STAT 6351: Assignment #5

Notes

This assignment is concerned with alternative approaches to volatility calculation (Lecture 6) and nonlinear models (Lecture 7). Relevant data can be downloaded from Yahoo via the quantmod package. For instance:

getSymbols('AMZN', from='2007-01-03', to='2017-04-28')

Because of possible stock splits, be sure to consider only the <u>adjusted closing prices</u>, whenever the question has you consider closing prices.

Questions

- 1. (Alcoa stock volatility). Consider the daily log returns of Alcoa stock (tick symbol AA), over the period from January 3, 2007 to April 28, 2017. Construct the (conditional) variance estimates $\sigma_{i,t}^2$ of Section 3.15.2 of AFTS for i = 0, 1, 2, 3, 5, 6. These conditional variances are for the stock price. Multiply by 252 and take the square-root to obtain annualized volatility. I.e., $\sigma_{i,t}(\text{annual}) = \sqrt{252\sigma_{i,t}^2}$. Obtain mean, median, maximum, and minimum values for each $\sigma_{i,t}(\text{annual})$ series.
- 2. Use the data and the script yz. R with n = 63 to compute the Yang and Zhang (2000) variance estimate σ_{yz}^2 of Section 3.15.2 of AFTS, for the Alcoa stock as in Question 1. Obtain a time plot of the estimated (and annualized) volatility series. (Note that this volatility is for the log returns, and is therefore not comparable to those of Question 1.)

3. (Amazon stock price movement). Consider the daily closing prices of Amazon stock (AMZN) and the S&P 500 index, from January 4, 2007 to April 28, 2017. Download the data using quantmod from 2007-01-03 to 2017-04-28, and then compute the two series of log returns $\{r_t\}$. Define the *price movement* of Amazon as follows:

$$A_t = \begin{cases} 1, & \text{if } r_t > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Similarly, let M_t denote the market movement of the S&P 500 index log returns.

- (a) Fit a (linear) logistic regression model for $P(A_t = 1)$ using $\{A_{t-1}, A_{t-2}, A_{t-3}, M_{t-1}, M_{t-2}, M_{t-3}\}$ as inputs. Write down the fitted model. Based on the model, can past price movements of either the stock or the market predict the future direction of Amazon price movement? Why?
- (b) Use $\{A_{t-1}, A_{t-2}, M_{t-1}, M_{t-2}\}$ to build a 4-3-1 feed-forward neural network with direct link to model $P(A_t = 1)$.
- (c) Divide the sample into training and testing subsamples, with the latter consisting of the last 250 observations. Use 1-step ahead predictions \hat{A}_t to compare the models in (a) and (b). Since predictions must be binary, let $\hat{A}_t = 1$ if $P(A_t = 1) > 0.5$, and zero otherwise. (Note: the script backnnet.R will perform backtesting for neural nets, just as backtest.R does for ARMA models.)