Lab 6. Linear time series models

There are several packages in R that have time series functionalities. Examples are "TimeSeries", "TSA" and "AER". You need to load them before using them. Below, we use the basic "stats" package (no installation needed).

1

Let us first play around acf, pacf, and the Ljung and Box test.

```
> x <- arima.sim(100, model=list(ar=c(.8,-.6), ma=c(0.7, 0.4)))
# simulate a TS with length 100
> par(mfrow = c(2,2))
> plot(x)                        # plotting the series
> ?acf                           # help on acf
> acf(x,40)                      # acf plot
> acf(x,40,"partial")# PACF plot
> xacf <- acf(x, 40, plot=F)     # computing ACF without plot
> xacfsq <- xacf$acf[2:11]^2     # square autocorrelation 1:10
> Q <- 100*102*sum(xacfsq/(100-1:10)) # Box-Ljung test
> 1-pchisq(Q,10)                 # computing P-value
[1] 0.00005613359
```

2

Let us now play around the efficient market hypothesis. First, let us load the .xls file for SP500.

```
> SP500 <- read.csv("SP500Daily.csv")
> head(SP500)
        Date    Open    High     Low   Close     Volume Adj.Close
1 2008-03-28 1327.02 1334.87 1312.95 1315.22 3686980000   1315.22
2 2008-03-27 1340.34 1345.62 1325.66 1325.76 4037930000   1325.76
3 2008-03-26 1352.45 1352.45 1336.41 1341.13 4055670000   1341.13
4 2008-03-25 1349.07 1357.47 1341.21 1352.99 4145120000   1352.99
5 2008-03-24 1330.29 1359.68 1330.29 1349.88 4499000000   1349.88
6 2008-03-20 1299.67 1330.67 1295.22 1329.51 2078450000   1329.51
```

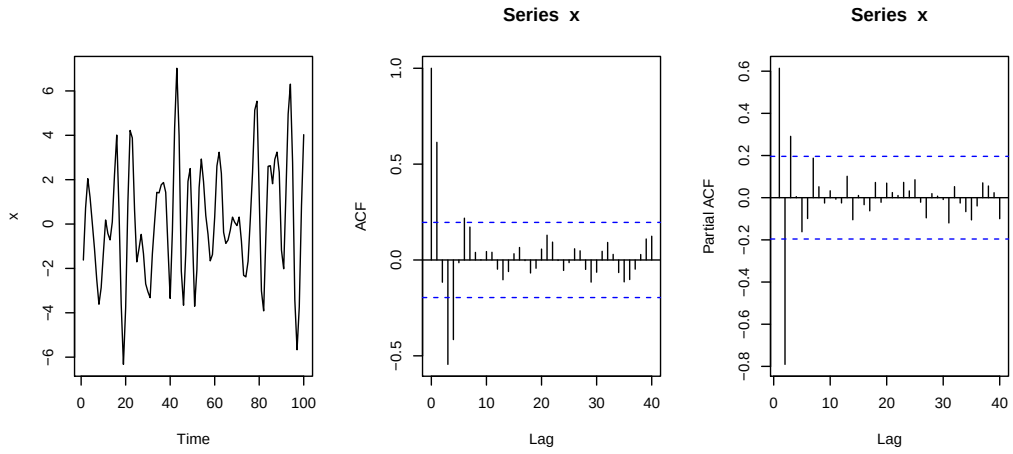

Figure 1: Arima simulation

```
> sp500 <- rev(SP500[1:1000,7])
# extracting the Adjusted Closing Prices for the last 4 years.
> head(sp500)
[1] 1140.53 1139.32 1145.20 1129.44 1128.17 1128.84
> sp500 <- log(sp500)       # log-prices
> acf(sp500, 400)           # examing log-prices
> n <- length(sp500)        # length of series

> ## calculating correlation without drift; here are two methods
> x <- sp500[1:(n-1)]
> y <- sp500[2:n]
> sum(x*y)/sum(x^2)       # correlation without drift
[1] 1.000019
> lsfit(x,y,intercept=F)$coef    # here is a short cut
       X
1.000019

> #### calcuating the correlation with drift
> lsfit(x,y)$coef
 Intercept          X
0.02953115 0.99589617
```

3

Here is an example for fitting an AMRA model to the daily log-returns of the SP500 stocks.

```
> return500 <- 100 * diff(sp500)
> mean(return500)     # mean daily return
[1] 0.01426535
> return500 <- return500 - mean(return500)    # subtract the mean out
> return500 <- as.ts(return500)   # create a time series structure
> acf(return500)      # acf plot
> acf(return500, lag.max=100, type="partial") # pacf plot
> #?arima
> return500.arima <- arima(return500, order=c(1,0,1))
> return500.arima

Call:
arima(x = return500, order = c(1, 0, 1))

Coefficients:
         ar1      ma1  intercept
      0.6223  -0.7106      5e-04
s.e.  0.1927   0.1740      2e-02

sigma^2 estimated as 0.6808:  log likelihood = -1225.49,  aic = 2458.98
> tsdiag(return500.arima)   # arima diagnostics
>
> fore500 <- predict(return500.arima,n.ahead=10)
> ## 10 step forward prediction
> fore500
$pred
Time Series:
Start = 1000
End = 1009
Frequency = 1
 [1] 0.114530973 0.071453873 0.044647472 0.027966146 0.017585543 0.011125808
 [7] 0.007105987 0.004604497 0.003047848 0.002079162

$se
Time Series:
Start = 1000
End = 1009
Frequency = 1
 [1] 0.8251118 0.8283234 0.8295637 0.8300435 0.8302292 0.8303011 0.8303290
 [8] 0.8303397 0.8303439 0.8303455
```

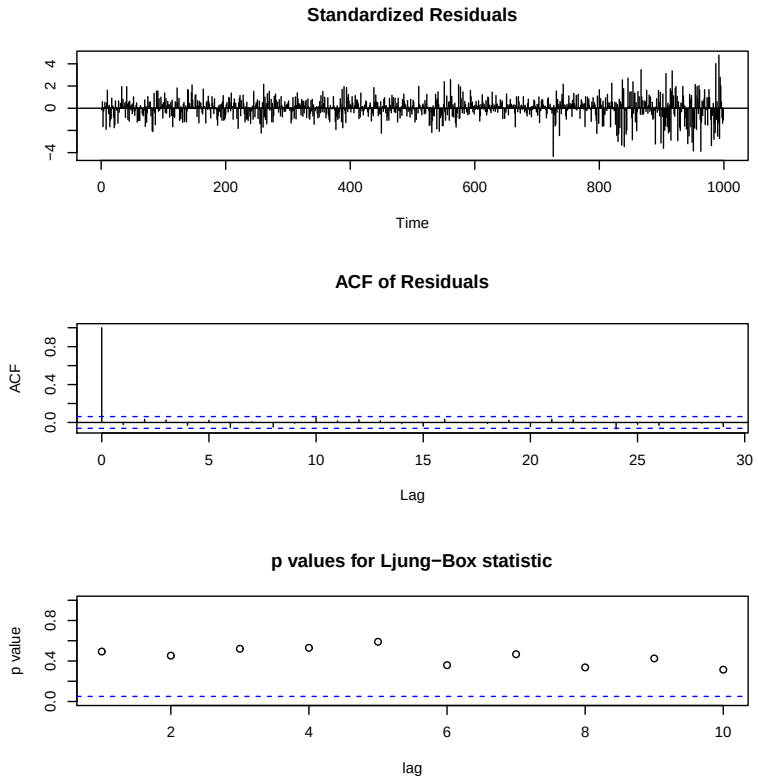

Figure 2: Arima fitting diagnostics

4

For the airline model, try the following example

```
> #fitting an airline model
> fit = arima(ship, order = c(0,1,1), seasonal = list(order = c(0,1,1), period = 4),
+ method = "ML" )
> fit

Call:
arima(x = ship, order = c(0, 1, 1), seasonal = list(order = c(0, 1, 1), period = 4),
    method = "ML")

Coefficients:
          ma1     sma1
      -0.8520  -0.2978
s.e.   0.0727   0.1158
```

```
sigma^2 estimated as 0.008241:  log likelihood = 76.45,  aic = -146.89
> fit$var.coef
              ma1         sma1
ma1   0.005278143 -0.002751018
sma1 -0.002751018  0.013409825
```