UNIVERSITY OF CAMBRIDGE
Judge Business School

Updated: February 11, 2012
Lecturer: Dr. Paul Kattuman
Lab Instructor: Michael Freeman
Contact: mef35@cam.ac.uk

Contest Quiz 4
Question Sheet

Due Date : 27th February 2012

In this quiz we will review concepts relating to ARIMA and ARCH-based models.

Each group has been assigned a specific stock from the Dow Jones 30 Constituents. In your data set is the daily closing value log returns from 2000-01-03 to 2009-02-03. Your objective is to try and produce a good model for the stock returns using methods learnt in the course.

Question 1

We start by reviewing the ARIMA model. Before continuing I suggest that you try fitting a ARIMA good model to your data by checking the ACFs and PACFs and manually constructing your model. You can then compare this to the model we develop in the rest of this question.

i) Perform an ADF test with 8 lags including both intercept and trend, and report the p-value. Convince yourself that the series is stationary by following through the full ADF procedure, if necessary.

ii) Test whether we have included enough lags in part (i) by considering only the first 20 lags of the residuals. Do you believe we have included enough lags in (i)?
(A) Yes (B) No

iii) Now we would like to find a good ARMA model for your stocks returns. Do this manually to start with, and then use the `auto.arima` function to compare your best model to the best automatically generated model. How many AR terms are included in the ARMA model calculated using `auto.arima`?

iv) How many MA terms are included in the ARMA model calculated using `auto.arima` from part (iii)?

v) Now check whether the model described in (iv) appears to suffer from conditional heteroskedasticity at the 5% level by considering the first 10 lags, and then every lag that is a multiple of 5 up to 40 lags. Based on this alone, does the model:
(A) Display conditional heteroskedasticity (B) Not display conditional heteroskedasticity

Question 2

Regardless of your answer to part (v) in Question 1, with financial time series it is often the case that the returns suffer from conditional heteroskedasticity, where the variance shows signs of banding. Thus, it is always a useful exercise to try and construct better models that can account

for this characteristic. In this question we will look to build an ARCH, GARCH, EGARCH, and EGARCH-M model for your stocks time series. The number of AR terms and MA terms in the mean model should be the same number as those given in parts (iii) and (iv) in Question 1. You should also specify that `include.mean=F`.

i) Produce a basic ARCH(7) model (i.e. set $q = 7$). Inspect the coefficients for the ARCH terms. How many of these are insignificant at the 5% level when using robust standard errors.

ii) Now check whether the model described in (i) appears to suffer from conditional heteroskedasticity at the 5% level by considering the first 10 lags, and then every lag that is a multiple of 5 up to 40 lags. Based on this alone, does the model:
(A) Display conditional heteroskedasticity (B) Not display conditional heteroskedasticity

iii) Regardless of your answer to (ii), now fit a GARCH(1,1) model (i.e. $q = 1$) to your time series. Test for remaining autocorrelation in the residuals by considering the first 50 lags. At the 5% level, are there:
(A) No signs of autocorrelation in the residuals (B) Signs of autocorrelation in the residuals

iv) Now fit an EGARCH(1,1) model to your time series. Determine whether this model appears to suffer from conditional heteroskedasticity at the 5% level by considering the first 10 lags, and then every lag that is a multiple of 5 up to 40 lags. Based on this alone, does the model:
(A) Display conditional heteroskedasticity (B) Not display conditional heteroskedasticity

v) Now include an in-mean term to the EGARCH model you constructed in (iv). At the 5% level, is there evidence that the in-mean term should be included in your model?
(A) Yes (B) No

vi) Now reconstruct the 4 models that you produced in parts (i), (iii), (iv), and (v), but this time leave our the last 50 entries from the model fit. Use each of your ARCH-based models to forecast these missing 50 entries. Compare the MSE of the 4 forecasts with the actual values, and report the best model based on this statistic. Report the numeric value of either 1, 2, 3 or 4 for the ARCH, GARCH, EGARCH and EGARCH-M models respectively.