Homework 8 – Due Tuesday, Apr. 25
STAT-GB.2302, STAT-UB.0018: Forecasting Time Series Data

In these problems, we will consider the `rupee` series (`rupee.csv`), the exchange rate for the Indian Rupee to 1 U.S. Dollar. The data is daily from July 1st, 2002 to April 8th, 2011 ($n = 2235$). We will work with the logs of the exchange rates

Problem 1

Plot the logs of `rupee`. Based on this plot, and the ACF and PACF of the logs and differenced logs, does the series appear to be stationary? Can you identify an ARIMA(p, d, q) model from these plots?

Problem 2

Using AIC_C, select an ARIMA$(p, 1, q)$ (without constant) with $0 \leq p \leq 2$, $0 \leq q \leq 2$. Compute the residuals and fitted values from this model. Also, compute the (ARIMA) one step ahead forecast and 95% forecast interval.

Problem 3

Plot the residuals, as well as ACF and PACF of both the residuals and the squared residuals. Use these plots to argue that the residuals, although approximately uncorrelated, are not independent; instead, they show evidence of conditional heteroscedasticity.

Problem 4

Using the residuals from the ARIMA model, find the log-likelihood and AIC_C values for ARCH(q) models where q ranges from 0 to 10. You will need to calculate the log likelihood for the ARCH(0) model by hand.

Next, fit a GARCH$(1, 1)$ model to the residuals. Evaluate AIC_C for the $GARCH(1, 1)$ model, using $k = 2$ in the formula for AIC_C. If the GARCH$(1, 1)$ model is preferred by AIC_C, use it as your selected model. Report the `summary` and `logLik` of your selected ARCH or GARCH model. Comment on the statistical significance of the parameter values of your selected model. Write the complete form of the ARCH or GARCH model you have selected.

Problem 5

Construct a 95% one step ahead forecast interval for the log exchange rate, based on your ARIMA-ARCH model. (If you decided to use a GARCH$(1, 1)$ model, you will need to first get the conditional variances from R. See Problem 6.) Compare this to the interval based on the ARIMA only model from Problem 2.

Problem 6

Plot the conditional variances, `ht`, for your fitted ARCH or GARCH model from Problem 4. Use this plot to locate bursts of high volatility. Do these highly volatile periods agree with those found form examination of the time series plot of the log exchange rates themselves?

Problem 7

Make a time series plot which simultaneously shows the log exchange rates, together with the ARIMA-ARCH one-step-ahead 95% forecast intervals based on information available in the previous day. Using this plot, comment on the accuracy and practical usefulness of the forecast intervals. Keep in mind that the performance may be somewhat better here than in an actual forecasting context, since the ARIMA-ARCH parameters are estimated from the entire data set, not just the observations up to the time at which the forecaast is to be constructed.

Problem 8

Compute the residuals from your ARIMA-ARCH model, that is, $e_t = \varepsilon_t \sqrt{h_t}$. If the ARIMA-ARCH model is adequate, these residuals should be normally distributed with mean zero and variance 1. Make a normal probability plot of the arch residuals. Does the model seem to have adequately described the leptokurtosis ("long-tailedness") in the data?

Problem 9

From the formula for the prediction intervals, it follows that the 95% prediction interval constructed yesterday fails to cover today's log exchange rate whenever today's residual exceeds 1.96 in absolute value. Use the `sum` function to count up how many failures there were. What percentage of the time did the intervals fail? (Keep in mind that even though $n = 2235$, there are not 2235 forecast intervals.)