

MATH 545: INTRODUCTION TO TIME SERIES ANALYSIS

WINTER 2025

COINTEGRATION, ERROR CORRECTION AND GRANGER CAUSALITY

Example: Consider the VARMA(1, 1) model

$$\Phi(B)\mathbf{X}_t = \Theta(B)\mathbf{Z}_t$$

with

$$\Phi(B) = \Phi_1(B) = \begin{bmatrix} 1 - 0.5B & B \\ 0.25B & 1 - 0.5B \end{bmatrix} \quad \Theta(B) = \Theta_1(B) = \begin{bmatrix} 1 - 0.2B & 0.4B \\ 0.1B & 1 - 0.2B \end{bmatrix}$$

with $\text{Var}[\mathbf{Z}_t] = \Sigma$. That is

$$\begin{bmatrix} X_{1,t} \\ X_{2,t} \end{bmatrix} = \begin{bmatrix} 0.5 & -1 \\ -0.25 & 0.5 \end{bmatrix} \begin{bmatrix} X_{1,t-1} \\ X_{2,t-1} \end{bmatrix} + \begin{bmatrix} Z_{1,t} \\ Z_{2,t} \end{bmatrix} + \begin{bmatrix} -0.2 & 0.4 \\ 0.1 & -0.2 \end{bmatrix} \begin{bmatrix} Z_{1,t-1} \\ Z_{2,t-1} \end{bmatrix}$$

or

$$\mathbf{X}_t = \Phi \mathbf{X}_{t-1} + \mathbf{Z}_t + \Theta \mathbf{Z}_{t-1}.$$

say. Then

$$\begin{bmatrix} 1 - 0.5B & B \\ 0.25B & 1 - 0.5B \end{bmatrix} \mathbf{X}_t = \begin{bmatrix} 1 - 0.2B & 0.4B \\ 0.1B & 1 - 0.2B \end{bmatrix} \mathbf{Z}_t$$

This process is nonstationary, as we can discover by finding the eigenvalues of the Φ matrix:

```
Phi<-matrix(c(0.5,-1,-0.25,0.5),2,2,byrow=T)
eigen(Phi)$values
## [1] 1 0
```

Filtering on both sides by

$$\begin{bmatrix} 1 - 0.5B & -B \\ -0.25B & 1 - 0.5B \end{bmatrix}$$

we obtain on the left and right hand sides respectively

$$\begin{bmatrix} 1 - 0.5B & -B \\ -0.25B & 1 - 0.5B \end{bmatrix} \begin{bmatrix} 1 - 0.5B & B \\ 0.25B & 1 - 0.5B \end{bmatrix} = \begin{bmatrix} 1 - B & 1 \\ 0 & 1 - B \end{bmatrix}$$

and

$$\begin{bmatrix} 1 - 0.5B & -B \\ -0.25B & 1 - 0.5B \end{bmatrix} \begin{bmatrix} 1 - 0.2B & 0.4B \\ 0.1B & 1 - 0.2B \end{bmatrix} = \begin{bmatrix} 1 - 0.7B & -0.6B \\ -0.15B & 1 - 0.7B \end{bmatrix}.$$

Therefore

$$\begin{bmatrix} 1 - B & 0 \\ 0 & 1 - B \end{bmatrix} \mathbf{X}_t = \begin{bmatrix} 1 - 0.7B & -0.6B \\ -0.15B & 1 - 0.7B \end{bmatrix} \mathbf{Z}_t$$

That is

$$\begin{bmatrix} (1 - B)X_{1,t} \\ (1 - B)X_{2,t} \end{bmatrix} = \begin{bmatrix} Z_{1,t} - 0.7Z_{1,t-1} - 0.6Z_{2,t-1} \\ Z_{2,t} - 0.15Z_{1,t-1} - 0.7Z_{2,t-1} \end{bmatrix}$$

Processes $X_{1,t}$ and $X_{2,t}$ have unit roots, specifically

$$X_{1,t}, X_{2,t} \sim \text{ARIMA}(0, 1, 1).$$


For example, with

$$\Sigma = \begin{bmatrix} 0.50 & 0.0 \\ 0.0 & 0.50 \end{bmatrix}$$

```

set.seed(132)
library(mvtnfast)
n<-1000
Sigma<-matrix(c(0.5,0,0,0.5),nrow=2,byrow=T)
Phi<-matrix(c(0.5,-1,-0.25,0.5),2,2,byrow=T)
Theta<-matrix(c(-0.2,0.4,0.1,-0.2),2,2,byrow=T)
Zt<-rmvn(n,c(0,0),Sigma)
Xt<-matrix(0,nrow=n,ncol=2)
for(i in 2:n){
  Xt[i,]<-Phi %*% Xt[i-1,] + Zt[i,] + Theta %*% Zt[i-1,]
}
par(mar=c(4,4,2,2))
plot(Xt[,1],type='l',ylim=range(Xt))
lines(Xt[,2],col='red')

```



We can check the ARIMA conclusion by inspecting the ACFs

```

cbind(acf(Xt[,1],plot=F,lag.max=10)$acf,acf(Xt[,2],plot=F,lag.max=10)$acf)
##           [,1]      [,2]
## [1,] 1.0000000 1.0000000
## [2,] 0.9784043 0.9481514
## [3,] 0.9689795 0.9389766
## [4,] 0.9624603 0.9340313
## [5,] 0.9544856 0.9245850
## [6,] 0.9474375 0.9197050
## [7,] 0.9397599 0.9080038
## [8,] 0.9312575 0.8987133
## [9,] 0.9227290 0.8935457
## [10,] 0.9137454 0.8845359
## [11,] 0.9047615 0.8785093
cbind(acf(diff(Xt[,1]),plot=F,lag.max=10)$acf,acf(diff(Xt[,2]),plot=F,lag.max=10)$acf)

```

```
##           [,1]      [,2]
## [1,]  1.000000000  1.000000000
## [2,] -0.329987357 -0.441782910
## [3,] -0.078731213 -0.039128029
## [4,]  0.045285246  0.048872130
## [5,] -0.021237843 -0.040325691
## [6,]  0.009397625  0.067059013
## [7,]  0.023253250 -0.035608664
## [8,] -0.008975341 -0.023806951
## [9,]  0.024083096  0.027023208
## [10,] -0.010672738 -0.027115645
## [11,] -0.031047795  0.007367744
```

However, let

$$\mathbf{L} = \begin{bmatrix} 1 & -2 \\ 0.5 & 1 \end{bmatrix}$$

and let $\mathbf{Y}_t = \mathbf{L}\mathbf{X}_t$ and $\mathbf{U}_t = \mathbf{L}\mathbf{Z}_t$ Then

$$\begin{aligned} \mathbf{Y}_t = \mathbf{L}\mathbf{X}_t &= \mathbf{L}\Phi\mathbf{X}_{t-1} + \mathbf{L}\mathbf{Z}_t + \mathbf{L}\Theta\mathbf{Z}_{t-1} \\ &= \mathbf{L}\Phi\mathbf{L}^{-1}(\mathbf{L}\mathbf{X}_{t-1}) + \mathbf{L}\mathbf{Z}_t + \mathbf{L}\Theta\mathbf{L}^{-1}(\mathbf{L}\mathbf{Z}_{t-1}) \\ &= \mathbf{L}\Phi\mathbf{L}^{-1}\mathbf{Y}_{t-1} + \mathbf{U}_t + \mathbf{L}\Theta\mathbf{L}^{-1}\mathbf{U}_{t-1} \end{aligned}$$

or equivalently

$$\mathbf{Y}_t - \Phi^*\mathbf{Y}_{t-1} = \mathbf{U}_t + \Theta^*\mathbf{U}_{t-1}$$

In this equation

$$\Phi^* = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad \Theta^* = \begin{bmatrix} -0.4 & 0 \\ 0 & 0 \end{bmatrix}$$

so that

$$\begin{aligned} Y_{1,t} - Y_{1,t-1} &= U_{1,t} - 0.4U_{1,t-1} \\ Y_{2,t} &= U_{2,t} \end{aligned}$$

Hence $\{Y_{1t}\} \sim \text{ARIMA}(0, 1, 1)$ and $\{Y_{2t}\}$ is a white noise process, and there is no dependence between the two series.

```
L<-matrix(c(1,-2,0.5,1),2,2,byrow=T)
Phistar<-L %*% Phi %*% solve(L)
Thetastar<- L %*% Theta %*% solve(L)
Sigmastar<-L %*% Sigma %*% solve(L)
Phistar
```

```
##           [,1] [,2]
## [1,]      1    0
## [2,]      0    0
```

```
Thetastar
```

```
##           [,1] [,2]
## [1,] -0.4    0
## [2,]  0.0    0
```

```
Sigmatstar
```

```
##           [,1] [,2]
## [1,]  0.5  0.0
## [2,]  0.0  0.5
```

```


Yt<-t(L %*% t(Xt))
acf.mat<-cbind(acf(Yt[,1],plot=F,lag.max=10)$acf,acf(Yt[,2],plot=F,lag.max=10)$acf,
              acf(diff(Yt[,1]),plot=F,lag.max=10)$acf)
colnames(acf.mat)<-c('Y1','Y2','diff(Y1)')
rownames(acf.mat)<-paste('h=',0:10)
acf.mat

##           Y1           Y2      diff(Y1)
## h= 0  1.0000000  1.000000000  1.000000000
## h= 1  0.9840648  0.010720175 -0.293021091
## h= 2  0.9755148 -0.032124119 -0.055959755
## h= 3  0.9686119  0.015712022  0.093498512
## h= 4  0.9596378  0.017948690 -0.067015648
## h= 5  0.9527233  0.056018549  0.045540369
## h= 6  0.9443686 -0.016105697  0.015150209
## h= 7  0.9351219 -0.009743926 -0.033179654
## h= 8  0.9274533  0.022332015 -0.015082854
## h= 9  0.9197033 -0.042702032  0.011439645
## h=10  0.9116823 -0.025228773 -0.001150626


par(mar=c(4,4,3,2))
acf(Yt)

```


Series 1




Series 1 & Series 2



Series 2 & Series 1



Series 2




Hence, the unit roots in $\{X_{1,t}\}$ and $\{X_{2,t}\}$ result from the unit root in $\{Y_{1,t}\}$ and a subsequent linear transformation. This phenomenon is known as *cointegration*.

Example: UK Consumption Data: In the UK consumption data, the three series are:

- Consumption, C_t
- Income, I_t
- Wealth, W_t

in millions of pounds, measured on the log scale.

```
library(tseries)
library(urca)
data(Raotbl3)
attach(Raotbl3)
lc <- ts(lc, start=c(1966,4),end=c(1991,2),frequency=4)
li <- ts(li, start=c(1966,4),end=c(1991,2),frequency=4)
lw <- ts(lw, start=c(1966,4),end=c(1991,2),frequency=4)
ukcons <- window(cbind(lc,li,lw), start=c(1967,2),end=c(1991,2))
par(mar=c(4,4,2,2))
ts.plot(cbind(lc,li,lw),ylim=range(10,14),col=c('black','blue','red'))
legend("left", legend = c('Consumption','Income','Wealth'),
      col = c('black','blue','red'), lty = 1)
```



The three series appear to vary together, but appear to be non-stationary.

Unit Root Testing: The following results show that there seems to be a cointegration relationship between the three series. We fit the models

$$C_t = \alpha_0 + \alpha_1 I_t + \alpha_2 W_t + Z_t^*$$

$$I_t = \alpha_0 + \alpha_1 C_t + \alpha_2 W_t + Z_t^*$$

$$W_t = \alpha_0 + \alpha_1 I_t + \alpha_2 C_t + Z_t^*$$

with the cointegrating vector of α s being estimated by OLS. If the residual series are white noise then there is evidence of cointegration as, for example

$$C_t - \alpha_0 - \alpha_1 I_t - \alpha_2 W_t$$

is stationary. The three lm fits below attempt to find the cointegrating vectors.

```
lc.eq <- lm(lc~li + lw, data=ukcons)
li.eq <- lm(li~lc + lw, data=ukcons)
lw.eq <- lm(lw~li + lc, data=ukcons)
round(summary(lc.eq)$coef,6)

+           Estimate Std. Error   t value Pr(>|t|)
+ (Intercept) -0.178458   0.103491  -1.724385 0.087925
+ li           0.910971   0.013457  67.697103 0.000000
+ lw           0.079761   0.007764  10.273688 0.000000

round(summary(li.eq)$coef,6)

+           Estimate Std. Error   t value Pr(>|t|)
+ (Intercept)  0.302998   0.109864   2.757944 0.00699
+ lc           1.075667   0.015889  67.697103 0.00000
+ lw          -0.077579   0.009331  -8.314451 0.00000

round(summary(lw.eq)$coef,6)

+           Estimate Std. Error   t value Pr(>|t|)
+ (Intercept)  1.068061   0.952109   1.121785 0.264811
+ li          -5.462487   0.656987  -8.314451 0.000000
+ lc           6.631514   0.645485  10.273688 0.000000
```

The residual series are computed by

```
error.lc <- ts(resid(lc.eq), start=c(1967,2),end=c(1991,2), frequency=4)
error.li <- ts(resid(li.eq), start=c(1967,2),end=c(1991,2), frequency=4)
error.lw <- ts(resid(lw.eq), start=c(1967,2),end=c(1991,2), frequency=4)
```

The code below looks for unit roots amongst the residual series:

```
ci.lc <- ur.df(error.lc, lags=1, type='none')
ci.li <- ur.df(error.li, lags=1, type='none')
ci.lw <- ur.df(error.lw, lags=1, type='none')

summary(ci.lc)$cval; summary(ci.lc)$teststat

+      1pct  5pct 10pct
+ tau1 -2.6 -1.95 -1.61
+      tau1
+ statistic -4.141447

summary(ci.li)$cval; summary(ci.li)$teststat

+      1pct  5pct 10pct
+ tau1 -2.6 -1.95 -1.61
+      tau1
+ statistic -4.059031

summary(ci.lw)$cval; summary(ci.lw)$teststat

+      1pct  5pct 10pct
+ tau1 -2.6 -1.95 -1.61
+      tau1
+ statistic -2.713707
```

The conclusion is that

- there appears to be strong evidence for stationarity in the residuals from the regression of C_t on (I_t, W_t) (the test statistic is -4.1414, where the 1% critical value of -2.6);
- there appears to be strong evidence for stationarity in the residuals from the regression of I_t on (C_t, W_t) (the test statistic is -4.059);
- evidence from the third regression (test statistic -2.7137) is less compelling, but still significant.

Error Correction Models: The three error correction models (ECMs)

```
#Construct the error correction models
#Columns contain differenced, and lag-1 differenced data
ukcons2 <- ts(embed(diff(ukcons), dim=2), start=c(1967,4), freq=4)
colnames(ukcons2) <- c('lc.d', 'li.d', 'lw.d', 'lc.d1', 'li.d1', 'lw.d1')

#Construct the ECM Error terms (lagged residuals from original regression)
error.ecm1 <- window(lag(error.lc, k=-1), start=c(1967,4), end=c(1991, 2))
error.ecm2 <- window(lag(error.li, k=-1), start=c(1967,4), end=c(1991, 2))
error.ecm3 <- window(lag(error.lw, k=-1), start=c(1967,4), end=c(1991, 2))
ecm.eq1 <- lm(lc.d ~ error.ecm1 + lc.d1 + li.d1 + lw.d1, data=ukcons2)
ecm.eq2 <- lm(li.d ~ error.ecm2 + lc.d1 + li.d1 + lw.d1, data=ukcons2)
ecm.eq3 <- lm(lw.d ~ error.ecm3 + lc.d1 + li.d1 + lw.d1, data=ukcons2)
```

examine the equilibrium relationships between differenced versions of the variables. The three models fitted are

$$(1 - B)C_t = \psi_0 + \gamma_1 \hat{z}_{1,t-1}^* + \psi_1(1 - B)C_{t-1} + \psi_2(1 - B)I_{t-1} + \psi_3(1 - B)W_{t-1} + Z_t$$

$$(1 - B)I_t = \psi_0 + \gamma_2 \hat{z}_{2,t-1}^* + \psi_1(1 - B)C_{t-1} + \psi_2(1 - B)I_{t-1} + \psi_3(1 - B)W_{t-1} + Z_t$$

$$(1 - B)W_t = \psi_0 + \gamma_3 \hat{z}_{3,t-1}^* + \psi_1(1 - B)C_{t-1} + \psi_2(1 - B)I_{t-1} + \psi_3(1 - B)W_{t-1} + Z_t$$

where $\hat{z}_{i,t-1}^*$ for $i = 1, 2, 3$ are the residuals from the earlier fitted cointegration models.

Estimates of the parameters $\gamma_1, \gamma_2, \gamma_3$ will give information about the “correction” of the system; for example γ_1 measures the magnitude of the effect of the (lagged) residual from the series

$$C_t - (\alpha_0 + \alpha_1 I_t + \alpha_2 W_t).$$

If this lagged residual $\hat{z}_{1,t-1}^*$ is positive, then $\gamma_1 < 0$ indicates that the expected **change** in C_t over the next time interval (as measured by the left-hand side of the ECM, $(1 - B)C_t$, will also be negative as $\gamma_1 \hat{z}_{1,t-1}^* < 0$, that is, C_t is “corrected” by the long-run stability of the system.

(i) From the first ECM, we obtain the following result:

```
summary(ecm.eq1)$coef
```

	Estimate	Std. Error	t value	Pr(> t)
+ (Intercept)	0.005769527	0.001496385	3.855643	0.0002163621
+ error.ecm1	0.062527706	0.098408076	0.635392	0.5267850247
+ lc.d1	-0.285563063	0.115824073	-2.465490	0.0155792467
+ li.d1	0.261417064	0.086361329	3.027015	0.0032209818
+ lw.d1	0.082694488	0.031687819	2.609662	0.0106140443

The coefficient of error.ecm1 is not significant, so the error correction from the lagged cointegration relationship is negligible in this case. Thus there appears to be no “causal” structure, that is, there is little additional predictive power in understanding the cointegration relationships between the variables.

(ii) For the second ECM, we have the following results

```
summary(ecm.eq2)$coef
```

```
+           Estimate Std. Error    t value    Pr(>|t|)
+ (Intercept)  0.006643443 0.001879559   3.53457508 0.000647320535
+ error.ecm2   -0.539482421 0.114210976  -4.72356019 0.000008490864
+ lc.d1        -0.149634949 0.146448055  -1.02176126 0.309632576090
+ li.d1        -0.006035003 0.108544106  -0.05559954 0.955784070767
+ lw.d1         0.062672652 0.039784363   1.57530867 0.118694046766
```

Here, the coefficient of `error.ecm2` is significantly different from zero, thus indicating a causal relationship; the estimate is negative as required, indicating a “correction” mechanism is present.

- (iii) For the third ECM, there appears to be no causal relationship as the coefficient of `error.ecm3` is not significantly different from zero.

```
summary(ecm.eq3)$coef
```

```
+           Estimate Std. Error    t value    Pr(>|t|)
+ (Intercept)  0.005549666 0.004975275   1.1154492 0.26762753
+ error.ecm3   -0.043771275 0.032810904  -1.3340466 0.18555451
+ lc.d1        -0.221625237 0.345405053  -0.6416387 0.52273864
+ li.d1         0.196085218 0.270032830   0.7261533 0.46962909
+ lw.d1         0.224392608 0.106155576   2.1138090 0.03729802
```

These findings are confirmed by looking at the correlation between the differenced responses (that is, $(1 - B)C_t$, $(1 - B)I_t$, and $(1 - B)W_t$) and the cointegration residuals $\hat{z}_{i,t-1}^*$:

```
cor(as.numeric(ukcons2[,1]),error.ecm1)
```

```
+ [1] -0.1624388
```

```
cor(as.numeric(ukcons2[,2]),error.ecm2)
```

```
+ [1] -0.4699282
```

```
cor(as.numeric(ukcons2[,3]),error.ecm3)
```

```
+ [1] -0.06952025
```