then let me divide this time series into multiple regimes. And due to the highly volatile nature of the time series or due to lots and lots of ups and downs overall in the entire time series, let us say hypothetically the time series behaves like this.

So, initially it might have a trend, but in the second regime it might be slightly volatile like that, in the third regime it might be stationary for example like that and towards the end again it might show some downward trend ok. So, again if somebody wants to put forward an overall sort of a model structure to capture the entire time series in one piece, then it might not be very easy, but the moment you divide the time series into multiple regimes, we can actually capture the individual non-linearities, let us say capturing the trend or volatility in the second regime or stationary in the third regime etcetera by dividing the entire time series into different states or different regimes. So, again just to summarize improved forecasting obviously interpretability And the third one is flexibility when it comes to regime modeling. But having said that, the regime models are not entirely advantageous always.

So, what exactly the challenges of regime models? So, what exactly the disadvantages when it comes to modeling using the regime idea? So, the first one is threshold selection. What do you mean by that? So, determining the appropriate threshold value let us say gamma or transition function parameters can be computationally intensive ok.

So, again probably let us say that again just to give you the same example let us say if you have a time series divided into multiple regimes, but where exactly should you place all these different threshold values or threshold levels right. So, that could be a challenge all right. So, one has to identify exactly as to where change is sort of occurring. So, one cannot simply put the threshold values randomly.

There has to be some differentiating points in all the regimes. So, threshold selection could be one challenge. So, determining the appropriate threshold value which is given by gamma or the transition function parameters can be computationally intensive. Now, the second problem could be overfitting. So, let us say adding too many regimes can lead to overly complex models which generalize poorly.

So, let us say if you have lots and lots of regimes—let us say 10 regimes, 15 regimes, or 20 regimes—that again might be a problem because that might lead to overfitting. So, adding too many regimes can lead to overly complex models, which generalize poorly. And then, thirdly, data sufficiency. So, each regime requires sufficient data points for reliable estimation. So, if you do not have sufficient data points inside each of the regimes, that might be a problem in itself.

So, again, these could be some brief challenges when it comes to, let us say, modeling using the regime idea. So, in the earlier slide, we talked about advantages, but then, at the same time, I should understand that regime modeling is not free from any challenges. And then, some of the challenges which exist are, let us say, threshold selection, overfitting, or data sufficiency, etcetera, okay. So, now we pay attention to the very first kind of model, which is called a threshold autoregressive model. So, in itself, we have a very wide umbrella of different kinds of threshold models, and then the very first one that we will discuss is called the TAR model in short, or the threshold autoregressive models.

So, firstly, just to give you a brief introduction as to what you mean by the STAR model. So, in short, we can call them STAR models, by the way. So, threshold autoregressive or TAR models are a class of nonlinear time series models where the behavior of the system switches between different regimes depending on whether the threshold variable exceeds a predefined value. So, I think even in the last session, we gave a small introduction of threshold autoregressive models, which make use of a threshold value as such. So, let us say again, just to give you an example—let us say if you have a time series and a single threshold value like that, and let us say the time series is behaving something like this before the threshold, and then, once the time series crosses the threshold, you can see a lot of volatility in the time series, okay.

So, a simple TAR model can be applied to this because it contains one threshold value, and then you transition from one of the thresholds to the other, basically. The second point is that this structure makes TAR models well-suited to capturing abrupt changes or regime shifts in time series data. So, I guess this entire slide is more from a basic point of view, just to give an intro about threshold autoregressive models in general or a small motivation as to where one can apply TAR models, under what situations one can apply TAR models, etc. Now, talking about some of the key features of TAR models, right? So, the very first key feature is regime switching, obviously.

So, the model assumes different dynamics in separate regimes. We have discussed this point a number of times so far—that each regime has some different dynamics when it comes to, let us say, different characteristics of the underlying time series. And the second important point is that each regime could be modeled by its own AR process or by its own autoregressive process, okay? So, let us say if the overall time series is divided into two regimes by deciding upon some threshold value, then on one hand, I can put forward a certain AR model, let us say AR(1), okay. And on the other hand of the threshold value, I can put forward another AR model, let us say AR(2), okay.

So, how you combine these piecewise AR models to form an overall kind of TAR model is the idea here. Then the second ingredient, or the second key feature of the TAR model, is the threshold variable itself. So, the threshold variable is what determines the actual regime. Commonly, this is a lag value of the time series itself—let us say y(t - d). So, y(t - d) is some lag value of y(t) itself, right? So, let us say y(t) is a time series.

$$y_t = \begin{cases} \phi_{1,0} + \phi_{1,1} y_{t-1} + \cdots + \phi_{1,p} y_{t-p} + e_t, & if \quad y_{t-d} \leq \gamma \\ \phi_{2,0} + \phi_{2,1} y_{t-1} + \cdots + \phi_{2,p} y_{t-p} + e_t, & if \quad y_{t-d} > \gamma \end{cases}$$

Where, y_t is the time series at time *t*, y_{t-d} is the threshold variable with lag *d*, and γ is the threshold value, and $\phi_{i,j}$ is the AR coefficient for regime i.

So, depending upon what value d takes let us say 1, 2, 3, 4 whatever. So, y t minus 1 or y t minus 2 or y t minus 3 any one of the lag values could be taken as a threshold variable.

Then the third idea is piecewise linearity. So, although non-linear overall the model is linear within each regime. So, I think we discussed this point even in the last session that if you delve deeper down into each of the individual states or individual regimes then I can actually fit some linear models inside each of the regimes.

And again in front of you even you can see the same thing in this example as well that both the autoregressive models are sort of linear models and I am trying to fit the individual AR models inside each of the different regimes. And lastly, flexibility. So, handles processes where the relationship between variables changes depending on the state of the system. So, these are I think some of the key features when it comes to TAR models or modeling using a threshold autoregressive kind of a structure. Now, we will pay some attention into the model structure itself.

So, again do not be bogged down by all the notations. So, again if you are not very comfortable by any of the coefficients any of the notations here again a suggestion would be to go back let us say not just this week, but there may be some of the elementary weeks as well and then try to understand as to what the individual notations mean. But in general, since we are discussing a very easy sort of a model to start with, which is based on the autoregressive nature, it should not be very difficult to understand that where do you stop the model, right? For example, what do you mean by P? So, again, even before discussing, I am very sure that most of you might remember what do you mean by P.

So, if you are discussing any AR model structure, then you should be very comfortable knowing what the meaning of P is or, for example, what you mean by ET, etc. So, what exactly is the model formulation? So, the simplest threshold model is where the autoregressive or AR process switches regimes based on some threshold. And then, this is the model formulation in general. So, y_t equals phi_1,0 plus phi_1,1 y_t minus 1 all the way up to phi_1,p y_t minus p plus e_t. Now, here the only difference is rather than simply writing down phi_1, phi_2, phi_3 up to phi_p, since we have two regimes or two states, I had to write it in this form, alright.

So, the first row is true. whenever y_t minus d is less than or equal to some threshold value, which is gamma. On the other hand, the second row is phi_2,0 plus phi_2,1 y_t minus 1 all the way up to phi_2,p y_t minus p plus e_t if this is true, which is y_t minus d is bigger than gamma. And again, just a small description of what the notations mean. So, y_t is a time series at time t.

And then, y_t minus d is the threshold variable with lag d, and gamma is the threshold value. The most important point in this slide is this notation, which is phi_i,j, which is nothing but the AR coefficient for regime i, ok. So, phi_1,1, for example, would be nothing but the AR coefficient for regime 1. Similarly, phi_2,1 is nothing but the AR coefficient for regime 2, etc. So, again, you can pause the video, ponder over this slide, and try to understand all the individual notations, such as phi_2,p or phi_1,p, etc.

Now, we will pay some attention to what exactly the steps in model building are. So, the very first step is to specify the threshold variable. So, often a lagged value of the series So, often a lagged value of the series itself, let us say yt minus d, but it could also be some exogenous variable. So, exogenous means some external variable.

Now, the second point is determining the lag structure. So, you should always select the number of lags, let us say p in this case, in the AR process for each regime. And thirdly, estimate the threshold. So, estimate gamma using techniques like grid search or minimizing the residual variance or using an information criterion, let us say AIC or BIC, etc. So, these are some broad steps when it comes to model building.

So, specifying the threshold variable, then determining the lag structure, and then thirdly estimating the threshold value. All right. And the next step in model building is estimating the parameters. So, once you estimate the threshold value, then inside the equations, you have to estimate or find out all the parameters. So, once the threshold is fixed, fix the AR processes for each regime using methods like least squares, etc.

And lastly, obviously, as always, diagnostic checking. So, check for residual autocorrelation, stationarity, and model adequacy, and conduct some hypothesis tests for non-linearity. Let us say Hansen's test, etc. So, these are broadly all the steps when it comes to model building using a regime sort of approach. So, estimating parameters, diagnostic checking, etc.

Now, toward the end, we will describe what the advantages of TAR models are. So, the first is interpretability, obviously. So, the piecewise linear nature makes these models easier to understand compared to any other non-linear models. For example, let us say neural networks. Then, the second one is flexibility.

So, the ability to model systems with regime shifts or asymmetric behaviors is one of the advantages of TAR models. And then, thirdly, stationarity within regime. So, each regime can be stationary even if the overall process is not stationary. And on the other hand, what are the challenges of TAR models? So, the first one could be threshold estimation, as discussed even a short while back, that selecting the correct threshold value can be computationally intensive.

Second is overfitting. So, let us say adding too many regimes or lags can lead to overfitting. And thirdly, boundary issues. So, sparse data near the threshold value gamma can make estimation very difficult. So, firstly, what do you mean by the third point here? I will try to explain in a bit more detail.

So, let us say if the overall structure has two regimes divided by one threshold value, and let us say you have a time series which behaves like that. In some cases, it might happen that just next to the threshold value, you do not have any data. So, let us say if you have some missing data here and missing data there, and again, let us say you have a highly volatile kind of structure here. So, this is called sparse data. So, sparse data means that you do not have many observations near the threshold.

This can make the boundary issue problem kind of predominant. So, boundary issue is exactly this problem: sparse data near the threshold value can make estimation very difficult in a way. And lastly, a few applications. So, first could be in economics, let us say analyzing business cycles, expansion versus contraction, or modeling inflation dynamics with thresholds like unemployment rates. Second one could be in climatology.

So, let us say representing shifts in climate variables, for example, temperature, rainfall, etc. based on some external drivers. Or thirdly, engineering. So, let us say modeling systems with operational thresholds, such as load-bearing structures or machinery, etc. So, I think the second session this week is more of an intro to one particular kind of threshold model called TAR or Threshold Autoregressive Regime Models.

So, initially, we started this session by describing in detail what the different kinds of regimes are. So, smooth transition regimes, Markov regimes, or threshold-based regimes, and then towards the end, we discussed one particular time series example or one particular time series model, which is widely applicable, called TAR or threshold autoregressive model.