Tutorial 9

1. Consider the monthly log stock returns of Intel Corporation from 1973 - 2008, found in the file `m-intc7308.txt`.

   (a) Transform the returns into log returns.

   (b) Make a time series plot. Does this suggest conditional heteroscedasticity (clusters of high volatility)?

   (c) Does the acf of the log stock returns indicate significant serial correlations?

   (d) How about the acf and pacf of the squared log returns?

   (e) Try a Ljung - Box test. For which m do the $Q(m)$ statistics indicate ARCH effects? Build a GARCH model for the transformed series and compute 1-step to 5-step ahead volatility forecasts at the forecast origin December 2008.

   The R script is as follows:

```
www =
"https://www.mimuw.edu.pl/~noble/courses/TimeSeries/data/m-intc7308.
txt"
mintc7308 = read.table(www,header=T)
da = mintc7308
intc = log(da[,2]+1)
plot(intc, type = "l")
acf(intc)
pacf(intc)
acf(intc^2)
pacf(intc^2)
```

   to get the acf and pacf of the series and the squared values. The command for the Ljung Box test is:

```
Box.test(intc,lag=12,type='Ljung')
```

   Now remove the mean and take the square:

```
at=intc-mean(intc)
Box.test(at^2,lag=12,type='Ljung')
```

   For garch fitting:

```
library(fGarch)
m1 = garchFit(intc~garch(1,0),data=intc,trace=F)
```

2. The file `m-mrk4608.txt` contains monthly simple returns of Merck stock from June 1946 to December 2008. The file has two columns denoting date and simple return. Transform the simple returns into log returns.

   (a) Is there any evidence of serial correlations in the log returns? Use autocorrelations and 5% significance level to answer the question. If yes, remove the serial correlations. This is done by fitting an ARMA model. One is then interested in ARCH-GARCH effects in the residuals.

   (b) Is there any evidence of ARCH effects in the log returns? Use the residual series if there are serial correlations in part (a). Use Ljung-Box statistics for the squared returns (or residuals) with 6 and 12 lags of autocorrelations and 5% significance level to answer the question.

   (c) Identify an ARCH model for the data and fit the identified model. Write down the fitted model.

3. The file `m-3m4608.txt` contains two columns. They are date and the monthly simple return for 3M stock. Transform the returns to log returns.

   (a) Is there any evidence of ARCH effects in the log returns? Use Ljung-Box statistics with 6 and 12 lags of autocorrelations and 5% significance level to answer the question.

   (b) use the PACF of the squared returns to identify an ARCH model. What is the fitted model?

   (c) There are 755 data points. Refit the model using the first 750 observations and use the fitted model to predict the folatilities for times 751 to 755 (the forecast origin is 750).

   (d) Build an EGARCH model for the log return series of 3M stocks using the first 750 observations. Use the fitted model to compute the 1-step to 5-step-ahead volatility forecasts at the forecast origin $h = 750$. For EGARCH, you may find the package `betategarch` useful. Look at the documentation.

4. The file `m-gmsp5008.txt` contains the dates and monthy simple returns of General Motors stock and the S&P index from 1950 - 2008.

   (a) Build a GARCH model with Gaussian innovations for the log returns of GM stock. Check the model and write down the fitted model.

   (b) Build a GARCH model with Student-t distribution for the log returns of GM stock including estimation of the degrees of freedom. Write down the fitted model. Let ν be the degrees of freedom. Test the hypothesis $H_0 : \nu = 6$ versus $H_1 : \nu \neq 6$ at a significance level of 5%.

   (c) Build an EGARCH model for the log returns of GM stock. What is the fitted model?

   (d) Obtain 1-step to 6-step-ahead volatility forecasts for all models obtained. Compare the forecasts.

5. Consider again the data in `m-gmsp5008.txt`.

   (a) Build a Gaussian GARCH model for the monthly log returns of the S&P 500 index.

(b) Is there are summer effect on the volatility of the index return? This requires writing an R script. Make a variable

$$u_t = \begin{cases} 1 & \text{month is: June, July, August} \\ 0 & \text{other months.} \end{cases}$$

Try fitting the model:

$$\sigma^2(t) = \alpha X^2(t-1) + \beta\sigma^2(t-1) + \gamma(1-u_t).$$

Then the coefficients are (α, β, γ) for September to May and $(\alpha, \beta, 0)$ for June, July and August. Is γ, which represents the difference, significant?

This is a GARCH(1,1) model with an exogenous variable in the volatility. This is dealt with by the package **rugarch**. You will find the syntax in the documentation. You need `ugarchspec` to specify the model and `ugarchfit` to fit the model.

(c) Are the lagged returns of GM stock useful for modelling the index volatility? Use your GARCH model as a baseline for comparison.

6. The purpose of this exercise is to simulate a GARCH(1,1) process

$$\begin{cases} u(t) = v(t)\sqrt{h(t)} \\ h(t) = a_0 + a_1 u(t-1)^2 + b_1 h(t-1) \\ a_0 = 0.1 \qquad a_1 = 0.4 \qquad b_1 = 0.2 \end{cases}$$

```
> v<-rnorm(1000)
> u<-array(0,1000)
> h<-array(0,1000)
> for(t in 2:1000){}
> for(t in 2:1000){}
> a0 = 0.1; a1 = 0.4; b1 = 0.2
> for(t in 2:1000){
+ h[t] <- a0 + a1*(u[t-1]^2)+b1*h[t-1]
+ u[t]<-v[t]*sqrt(h[t])
+ }
> plot(u,type="l")
> acf(u)
> acf(u^2)
> library("fGarch")
> u.garch <- garchFit(~garch(1,1),u,trace=F)
> u.garch@fit$matcoef
```

How do the estimated coefficients compare with the true coefficients?