# ARMA-GARCH modelling and white noise tests

**James Proberts**
University of Manchester

**Georgi N. Boshnakov**
University of Manchester

**Abstract**

This vignette illustrates applications of white noise tests in GARCH modelling. It is based on an example from an MMath project by the first author.

*Keywords*: autocorrelations, white noise tests, IID tests, GARCH models, time series.

## 1. The data

In this example we consider data from Freddie Mac, a mortgage loan company in the USA. This stock is an interesting case for study. In the financial crash of 2008 it dropped from roughly \$60 to \$0.5 over the course of a year. It is now (April 2017) majority owned by the government and has all its profits and dividends sweeped. There has been speculation on this stock being returned to private ownership for years making it prone to clusters of volatility.

We import weekly data from Yahoo Finance covering the period from 10/05/2006 to 22/04/2017, and calculate the weekly simple log returns.

```
R> ## using a saved object, orginally imported with:
R> ##  FMCC <- yahooSeries("FMCC", from = "2006-05-10", to = "2017-04-22",
R> ##                           freq = "weekly")
R> FMCC <- readRDS(system.file("extdata", "FMCC.rds", package = "sarima"))
R> ###### JAMIE (23/08/2018): above command wouldn't work for me!
R> ##  @JAMIE (2018-10-11): should work now. If it doesn't install the package.
R> ##      (.gitignore had a rule that prevented pushing the file on bibbucket)
R> ##FMCC <- readRDS("FMCC.rds")
R> logreturns <- diff(rev(log(FMCC$FMCC.Close)))
```

A plot of the log-returns. is given in Fig. 1. We also calculate the autocorrelations and partial autocorrelations for the log returns.

```
R> FMCClr.acf <- autocorrelations(logreturns)
R> FMCClr.pacf <- partialAutocorrelations(logreturns)
```

## 2. Autocorrelations

We now produce a plot of the autocorrelations to assess whether the series is autocorrelated, see Fig. 2. There are two bounds plotted on the graph. The straight red line represents

```
R> plot(logreturns, type="l", main="Log-returns of FMCC")
```

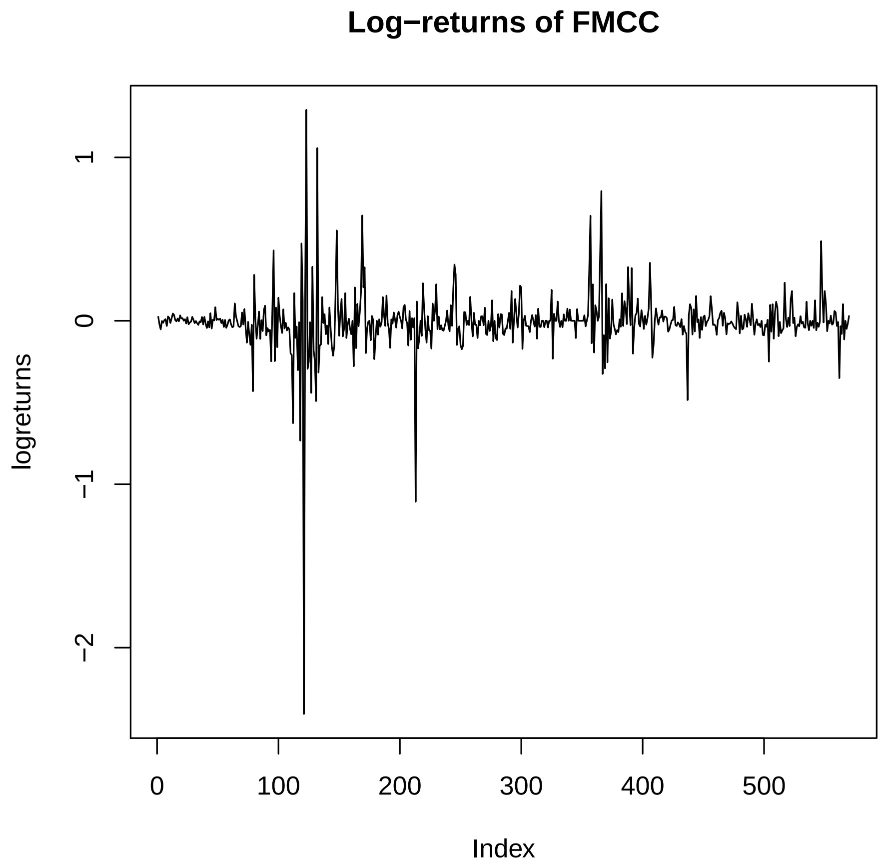


Figure 1: Log-returns of weekly log-returns of FMCC from 10 May 2006 to 22 Apr 2017.

the standard bounds under the strong white noise assumption. The second line is under the hypothesis that the process is GARCH.

Several autocorrelations seem significant under the iid hypothesis. This may lead us to fitting an ARMA or ARMA-GARCH model. On the other hand, the autocorrelations are well into the bands produced under the GARCH hypothesis, suggesting a pure GARCH model, without any ARMA terms. So, it matters on which test we base our decision.

The partial autocorrelation function can be used instead of the autocorrelations, with similar inferences, see Fig. 3.

## 3. Pormanteau tests

Routine portmanteau tests, such as Ljung-Box, also reject the IID hypothesis. Here we carry out IID tests using the method of Li-McLeod:

```
R> wntLM <- whiteNoiseTest(FMCClr.acf, h0 = "iid", nlags = c(5,10,20),
```

```
R> plot(FMCClr.acf, data = logreturns)
```

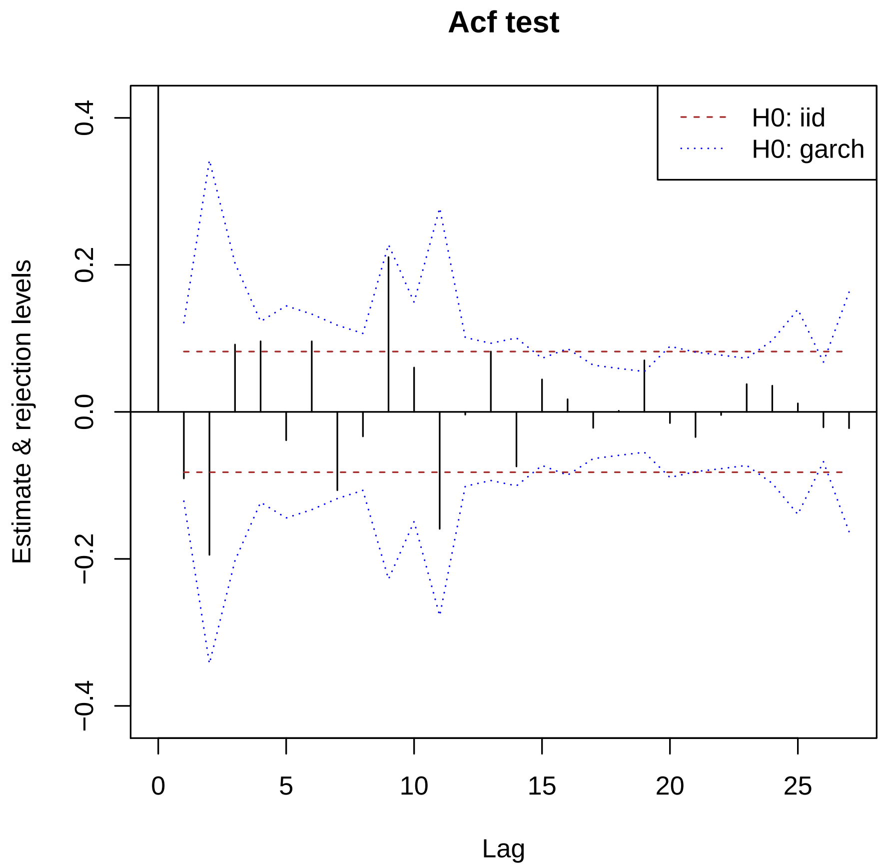


Figure 2: Autocorrelation test of the log returns of FMCC

```
+                          x = logreturns, method = "LiMcLeod")
R> wntLM$test
```

```
         ChiSq DF       pvalue
[1,]  37.18469  5 5.499929e-07
[2,]  76.99131 10 1.946524e-12
[3,] 103.19392 20 3.363466e-13
attr(,"method")
[1] "LiMcLeod"
```

Small p-values lead to rejection of the null hypothesis at reasonable levels. Rejection of the null hypothesis is often taken to mean that the data are autocorrelated.

Let us test for fitting a GARCH-type model by using the following code which has the weaker assumption that the log returns are GARCH. Let us change the null hypothesis to "garch" (one possible weak white noise hypothesis):

```
R> plot(FMCClr.pacf, data = logreturns,
+ main="Partial Autocorrelation test of the log returns of FMCC")
```

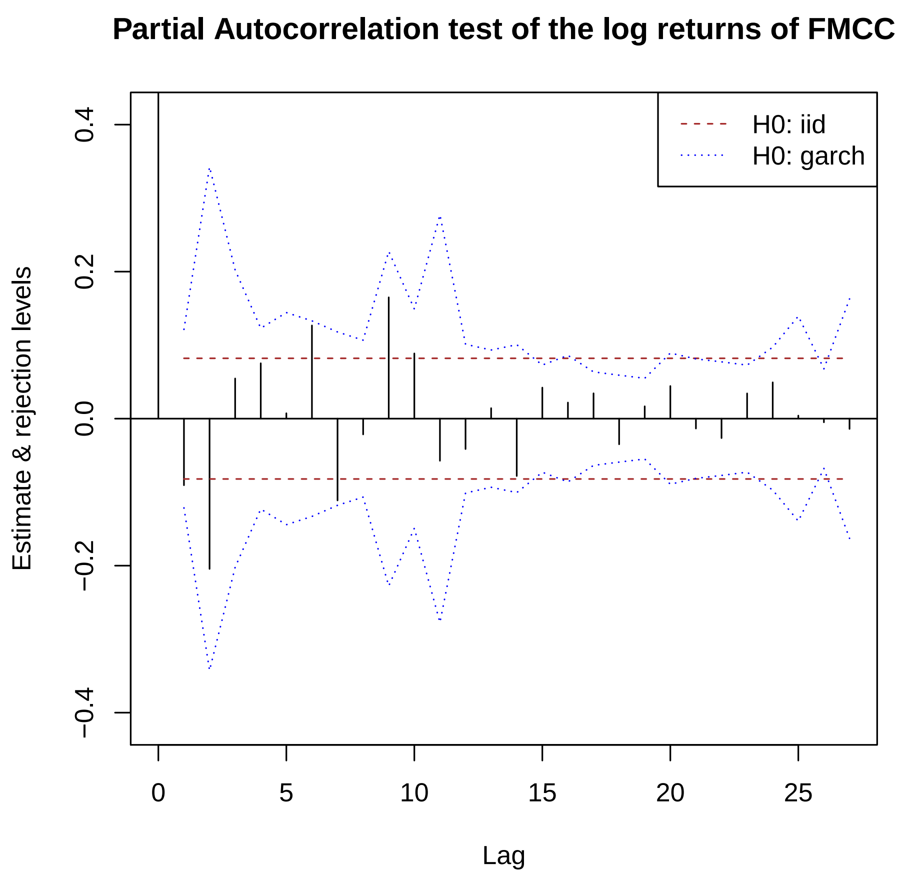

```
R> wntg <- whiteNoiseTest(FMCClr.acf, h0 = "garch", nlags = c(5,10,15), x = logreturns)
R> wntg$test
```

```
      h         Q      pval
[1,]  5  5.115077 0.4019983
[2,] 10 12.667143 0.2428821
[3,] 15 18.036831 0.2607342
```

The high p-values give no reason to reject the hypothesis that the log-returns are a GARCH white noise process. In other words, there is no need to ARMA modelling.

# 4. Fitting GARCH(1,1) models and their variants

Based on the discussion above, we go on to fit GARCH model(s), starting with a GARCH(1,1) model with Gaussian innovations.

```
R> fit1 <- garchFit(~garch(1,1), data = logreturns, trace = FALSE)
R> summary(fit1)
```

```
Title:
 GARCH Modelling

Call:
 garchFit(formula = ~garch(1, 1), data = logreturns, trace = FALSE)

Mean and Variance Equation:
 data ~ garch(1, 1)
<environment: 0x0000000010706798>
 [data = logreturns]

Conditional Distribution:
 norm

Coefficient(s):
         mu        omega       alpha1        beta1
-6.3541e-05   2.9206e-03   4.3649e-01   5.8992e-01

Std. Errors:
 based on Hessian

Error Analysis:
         Estimate  Std. Error  t value Pr(>|t|)
mu     -6.354e-05   5.006e-03   -0.013     0.99
omega   2.921e-03   6.982e-04    4.183 2.87e-05 ***
alpha1  4.365e-01   7.623e-02    5.726 1.03e-08 ***
beta1   5.899e-01   5.427e-02   10.869  < 2e-16 ***
```

```
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Log Likelihood:
 341.7229    normalized:  0.5995139

Description:
 Sun May 12 20:52:49 2019 by user: mcbssgb2


Standardised Residuals Tests:
                                Statistic p-Value
 Jarque-Bera Test   R    Chi^2  24230.08  0
 Shapiro-Wilk Test  R    W      0.7433933 0
 Ljung-Box Test     R    Q(10)  9.525801  0.4830325
 Ljung-Box Test     R    Q(15)  12.92386  0.6081792
 Ljung-Box Test     R    Q(20)  14.75224  0.7904048
 Ljung-Box Test     R^2  Q(10)  0.7315935 0.9999597
 Ljung-Box Test     R^2  Q(15)  0.9445704 0.9999998
 Ljung-Box Test     R^2  Q(20)  1.338934  1
 LM Arch Test       R    TR^2   0.8791397 0.9999931

Information Criterion Statistics:
      AIC       BIC       SIC      HQIC
-1.184993 -1.154497 -1.185090 -1.173094
```

The diagnostics suggest that the standardised residuals and their squares are IID and that the ARCH effects have been accommodated by the model. Their distribution is clearly not Gaussian however (see the p-values for Jarque-Bera and Shapiro-Wilk Tests), so another conditional distribution can be tried.

Another possible problem is that $\alpha_1 + \beta_1 > 0$.

```
R> fit2 <- garchFit(~garch(1,1), cond.dist = c("sstd"), data = logreturns, trace = FALSE)
R> summary(fit2)


Title:
 GARCH Modelling

Call:
 garchFit(formula = ~garch(1, 1), data = logreturns, cond.dist = c("sstd"),
    trace = FALSE)

Mean and Variance Equation:
 data ~ garch(1, 1)
<environment: 0x000000001217e4a8>
 [data = logreturns]
```

```
Conditional Distribution:
 sstd

Coefficient(s):
        mu       omega      alpha1       beta1        skew       shape
0.00024523  0.00277227  0.99999999  0.73057510  1.16531856  2.14375224

Std. Errors:
 based on Hessian

Error Analysis:
        Estimate  Std. Error  t value Pr(>|t|)
mu     0.0002452   0.0033295    0.074   0.9413
omega  0.0027723   0.0017142    1.617   0.1058
alpha1 1.0000000   0.5302728    1.886   0.0593 .
beta1  0.7305751   0.0763615    9.567   <2e-16 ***
skew   1.1653186   0.0576821   20.202   <2e-16 ***
shape  2.1437522   0.0969271   22.117   <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Log Likelihood:
 555.2528    normalized:  0.9741278

Description:
 Sun May 12 20:52:49 2019 by user: mcbssgb2


Standardised Residuals Tests:
                                Statistic p-Value
 Jarque-Bera Test   R    Chi^2  27547.01  0
 Shapiro-Wilk Test  R    W      0.7324028 0
 Ljung-Box Test     R    Q(10)  7.820836  0.6463324
 Ljung-Box Test     R    Q(15)  10.34984  0.7971759
 Ljung-Box Test     R    Q(20)  11.87712  0.9202405
 Ljung-Box Test     R^2  Q(10)  0.7097748 0.9999651
 Ljung-Box Test     R^2  Q(15)  1.089078  0.9999995
 Ljung-Box Test     R^2  Q(20)  1.449253  1
 LM Arch Test       R    TR^2   0.9024198 0.999992

Information Criterion Statistics:
      AIC       BIC       SIC      HQIC
-1.927203 -1.881459 -1.927421 -1.909355
```

The qq-plot of the standardised residuals, suggests that the fitted standardised skew-t conditional distribution is not good enough.

```
R> plot(fit2, which = 13)
```

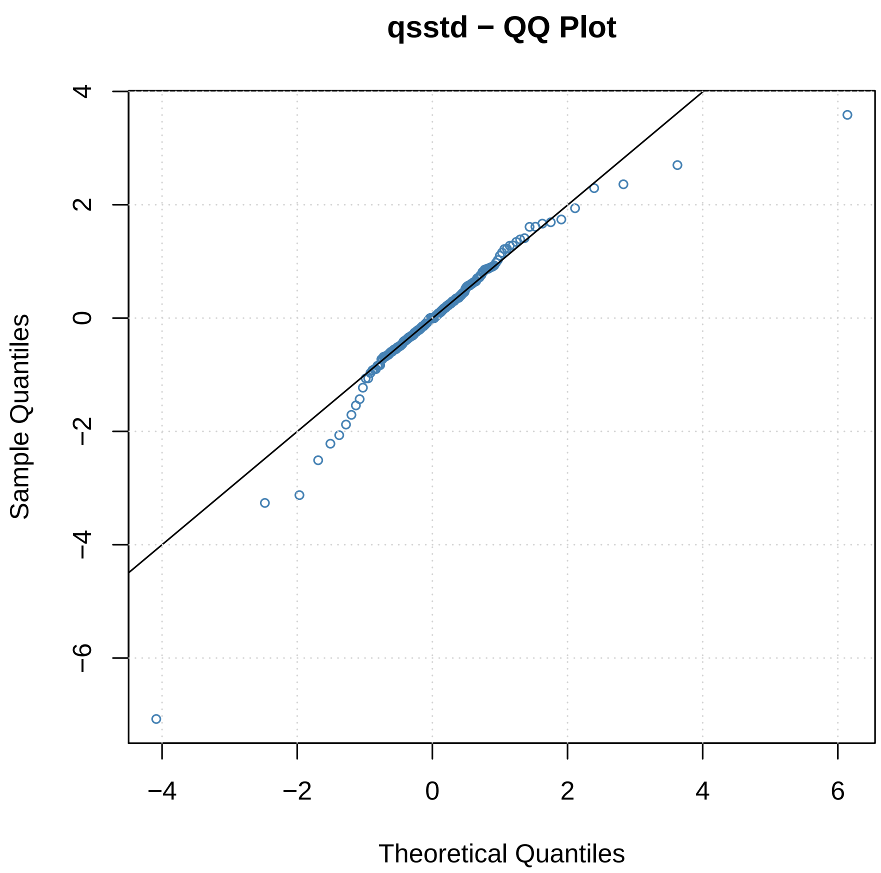


```
R> fit3 <- garchFit(~aparch(1,1), cond.dist = c("sstd"), data = logreturns, trace = FALSE)
R> summary(fit3)

Title:
 GARCH Modelling

Call:
 garchFit(formula = ~aparch(1, 1), data = logreturns, cond.dist = c("sstd"),
    trace = FALSE)

Mean and Variance Equation:
 data ~ aparch(1, 1)
<environment: 0x0000000010469808>
 [data = logreturns]

Conditional Distribution:
 sstd
```

```
Coefficient(s):
       mu      omega     alpha1     gamma1      beta1      delta       skew
0.0034733  0.0425723  1.0000000  0.2025020  0.7970743  0.7374463  1.2050209
    shape
2.0099078

Std. Errors:
 based on Hessian

Error Analysis:
        Estimate  Std. Error  t value Pr(>|t|)
mu      0.003473    0.002574    1.349 0.177211
omega   0.042572    0.020910    2.036 0.041752 *
alpha1  1.000000    0.520738    1.920 0.054814 .
gamma1  0.202502    0.147550    1.372 0.169930
beta1   0.797074    0.045315   17.590  < 2e-16 ***
delta   0.737446    0.210335    3.506 0.000455 ***
skew    1.205021    0.052296   23.042  < 2e-16 ***
shape   2.009908    0.004466  450.084  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Log Likelihood:
 557.6893    normalized:  0.9784022

Description:
 Sun May 12 20:52:50 2019 by user: mcbssgb2


Standardised Residuals Tests:
                                Statistic p-Value
 Jarque-Bera Test   R    Chi^2  17899.48  0
 Shapiro-Wilk Test  R    W      0.7489928 0
 Ljung-Box Test     R    Q(10)  10.27713  0.4165255
 Ljung-Box Test     R    Q(15)  13.3245   0.5772494
 Ljung-Box Test     R    Q(20)  15.23641  0.7627242
 Ljung-Box Test     R^2  Q(10)  2.251332  0.9940257
 Ljung-Box Test     R^2  Q(15)  2.625092  0.9998262
 Ljung-Box Test     R^2  Q(20)  3.271614  0.9999914
 LM Arch Test       R    TR^2   2.218819  0.9989894

Information Criterion Statistics:
      AIC       BIC       SIC      HQIC
-1.928734 -1.867743 -1.929121 -1.904937
```

The qq-plots of the standardised results for all models fitted above suggest that the chosen conditiional distributions are unsatisfactory. Moreover, the fitted standardised-t distributions have shape parameters (degrees of freedom) slightly over two. Suggesting extremely heavy tails, maybe even the need for stable distributions.

Note also that in all models above $\alpha_1 + \beta_1$ is greater than one, a possible violation of any form of stationarity.

Or maybe, it is just that the GARCH models tried here are not able to accomodate varying behaviour before, during and after the financial crisis.

**Affiliation:**

Georgi N. Boshnakov
School of Mathematics
The University of Manchester
Oxford Road, Manchester M13 9PL, UK
URL: http://www.maths.manchester.ac.uk/~gb/