Financial Risk Forecasting
Seminar 6

Jon Danielsson
London School of Economics

Version 4.0 August 2024

6 Univariate volatility

In this seminar, we will use the `Returns.RData` file that we created in last session. We focus on coding the univariate GARCH model using `rugarch` package as well as `car` for QQ plots. You need to install both GARCH packages before proceeding. Please review the theory of GARCH that we discussed in the lecture.

6.1 The plan for this week:

1. Load data and packages and do simple statistics
2. Build and run univariate GARCH models
3. Print and plot the outputs
4. Work with various GARCH specifications
5. Assess model quality using likelihood ratio tests and residual analysis
6. GARCH simulations

6.2 Loading data and libraries

```
library(rugarch)
library(car)
```

```
load('Returns.RData')
```

6.3 Plot the ACF

```
y=Returns$JPM
plot(y,type='l')
acf(y,lag=10)
plot(y^2)
acf(y^2,lag=20)
```

6.4 Build and run univariate GARCH models

Here we have two specifications, both the default and one we have modified. What is the difference between these two? Note the `solver = "hybrid"` argument. Why is that necessary? You can find the answer in the notebook.

```
spec0 = ugarchspec()
spec0
spec1=ugarchspec(mean.model =
  list(armaOrder=c(0,0),include.mean=FALSE))
spec1
output0=ugarchfit(spec0,y,solver = "hybrid")
output1=ugarchfit(spec=spec1,
  data=y,solver = "hybrid")
```

6.5 Print and plot the outputs

When you run `plot(output1)`, you get a menu of options for making various plots.

```
plot(output1)
names(output1@fit)
output1@fit$coef
l1=likelihood(output1)
coef(output1)
```

6.6 Work with various GARCH specifications

There are many alternatives you could use, here we show the Student-t.

```
spec2=ugarchspec(
  variance.model=list(garchOrder=c(1,0)),
  mean.model = list(armaOrder=c(0,0),include.mean=FALSE))
output2=ugarchfit(spec=spec2,data=y)
output2
spec3=ugarchspec(
  distribution.model='std',
  mean.model = list(
    armaOrder=c(0,0),
    include.mean=FALSE)
  )
output3=ugarchfit(spec=spec3,data=y)
```

6.7 Assess model quality using likelihood ratio tests and residual analysis

We can use the likelihood ratio test to compare the effect in **nested** specifications.

```
LR_statistic=2*(likelihood(output3)-likelihood(output1))
p_value  = 1 - pchisq(LR_statistic, df = 1)
```

The residuals provide information on whether the model is correctly specified.

```
residuals=y/output3@fit$sigma
plot(residuals)
acf(residuals)
acf(residuals^2)
qqPlot(residuals)
```

6.8 GARCH simulations

It can be useful to simulate from a volatility model. Perhaps to identify whether a model describes the data properly, to run a Monte Carlo test with known stochastic process as input or in derivative pricing.

```
ugarchsim(output1,10)
```

6.9 Recap

In this seminar we have covered:

- Making univariate volatility models

Some new functions used:

- `ugarchspec()`
- `ugarchfit()`
- `ugarchsim()`

6.10 Optional exercises

1. Run all the volatility models discussed in the lectures, put the one day ahead forecast for each into a vector and plot them together.
2. Make a function that takes volatility specifications as input (like `GARCH11`) and runs the models corresponding to them, putting all necessary output into a data frame.
3. Evaluate the quality of a fitted volatility model as the sample size increases from 500, using all of the models discussed in the course.
4. Use simulations from the tAPARCH model to evaluate the quality of risk forecasts for all the univariate volatility models we discussed as the sample size increases from 500 to 1,000,000.