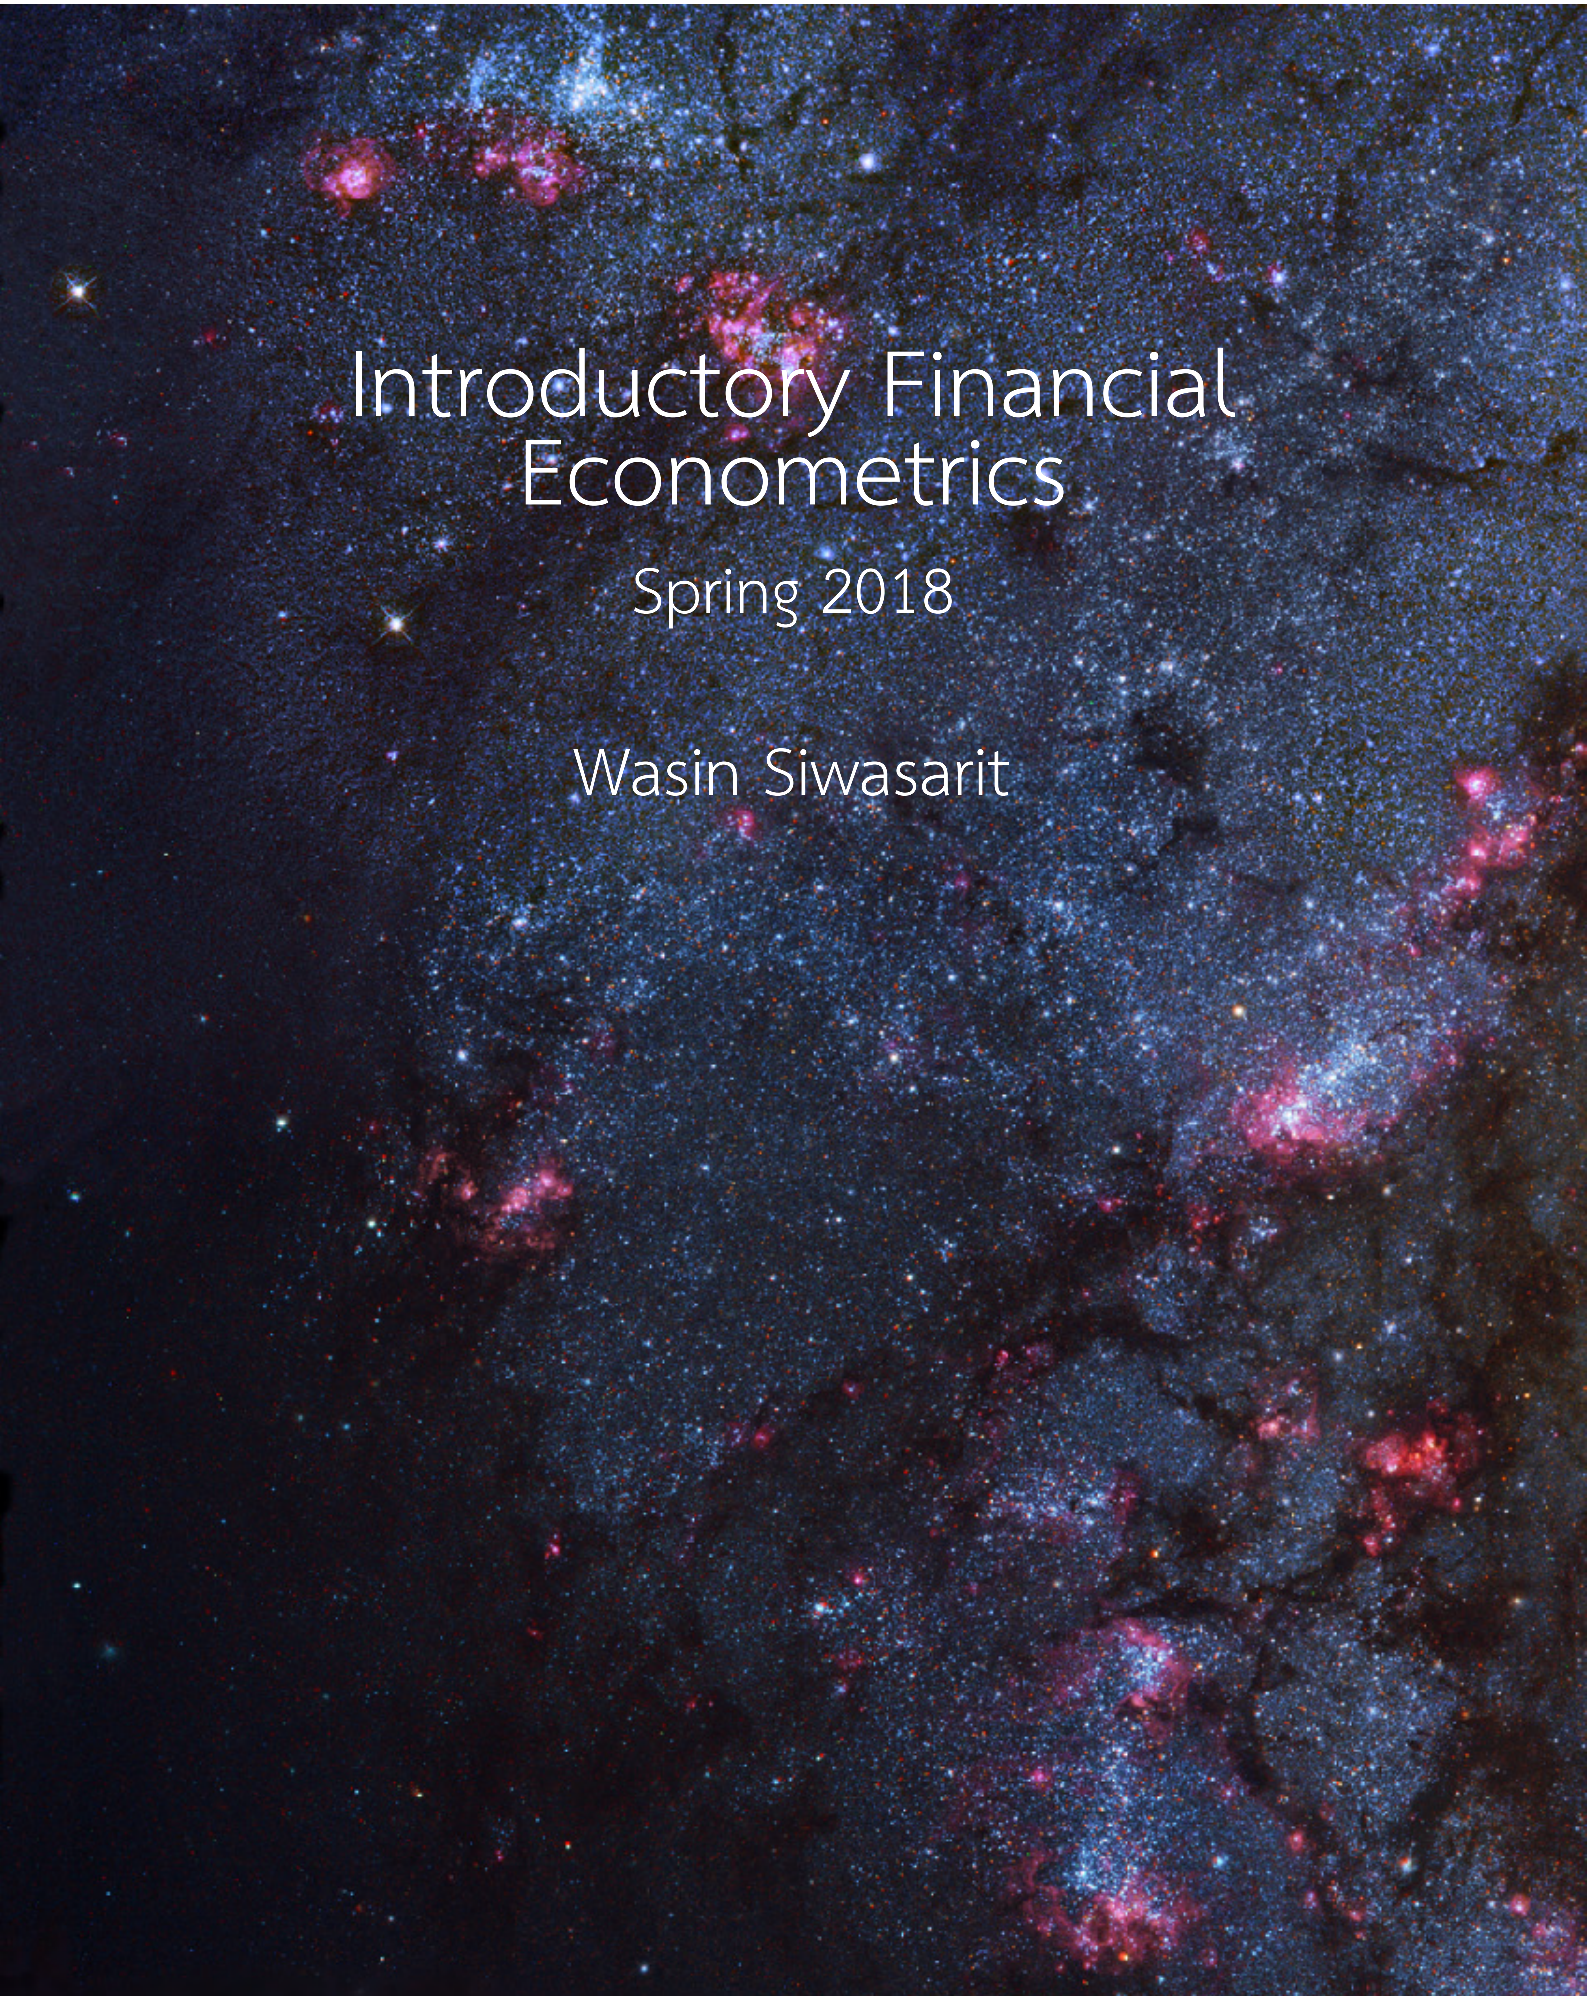

Introductory Financial Econometrics

Spring 2018

Wasin Siwasarit

Project Spring 2018, Faculty of Economics, Thammasat University.

Academic Year Spring 2018

Road Map of this class:

1. Financial Time Series and Their Characteristics

Financial time series (FTS) analysis

Financial time series (FTS) analysis is concerned with theory and practice of asset valuation over time.

What is the difference, if any, from traditional time series analysis?

Two topics are highly related, but FTS has added uncertainty, because it must deal with the ever-changing business & economic environment and the fact that volatility is not directly observed.

1.1 The Objectives of this chapter

1. to access financial data online and to process the embedded information
2. to provide basic knowledge of FTS data such as skewness, heavy tails, and measure of dependence
between asset returns
3. to introduce statistical tools econometric models useful for analyzing these series.
4. to gain experience in analyzing FTS

1.2 Examples of financial time series

1. Daily log returns of Apple stock: 2007 to 2018 (12 years). Data downloaded using quantmod

2. The VIX index.

3. CDS spreads: Daily 3-year CDS spreads of JP Morgan from July 20, 2004 to September 19, 2018.

4. Quarterly earnings of Coca-Cola Company: 1983-2009 Seasonal time series useful in

• earning forecasts

• pricing weather related derivatives (e.g. energy) • modeling intraday behavior of asset returns

5. US monthly interest rates (3m & 6m Treasury bills)

Relations between the two asset returns? Term structure of interest rates.

6. Exchange rate between US Dollar vs Euro Fixed income, hedging, carry trade.

7. Size of insurance claims.

8. High-frequency financial data:
Tick-by-tick data of Caterpillars stock: January 04, 2010.

1.3 Asset Returns

Let P_t be the price of an asset at time t, and assume no dividend. One-period simple return: Gross return

$$1+R_t = \frac{P_t}{P_{t-1}}$$

One-Period Simple Net Return or Simple Return:

Multiperiod simple return: Gross return)

Example:Table below gives six daily (adjusted) closing prices of Apple stock in December 2015.

Date	Price
12/23	108.02
12/24	107.45
12/28	106.24
12/29	108.15
12/30	106.74
12/31	104.69

what is one-day gross return of holding the stock from 12/28 to 12/29 and the daily simple return?

Time interval is important! Default is one year. Annualized (average) return:

Besides the simple return, we can also compute the continuously compounding interest rate where r is the interest rate per annum, C is the initial capital, n is the number of years, and exp is the exponential function.

$$A = C \times exp(r \times n)$$

Continuously compounded (or log) return

$$r_t = ln(1+R_t) = ln\frac{P_t}{P_{t-1}} = p_t - p_{t-1}$$

where $p_t = ln(P_t)$

Multiperiod log return:

Continuously compounding: Illustration of the power of compounding (int. rate 10 % per annum)

Type	#(payment)	Int.	Net
Annual	1	0.1	\$1.10000
Semi-Annual	2	0.05	\$1.10250
Quarterly	4	0.025	\$1.10381
Monthly	12	0.0083	\$1.10471
Weekly	52	$\frac{0.1}{52}$	\$1.10506
Daily	365	$\frac{0.1}{365}$	\$1.10516
Continuously	∞		\$1.10517

Portfolio return: N assets

Dividend payment:

Excess Returns (adjusting for risk)

Example

1. Whai is the log return from 12/23 to 12/24?

Remarks:

Example If the monthly log returns of an asset are 4.46 %, −7.34 % and 10.77 %, then what is the corresponding quarterly log return?

Example If the monthly simple returns of an asset are 4.46 %, −7.34 %, and 10.77 %, then what is the corresponding quarterly simple return?

1.4 Distributional Properties of Returns

What is the distribution of
r_{it} where $i = 1,...,N$; and $t = 1,...,T$

Some theoretical properties:

Moments of a random variable X with density f(x): l-th moment

$$m_l' = E(X^l) = \int_{\infty}^{\infty} x^l f(x) dx$$

First Moment: mean or expectation of X.

l-th central moment

$$m_l = E[(X - \mu_x)^l] = \int_{\infty}^{\infty} (x - \mu_x)^l f(x) dx$$

2nd central moment.

Standard deviation: square-root of variance

Skewness (Symmestry)

$$S(x) = E\left[\frac{(X-\mu_x)^3}{\sigma_x^3}\right]$$

Kurtosis (Fat-taills)

$$K(x) = E\left[\frac{(X-\mu_x)^4}{\sigma_x^4}\right]$$

Q1: Why study the mean and variance of returns?
They are concerned with long-term return and risk, respectively.
Q2: Why is symmetry important?
Symmetry has important implications in holding short or long financial positions and in risk management.
Q3: Why is kurtosis important?
Related to volatility forecasting, efficiency in estimation and tests High kurtosis implies heavy (or long) tails in distribution.

Estimation

Sample mean, Sample Variance, Sample Skewness and Sample Kurtosis

1.5 Hypothesis Testing

A random variable under the normal distribution

$$\widehat{S(x)} \sim N(0, \frac{6}{T})$$

$$\widehat{K(x)} - 3 \sim N(0, \frac{24}{T})$$

Test for symmetry

Test for tail thickness

Test for normality :(Jarque-Bera test)

1.6 Empirical work using R program

FE Toolbox

```
   #EE435 Wasin Siwasarit  Lecture1  Spring/2018
setwd("/Users/wasinsiwasarit/Desktop/EE435")
cat(rep("\n",50))  #clear R Console
#install.packages("quantmod")
#install.packages("fBasics")
#install.packages("sn")
#install.packages("PerformanceAnalytics")
#install.packages("car")
#install.packages("tseries")
#install.packages("forecast")
library(quantmod)
library(fBasics)
library(sn)
library(PerformanceAnalytics)
library(car)
library(tseries)
library(forecast)

getSymbols("^GSPC",from="2000-01-03",to="2017-01-28")
dim(GSPC)
head(GSPC)
tail(GSPC)
da=GSPC
chartSeries(GSPC,theme="white")
price=da[,6]
plot(price,type='l')
```

```r
logprice=log(price)
plot(logprice,type='l')
logreturn=diff(log(price))
simplereturn <-exp(logreturn)-1
#1 Plot the series of log return and simple return

par(mfrow=c(1,1))
plot(logreturn,type='l')
plot(simplereturn)

newlogreturn <- logreturn[2:nrow(logreturn),]
newsimplereturn <- simplereturn[2:nrow(logreturn),]

#2 Histogram and sample statistics
hist(logreturn, breaks=100, col="slateblue")
chart.Histogram(logreturn,methods = c("add.normal"))
table.Stats(logreturn)

#3 QQ-plots and tests for normality
#
# use qqnorm function
par(mfrow=c(1,1))
qqnorm(newlogreturn)
qqline(newlogreturn, col = 2)
jarque.bera.test(newlogreturn)
```

FE Analysis

```
> table.Stats(logreturn)
                GSPC.Adjusted
Observations        4377.0000
NAs                    1.0000
Minimum               -0.0947
Quartile 1            -0.0051
Median                 0.0005
Arithmetic Mean        0.0001
Geometric Mean         0.0000
Quartile 3             0.0058
Maximum                0.1096
SE Mean                0.0002
LCL Mean (0.95)       -0.0002
UCL Mean (0.95)        0.0005
Variance               0.0002
Stdev                  0.0123
Skewness              -0.1989
Kurtosis               8.3611
> par(mfrow=c(1,1))
> qqnorm(newlogreturn)
> qqline(newlogreturn, col = 2)
> jarque.bera.test(newlogreturn)

        Jarque Bera Test

data:  newlogreturn
X-squared = 12778, df = 2, p-value < 2.2e-16
```

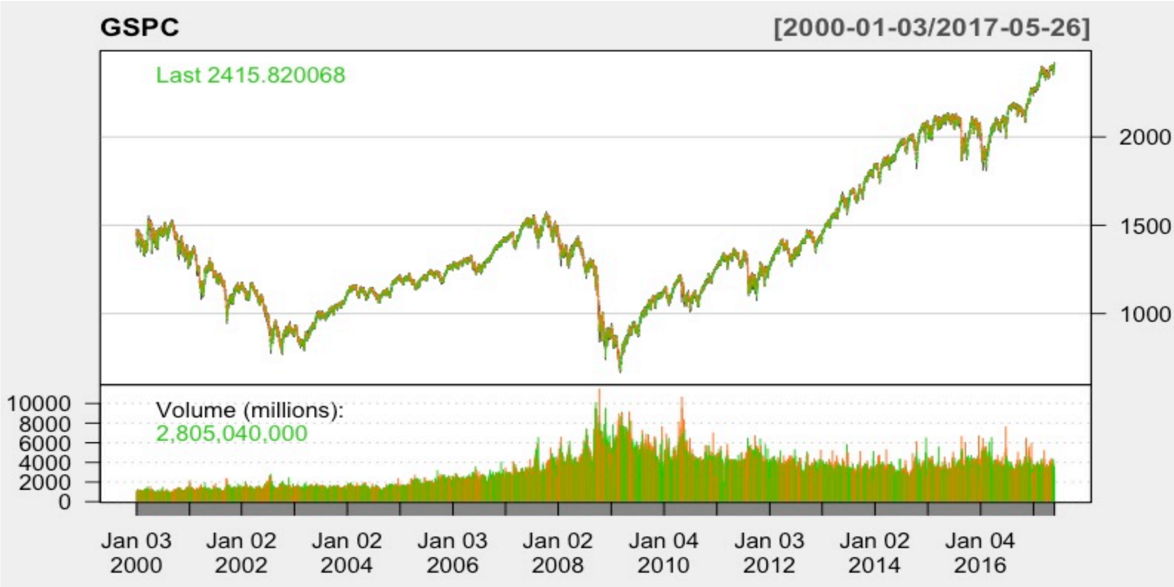
GSPC
[2000-01-03/2017-05-26]
Last 2415.820068
2000
1500
1000
Volume (millions):
2,805,040,000
10000
8000
6000
4000
2000
0
Jan 03 2000
Jan 02 2002
Jan 02 2004
Jan 03 2006
Jan 02 2008
Jan 04 2010
Jan 03 2012
Jan 02 2014
Jan 04 2016

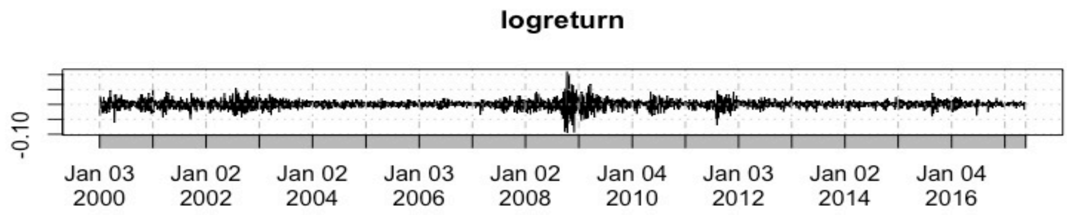
logreturn
-0.10
Jan 03 2000
Jan 02 2002
Jan 02 2004
Jan 03 2006
Jan 02 2008
Jan 04 2010
Jan 03 2012
Jan 02 2014
Jan 04 2016

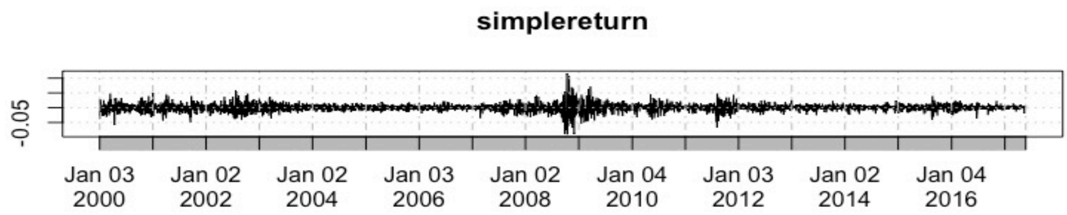
simplereturn
-0.05
Jan 03 2000
Jan 02 2002
Jan 02 2004
Jan 03 2006
Jan 02 2008
Jan 04 2010
Jan 03 2012
Jan 02 2014
Jan 04 2016

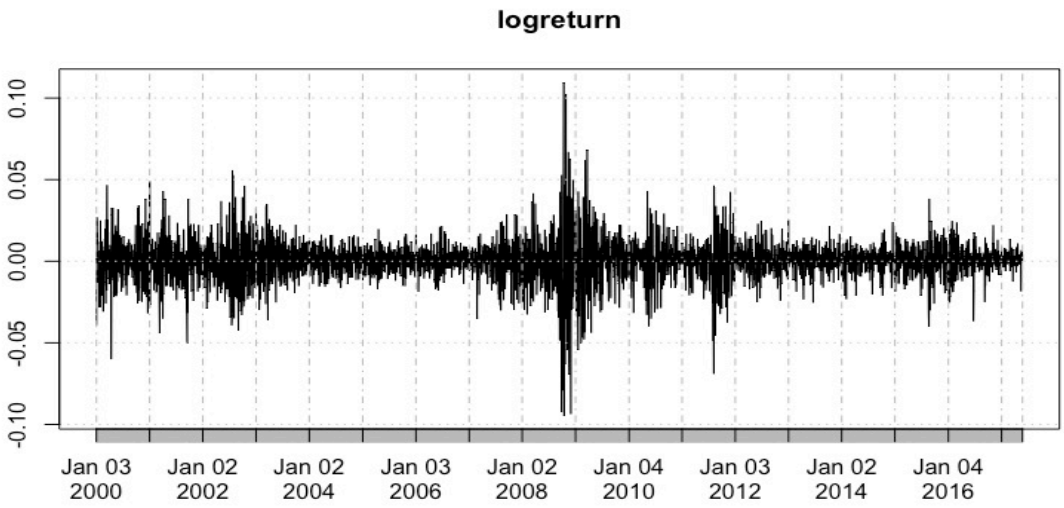
logreturn
0.10
0.05
0.00
-0.05
-0.10
Jan 03 2000
Jan 02 2002
Jan 02 2004
Jan 03 2006
Jan 02 2008
Jan 04 2010
Jan 03 2012
Jan 02 2014
Jan 04 2016

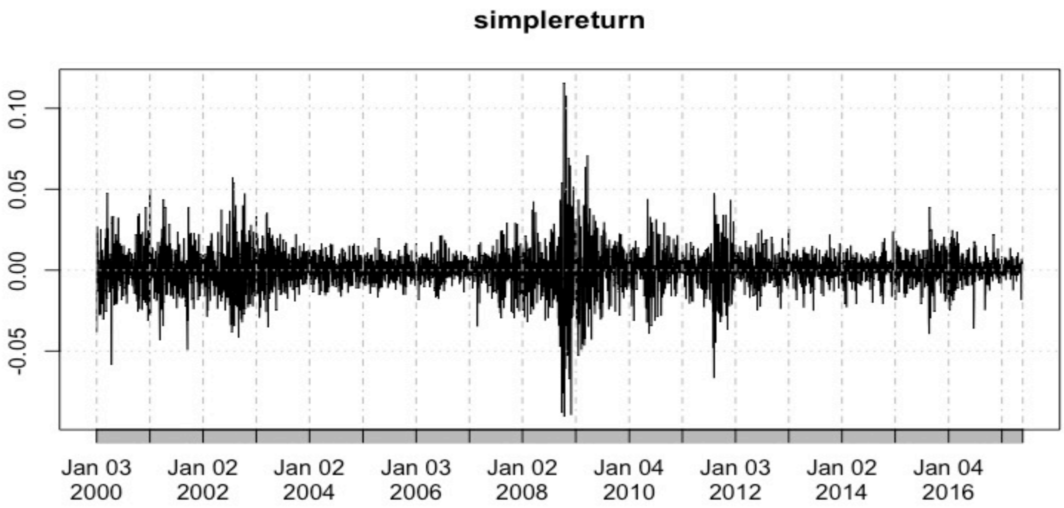
simplereturn
0.10
0.05
0.00
-0.05
Jan 03 2000
Jan 02 2002
Jan 02 2004
Jan 03 2006
Jan 02 2008
Jan 04 2010
Jan 03 2012
Jan 02 2014
Jan 04 2016

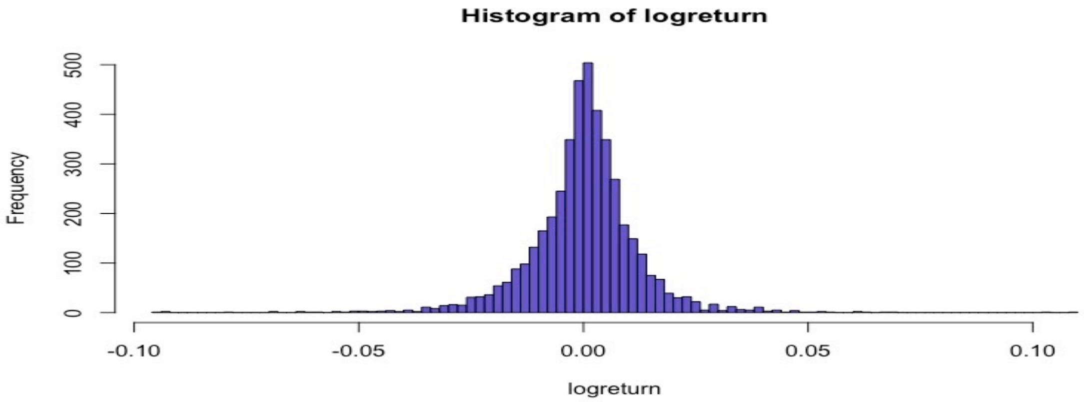
Histogram of logreturn
Frequency
0
100
200
300
400
500
-0.10
-0.05
0.00
0.05
0.10
logreturn

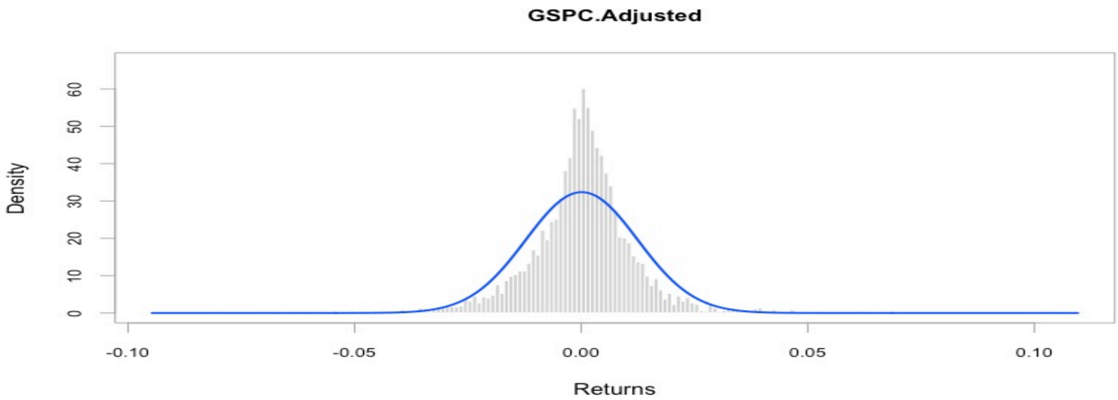
GSPC.Adjusted
Density
0
10
20
30
40
50
60
-0.10
-0.05
0.00
0.05
0.10
Returns

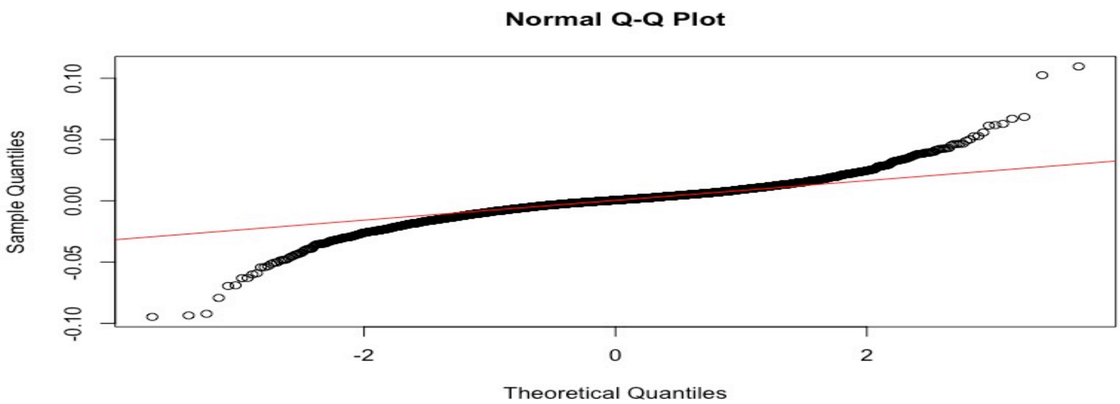
Normal Q-Q Plot
Sample Quantiles
-0.10
-0.05
0.00
0.05
0.10
-2
0
2
Theoretical Quantiles

FE toolbox (Cont.)

```
#4 Test mean = 0
t.test(newlogreturn)

#5 Test Skewness = 0
T=length(newlogreturn)
s3=skewness(newlogreturn)
tst = s3/sqrt(6/T)
tst
pv = 2*pnorm(tst)
pv

#6 Test excess kurtosis =0
k4 = kurtosis(newlogreturn)
tst = k4/sqrt(24/T)
tst
pv = 2*(1-pnorm(tst))
pv
```

FE Analysis (Cont.)

```
> t.test(newlogreturn)

        One Sample t-test

data:  newlogreturn
t = 0.62168, df = 4376, p-value = 0.5342
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 -0.0002493954  0.0004810069
sample estimates:
   mean of x
0.0001158057

> T=length(newlogreturn)
> s3=skewness(newlogreturn)
> tst = s3/sqrt(6/T)
> tst
[1] -5.370912
> pv = 2*pnorm(tst)
> pv
[1] 7.833935e-08
> k4 = kurtosis(newlogreturn)
> tst = k4/sqrt(24/T)
> tst
```

```
[1] 112.913
> pv = 2*(1-pnorm(tst))
> pv
[1] 0
>
```

2. Linear Time Series (TS) Models

2.1 Basic Concepts

Financial TS: collection of a financial measurement over time.

Example: log return of apple r_t

Data: $\{r_1, r_2, ..., r_T\}$

Purpose What is the information contained in series of r_t

Definition: Stationarity

-Strict: Distributions are time-invariant

-Weak: First 2 moments are time-invariant

What does weak stationarity mean in practice?

Past: time plot of r_t varies around a fixed level within a finite range!

Future:the first 2 moments of future r_t are the same as those of the data so that meaningful inferences can be made.

Mean (or expectation) of returns

$$\mu = E(r_t)$$

Variance (variability) of returns

$$Var(r_t) = E[(r_t - \mu)^2]$$

Sample mean and Sample Variance are used to estimate the mean and variance of returns.

$$\bar{r} = \frac{1}{T}\sum_{t=1}^{T} r_t$$

$$Var(r_t) = \frac{1}{T-1}\sum_{t=1}^{T}(r_t - \bar{r})^2$$

testing the mean of r_t is different from zero or not
$H_0 : \mu = 0$
$H_a : \mu \neq 0$

$$t_{cal} =$$

Decision rule: Reject H_0 if $|t| > Z_{\frac{\alpha}{2}}$ or p-value is less than α

Lag-kk autocovariance:

$$\gamma_k = Cov(r_t, r_{t-k}) = E[(r_t - \mu)(r_{t-k} - \mu)]$$

Serial (or-auto) correlations:

$$\rho_l = \frac{cov(r_t, r_{t-l})}{var(r_t)}$$

Remark The existence of serial correlation in r_t implies that..

Sample Autocorrelation function (AFC) can be computed by:

$$\hat{\rho}_l = \frac{\sum_{t=1}^{T-l}(r_t - \bar{r})(r_{t+l} - \bar{r})}{\sum_{t=1}^{T}(r_t - \bar{r})^2}$$

Test Zero Serial Correlations (Market Efficiency)

1. Individual Test

$H_0 : \rho_1 = 0$

$H_a : \rho_1 \neq 0$

$$t_{cal} =$$

Decision rule: Reject the null hypothesis when $|t| > Z_{\frac{\alpha}{2}}$ or the p-value has the value less than α

2. Joint Test (Ljung-Box Statistics):

$H_0 : \rho_1 = ... = \rho_m = 0$

$H_a : \rho_i \neq 0$

$$Q(m) = T * (T+2) \sum_{l=1}^{m} \frac{\rho_l^2}{T-l}$$

Decision rule: Reject the null hypothesis when $Q(m) > \chi_m^2(\alpha)$ or the p-value has the value less than α

FE toolbox 2

```
#EE435 Wasin Siwasarit
setwd("/Users/wasinsiwasarit/Desktop/EE435")
library(fBasics)
cat(rep("\n",50))  #clear R Console
da <- read.table("CRSP.txt")
log_return =da[,1]
par(mfcol=c(1,1))
length(log_return)
tdx = c(1:456)/12+1961
plot(tdx,log_return,xlab='year',ylab='log_return',type='l')
basicStats(log_return)
normalTest(log_return,method="jb")
t.test(log_return)
tt1 = skewness(log_return)/sqrt(6/546)
tt1
pv =  2*pnorm(tt1)
pv
tt2 = kurtosis(log_return)/sqrt(24/546)
tt2
```

```
pv = 2*(1-pnorm(tt2))
pv
m1=acf(log_return)
names(m1)
m1$acf
m2=pacf(log_return)
names(m2)
m2$acf
Box.test(log_return,lag=12,type='Ljung')
```

FE Analysis 2

```
> da <- read.table("CRSP.txt")
> log_return =da[,1]
> par(mfcol=c(1,1))
> length(log_return)
[1] 456
> tdx = c(1:456)/12+1961
> plot(tdx,log_return,xlab='year',ylab='log_return',type='l')
> basicStats(log_return)
             log_return
nobs         456.000000
NAs            0.000000
Minimum      -31.588000
Maximum       26.175000
1. Quartile   -1.860000
3. Quartile    4.268250
Mean           1.059511
Median         1.494500
Sum          483.137000
SE Mean        0.262245
LCL Mean       0.544149
UCL Mean       1.574873
Variance      31.360270
Stdev          5.600024
Skewness      -0.673271
Kurtosis       4.122884
> normalTest(log_return,method="jb")

Title:
 Jarque - Bera Normalality Test

Test Results:
  STATISTIC:
    X-squared: 362.5726
```

```
  P VALUE:
    Asymptotic p Value: < 2.2e-16

Description:
 Sat Aug 19 22:51:45 2017 by user:

> t.test(log_return)

        One Sample t-test

data:  log_return
t = 4.0402, df = 455, p-value = 6.27e-05
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 0.544149 1.574873
sample estimates:
mean of x
 1.059511

> tt1 = skewness(log_return)/sqrt(6/546)
> tt1
[1] -6.443776
> pv =  2*pnorm(tt1)
> pv
[1] 1.165367e-10
> tt2 = kurtosis(log_return)/sqrt(24/546)
> tt2
[1] 19.81441
> pv = 2*(1-pnorm(tt2))
> pv
[1] 0
> m1=acf(log_return)
> names(m1)
[1] "acf"     "type"    "n.used" "lag"     "series" "snames"
> m1$acf
, , 1

              [,1]
 [1,]  1.000000000
 [2,]  0.226356832
 [3,] -0.009975215
 [4,] -0.038128697
 [5,] -0.015760585
 ...
[26,] -0.012220663
[27,]  0.064944965
```

```
> m2=pacf(log_return)
> names(m2)
[1] "acf"     "type"    "n.used" "lag"     "series" "snames"
> m2$acf
, , 1

               [,1]
 [1,]  0.2263568318
 [2,] -0.0645183854
 [3,] -0.0223545572
 [4,] -0.0022849724
 [5,]  0.0116224218
 ...
[25,] -0.0142841594
[26,]  0.0788941527

> Box.test(log_return,lag=12,type='Ljung')

        Box-Ljung test

data:  log_return
X-squared = 37.302, df = 12, p-value = 0.0001995
```

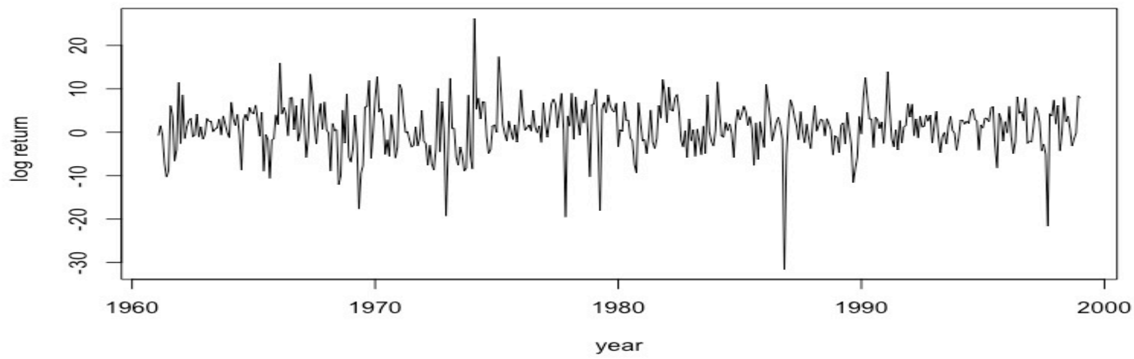
log return
year
1960
1970
1980
1990
2000

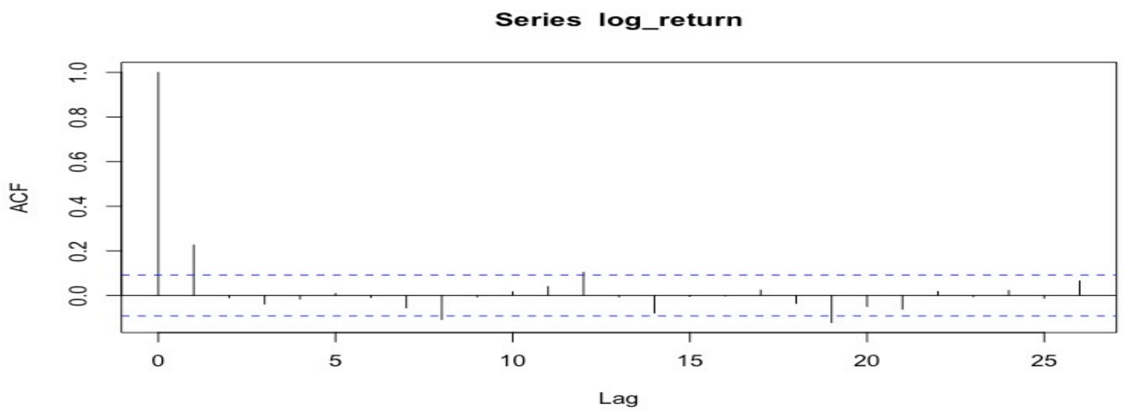
Series log_return
ACF
Lag

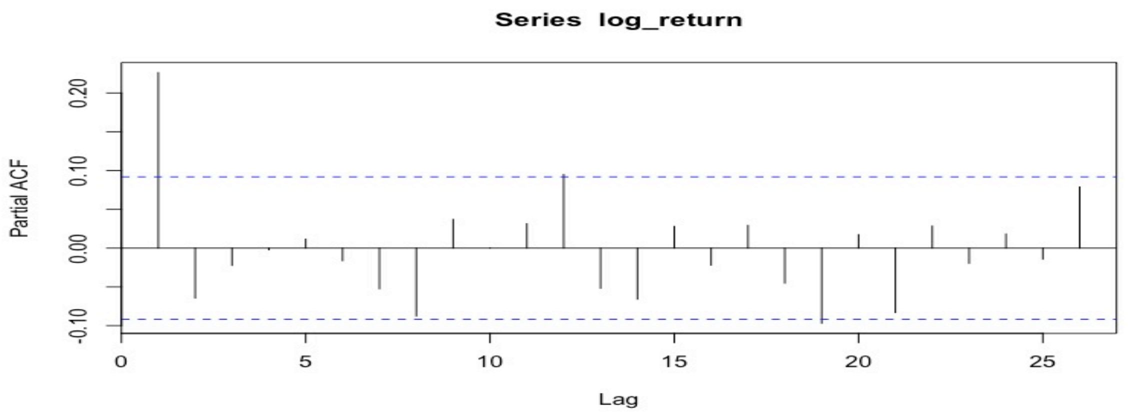
Series log_return
Partial ACF
Lag

2.2 Back-Shift (lag) operator

Definition $Br_t = r_{t-1}$ or $Lr_t = r_{t-1}$
$B^2 r_t = B(Br_t) = Br_{t-1} = r_{t-1}$
B or L means Time Shift
For example Br_t is the value of the series at time $t-1$
For example
The table of log return:

Date	r_t
1	0.025
2	0.013
3	-0.003
4	0.035

What is the following value of
r_2 =
Br_3 =
$B^2 r_5$ =

A proper perspective: at a time point t

Available data: $\{r_1, r_2, ...r_{t-1}\}$ ≡F$_{t-1}$

The return is decomposed into two parts:

$r_t =$ predictable part + not predictable part

Traditional TS modeling is concerned with μ_t:
Model for μ_t : mean equation
Volatility modeling concerns σ_t
Model for σ_t^2: volatility equation

2.3 Linear Time Series

r_t is linear if
. the predictable part is a linear function of F_{t-1}
. $\{a_t\}$ are independent and have the same distribution (iid)

Mathematically, it means r_t can be written as

$$r_t = \mu + \sum_{i=1}^{\infty} \psi_i a_{t-1}$$

where μ is constant and $\psi_0 = 1$ and a_t is an iid sequence with mean 0 and well-defined distribution.

In the economic itterature a_t is the shock or innovation at time t and ψ_i are the impulse responses of r_t.

White noise: iid sequence (with finite variance), which is the build- ing block of linear TS models. White noise is not predictable, but has zero mean and finite variance.

In EE435 we will study the (Univariate linear time series models) as follow:

1. autoregressive (AR) models
2. moving-average (MA) models
3. mixed ARMA models
4. seasonal models

Important properties of a model

- Stationarity condition
- Basic properties: mean, variance, serial dependence
- Empirical model building: specification, estimation, & checking
- Forecasting

2.4 Autoregressive(AR) model

If the series r_t and r_{t-1} are correlated, we might be able to use the series r_{t-1} in forecasing r_t. Thus the linear model can be:

$$r_t = \phi_0 + \phi_1 r_{t-1} + a_t$$

wherea_t are white noise series with mean equal to 0 and variance equals to σ_a^2

The above model is known as AR(1) since the variation of r_t can be explained by the variation of r_{t-1} . From this model we can calculate the conditional mean and conditional variance as the follow:

In general, we can write down the model of AR(p) as

$$r_t = \phi_0 + \phi_1 r_{t-1} + ... + \phi_p r_{t-p} + a_t$$

2.4.1 Properties of AR models

AR (1) model

the necessary condition for AR model is that series r_t have to be weak stationarity.

Unconditional mean

Unconditional variance

Unconditional autocorrelations

Figure The autocorrelation function of an AR(1) model: (a) for ϕ_1 = 0.8, and (b) for ϕ_1 = −0.8.

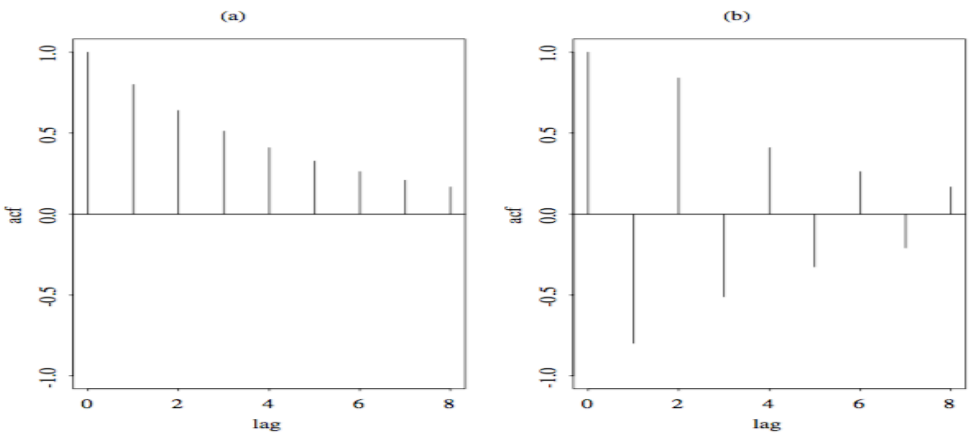

AR (2) model

$$r_t = \phi_0 + \phi_1 r_{t-1} + \phi_2 r_{t-2} + a_t$$

Unconditional mean

Unconditional variance

Unconditional autocorrelations

Figure The autocorrelation function of an AR(2) model: (a) ϕ_1 = 1.2 and ϕ_2 = −0.35, (b) ϕ_1 =0.6 and ϕ_2 =−0.4,(c) ϕ_1 =0.2and ϕ_2 =0.35,(d) ϕ_1 =−0.2 and ϕ_2 =0.35.

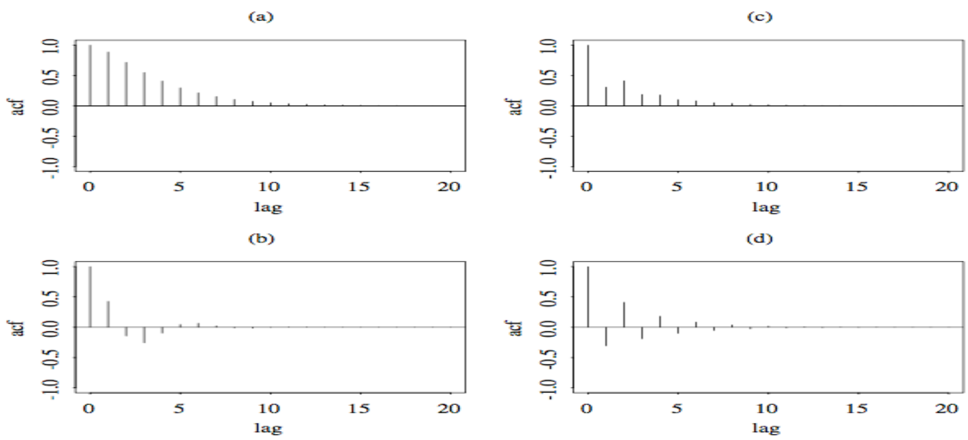

2.4.2 Stationarity of AR(p) Model

In case of AR(p)

$$r_t = \phi_0 + \phi_1 r_{t-1} + ... + \phi_p r_{t-p} + a_t$$

We can write down the Unconditional Mean as :

$$E(r_t) = \frac{\phi_o}{1 - \phi_1 - ... - \phi_p},$$

which the Polynomial equation can be expressed as

$$x^p - \phi_1 x^{p-1} - \phi_2 x^{p-2} - ... - \phi_p = 0$$

The above equation is also known as Equation in which if Characteristic roots has the value less than 1 in modulus, we can say that the model is stationary.

Moreover, AR(p) , the ACF can be written as difference equation

$$(1 - \phi_1 B - \phi_2 B^2 - ... - \phi_p B^P)\rho_l = 0$$

where $l > 0$

The graph of ACF of AR(p) has the pattern as the graph of sine and cosine.

2.4.3 Identifying AR Models

Partial Autocorrelation Function (PACF)

PACF is considered to be a tool to determine the order of AR(p). We can calculate the PACF from the following equations:

$$r_t = \phi_{0,1} + \phi_{1,1} r_{t-1} + e_{1t}$$

$$r_t = \phi_{0,2} + \phi_{1,2} r_{t-1} + \phi_{2,2} r_{t-2} + e_{2t}$$

$$r_t = \phi_{0,3} + \phi_{1,3} r_{t-1} + \phi_{2,3} r_{t-2} + \phi_{3,3} r_{t-3} + e_{3t}$$

In case of lag-2 PACF $\phi_{2,2}$ shows the marginal effect of r_{t-2} on r_t
In case of lag-3 PACF $\phi_{3,3}$ shows the marginal effect of r_{t-3} on r_t

Theefore, If the AR(p) is the optimal model, we then have the the lag-p PACF have to significantly different from 0, but $\phi_{j,j}$ have to be insignificant when $j > p$

Information Criteria

There are two methods to select the optimal lag AR(p)

1. Akaike Information Criterion

$$AIC(l) = ln(\tilde{\sigma}_l^2) + \frac{2l}{T}$$

For AR(l), $\tilde{\sigma}_l^2$ is the MLE of residual variance

We select the AR(l) model that provides the minimum AIC for all $l \in [0, ..., P]$

2. BIC Criterion

$$BIC(l) = ln(\tilde{\sigma}_l^2) + \frac{l * ln(T)}{T}$$

For AR(l), $\tilde{\sigma}_l^2$ is the MLE of residual variance

We select the AR(l) model that provides the minimum BIC for $l \in [0, ..., P]$

Example: GDP Growth

```
#EE 435 Wasin Siwasarit
setwd("/Users/wasinsiwasarit/Desktop/EE435")
cat(rep("\n",50))  #clear R Console
library(fBasics)
library(quantmod)
library(sn)
library(PerformanceAnalytics)
library(car)
library(tseries)
library(forecast)
library(Matrix)
da=read.table("dgnp82.txt")
x=da[,1]
par(mfcol=c(2,2))
plot(x,type='l')
plot(x[1:175],x[2:176])
plot(x[1:174],x[3:176])
acf(x,lag=12)
par(mfcol=c(1,1))
pacf(x,lag.max=12)
```

Figure: GDP growth, ACF and PACF

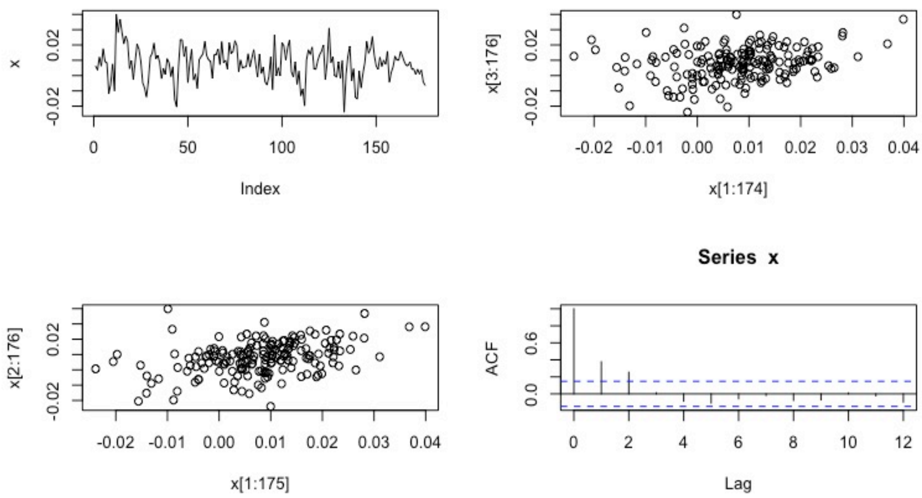

Figure: GDP growth, ACF and PACF (cont.)

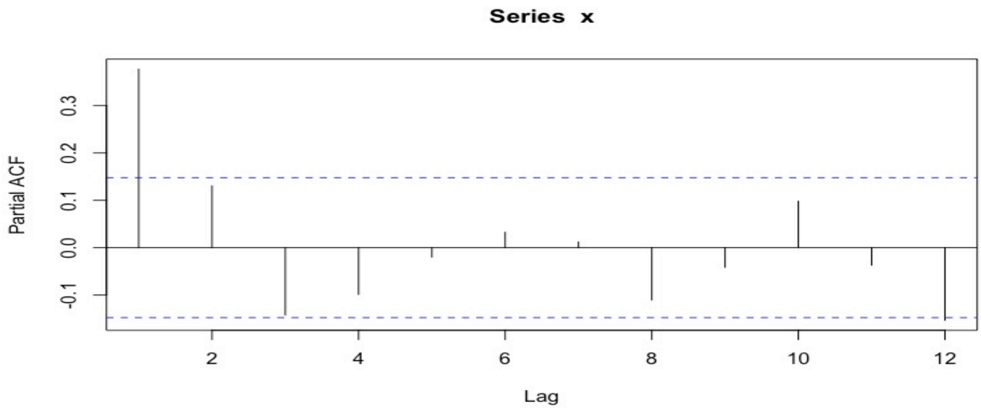

2.5 AR(P) in Lag Operator Notation

AR(1) in Lag Operator Notation

$$(r_t - \mu) = \phi_1 (r_{t-1} - \mu) + a_t$$

if $|\phi_1| < 1$ then,

AR(P) model

From the Mean-Adjusted Form:

$$(r_t - \mu) = \phi_1(r_{t-1} - \mu) + \ldots + \phi_p(r_{t-p} - \mu) + a_t$$

Stability and Stationarity Condition

$$\begin{bmatrix} r_t \\ r_{t-1} \\ \vdots \\ r_{t-p+1} \end{bmatrix} = \begin{bmatrix} \phi_1 & \phi_2 & \phi_3 & \ldots & \phi_p \\ 1 & 0 & \ldots & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 1 & 0 \end{bmatrix} \begin{bmatrix} r_{t-1} \\ r_{t-2} \\ \vdots \\ r_{t-p} \end{bmatrix} + \begin{bmatrix} a_t \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

we can write it as

$$\boldsymbol{\xi_t} = \mathbf{F}\boldsymbol{\xi_{t-1}} + \mathbf{vt}$$

Example: AR(2)

$$r_t = \phi_1 r_{t-1} + \phi_2 r_{t-2} + a_t$$

Results: The AR(p) model is weakly stationary and has Wold representation

$$r_t = \mu + \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j}$$

with $\psi_j = (1,1)$ element of $\mathbf{F^j}$ provided all of the eigenvalues of $\mathbf{F}$ have modulus less than 1.

2.6 Finding Eigenvalue

λ is an eigenvalue of $\mathbf{F}$ and $\mathbf{x}$ is the eigenvector if

$$\mathbf{Fx} = \boldsymbol{\lambda}\mathbf{x}$$

Example: AR(2)

The eigenvalues of $\mathbf{F}$ solve the ``reverse" characteristic equation

$$\lambda^2 - \phi_1 \lambda - \phi_2 = 0$$

Using the quadratic equation, the roots satisfy

$$\lambda_i = \frac{\phi_1 \pm \sqrt{\phi_1^2 + 4\phi_2}}{2}$$

These root may be real or complex. Complex roots induce periodic behavior in y_t. If λ_i is complex then

$$\lambda_i = a + bi$$

$$a = Rcos(\theta), b = Rsin(\theta)$$

$$R = \sqrt{a^2 + b^2}$$

Remark: R=modulus

Example 1: AR(2)

$$Y_t = 0.6Y_{t-1} + 0.2Y_{t-2} + \epsilon_t$$

Example 1

```
> Re(polyroot(c(-0.2, -0.6, 1)))
[1] -0.2385165  0.8385165
> Im(polyroot(c(-0.2, -0.6, 1)))
[1] -1.29247e-26  1.29247e-26
```

Example 2: AR(2)

$$Y_t = 0.5Y_{t-1} - 0.8Y_{t-2} + \epsilon_t$$

Example 2

```
> Re(polyroot(c(0.8, -0.5, 1)))
[1] 0.25 0.25
> Im(polyroot(c(0.8, -0.5, 1)))
[1]  0.8587782 -0.8587782
```

2.7 Parameter Estimation

For a specified AR(p) model , the conditional least squares method, which starts with the (p + 1)th observation, is often used to estimate the parameters. Specifically, conditioning on the first p observations, we have

$$r_t = \phi_0 + \phi_1 r_{t-1} + \ldots + \phi_p r_{t-p} + a_t$$

t=p+1,...,T

which can be estimated by the least squares method. Denote the estimate of ϕ by $\widehat{\phi}$. The fitted model is

$$r_t = \widehat{\phi}_0 + \widehat{\phi}_1 r_{t-1} + \ldots + \widehat{\phi}_p r_{t-p} + a_t$$

where the residual term is

$$\widehat{a_t} = r_t - \widehat{r_t}$$

$$\widehat{\sigma}_a^2 = \frac{\sum_{t=p+1}^{T} \widehat{a}_a^2}{T-2p-1}$$

2.8 Model Checking

A fitted model must be examined carefully to check for possible model inadequacy. If the model is adequate, then the residual series should behave as a white noise. The ACF and the Ljung–Box statistics of the residuals can be used to check the closeness of a_t to a white noise. For an AR(p) model, the Ljung–Box statistic Q(m) follows asymptotically a chi-squared distribution with m-p degrees of freedom. Here the number of degrees of freedom is modified to signify that p AR coefficients are estimated. If a fitted model is found to be inadequate, it must be refined.

2.9 Forecasting

Forecasting is an important application of time series analysis. For the AR(p) model, suppose that we are at the time index h and are interested in forecasting r_{h+l}, where $l \geq 1$. The time index h is called the forecast origin and the positive integer l is the forecast horizon. Let $\hat{r}_h(l)$ be the forecast of r_{h+l},

2.9.1 1-Step Ahead Forecast

2.9.2 2-Step Ahead Forecast

2.9.3 3-Step Ahead Forecast

Example: Analysis of U.S. GNP growth rate series.

```
#EE435
setwd("/Users/wasinsiwasarit/Desktop/EE435")
cat(rep("\n",50))  #clear R Console
da=read.table("dgnp82.txt")
x=da[,1]
> da=read.table("dgnp82.dat")
> x=da[,1]
> par(mfcol=c(2,2))  % put 4 plots on a page
plot(x,type'='l)  % first plot
plot(x[1:175],x[2:176]) % 2nd plot
plot(x[1:174],x[3:176]) % 3rd plot
acf(x,lag=12)  % 4th plot
pacf(x,lag.max=12) % Compute PACF
Box.test(x,lag=10,type'='Ljung) % Compute Q(10) statistics
Box-Ljung test
data:  x
X-squared = 43.2345, df = 10, p-value = 4.515e-06
m1=ar(x,method'='mle) % Automatic AR fitting using AIC criterion.
m1
Call: ar(x = x, method = "mle")
Coefficients:
1         2         3      % An AR(3) is specified.
0.3480    0.1793   -0.1423
Order selected 3  sigma^2 estimated as  9.427e-05
names(m1)
[1]  "order"   "ar"   "var.pred"   "x.mean"   "aic"
[6]  "n.used"    "order.max"      "partialacf"    "resid"         "method"
[11]  "series"     "frequency"    "call"           "asy.var.coef"

plot(m1$resid,type'='l)  % Plot residuals of the fitted model (not shown)

Box.test(m1$resid,lag=10,type'='Ljung) % Model checking
Box-Ljung test
data:  m1$resid
X-squared = 7.0808, df = 10, p-value = 0.7178


m2=arima(x,order=c(3,0,0))  % Another approach with order given.
 m2
Call:  arima(x = x, order = c(3, 0, 0))
Coefficients:
         ar1      ar2       ar3   intercept  % Fitted model is
0.3480  0.1793  -0.1423     0.0077  % y(t)=0.348y(t-1)+0.179y(t-2)
s.e.  0.0745  0.0778   0.0745     0.0012  %      -0.142y(t-3)+a(t),
```

```
% where y(t) = x(t)-0.0077
sigma^2 estimated as 9.427e-05:  log likelihood = 565.84,  aic = -1121.68
> names(m2)
[1] "coef"      "sigma2"    "var.coef"  "mask"    "loglik"    "aic"
[7] "arma"      "residuals" "call"      "series"   "code"      "n.cond"
[13] "model"
Box.test(m2$residuals,lag=10,type'='Ljung)
Box-Ljung test
data:  m2$residuals
X-squared = 7.0169, df = 10, p-value = 0.7239
ts.plot(m2$residuals)  % Residual plot
tsdiag(m2) % obtain 3 plots of model checking (not shown in handout).
p1=c(1,-m2$coef[1:3])  % Further analysis of the fitted model.
roots=polyroot(p1)
roots
[1]  1.590253+1.063882e+00i -1.920152-3.530887e-17i  1.590253-1.063882e+00i
Mod(roots)
[1] 1.913308 1.920152 1.913308
predict(m2,8)  % Prediction 1-step to 8-step ahead.
$pred
Time Series:
Start = 177
End = 184
Frequency = 1
[1] 0.001236254 0.004555519 0.007454906 0.007958518
[5] 0.008181442 0.007936845 0.007820046 0.007703826
$se
Time Series:
Start = 177
End = 184
Frequency = 1
[1] 0.009709322 0.010280510 0.010686305 0.010688994
[5] 0.010689733 0.010694771 0.010695511 0.010696190
```

Figure: GDP growth and PACF of GDP growth

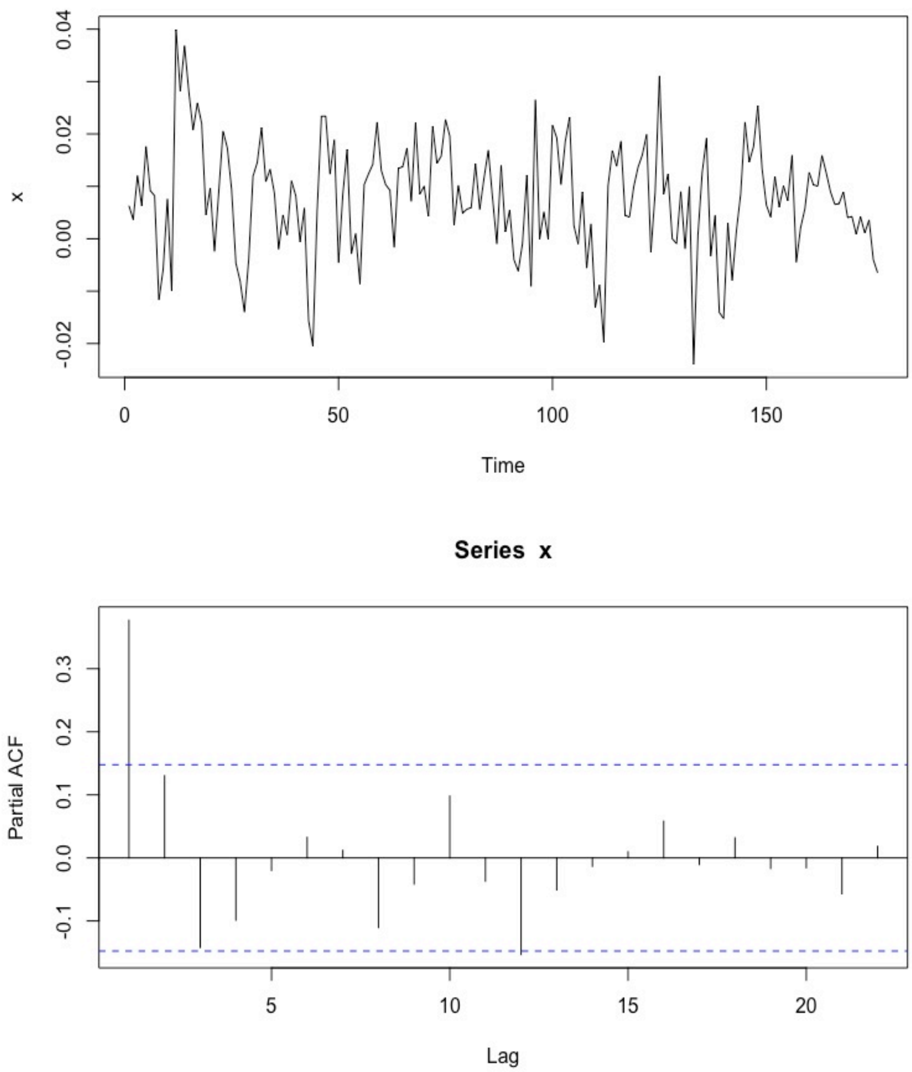

Figure: Residual Term from the estimated model and ACF of Residual Term

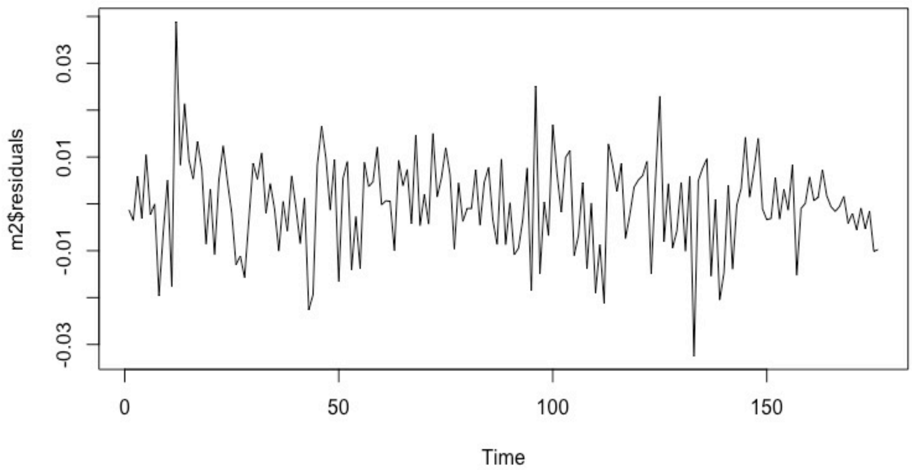

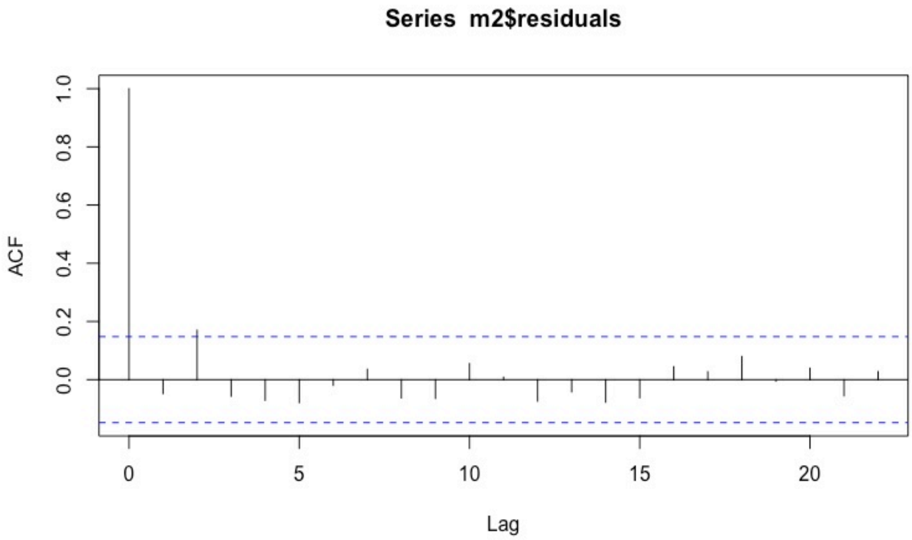

2.10 Moving-Average Models (MAs)

An AR model with infinite order can be written as:

$$r_t = \phi_0 + \phi_1 r_{t-1} + \phi_2 r_{t-2} + \ldots + a_t$$

The above AR model is not realistic because it has infinite many parameters. One way to make the model practical is to assume that the coefficients satisfy some constraints so that they are determined by a finite number of parameters. A special case of this idea is

$$r_t = \phi_0 - \theta_1 r_{t-1} - \theta_1^2 r_{t-2} - \theta_1^3 r_{t-3} - \ldots + a_t$$

where $\phi_i = -\theta_1^i$ for all i

2.10.1 Properties of MA Models

Stationarity MA models are always weakly stationary because they are finite linear combinations of a white noise sequence for which the first two moments are time invariant.

$E(r_t)$

$Var(r_t)$

Autocorrelation Function

2.10.2 Identifying MA order

The ACF is useful in identifying the order of an MA model. For the MA(q), we find out that $\rho_q \neq 0$ but $\rho_l = 0$ for all $l > q$

The example of MA(5) can be depicted as follwing:

Figure of the MA(5) model and its ACF

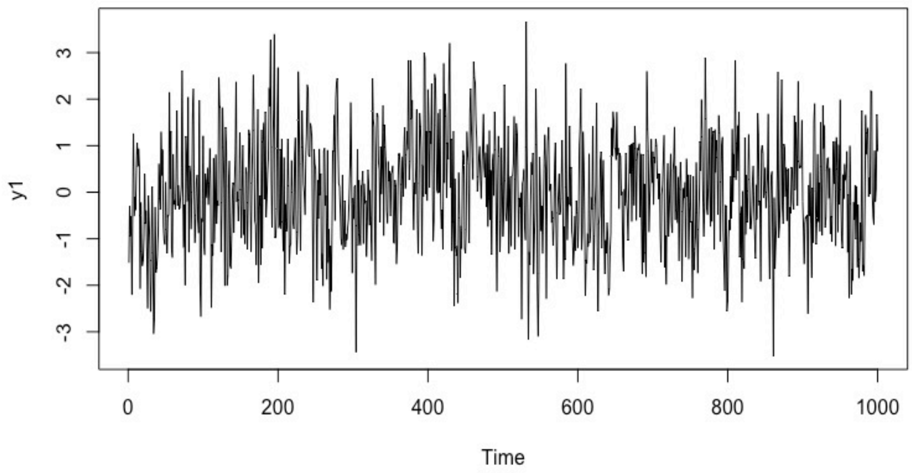

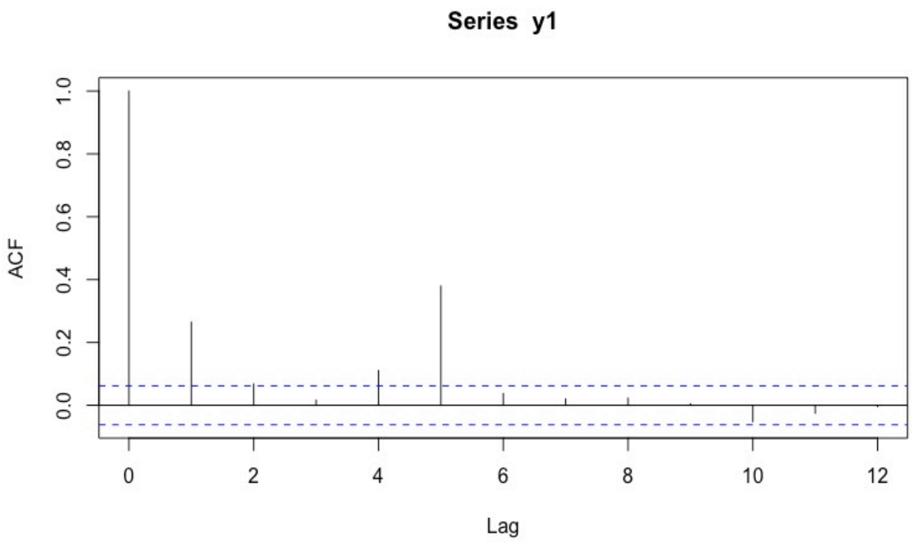

2.11 Parameter Estimation

For the MA(q), we can apply the Conditional Maximum Likelihood Estimation to estimate by starting from the obsevation (q+1). It can write down the model as following:

$$r_t = c_0 + a_t - \theta_1 a_{t-1} - \ldots - \theta_q a_{t-q}$$

t=q+1,...,T

The fitted model can be written as:

$$r_t = \widehat{c}_0 - \widehat{\theta}_1 a_{t-1} + \ldots + \widehat{\theta}_q a_{t-q} + \widehat{a}_t$$

We can define the residual terms from the below equation:

$$\widehat{a_t} = r_t - \widehat{r_t}$$

$$\widehat{\sigma}_a^2 = \frac{\sum_{t=q+1}^{T} \widehat{a}_a^2}{T - q - 1}$$

2.12 Model Checking

A fitted model must be examined carefully to check for possible model inadequacy. If the model is adequate, then the residual series should behave as a white noise. The ACF and the Ljung–Box statistics of the residuals can be used to check the closeness of a_t to a white noise. For an MA(q) model, the Ljung–Box statistic Q(m) follows asymptotically a chi-squared distribution with m-q degrees of freedom. Here the number of degrees of freedom is modified to signify that q MA coefficients are estimated. If a fitted model is found to be inadequate, it must be refined.

2.13 Forecasting

Forecasts of an MA model can easily be obtained. Because the model has finite memory, its point forecasts go to the mean of the series quickly. To see this, assume that the forecast origin is h. For the 1-step ahead forecast of an MA(1) process which is defined as

$$\hat{r}_h(l)$$

2.13.1 1-Step Ahead Forecast

2.13.2 2-Step Ahead Forecast

2.13.3 Multi-Step Ahead Forecast

The example of MA(q) model

```
> m2=arima(y1,order=c(0,0,5),include.mean = TRUE)
> m2

Call:
arima(x = y1, order = c(0, 0, 5), include.mean = TRUE)

Coefficients:
         ma1     ma2      ma3      ma4     ma5  intercept
      0.3172  0.0513  -0.0173  -0.0071  0.5308     0.0167
s.e.  0.0286  0.0285   0.0277   0.0292  0.0282     0.0580

sigma^2 estimated as 0.9601:  log likelihood = -1399.62,  aic = 2813.23
> names(m2)
 [1] "coef"      "sigma2"    "var.coef"  "mask"      "loglik"    "aic"       "
    arma"      "residuals" "call"
[10] "series"    "code"      "n.cond"    "nobs"      "model"
> Box.test(m2$residuals,lag=10,type='Ljung')

        Box-Ljung test

data:  m2$residuals
X-squared = 5.6582, df = 10, p-value = 0.8431

> ts.plot(m2$residuals)
> predict(m2,5)
$pred
Time Series:
Start = 1001
End = 1005
Frequency = 1
[1]  0.3016739 -0.4146540  0.1831065  1.0182231  0.1933976

$se
Time Series:
Start = 1001
End = 1005
Frequency = 1
[1] 0.9798251 1.0279423 1.0291705 1.0293095 1.0293327
```