Predictive Analytics - Regression and Classification
Prof. Sourish Das
Department of Mathematics
Chennai Mathematical Institute

Lecture - 33
Hands on with R for Bootstrap Regression

Hello all welcome back to the last part or the final part of lecture 9. In this video, we will be doing some hands on, on how to do Bootstrap Regression for capital asset pricing model.

(Refer Slide Time: 00:31)

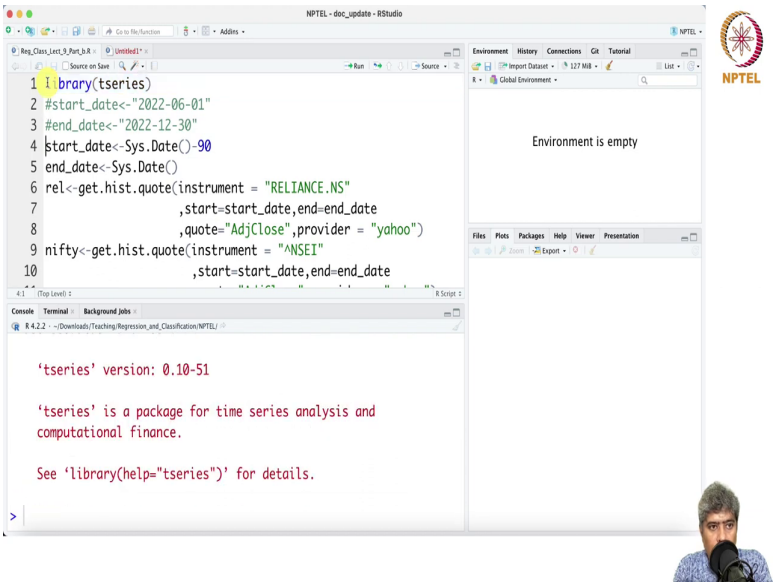

(Refer Slide Time: 00:48)

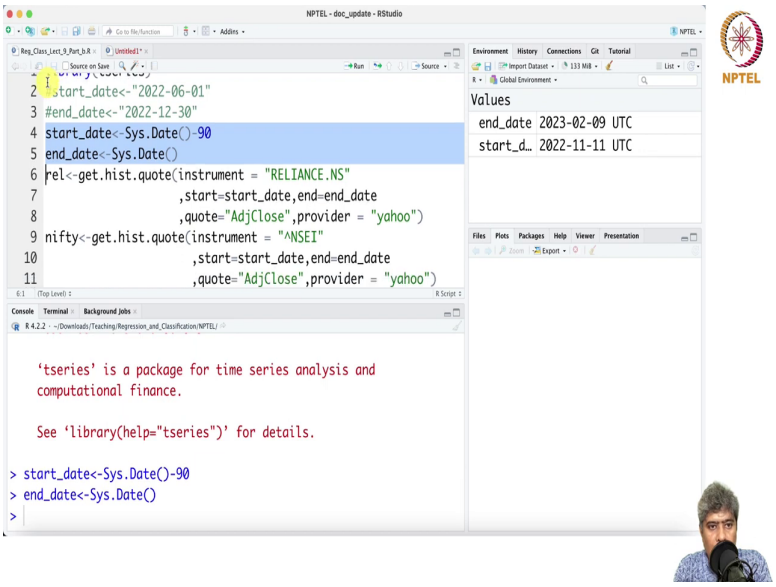

So, here in our this piece of code you have already seen, I will quickly remind you that you know we are going to use a T series package, then we are going to use system date or start date. Here we are then now this piece of code 9678 is going to download the data from internet from Yahoo Finance.

(Refer Slide Time: 01:01)

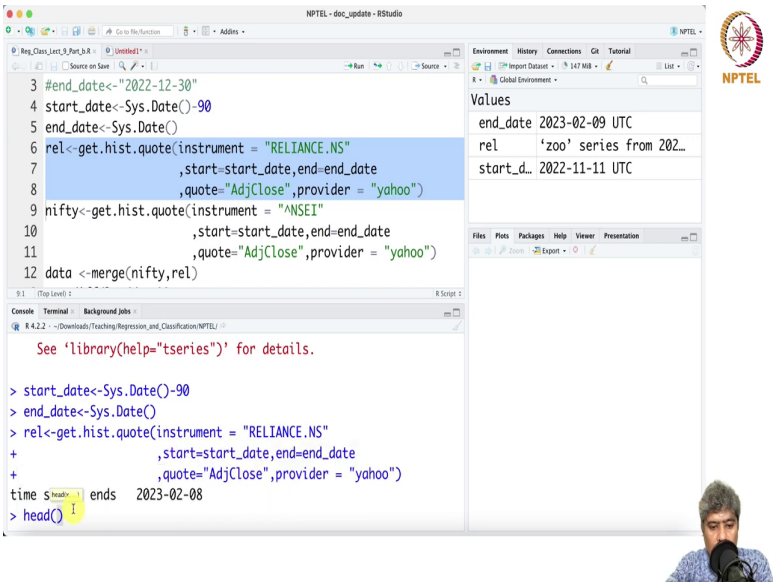

(Refer Slide Time: 01:12)

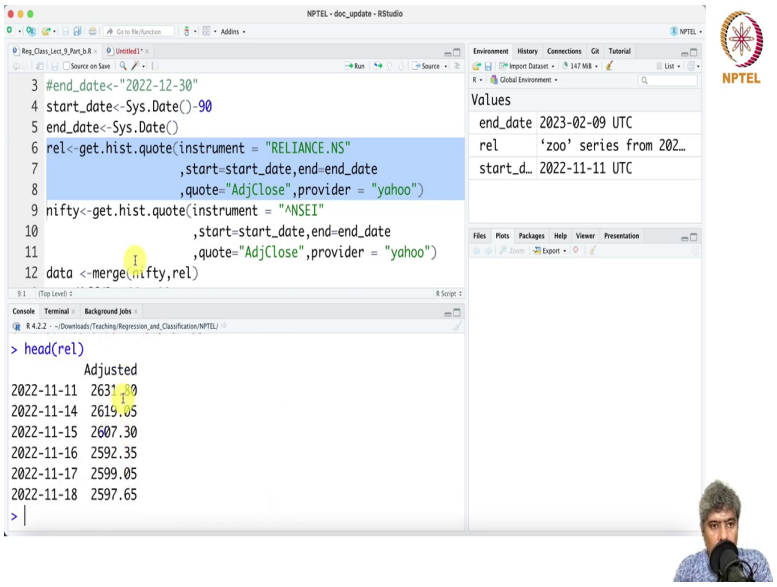

(Refer Slide Time: 01:20)

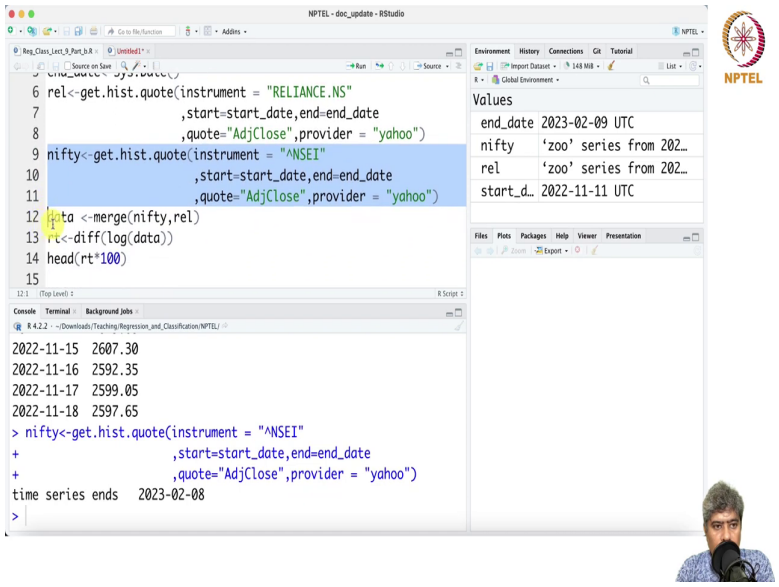

So, here we have data is downloaded from Yahoo Finance head rel. So, here we can see then NIFTY 50 data will be downloaded adjusted close prices as of 8th February 2023 all the data has come.

(Refer Slide Time: 01:28)

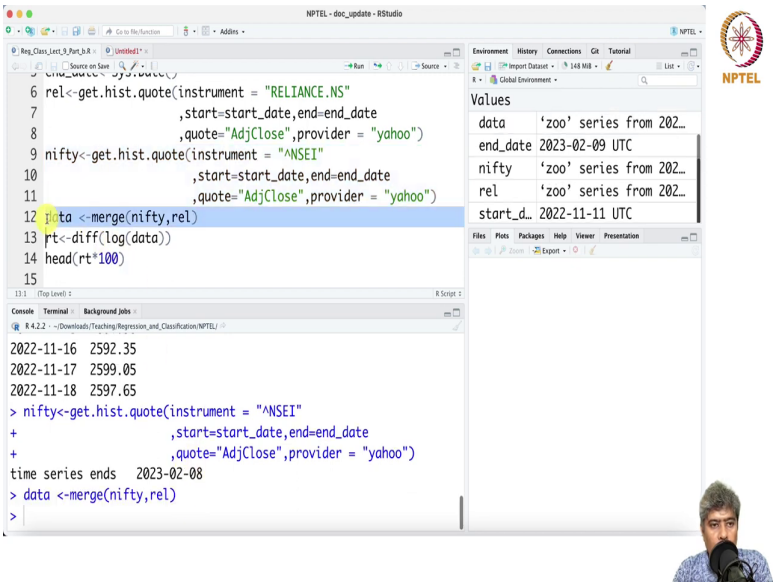

(Refer Slide Time: 01:37)

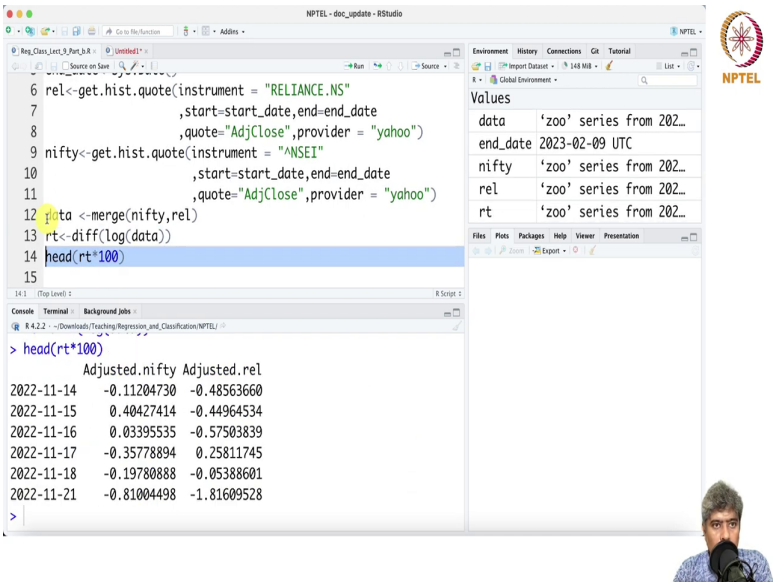

Then we are going to merge these data here, in line 13 we are going to calculate the log return. Here is the log returns in percentage term.

(Refer Slide Time: 01:14)

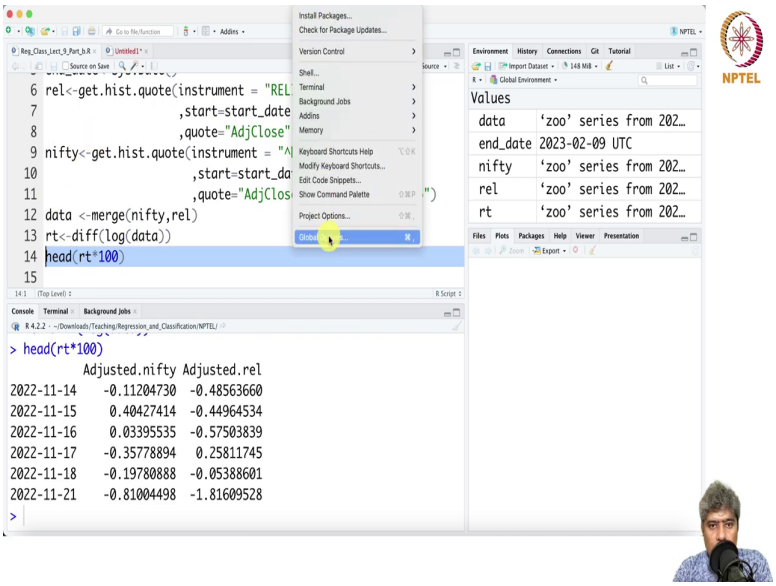

(Refer Slide Time: 01:46)

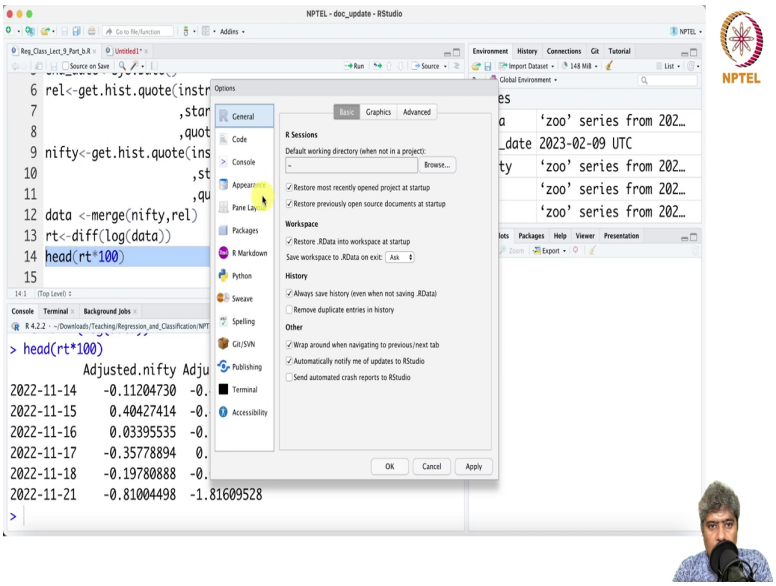

(Refer Slide Time: 01:47)

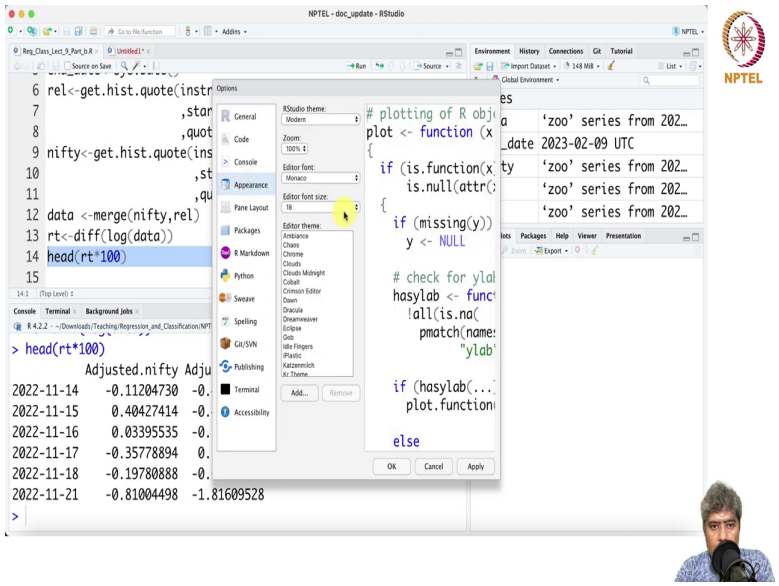

Let me increase the font size bit maybe I will use 24 apply ok; alright.

(Refer Slide Time: 01:58)

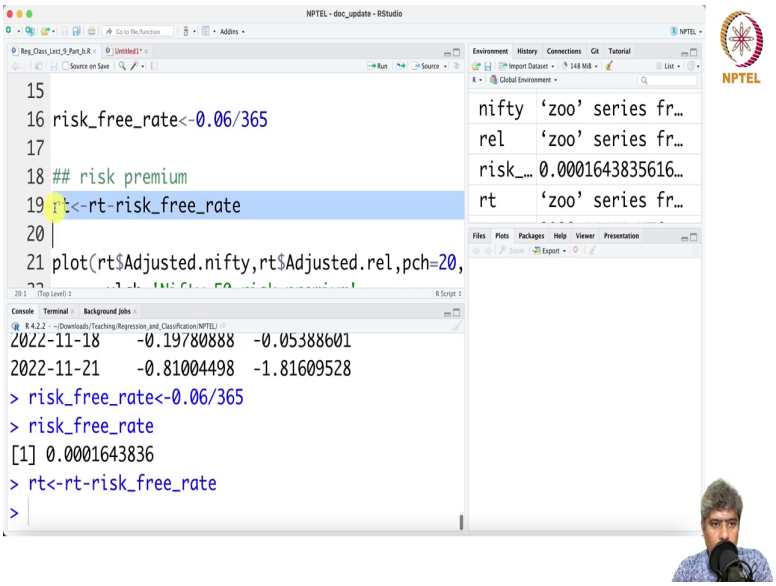

And now I am using risk free rate 6 percentage, annualized 6 percentages. So, I am dividing it to by 365, so this is the daily risk-free rate. Then I am calculating subtracting risk free rate from the return log return.

(Refer Slide Time: 02:24)

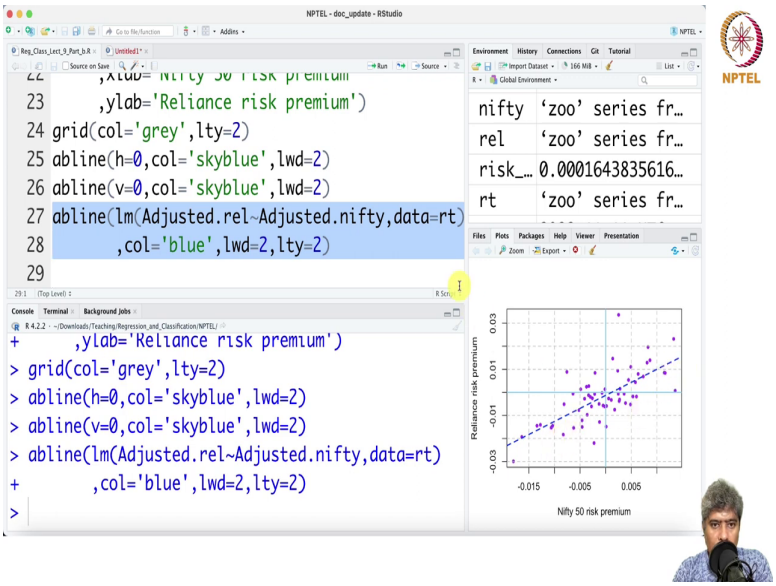

So, calculating risk premium and then drawing the plots of risk premium grids are being drawn. And then here are some ab lines and then. So, the fitted line this is the fitted line using OLS method.

(Refer Slide Time: 02:44)

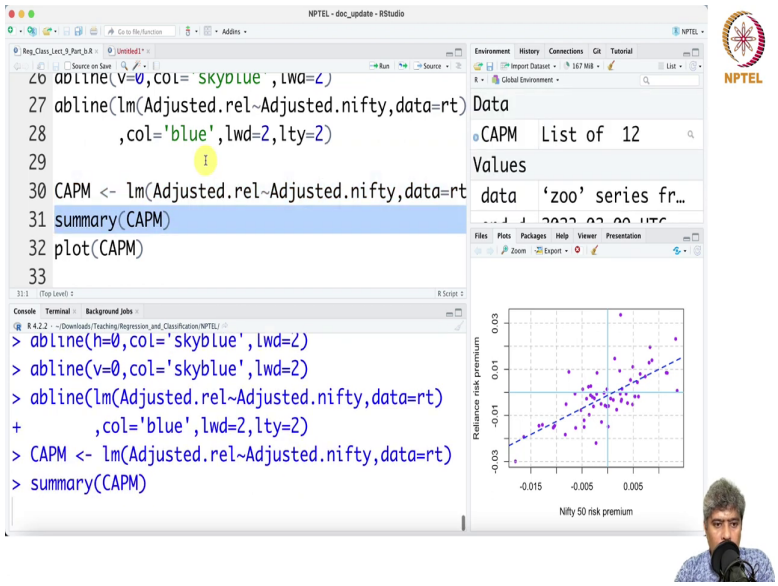

So, here I am fitting CAPM using OLS method and here is the summary statistics, so ok.

(Refer Slide Time: 02:52)

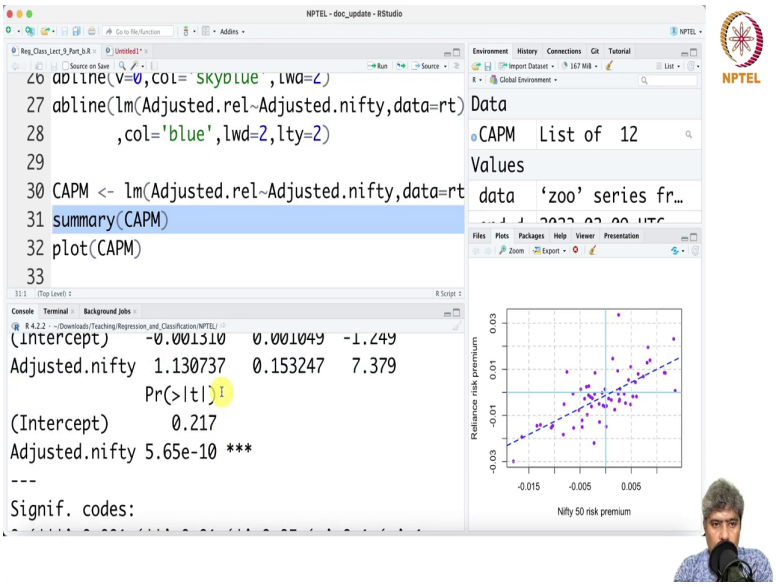

(Refer Slide Time: 02:57)

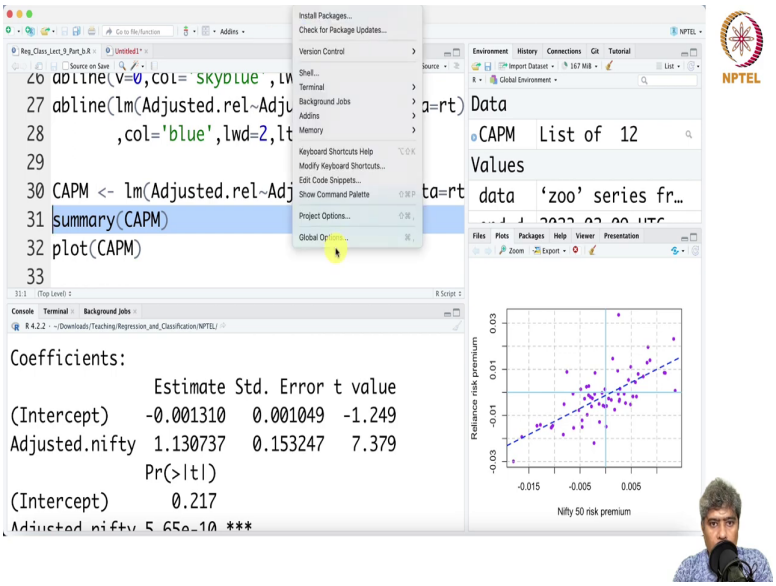

(Refer Slide Time: 03:01)

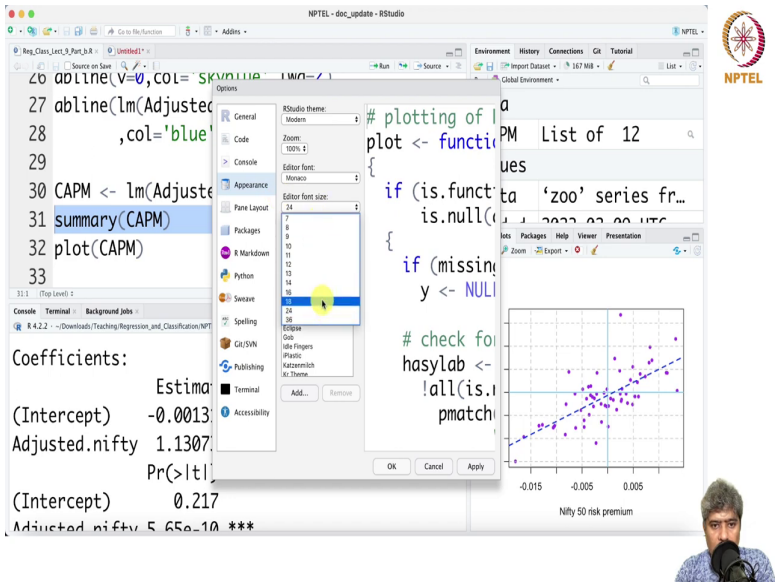

So, one problem is if we make the font size too big then it will not come as nicely as we would like to see it. So, now it like comes a bit nicely.

(Refer Slide Time: 03:10)

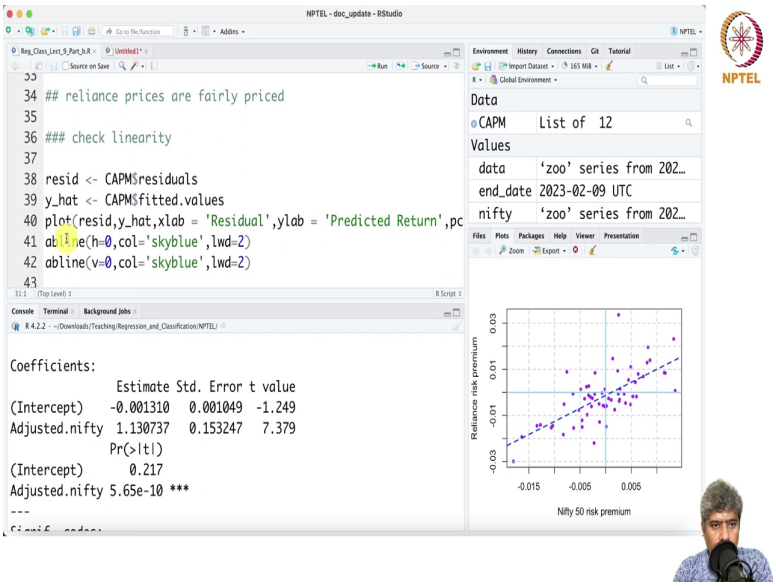

(Refer Slide Time: 03:17)

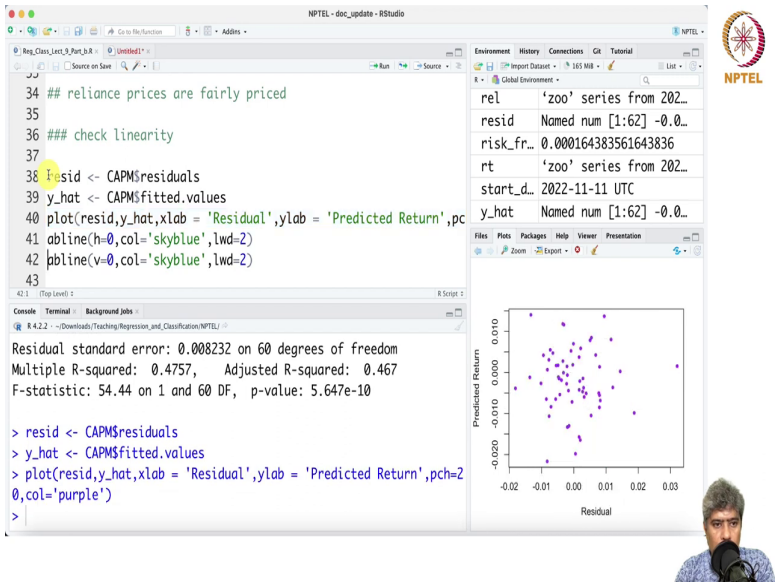

And then we can take the residuals and then plot the residuals and then can do the Bartels rank test and then we did Breusch Pagan test, rank test and Breusch Pagan test everything was quite good.

(Refer Slide Time: 03:20)

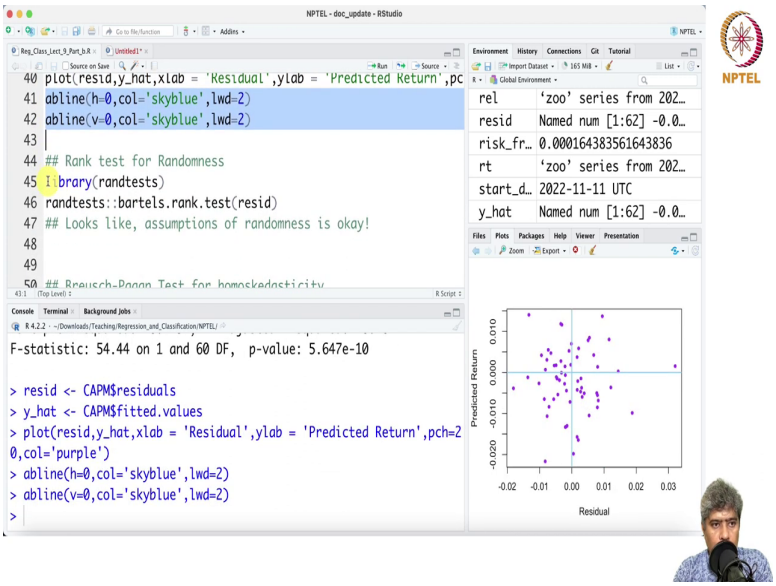

(Refer Slide Time: 03:28)

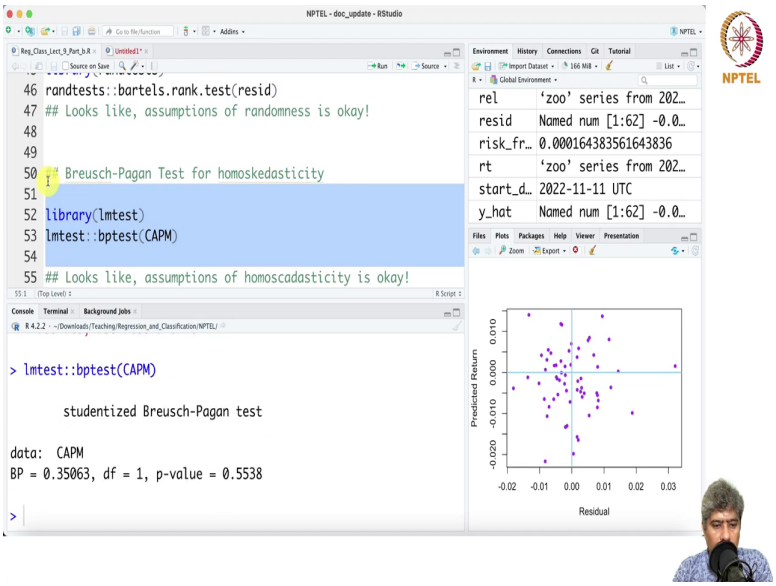

(Refer Slide Time: 03:34)

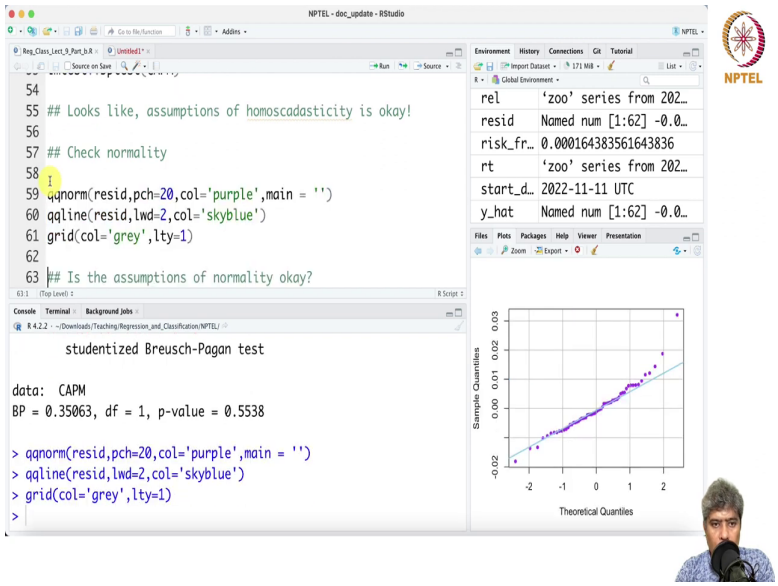

Then when we did a test for normality it was kind of get rejected and that is where the whole problem starts.

(Refer Slide Time: 03:44)

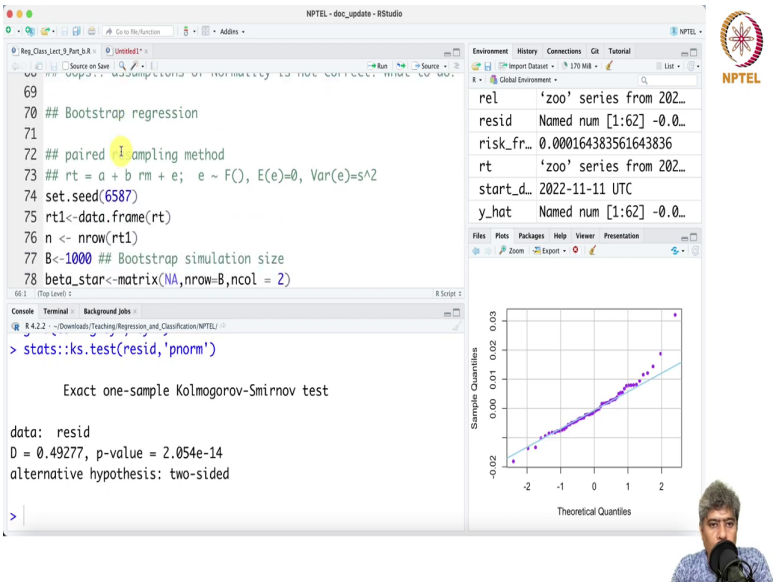

So, we decided to do a bootstrap regression first we are here we are going to apply paired resampling method or paired bootstrap regression method. So, here we are assuming risk premium equal to a alpha plus beta a plus b times market premium plus error. And we are assuming this idiosyncratic return or E is following some distribution. We do not know what is that distribution?

Some distribution F with expected value of E to be 0 and variance of E to be sigma square, but we do not know what is that distribution is this could be any distribution.

(Refer Slide Time: 04:27)

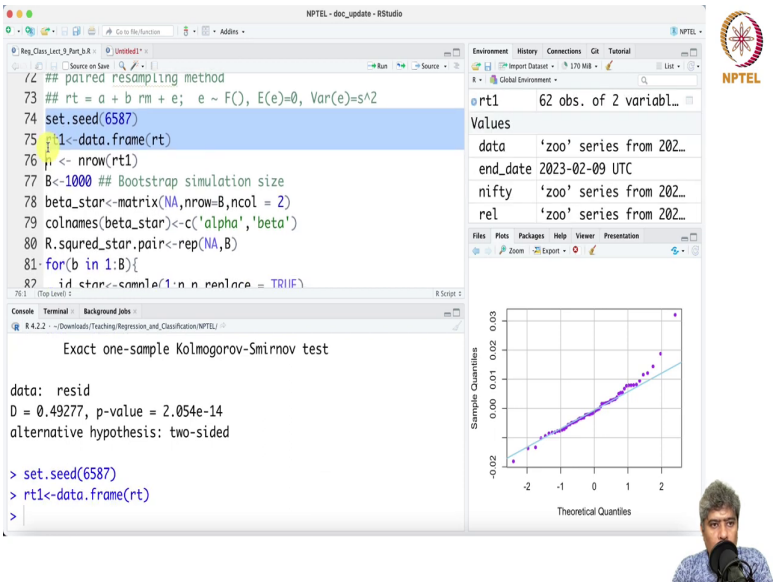

(Refer Slide Time: 04:34)

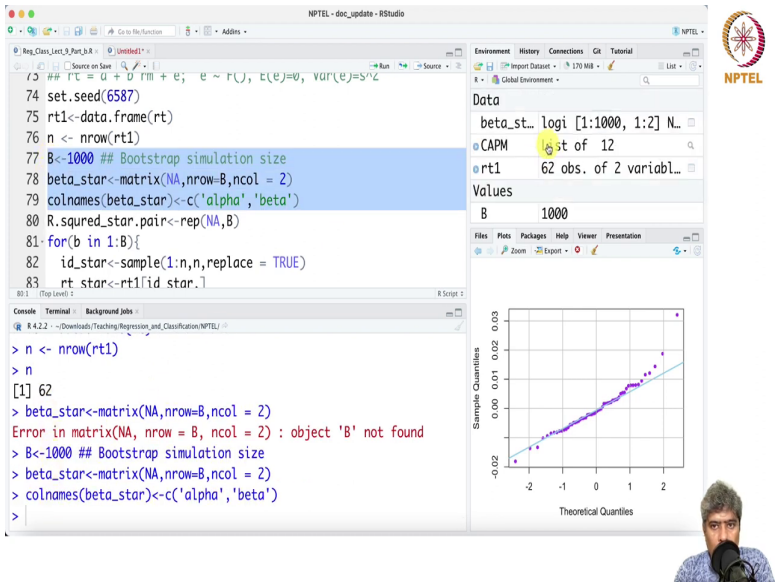

So, we put it as a data frame here is the number of samples is 62, we have taken bootstrap simulation size to be 1000, you can increase it to any number 10,000 or anything. And then I am defining a beta star a matrix of sample simulation size n and number of columns is beta 2, because I am going to estimate alpha and beta. So, oh sorry I have to run this ok.

(Refer Slide Time: 05:03)

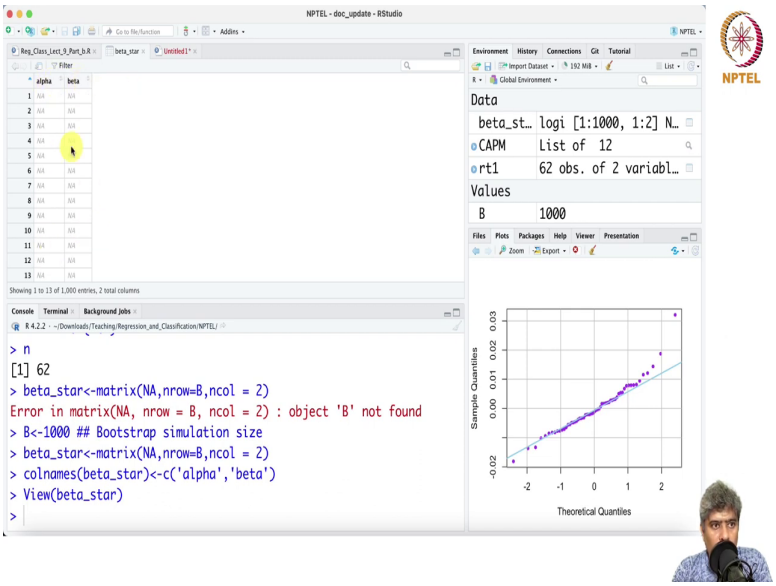

(Refer Slide Time: 05:15)

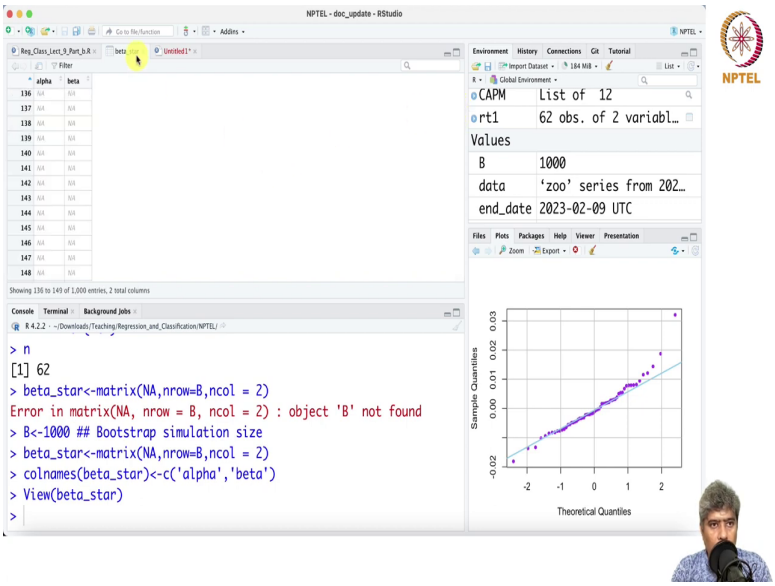

Now, if you see I have created this matrix with column name alpha beta and it has size of it has basically if you just go up 10th about 1000 rows number of simulation.

(Refer Slide Time: 05:22)

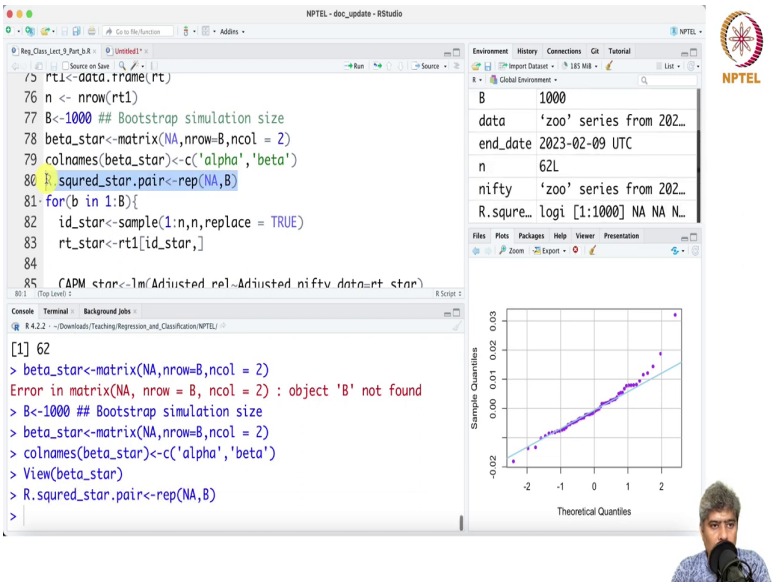

(Refer Slide Time: 05:33)

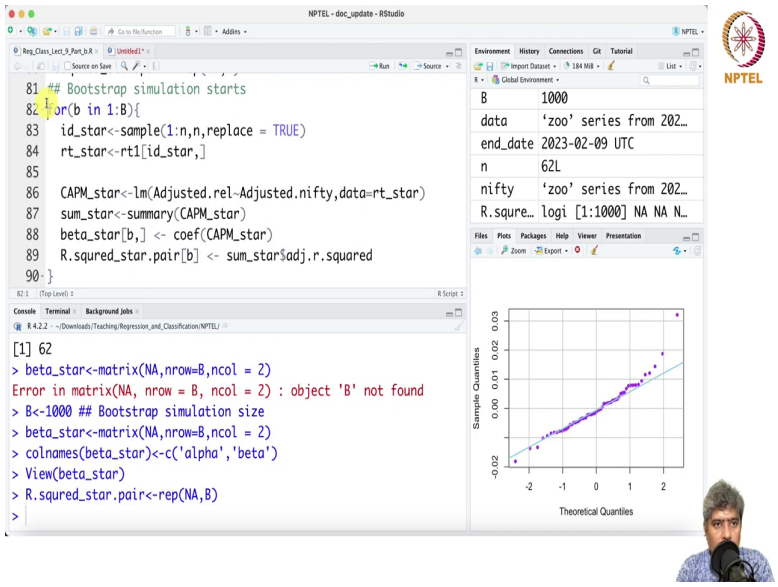

Now, I am going to also create another place for R square I am calling it R square star. Then here my loop start bootstrap regression bootstrap simulation starts bootstrap simulation starts. So, first from 1 is to n id I am going to draw ransom's id I am going to draw the random ids and those ids I am going to take the samples random samples from the samples itself.

So, rt 1 is the sample itself from there I am going to draw random samples and those are that is the rt star and from the rt star I am going to you know fit the model and from the model I am going to calculate the summary and from the summary I am going to I am extracting the coefficients alpha beta ok and the adjusted R squared I am keeping it here.

(Refer Slide Time: 06:40)

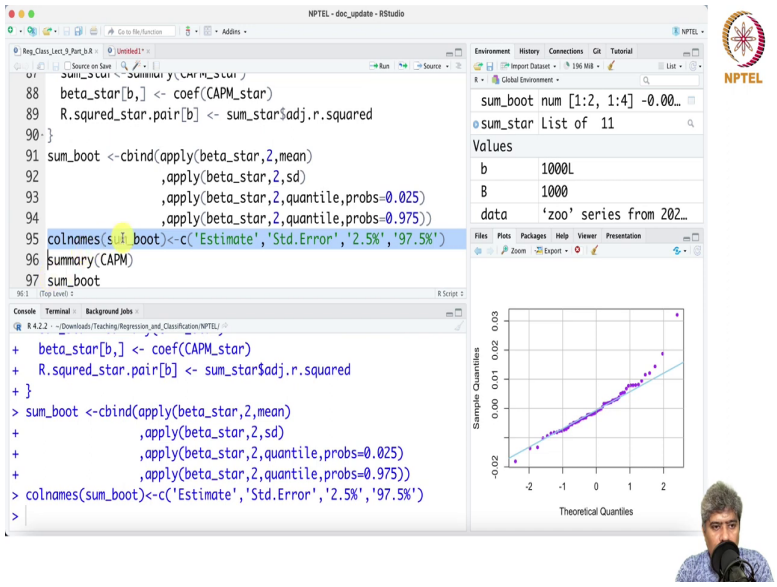

So, let me just run this piece of code and it is run in almost no time and then I am going to. So, summarize the bootstrap statistics here so few and then estimate standard error and 25 percent.

(Refer Slide Time: 06:57)

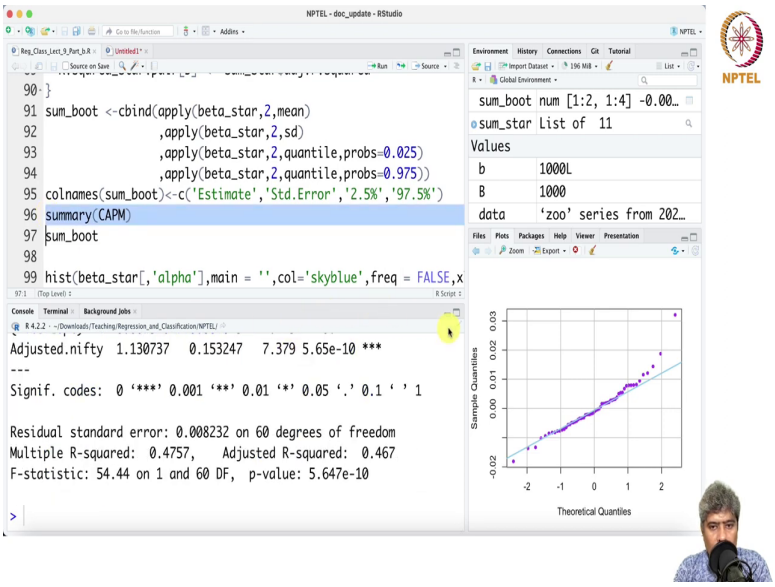

So, now if you run this, so estimates are like negative 0.0012 and 0.12. So, now you can see the alpha is containing the 0 the confidence interval this is the confidence interval for alpha and it contain 0. And similarly, this is the confidence interval for beta it contains 1. So, for this is for reliance and remember that summary statistics for CAPM this is the OLS thing.

(Refer Slide Time: 07:28)

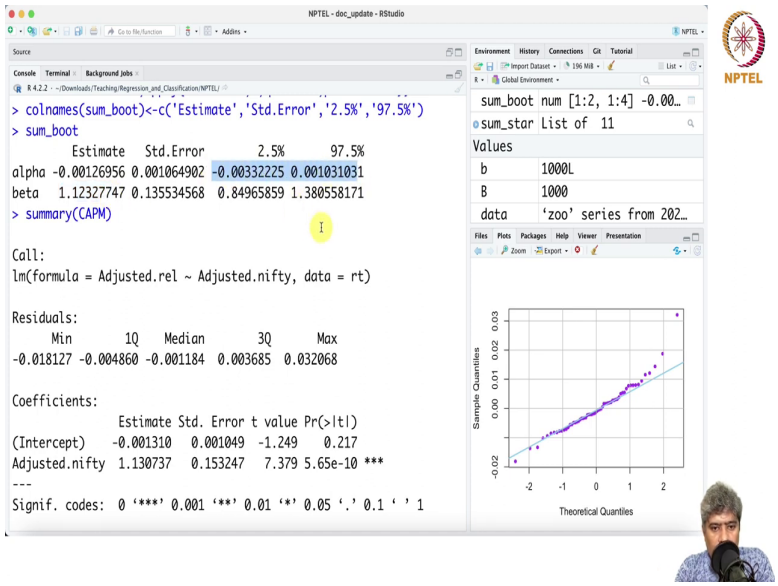

So, if you now carefully look into it that OLS estimate is 0.0013 and if you round it of to 4 decimal places it is actually negative 0.0013, beta is 1.1307 1 this is beta here 1.12327 the this is the coefficient.

(Refer Slide Time: 07:56)

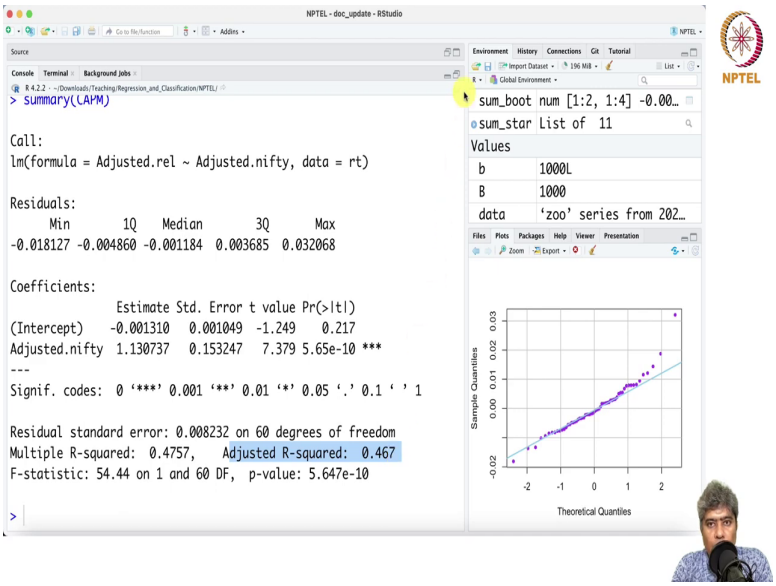

So, now since the underlying distribution does not or underlying model assumption does not hold good, so these inference based on t value and p value is not valid inference. But here I am assuming that the underlying distribution is just unknown and I am doing a non-parametric test. So, these inference is valid though there is not much change in my inference, but this inference is valid.

(Refer Slide Time: 08:41)

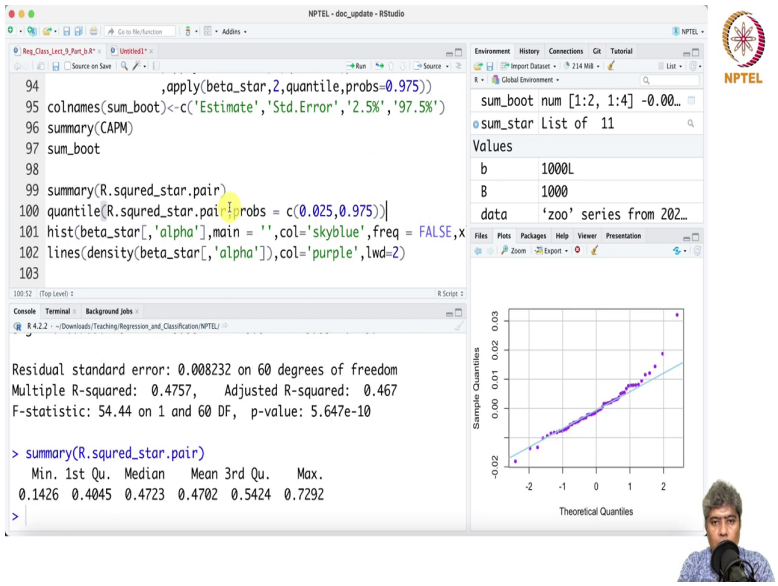

Here the adjusted R square is given 0.467, but what we can do is essentially we can just say summary of this. So, this is the bootstrap statistics for the R square adjusted R square and we can quantile say we can create a quantile of R squared with probability equal to 0.025 comma 0.975.

(Refer Slide Time: 09:10)

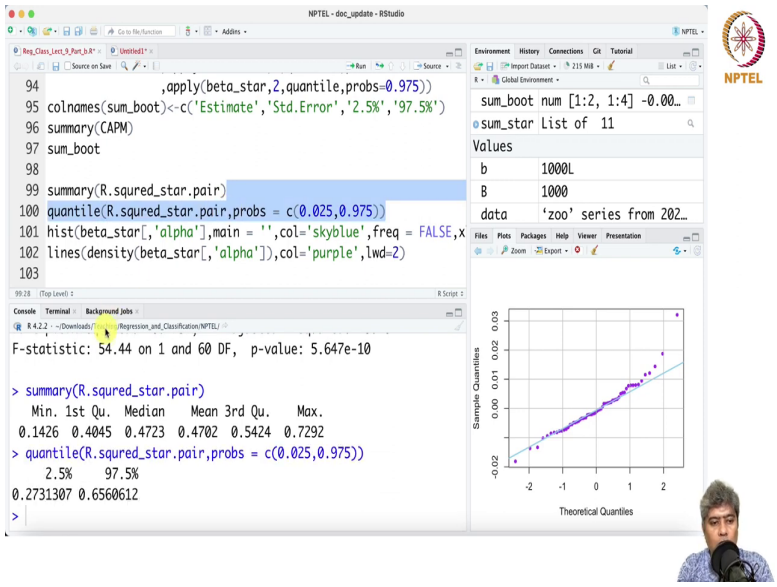

Now, this gives me a 95 percent confidence interval for adjusted R square. Now, in typically it is very difficult to get a any statistical inference for R square, but using bootstrap we can calculate a 95 percent confidence interval for adjusted R square.

(Refer Slide Time: 09:34)

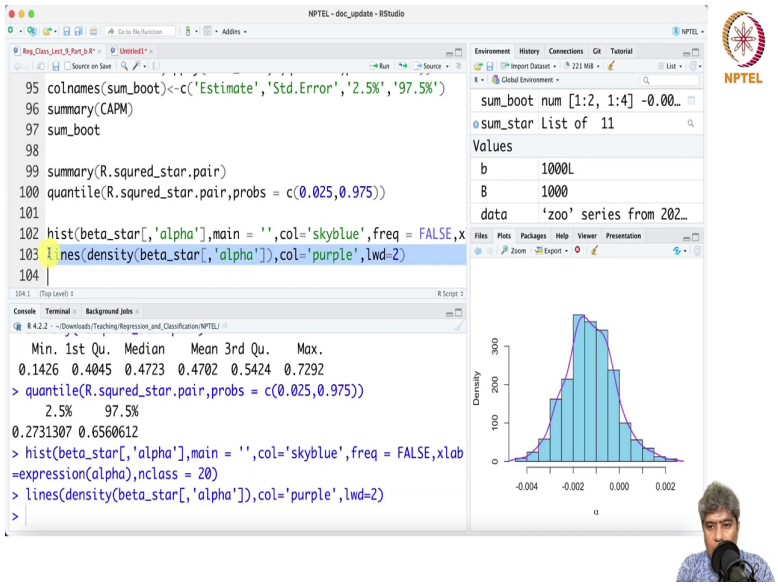

So, this is a very interesting thing. So, you can here we have the alpha the sampling this is bootstrap sampling distribution of alpha and here is the density that we are drawing and then if you would have lines.

(Refer Slide Time: 09:57)

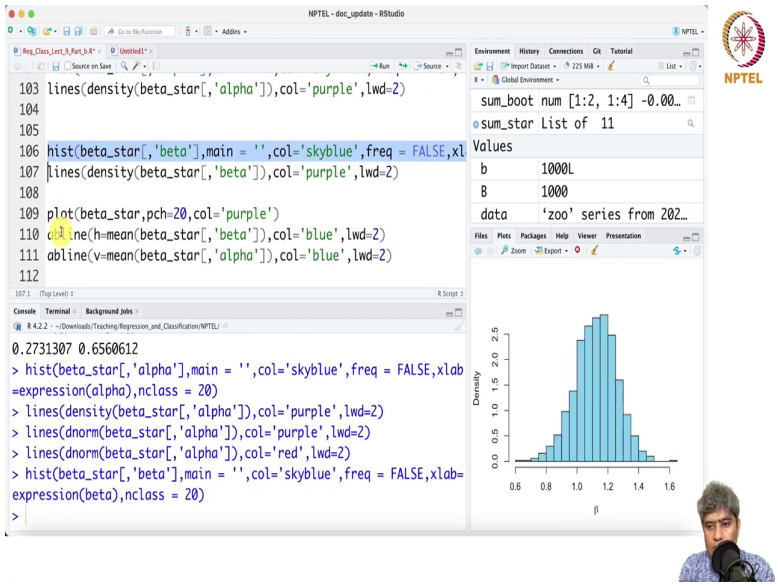

(Refer Slide Time: 10:03)

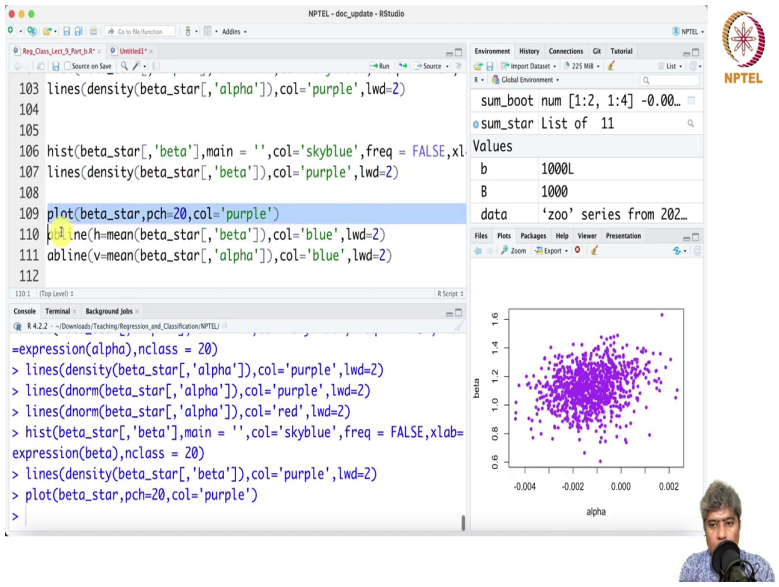

And similarly, we can do the histogram of the this is the beta and this is the lines.

(Refer Slide Time: 10:15)

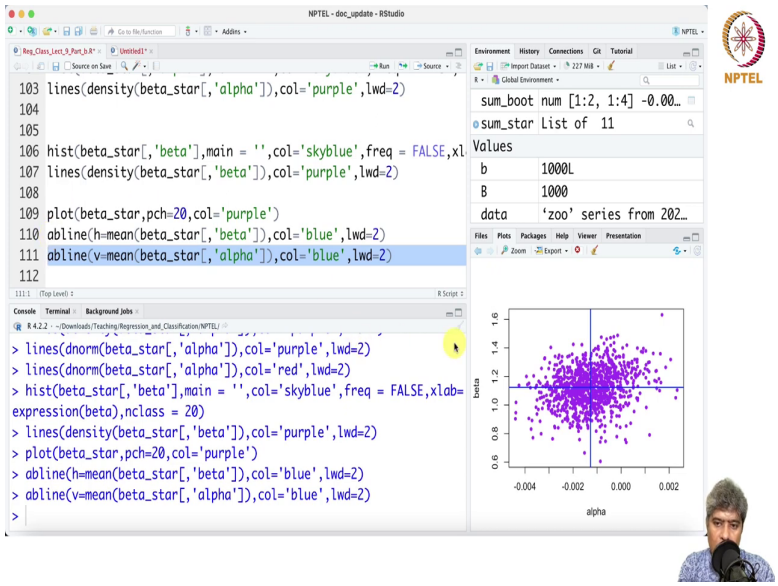

If we do the beta star this is the distribution between joint distribution of sampling distribution of alpha and beta. And we see that there is not much correlation that we are seeing between alpha and beta which is a good actually.

(Refer Slide Time: 10:26)

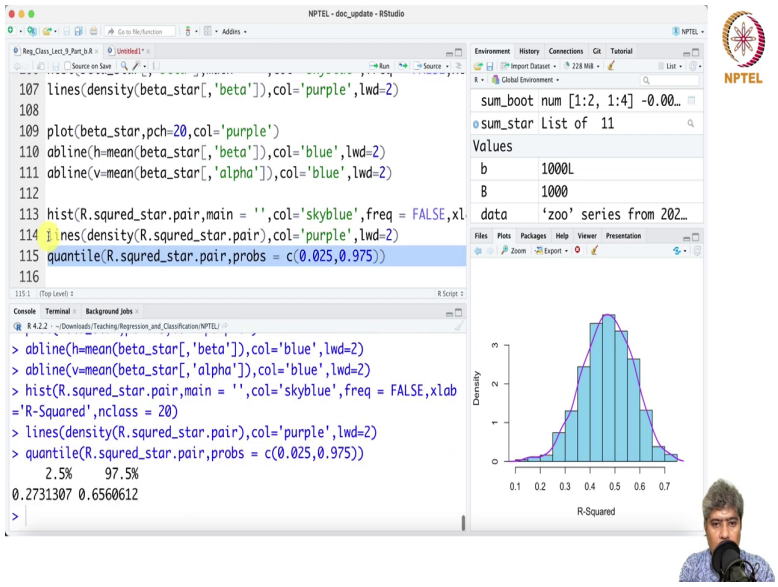

So, here is the sampling distribution of bootstrap sampling distribution of R square here is the lines density; kernel density here is the R square.

(Refer Slide Time: 10:41)

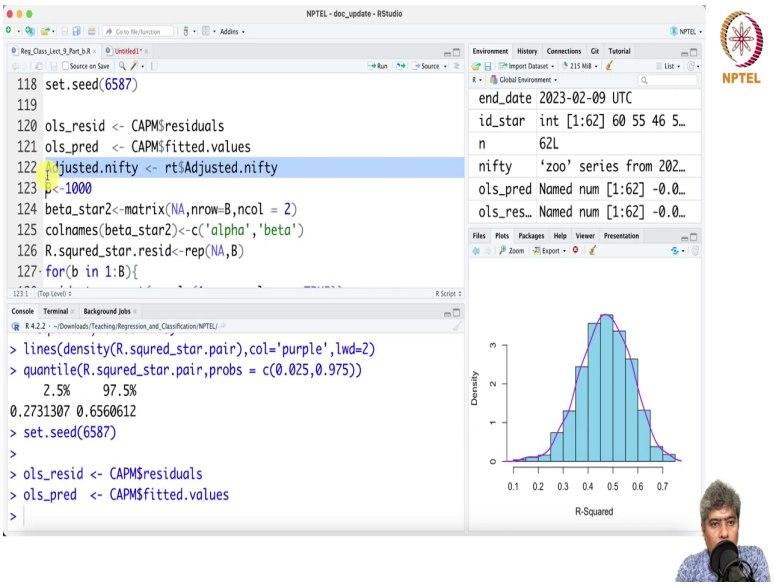

Now, we are going to do the residual resampling or Residual Bootstrap Pregression ok and from the CAPM we just take the OLS residual and OLS predicted values and then this is the adjusted nifty and same way we define the matrixes.

(Refer Slide Time: 11:11)

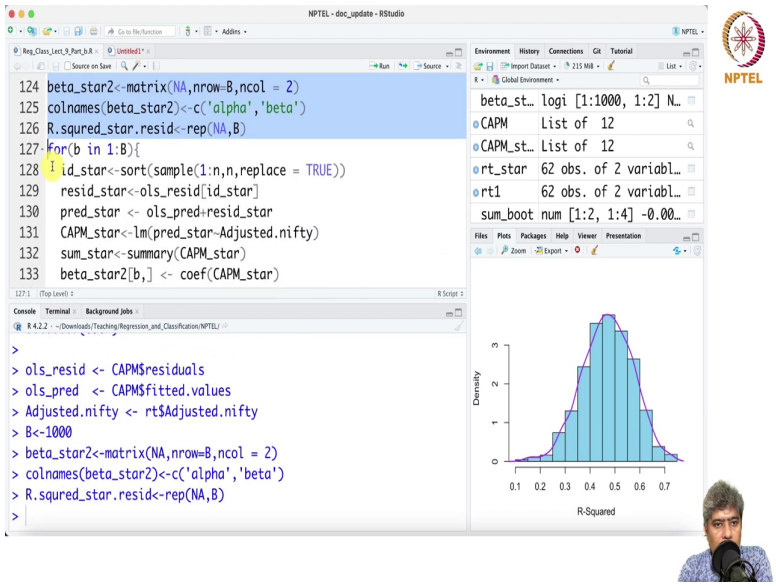

Now, here we have starting the bootstrap statistic Bootstrap simulation.

(Refer Slide Time: 11:18)

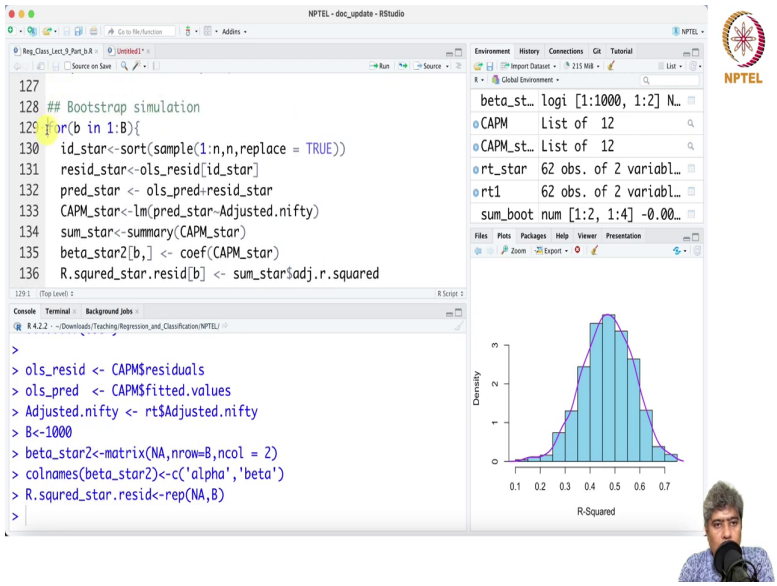

Now, here is the id star I am from the OLS residual I am just simulating the residuals, then from the predicted ols predicted value you just take the predicted star predicted value, because X beta hat is nothing but the ols predicted value right and then you just add the error and that will give you the Y star. And then based on the Y star you compute the again fitted value and CAPM star and then all you do from there you compute the coefficients and from the and adjusted R square.

(Refer Slide Time: 12:20)

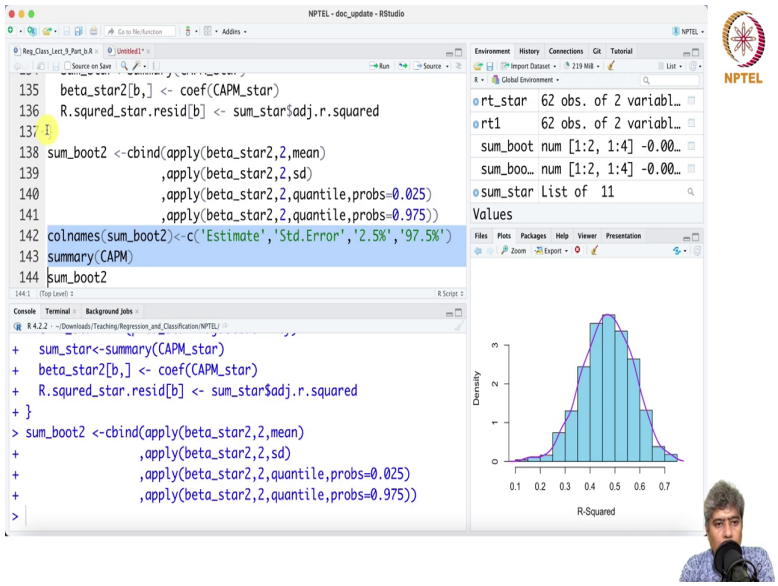

So, same way you just do that it is very fast for 1000 simulation it does very quickly actually.

(Refer Slide Time: 12:33)

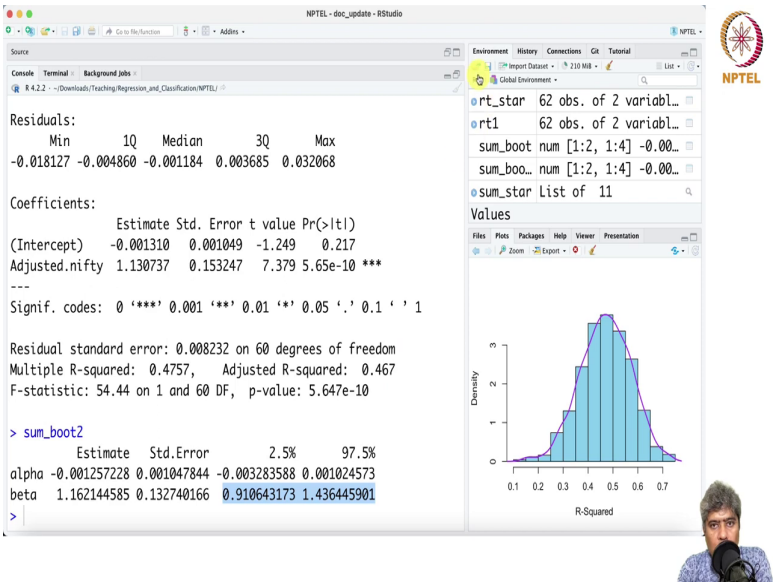

So, here is the estimates are pretty much very close to similar previously and confidence interval also does not change much.

(Refer Slide Time: 12:48)

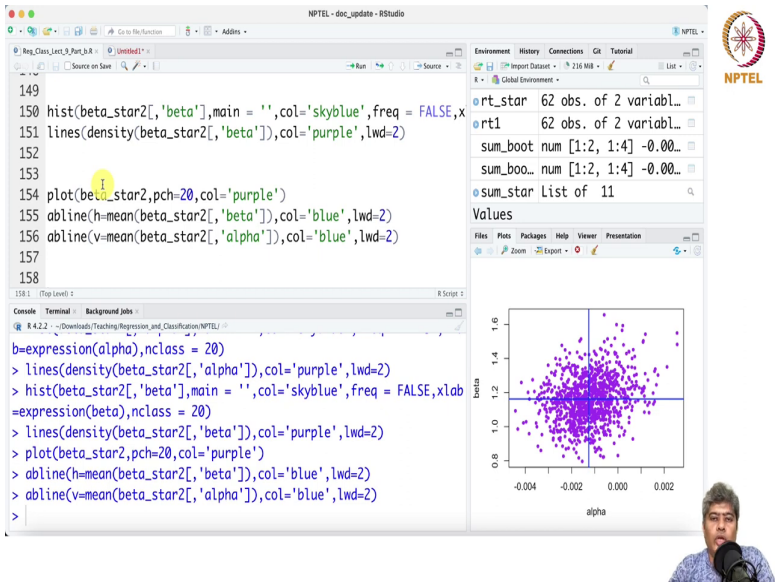

So, residual bootstrap holds pretty much similar not much change because and this is the alpha distribution of alpha, this is the distribution of beta, this is the joint distribution of alpha and beta when we are doing residual bootstrap regression. So, that is how we do paired bootstrap statistics or bootstrap regression and residual bootstrap regression.

If sometimes what happens as you know that you fit the model and what happens is you fit a model you did a very good job with modeling and do a reasonable predict prediction also. Suppose you have achieved a descent R squared very reasonably low RMSE, but and most likely you are ok to you know accept the model to the production.

Now, people can still re-question that if the model assumption does not hold good. Now, often time what I have seen that if this is the if the case is that randomness is ok and homoscedasticity is ok. Then and but the it is the distribution that is the problem because

distribution is always a very strong assumption; then I will always recommend you to resort to bootstrap.

So, if you do the bootstrapping what happens is you still work with the same model only the statistical testing of hypothesis and the inferences. Now, you are doing with the bootstrap distribution instead of assuming residual is normal. But rest of the assumptions is perfectly fine and if even the homoscedasticity does not hold good you can still work with paired bootstrap statistics and it will still work.

So, bootstrap regression is a big help for the statistician in case you have to do a statistical inference, that particularly if the model prediction is good. If the model prediction is not good then of course, you have to figure out a better model in that case.

But if the model prediction is good if the model is doing reasonable prediction in the out of the sample, if the model is doing if the RMSE if pretty low, if the R square adjusted R square is reasonably acceptable range for the given domain. Then just for inference you may find that the normality assumption is too strong assumption.

In that case you can do a statistical inference using bootstrap statistics and still you can go on with that prediction of the model, that predictive model because that predictive model do reason if it does do if it does good with prediction good prediction, then inference can be done using bootstrap sampling distribution. So, you will be good in that sense. You can salvage that situation when the model check assumption is not holding good.

As long as the IID sample assumption is working fine. You can salvage that situation using simple bootstrap statistic technique. So, with that I will stop here for this lecture and looking forward to the next lecture. I hope you enjoy this hands on session it is a very useful I find it in my many real life project I find it very useful this technique. So, I hope you will enjoy it and you will be able to use it and for your project experience.

So, thank you very much and see you in the next topic and next video next lecture.