TIME VARYING HETEROSKEDASTICITY MODELS: ARCH-GARCH

TIME VARYING HETEROSKEDASTICITY MODELS

In this section, we present illustration examples of estimating GARCH-type models to financial time series using R. To estimate time-varying conditional heteroskedasticity models, the command garchFit(formula=~arma(r,s)+garch(p,q),y) is used, where y is the analyzed return series. This command allows modeling of the mean equation using ARMA(r,s) models, and/or modeling the variance equation using GARCH(p,q)-type models[1]. The default conditional distribution is the normal (norm), however alternative conditional distributions that allow for fat tails can be considered, i.e. the student-t distribution (std), and the generalized error distribution (ged) among others. In order to run the command garchFit(), we need to install the package fGarch. To check the fitted model, say 'm', we can use the command plot(m), where 'm' is the name of the model output from the garchFit() command.

```
# Load library for time-varying volatility models
library(fGarch)
```

Example 1: GARCH modeling of the Intel stock returns

In this example, we fit a GARCH-type model to Intel stock monthly returns for the period 1973:01-2003:12. We will also compute several summary statistics and plots. First, load the data:

```
# Load data: Intel stock monthly returns (1973:01 - 2003:12)
data1<-read.table("…..intel-example.txt",header=T)
data1
head(data1)
y <- data1$rtn
y
```

Compute summary statistics and plots:

[1] Note that the R command garch() can be used to estimate GARCH models, however does not allow for exogenous variables or ARMA terms in the mean equation.

```
# Summary Statistics and plots
par(mfrow=c(3,2))
plot(y, type="l", main="intel stock monthly returns")
hist(y, main="histogram of returns")
acf(y, 12, main="ACF of returns")
pacf(y, 12, main="PACF of returns")
acf(y^2,12, main="ACF of squared returns")
pacf(y^2, 12, main="PACF of squared returns")
```

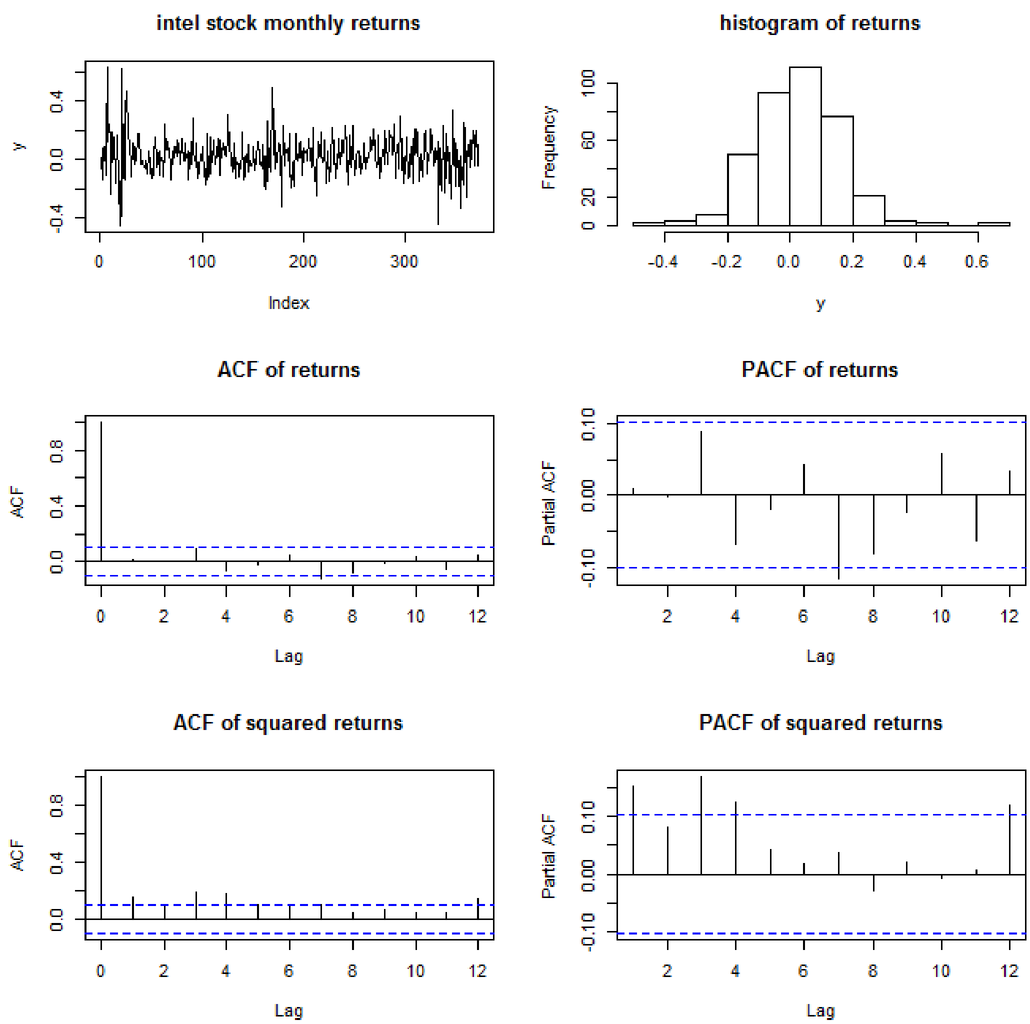


```
Box.test(y,lag=12,type="Ljung")
```

Box-Ljung test: data: y, X-squared = 16.853, df = 12, p-value = 0.1552

```
Box.test(y^2,lag=12,type="Ljung")
```

Box-Ljung test: data: y^2, X-squared = 58.215, df = 12, p-value = 4.765e-08

A brief discussion on the characteristics of the analyzed returns follows. From the upper left plot of the return series we observe that the volatility of the return series is not constant over time. There is evidence of volatility clustering phenomenon, since there are periods of high and low volatility. By applying Ljung-Box tests based on 12 lags on the return series, it seems that there is no evidence for autocorrelation in the return series, see also the autocorrelation and partial autocorrelation plots. However, applying Ljung-Box test on the squared returns shows that the null hypothesis of no autocorrelation is rejected, indicating heteroscedastic effects, i.e. the volatility of the return series is time-varying. To take into account for these characteristics, we will estimate ARCH/GARCH models.

First, we will fit an ARCH(1) model:

```
# Estimate ARCH(1) model
m1arch=garchFit(~garch(1,0),data=y,trace=F)      # trace = F   reduces the summary
summary(m1arch)
```

```
Title:  GARCH Modelling
Call:  garchFit(formula = ~garch(1, 0), data = y, trace = F)
Mean and Variance Equation:   data ~ garch(1, 0),  [data = y]
Conditional Distribution:  norm
Coefficient(s):
      mu        omega      alpha1
0.023985  0.012250  0.368317
Std. Errors:  based on Hessian Error Analysis:
            Estimate    Std. Error    t value Pr(>|t|)
mu        0.023985    0.006127    3.915 9.05e-05 ***
omega   0.012250    0.001474    8.308  < 2e-16 ***
alpha1   0.368317    0.121678    3.027  0.00247 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
Log Likelihood:  231.3696    normalized:  0.6219614
Standardised Residuals Tests:
                                    Statistic    p-Value
```

```
 Jarque-Bera Test    R    Chi^2  46.00786  1.022161e-10
 Shapiro-Wilk Test   R    W      0.9825606 0.0001809205
 Ljung-Box Test      R    Q(10)  13.23168  0.2110055
 Ljung-Box Test      R    Q(15)  20.35255  0.1588215
 Ljung-Box Test      R    Q(20)  21.34622  0.3770122
 Ljung-Box Test      R^2  Q(10)  10.35347  0.4100484
 Ljung-Box Test      R^2  Q(15)  26.08889  0.0370987
 Ljung-Box Test      R^2  Q(20)  28.45493  0.09906006
 LM Arch Test        R    TR^2   22.63579  0.03098202
Information Criterion Statistics:
     AIC       BIC       SIC      HQIC
-1.227794 -1.196190 -1.227922 -1.215243
```

The command m1arch=garchFit(~garch(1,0),data=y,trace=F) returns the m1arch object. A summary of this object is obtained through the command summary(m1arch). We obtain the parameter estimates, the corresponding standard errors, the t-statistics and the associated p-values, in order to examine the statistical significance of the model parameters. The ARCH(1) model is written:

$$y_t = \mu + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_{t,} \quad v_t \sim N(0,1)$$
$$\sigma_t^2 = a_0 + a_1 \varepsilon_{t-1}^2$$

Thus, the estimated ARCH(1) model is:

$$y_t = 0.0239 + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_{t,} \quad v_t \sim N(0,1)$$
$$\sigma_t^2 = 0.012 + 0.368 \varepsilon_{t-1}^2$$

We also obtain the log-likelihood and several information criteria to evaluate the fit of the estimated model. Finally, there are different diagnostic tests performed on the residuals [R] and/or the squared residuals [R^2]. The Jarque-Bera and the Shapiro-Wilk tests are used in order to examine the normality assumption of the model residuals. It seems that the null hypothesis of normality is rejected, i.e. there is evidence of non-normality of residuals. Most of the remaining tests are based on the Ljung-Box Q-statistic to examine the autocorrelation of the residuals and/or the squared residuals. The results indicate that there is no autocorrelation in the residuals series, however there is some autocorrelation in the squared residuals. That is, the ARCH(1) model did not capture adequately the time-varying characteristics of the return series.

Similar results are obtained by analyzing the returns series in Eviews:

Dependent Variable: Y
Method: ML - ARCH (Marquardt) - Normal distribution
Sample: 1 372
Included observations: 372
Convergence achieved after 9 iterations
Presample variance: backcast (parameter = 0.7)
GARCH = C(2) + C(3)*RESID(-1)^2

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.023991	0.005987	4.007148	0.0001
Variance Equation				
C	0.012264	0.001337	9.173166	0.0000
RESID(-1)^2	0.366818	0.102256	3.587266	0.0003

R-squared	-0.000526	Mean dependent var	0.027067
Adjusted R-squared	-0.005949	S.D. dependent var	0.134243
S.E. of regression	0.134642	Akaike info criterion	-1.227531
Sum squared resid	6.689421	Schwarz criterion	-1.195928
Log likelihood	231.3209	Hannan-Quinn criter.	-1.214981
Durbin-Watson stat	1.980747		

Note, that we can have a deeper access to the results of m1arch object by using the @ operator. For example, the command coef <- m1arch@fit$coef returns to coef the model estimates, while the command fittedvar <- m1arch@fit$series$h returns to fittedvar the estimated sigma squared (σ^2) process. For details, see the command m1arch@fit.

A useful command to evaluate the adequacy of the model fit via several graphs is the following:

```
plot(m1arch)
```

which returns a list of thirteen plots:

```
Make a plot selection (or 0 to exit):
 1:   Time Series
 2:   Conditional SD
 3:   Series with 2 Conditional SD Superimposed
 4:   ACF of Observations
 5:   ACF of Squared Observations
 6:   Cross Correlation
 7:   Residuals
 8:   Conditional SDs
 9:   Standardized Residuals
10:   ACF of Standardized Residuals
11:   ACF of Squared Standardized Residuals
```

12:	Cross Correlation between r^2 and r
13:	QQ-Plot of Standardized Residuals

For example, by selecting the option (10) and (11) we obtain the autocorrelation (ACF) plot of standardized residuals and the autocorrelation (ACF) plot of squared standardized residuals, confirming that there is no autocorrelation in the residuals series, however there is some autocorrelation in the squared residuals.

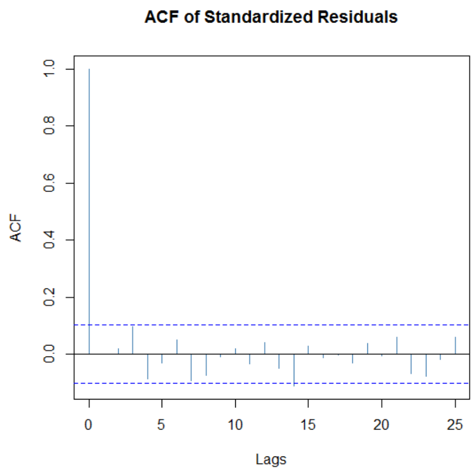

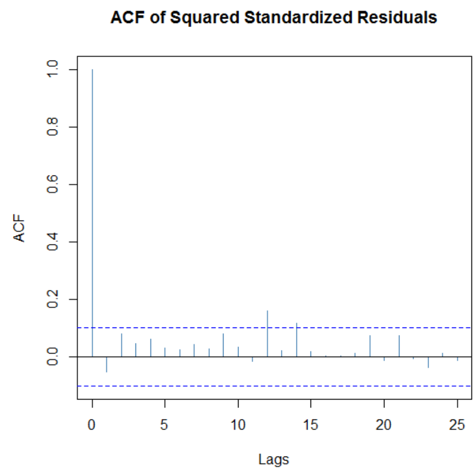

By selecting the last option (13) we obtain the normal QQ-plot of standardized residuals, indicating that the standardized residuals deviate from normality especially in the tails of the distribution.

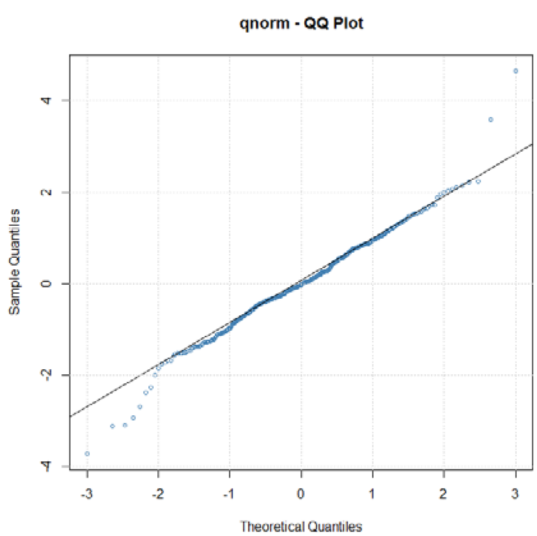

Predictions based on the ARCH(1) model can be obtained through the command predict(). For example, the command

```
predict(m1arch,6)
```

returns predictions of the mean and the standard deviation of the return series:

```
  meanForecast meanError standardDeviation
1  0.02398497 0.1182073        0.1182073
2  0.02398497 0.1318939        0.1318939
3  0.02398497 0.1365898        0.1365898
4  0.02398497 0.1382792        0.1382792
5  0.02398497 0.1388962        0.1388962
6  0.02398497 0.1391228        0.1391228
```

We can estimate an ARCH(1) model based on a student-t distribution for the error process. The following set of commands estimates the model, presents a summary of the obtained results, gives several graphs to evaluate the model fit, and computes predictions:

```
m1archst=garchFit(~garch(1,0),data=y,cond.dist="std",trace=F)
summary(m1archst)
plot(m1archst)
predict(m1archst,6)
```

The results are:

```
Title:  GARCH Modelling
Call:  garchFit(formula = ~garch(1, 0), data = y, cond.dist = "std",   trace = F)
Mean and Variance Equation:  data ~ garch(1, 0),  [data = y]
Conditional Distribution: std
Coefficient(s):
     mu       omega      alpha1     shape
0.024591  0.012917  0.295035  7.544325
Std. Errors:  based on Hessian Error Analysis:
            Estimate   Std. Error    t value  Pr(>|t|)
mu         0.024591    0.006073    4.049 5.15e-05 ***
omega    0.012917    0.001791    7.213 5.47e-13 ***
alpha1   0.295035    0.118014    2.500   0.0124 *
shape     7.544325    2.583299    2.920   0.0035 **
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
Log Likelihood: 238.8183    normalized:  0.6419846
Standardised Residuals Tests:
                        Statistic p-Value
```

```
 Jarque-Bera Test   R    Chi^2  49.00158  2.287937e-11
 Shapiro-Wilk Test  R    W      0.9819735 0.0001337123
 Ljung-Box Test     R    Q(10)  13.56394  0.1938264
 Ljung-Box Test     R    Q(15)  20.99709  0.1369222
 Ljung-Box Test     R    Q(20)  21.94734  0.3433748
 Ljung-Box Test     R^2  Q(10)  12.06373  0.2808139
 Ljung-Box Test     R^2  Q(15)  30.03666  0.01179
 Ljung-Box Test     R^2  Q(20)  32.81887  0.03531708
 LM Arch Test       R    TR^2   24.86223  0.01548936
Information Criterion Statistics:
     AIC       BIC       SIC      HQIC
-1.262464 -1.220325 -1.262692 -1.245729
```

The estimated ARCH(1) model based on Student-t errors is:

$$y_t = 0.0245 + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_t, \quad v_t \sim Student - t_{7.54}$$
$$\sigma_t^2 = 0.0129 + 0.295\varepsilon_{t-1}^2$$

The autocorrelation (ACF) plot of standardized residuals (option 10) and the autocorrelation (ACF) plot of squared standardized residuals (option 11), presented below, indicate that there is no autocorrelation in the residuals series, however there is some autocorrelation in the squared residuals.

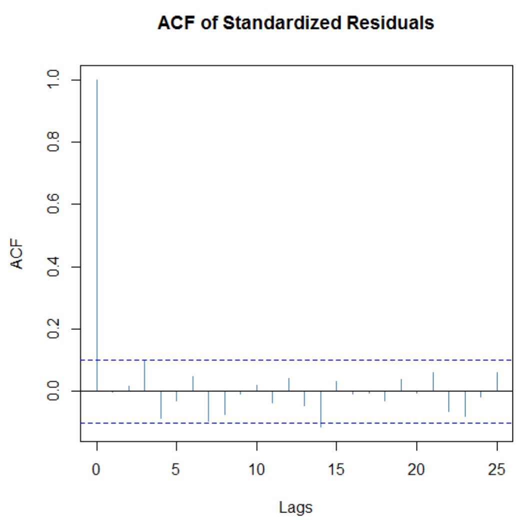

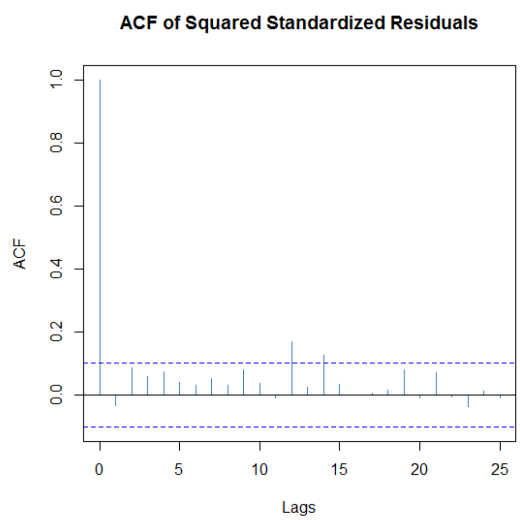

The QQ-plot of standardized residuals based on the Student-t distribution, indicate a better fit than those based on the Normal distribution. Obviously, the fit of the Intel returns based on the ARCH(1) model with Student-t errors provides a more appropriate modeling approach than that of ARCH(1) model based on normal errors.

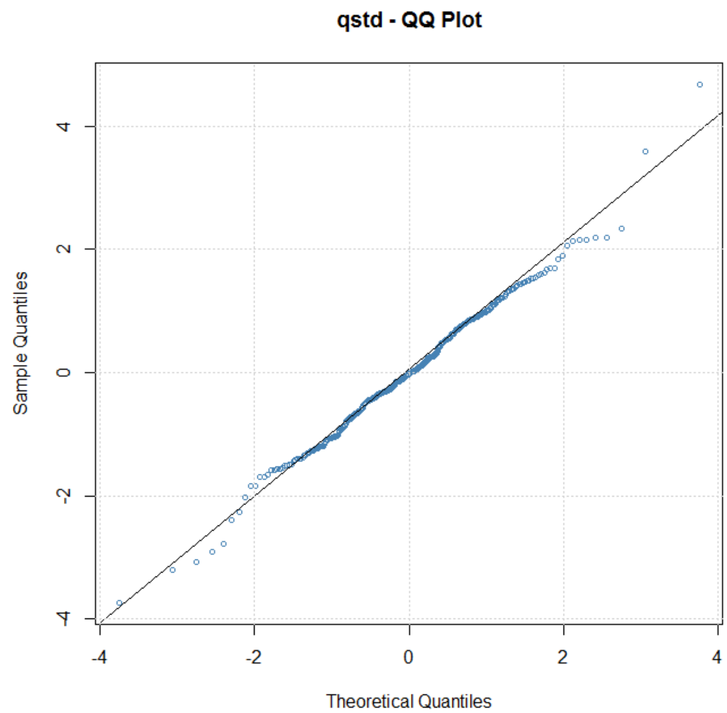

The predictions based on the Student-t ARCH(1) model are:

```
meanForecast meanError standardDeviation
1  0.02459108 0.1196754        0.1196754
2  0.02459108 0.1309299        0.1309299
3  0.02459108 0.1340700        0.1340700
4  0.02459108 0.1349825        0.1349825
5  0.02459108 0.1352506        0.1352506
6  0.02459108 0.1353295        0.1353295
```

Similar results are obtained by estimating the ARCH(1) model with Student-t errors in Eviews:

Dependent Variable: Y
Method: ML - ARCH (Marquardt) - Student's t distribution
Sample: 1 372
Included observations: 372
Convergence achieved after 7 iterations
Presample variance: backcast (parameter = 0.7)
GARCH = C(2) + C(3)*RESID(-1)^2

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.024598	0.006051	4.065034	0.0000
Variance Equation				
C	0.012934	0.001820	7.108712	0.0000
RESID(-1)^2	0.293488	0.110641	2.652614	0.0080
T-DIST. DOF	7.539362	2.490080	3.027760	0.0025

R-squared	-0.000339	Mean dependent var	0.027067
Adjusted R-squared	-0.008494	S.D. dependent var	0.134243
S.E. of regression	0.134812	Akaike info criterion	-1.262246
Sum squared resid	6.688170	Schwarz criterion	-1.220107
Log likelihood	238.7778	Hannan-Quinn criter.	-1.245512
Durbin-Watson stat	1.981118		

We can estimate a GARCH(1,1) model based on normal distribution for the error process. The following set of commands estimates the model, presents a summary of the obtained results, gives several graphs to evaluate the model fit, and computes predictions:

```
m2garch=garchFit(~garch(1,1),data=y,trace=F)
summary(m2garch)
plot(m2garch)
predict(m2garch,6)
```

The GARCH(1,1) model is written:

$$y_t = \mu + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_t, \quad v_t \sim N(0,1)$$
$$\sigma_t^2 = a_0 + a_1 \varepsilon_{t-1}^2 + b_1 \sigma_{t-1}^2$$

Based on the results presented below, the estimated GARCH(1,1) model is:

$$y_t = 0.0257 + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_{t,} \quad v_t \sim N(0,1)$$
$$\sigma_t^2 = 0.0012 + 0.088\varepsilon_{t-1}^2 + 0.837\sigma_{t-1}^2$$

The results obtained in R are:

```
Title: GARCH Modelling
Call:  garchFit(formula = ~garch(1, 1), data = y, trace = F)
Mean and Variance Equation:  data ~ garch(1, 1),  [data = y]
Conditional Distribution:  norm

Coefficient(s):
      mu         omega       alpha1      beta1
0.0257647  0.0012353  0.0883214  0.8379828
Std. Errors:  based on Hessian Error Analysis:
          Estimate      Std. Error     t value Pr(>|t|)
mu       0.0257647   0.0062869    4.098 4.16e-05 ***
omega  0.0012353   0.0006171    2.002   0.0453 *
alpha1  0.0883214   0.0328054    2.692   0.0071 **
beta1    0.8379828   0.0551283   15.201  < 2e-16 ***

Log Likelihood:  237.9957    normalized:  0.6397733

Standardised Residuals Tests:
                        Statistic p-Value
 Jarque-Bera Test   R    Chi^2  64.15007  1.176836e-14
 Shapiro-Wilk Test  R    W      0.9819539 0.0001323769
 Ljung-Box Test     R    Q(10)  9.972497  0.4429095
 Ljung-Box Test     R    Q(15)  16.40237  0.3558257
 Ljung-Box Test     R    Q(20)  17.24923  0.6367361
 Ljung-Box Test     R^2  Q(10)  0.5044128 0.9999931
 Ljung-Box Test     R^2  Q(15)  10.38199  0.7950486
 Ljung-Box Test     R^2  Q(20)  12.24078  0.9075393
 LM Arch Test       R    TR^2   10.59664  0.5637653

Information Criterion Statistics:
     AIC      BIC      SIC      HQIC
-1.258041 -1.215903 -1.258269 -1.241307
```

The estimated output obtain in Eviews for the GARCH(1,1) model with Normal errors is:

Dependent Variable: Y
Method: ML - ARCH (Marquardt) - Normal distribution
Sample: 1 372
Included observations: 372
Convergence achieved after 9 iterations
Presample variance: backcast (parameter = 0.7)
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.025786	0.005851	4.407438	0.0000
		Variance Equation		
C	0.001026	0.000516	1.987842	0.0468
RESID(-1)^2	0.078929	0.027446	2.875790	0.0040
GARCH(-1)	0.857852	0.046998	18.25275	0.0000

R-squared	-0.000091	Mean dependent var	0.027067
Adjusted R-squared	-0.008244	S.D. dependent var	0.134243
S.E. of regression	0.134796	Akaike info criterion	-1.263722
Sum squared resid	6.686512	Schwarz criterion	-1.221584
Log likelihood	239.0523	Hannan-Quinn criter.	-1.246988
Durbin-Watson stat	1.981609		

The autocorrelation (ACF) plot of standardized residuals (option 10) and the autocorrelation (ACF) plot of squared standardized residuals (option 11), based on the Normal-GARCH(1,1) model are:

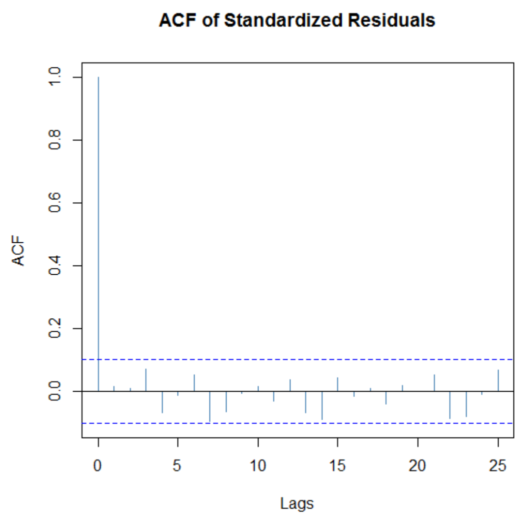

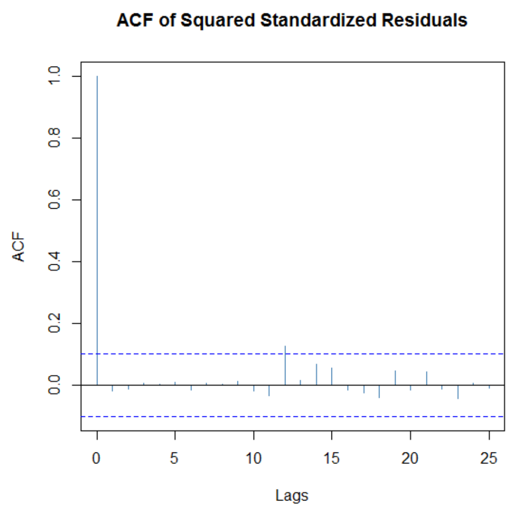

The Normal-QQplot (option 13) is:

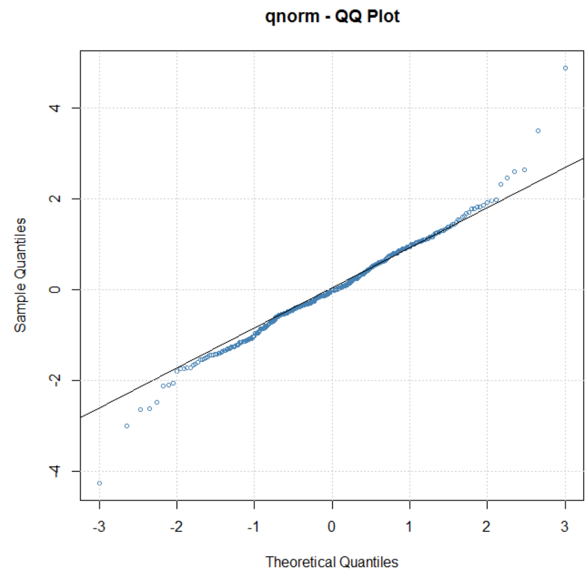

	meanForecast	meanError	standardDeviation
1	0.02576473	0.1236512	0.1236512
2	0.02576473	0.1240893	0.1240893
3	0.02576473	0.1244937	0.1244937
4	0.02576473	0.1248672	0.1248672
5	0.02576473	0.1252121	0.1252121
6	0.02576473	0.1255308	0.1255308

Finally, we estimate a GARCH(1,1) model based on a Student-t distribution for the error process. The following set of commands estimate the model, present a summary of the obtained results, give several graphs to evaluate the model fit, and compute predictions:

```
m2garchst=garchFit(~garch(1,1),data=y, cond.dist="std",trace=F)
summary(m2garchst)
plot(m2garchst)
predict(m2garchst,6)
```

The estimated GARCH(1,1) model based on Student-t errors is:

$$y_t = 0.025 + \varepsilon_t,$$
$$\varepsilon_t = \sigma_t v_{t,} \quad v_t \sim Student - t_{7.96}$$
$$\sigma_t^2 = 0.0014 + 0.1037\varepsilon_{t-1}^2 + 0.8119\sigma_{t-1}^2$$

The results obtained in R are:

```
Title: GARCH Modelling
Call:  garchFit(formula = ~garch(1, 1), data = y, cond.dist = "std", trace = F)
Mean and Variance Equation: data ~ garch(1, 1),  [data = y]
Conditional Distribution: std

Coefficient(s):
      mu        omega      alpha1     beta1      shape
0.0253871  0.0014384  0.1037353  0.8119818  7.9642135
Std. Errors: based on Hessian Error Analysis:
            Estimate   Std. Error    t value  Pr(>|t|)
mu       0.0253871   0.0060314    4.209 2.56e-05 ***
omega  0.0014384   0.0008013    1.795  0.07264 .
alpha1  0.1037353   0.0417075    2.487  0.01288 *
beta1    0.8119818   0.0705286   11.513  < 2e-16 ***
shape    7.9642135   2.7249221    2.923  0.00347 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Log Likelihood: 245.9957    normalized:  0.6612786
Standardised Residuals Tests:
                        Statistic p-Value
 Jarque-Bera Test   R    Chi^2  67.32763  2.442491e-15
 Shapiro-Wilk Test  R    W      0.9815083 0.0001055278
 Ljung-Box Test     R    Q(10)  9.849992  0.4537512
 Ljung-Box Test     R    Q(15)  16.36013  0.3585296
 Ljung-Box Test     R    Q(20)  17.27417  0.6351056
 Ljung-Box Test     R^2  Q(10)  0.6379137 0.9999789
 Ljung-Box Test     R^2  Q(15)  10.73515  0.7711331
 Ljung-Box Test     R^2  Q(20)  12.49515  0.8979801
 LM Arch Test       R    TR^2   11.31707  0.5019586

Information Criterion Statistics:
     AIC      BIC      SIC      HQIC
-1.295676 -1.243002 -1.296031 -1.274758
```

The estimated output obtain in Eviews for the GARCH(1,1) model with Student-t errors is:

Dependent Variable: Y
Method: ML - ARCH (Marquardt) - Student's t distribution
Sample: 1 372
Included observations: 372
Convergence achieved after 27 iterations
Presample variance: backcast (parameter = 0.7)
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.025425	0.006000	4.237553	0.0000
Variance Equation				
C	0.001335	0.000912	1.463865	0.1432
RESID(-1)^2	0.099002	0.052391	1.889669	0.0588
GARCH(-1)	0.821602	0.087544	9.384964	0.0000
T-DIST. DOF	8.144813	2.793239	2.915902	0.0035

R-squared	-0.000150	Mean dependent var	0.027067
Adjusted R-squared	-0.011051	S.D. dependent var	0.134243
S.E. of regression	0.134983	Akaike info criterion	-1.296606
Sum squared resid	6.686905	Schwarz criterion	-1.243933
Log likelihood	246.1688	Hannan-Quinn criter.	-1.275688
Durbin-Watson stat	1.981493		

The autocorrelation (ACF) plot of standardized residuals (option 10) and the autocorrelation (ACF) plot of squared standardized residuals (option 11), based on the Student-t GARCH(1,1) model are:

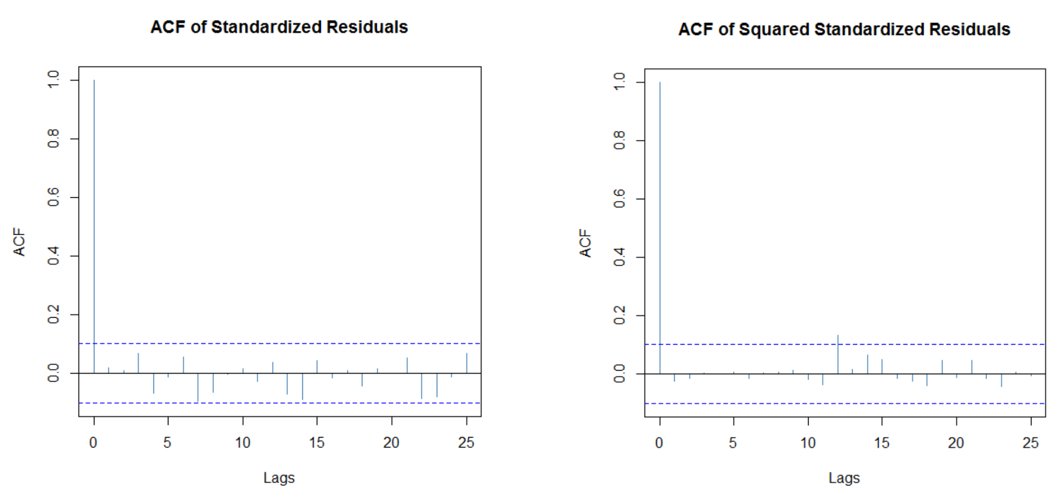

The QQplot of the residuals of the Student-t GARCH(1,1) model (option 13) and a plot of the return series together with the estimated conditional standard deviations (x2) obtained by option (3) are illustrated below:

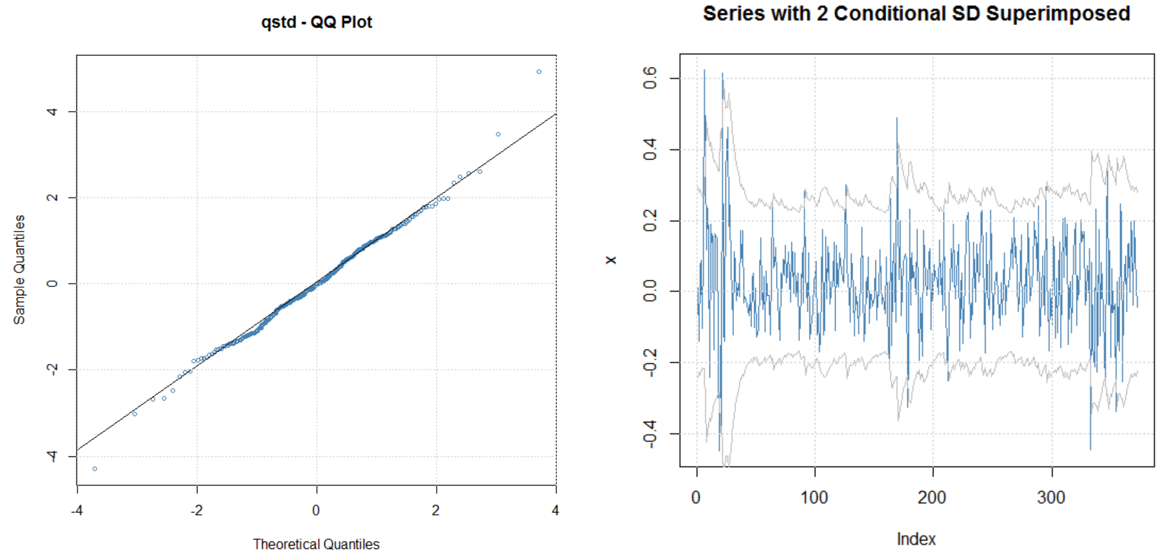

Finally, the predictions based on this Student-t GARCH(1,1) model are presented:

	meanForecast	meanError	standardDeviation
1	0.02538713	0.1216621	0.1216621
2	0.02538713	0.1224442	0.1224442
3	0.02538713	0.1231560	0.1231560
4	0.02538713	0.1238042	0.1238042
5	0.02538713	0.1243948	0.1243948
6	0.02538713	0.1249332	0.1249332

Example 2: GARCH modeling of the S&P500 index

In this example, we fit GARCH-type models to monthly excess returns of S&P500 index. Below are presented the R commands for this example:

```
#==========
# EXAMPLE 2
#==========
# Monthly excess returns of S&P 500 index starting from 1926 for 792 observations.
y=scan("…/sp500-example.txt")
y

# Summary Statistics and plots
par(mfrow=c(3,2))
plot(y, type="l")
hist(y, nclass=15)
acf(y)
pacf(y)
acf(y^2)
pacf(y^2)
```

```
Box.test(y,lag=10,type="Ljung")
Box.test(y^2,lag=10,type="Ljung")
```

```
m1=arima(y,order=c(3,0,0))
m1
acf(m1$residuals)
pacf(m1$residuals)
```

```
m2=garchFit(~arma(3,0)+garch(1,1),data=y,trace=F)
summary(m2)
```

```
m3=garchFit(~garch(1,1),data=y,trace=F)
summary(m3)
plot(m3)
predict(m3,6)
```

```
m4=garchFit(~garch(1,1),data=y,cond.dist="std",trace=F)
summary(m4)
plot(m4)
predict(m4,6)
```