STAT 6351: Assignment #2

Notes

- Use 5% significance level in all tests.
- ρ_h denotes lag-*h* autocorrelation.
- In some of the problems, I provide guidance in specifying a time series model. This is to help you gain experience in empirical data analysis. You can try your own models to gain further experience. The assignments show that multiple models can fit a given data set well and seasonally adjusted data might still have some residual seasonality.

Questions

1. (Gold price). Consider the daily gold fixing price 10:30 am (London time) in London Bullion Market in US dollars per Troy ounce from January 3, 1995 to November 10, 2016. Download the data from http://www.math.ttu.edu/~atrindad/tsdata. Read in the data, let x_t denote the log(gold price) and $r_t = 100(x_t - x_{t-1})$ the percent returns, as follows:

```
gold=read.csv("Gold.csv", header=T)
xt = log(as.numeric(gold$VALUE))
rt <- 100*diff(xt)</pre>
```

- (a) Obtain the time plots of x_t and r_t (in one page, using the command par(mfcol=c(2,1)).
- (b) Compute the first 12 lags of the ACF of x_t . Do these ACF values suggest the presence of a unit root in x_t ? Why?
- (c) Consider only the r_t series. Test $H_0: \rho_1 = \cdots = \rho_{12} = 0$ vs. $H_a: \rho_i \neq 0$ for some $1 \leq i \leq 12$, and draw your conclusion.
- (d) Use the command ar(rt,method=''mle'',order.max=20) to find an AR model for r_t .
- (e) Build an AR model for r_t , including model checking. Refine the model by excluding all estimates with t-ratio less than 1.645. Write down the fitted model.
- (f) Use the fitted AR model to compute 1-step to 4-step ahead forecasts of r_t . Also, compute the corresponding 95% interval forecasts.

- 2. (Model comparison). Consider, again, the log return series of gold price of Problem 1.
 - (a) Build an MA model for r_t . Refine the model by removing coefficient estimates with t-ratio less than 1.645. Write down the fitted model.
 - (b) Compute the Ljung-Box statistic Q(10) of the residuals of the fitted MA model. Is there serial correlation in the residuals? Why?
 - (c) Consider the in-sample fits of the AR model of Problem 1 and the MA model. Which model is preferred? Why?
 - (d) Use the function backtest (download from https://www.math.ttu.edu/~atrindad/ stat6351/Rprograms/) at the forecast origin t = 5225 with horizon h = 1 to compare the two models. Which model is preferred? Why?
- 3. (Volatility modeling). Consider the CBOE daily volatility index (VIX) from January 2, 1990 to March 31, 2017. The data are available from FRED via quantmod. Again, there are some missing values, which need to be removed. Use the commands below:

```
require(quantmod)
getSymbols(''VIXCLS'',src=''FRED'')
VIXCLS <- VIXCLS[1:7109]
idx <- c(1:7109)[is.na(VIXCLS)]
VIXCLS <- VIXCLS[-idx]
vix <- as.numeric(VIXCLS)</pre>
```

- (a) Find an AR model for the VIX series. Remove insignificant coefficient estimates (based on t-ratio 1.645). Provide model checking to confirm that the model is adequate. Write down the model.
- (b) Use the fitted model to obtain 1-step to 10-step ahead predictions at the forecast origin March 31, 2017.
- 4. (GDP revisited). Consider, again, the U.S. quarterly real GDP growth rate from 1947.II to 2016.IV. The data are available from FRED. See the command below.

```
getSymbols(''A191RL1Q225SBEA'',src=''FRED'')
gdp <- as.numeric(A191RL1Q225SBEA)</pre>
```

- (a) Obtain time-series plot of the real GDP growth rates.
- (b) Find an AR model for the real GDP growth rate, including model checking. Write down the fitted model.
- (c) Does the model imply existence of business cycle? Why?
- (d) If business cycles are present, compute the average length of the cycles. Otherwise, the length is infinity.
- (e) Obtain 95% interval forecasts of 1-step to 4-step ahead GDP growth rates at the forecast origin 2016.IV.

- 5. (Seasonal model). Consider the monthly US unemployment rates used in the lectures. Data are available from FRED and is entitled UNRATE. Let r_t be the unemployment at time t. Ignore any possible outliers in the data.
 - (a) Fit a seasonal model to r_t as follows:

```
getSymbols("UNRATE",src="FRED")
rt = as.vector(UNRATE[1:864])
m1 <- arima(rt,order=c(3,0,1),seasonal=list(order=c(1,0,1),period=12))</pre>
```

Perform model checking.

- (b) Fit an AR(11) model to r_t . Refine the model by excluding insignificant coefficients, similarly to what was done in Problem 1(e).
- (c) Use backtest at the forecast origin t = 770 and horizon h = 1 to compare the two models. Which model is preferred? Why?