Chapter 8 Solutions

8.2.3 Exercises

Exercise 8.1

(a) The two time series appear to be fluctuating around a constant mean with a few outbursts in variance.

(b) See program to obtain the graph. Based on the rules given in SEction 8.2.1, the model identified is the following:

$$Act_t = \alpha_1 app_{t-1} + \alpha_2 act_{t-1} + \alpha_3 act_{t-2} + w_t$$
$$app_t = w_t$$

(c) We identify the same model using the random term of both time series. The code for this is not in the program.

□

Exercise 8.2

See section 8.2.2 in the book.

(a) The theoretical ccf does not display the small vertical bars that the sample ccf does. The sample ccf contains 5% autocorrelations that are significant just by chance. The simulated ccf contains significant autocorrelation at lag 0.

(b) There is some difference but this one is due to chance variability. We are sampling from the same stochastic model, but the random samples are different but not different in the main features, which are that there only lag 0 containing significant cross-correlation. The other lags, do not show significant autocorrelation.

(c) Repeat the simulation several times. What are the differences in the images obtained? What are those differences due to? Sampling variability is the reason why the CCFs look different in each case.

□

8.3.4 Exercises

Exercise 8.4

This is the same model seen in Example 8.5.

(a) $\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.6 & -0.3 \\ -0.3 & 0.6 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} w_{x,t} \\ w_{y,t} \end{pmatrix}$

$\begin{vmatrix} 1-0.6B & 0.3B \\ 0.3B & 1-0.6B \end{vmatrix}$

The roots for the polynomial resulting from this determinant equal to 10/9 = 1.111 and 10/3 = 3.333. So the model is stable.

(b) The coefficients matrix is

$A = \begin{pmatrix} 0.6 & -0.3 \\ -0.3 & 0.6 \end{pmatrix}$. The eigenvalues of the coefficients of A are 0.9 and 0.3. We find them as follows in R.

```
A=matrix(c(0.6,-0.3, -0.3, 0.6), ncol=2)
 eigen(A)$values
```

(c) The VAR(1) is stationary. The roots of the determinant of the backshift polynomial are larger than one, and the eigenvalues of the matrix of coefficients are less than one. Both are equivalent conditions for stationarity.

(d) We observe that the eigenvalues are the reciprocal of the roots. That is,

0.9 = 1/1.111 and 0.3= 1/3.333

□

Exercise 8.5

We rewrite the model in backshift polynomial notation first.

$$(I-AB)\begin{pmatrix} X_t \\ Y_t \end{pmatrix} = \begin{pmatrix} 1-\theta_{11}B-\theta_{12}B^2-\theta_{13}B^3 & -\theta_{14}B-\theta_{15}B^2-\theta_{16}B^3 \\ -\theta_{21}B-\theta_{22}B^2-\theta_{23}B^3 & 1-\theta_{24}B-\theta_{25}B^2-\theta_{26}B^3 \end{pmatrix}\begin{pmatrix} X_t \\ Y_t \end{pmatrix} = \begin{pmatrix} W_{1t} \\ W_{2t} \end{pmatrix}$$

Then we extract the matrix whose determinant will give us the roots. This is

$$\begin{pmatrix} 1-\theta_{11}B-\theta_{12}B^2-\theta_{13}B^3 & -\theta_{14}B-\theta_{15}B^2-\theta_{16}B^3 \\ -\theta_{21}B-\theta_{22}B^2-\theta_{23}B^3 & 1-\theta_{24}B-\theta_{25}B^2-\theta_{26}B^3 \end{pmatrix}$$

The two matrix for which we must find the eigenvalues are:

$$\begin{pmatrix} \theta_{11} & \theta_{14} \\ \theta_{21} & \theta_{24} \end{pmatrix}, \quad \begin{pmatrix} \theta_{12} & \theta_{15} \\ \theta_{22} & \theta_{25} \end{pmatrix}, \text{ and } \begin{pmatrix} \theta_{13} & \theta_{16} \\ \theta_{23} & \theta_{26} \end{pmatrix}$$

□

8.6.2 Exercises

Exercise 8.6

See Program *chapter8var2.R*. We select automatically a model for the training period ending in 1995 quarter 4. The model automatically chosen was a VAR(1). We then fitted an AR(1) model. The time series were not differenced, as requested. The residuals are not white noise.

Looking at the impulse response function, the plots show that the system returns to stability relatively soon after a shock to U (within 10 quarters) but that is not the case after a shock in the other variables. The response of some of the variables in those other shocks is to not return to stability. The following R code was used.

□

Exercise 8.7

Use Program *chapter8var1.R* and revisit Example 8.6. The difference with that example is that now you will need to use the undifferenced time series. That is, instead of using the differenced GNP and M1, as in that example, do not difference. See the code at the bottom of that R program.

You will find that the residuals from the VAR(2) model fitted have the CCF of bivariate white noise, as in the differenced case, but the model is not stable. The eigenvalues are not all less than one. The impulse response functions do not return to steady state, but rather diverge farther and farther from it as the lag increases. The system explodes after a shock.

□

Exercise 8.8

See program *Chapter8tvar1.R* and adapt it to the Canadian data set.

□

8.7.3 Exercises

Exercise 8.9

The IRF for this new system rests much later than the one in Example 8.12. That one rested before time 10, whereas this one is still fluctuating past lag 15. That one also had more inertia in the decay, decaying soon but at a slow rate, whereas the IRF of the new system changes direction every two lags, due to the negative coefficient of the AR model of the system at lag 2. Thus an AR(2) process has very different dynamic response to a shock depending on the value of the coefficients of the AR model representing the process.

To learn about the true dynamic behavior of the model we would have to simulate different time series from the model

$$x_t = 0.2x_{t-1} - 0.8x_{t-2} + w_t$$

and for each of them do the impulse response function. Then take the average of all impulse response functions.

time(t)	w_t	ψ_j
-2	0	0
-1	0	0
0	1	$x_0 = 0.2(0) - 0.8(0) + 1 = 1$
1	0	$x_1 = 0.2(1) - 0.8(0) + 0 = 0.2$
2	0	$x_2 = 0.2(0.2) - 0.8(1) + 0 = -0.76$
3	0	$x_3 = 0.2(-0.76) - 0.8(0.2) + 0 = -0.312$
4	0	0.5456
5	0	0.35872
6	0	-0.3647
7	0	-0.3599
8	0	0.2198
9	0	0.3319
10	0	-0.1095
11	0	-0.2874
12	0	0.0301
13	0	0.2359
14	0	0.02312
15	0	$x_{15} = 0.2(0.02312) - 0.8(0.2359) + 0 = -0.1841$
16	0	-0.055316
17	0	0.1362168
18	0	-0.07613242
19	0	-0.07613242

To obtain the plot of the impulse response function in discrete and continuous form. you may run the following R code.

```
#Plot discrete IRF
lag=c(0,0, 1:16)
    psi= c(0,0,1,0.2, -0.76, -0.312, 0.5456, 0.35872, -0.3647, -0.3599, 0.2198, 0.3319,
     -0.1095, -0.2874, 0.0301, 0.2359, 0.02312, -0.1841)

    par(mfrow=c(1,2))
    plot(lag,psi,type="h", ylab="Impulse response",xlab="lag",main="IRF of
     x_t = 0.2 x_t-1 - 0.8 x_t-2 + w_t ")
    abline(h=0)

  ### Plot continuous IRF

 plot(lag,psi,type="l", ylab="Impulse response",xlab="lag",main="IRF of
      x_t = 0.2 x_t-1 - 0.8 x_t-2 + w_t ")
  abline(h=0)
  dev.off()
```

If the reader checks the roots of the model's backshift polynomial, it will be seen that there is a unit root. The system in this exercise is not stationary, hence the behavior of the IRF.

□

Exercise 8.10

- With forecasts, we use the past values already existing for the time series to forecast the future according to the equation, e.g., for t+1,

$$\hat{x}_{t+1} = 0.9824x_t - 0.3722x_{t-1}$$

- Simulation is a different experience. We generate a white noise series, use its first two values as initial values for the simulation of x_t, and then use the model to simulate values of x_t. For example, a simulation to obtain data from the process given could be

```
x<- w <- rnorm(100)
 for(t in 3:100) x[t] = 0.9824*x[t-1] -0.3722*x[t-2] + w[t]
 plot(x,type="l")
 par(mfrow=c(1,2))
  acf(x)
  pacf(x)
  dev.off()
```

  Run the chunk of code several times to obtain different realizations of the ensemble and their ACF and PACF plots.

- With IRF, we simulate the values of the time series after a one standard error shock is inflicted to the w_t. After that, there are no more error terms added. Rather, we look at how the deterministic values of the time series are dynamically self-generated after that initial shock.

□

□

8.8.1 Exercises

Exercise 8.12

(a) For every 10 repetitions about 2 out of 10 on average might have correlation larger than 0.7.

(b) When doing a regular scatter plot of the two variables, we do not see relationship between them, except in the cases where the correlation is relatively higher than 0.6, for example.

(c) Save the last x and y generated.

□

8.8.3 Exercises

Exercise 8.13

(a) No matter what simulated set of two random walks we obtain, the unit root test for either one of them does not reject the null hypothesis that there is a unit root. Makes sense, since the time series are random walks, for which the root is 1.

(b) The correlograms of the first regular difference are the correlograms of white noise, for both variables, which makes sense, since the x and y are random walks. See Program *chapter8manirf.R*. Notice though that depending on the random walks simulated, the acf might have some extras, the 5 percent autocorrelations that are scattered randomly in the acf just by chance.

(c) The ACF is reliable, since the theory says that if you difference the random walk model,

$$y_t = y_{t-1} + w_t$$

then you are doing

$$y_t - y_{t-1} = w_t.$$

So you are left with just white noise, and therefore, in theory the ACF of white noise is such that no r_k is statistically significant. If there is any that are, this is because of chance variability.

□

8.8.6 Exercises

Exercise 8.14

See Program *chapter8manirf.R*, section corresponding to this exercise. For the two unit roots (random walks) time series that I generated (notice that there is no seed set, so each reader will have a different, numerically, answer, but same conclusions), the unit root test concludes that there is a unit root in both. The cointegration test's conclusion is that the two time series are not cointegrated.

However, if we had generated two random walks with a common trend, i.e., two cointegrated series, the test would conclude that they are cointegrated (we do that also in the program) we find that the two cointegrated)

```
Phillips-Ouliaris Cointegration Test

data:  cbind(x1, y1)
Phillips-Ouliaris demeaned = -986.45, Truncation lag parameter = 9, p-value = 0.01
```

The p-value is less than 0.05 so we reject the null hypothesis at that level of significance and conclude that the two cointegrated time series are cointegrated. See Section 8.8.5 for the null and alternative hypotheses specifications for this test.

Notice that answers will vary. □

Exercise 8.16

We modify slightly the beginning code of Program *chapter8var2.R* to conduct the unit root tests and the cointegration tests. Although we can not reject the null hypothesis that the time series each has a unit root, we do not have evidence that the variables are cointegrated. So differencing is a good idea. The program already runs the VAR with the differenced data.

```
#install.packages("vars")
library(vars)
data(Canada)
adf.test(Canada[,1])
adf.test(Canada[,2])
adf.test(Canada[,3])
adf.test(Canada[,4])

po.test(cbind(Canada[,1],Canada[,2]))
po.test(cbind(Canada[,1],Canada[,3]))
po.test(cbind(Canada[,1],Canada[,4]))

po.test(cbind(Canada[,2],Canada[,4]))
po.test(cbind(Canada[,2],Canada[,3]))
po.test(cbind(Canada[,3],Canada[,4]))
```

□

8.10 Problems

Problem 8.2

Let Y_t be the log of the area planted with sugar cane in Bangladesh in year t and let X_t be the log of the price of sugar cane in year t. We fitted a simple linear regression model

$$\hat{y}_t = 6.111 + 0.974x_t$$

$$s.e. = (0.169) \quad (0.111)$$

The ACF of the residuals shows that there is an AR structure left in the residuals, $w_t = 0.8w_{t-1} + v_t$, where v_t is white noise.

Do we have enough information to say that the log area planted and log price are co-integrated? Explain.

Two time series are co-integrated if there is a unit root in each of them and there is a linear combination of them that gives a stationary time series.

We do not know whether the above two time series, log area and log price have unit roots, so we do not have all the information. But we have information that a linear combination of them has stationary residuals, because we identified a stationary AR(1) model for the residuals of this model. That is,

$$\hat{y}_t - 6.111 - 0.974x_t = \alpha_1 w_{t-1} + v_t$$

□

Problem 8.3

(a) Y leads X, because Y_{t-1} has a statistically significant (p-value =0.000) effect on X_t, but X_{t-1} has no statistically significant effect on Y_t. So X_t lags Y. Or Y leads (affects X).

(b) NO. The ACF of residuals of X, the ACF of residuals of Y and the CCF of residuals of X and Y are not the ACF and CCF of bivariate white noise. A good VAR model would have to have bivariate white noise residuals.

□

Problem 8.4

(a) I conducted the augmented Dickey-Fuller Test of the null hypothesis that there is a unit root against the alternative that there is no unit root. The null is not rejected (p-value =0.7303 for the London series, p-value=0.99 for the New York series). Therefore, there is statistically significant evidence that there is a unit root in each of the series. From this it is not clear whether the process is a random walk or some other ARIMA model, since the Augmented Dickey Fuller test allows for other models with unit roots.

(b)

$$\begin{aligned} \hat{x}t(London) &= 8.386 + 0.985x_{t-1} + 0.0749y_{t-1} \\ (P - value) &= (0.000) \quad (0.000) \quad\quad (0.000) \\ \\ \hat{y}t(NY) &= 0.078 + -0.0003x_{t-1} + 1.002y_{t-1} \\ (P - value) &= (0.830) \quad (0.477) \quad\quad (0.000) \end{aligned}$$

From the fitted VAR given above, we can conclude that while the stock market in London at time t gets affected by the Market in New York at time $t-1$ (p-value=0.000), the market in London at time $t-1$ does not affect the NY market at time t (p-value=0.477). But both series get affected at time t by what happens in their own markets at time $t-1$ (p-values 0.000 for London and 0.000 for NY). The ACF of the residuals of each series, however, reveals that there might be some structure left over (small significant autocorrelation at lag 1 for the London residuals–I am not worried about the others at higher lags because they are too small and too far and they don't indicate any seasonality of any kind). The cross-correlation function reveals that there is some cross-correlation between the residuals series at positive lag 1(see Figure 21. All this indicates that the residuals may not be bivariate white noise, hence the model is not good as is.

Notice that in the R program, we use acf() to find the acf and ccf of the residuals in one plot. The reader will find that one easier to interpret based on our explanation in Section 8.2.

(c) Based on the results seen in part (b), we can say that NY affects London the most. First, because NY significantly affects London while viceversa is not happening. We know that financial disasters in the US have a large impact in the rest of the world, because of the US being an important economy in the world, so it is not surprising that NY leads the chaos or the bliss that may pervade the world's financial markets.

(d) I conducted the Phillips-Ouliaris (PO) test of the null hypothesis that the series are not cointegrated against the alternative that they are cointegrated. The null hypothesis is rejected (p-value=0.01) at the 5% level of significance. This means that there is statistically significant evidence that the series are cointegrated, they share a common stochastic trend, and therefore it is ok to have conducted regression of one against the other. (Note: the cointegration definition is that some linear combination of the two of them should give us a stationary series, not necessarily white noise. Cointegration implies that the two series share common stochastic trend).

(e)

$$x_t = a_0 + a_1 y_t + w_t$$

where x is for London's and y for New York The fitted model is

$$\begin{aligned} \hat{x}_t &= 629.866 + 4.827 y_t \\ (p - value) &= (0.000) \qquad (0.000) \end{aligned}$$

Both coefficients are statistically significant (different from 0). The R^2 says that 96.79% of the variability in the London series at time t is explained by the NY series at time t. The F test (p-value=0.00) is also statistically significant at 5% level or any level of significance. So the model is good. The ACF of the residuals however shows a very smooth decay typical of a nonstationary series, so they are not stationary. They could be very close to having a unit root or have one.

(f) Test the residuals of the fitted model for unit roots. Does the test support or contradict the result in part (d)? Explain your answer.

I conduct again an augmented Dickey Fuller test, this time for the residuals of the model fitted in part (e) and I reject the null that there is a unit root in these residuals (p-value=0.0178). This implies that the residuals are stationary.

The definition of cointegrated is that some linear combination of the two series is stationary. The residuals of the linear regression is a linear combination of the two series (y-linear combination of x), so we should be happy.

However, it is interesting to see that the ACF gives us the opposite information. If you fit an AR(1) model to the residuals of part (e) you will find that the regression coefficient is 0.98 ($res_t = 0.98 res_{t-1}$), suggesting that the residuals are borderline to having a unit root, but not quite. They are an AR process with close to having a unit root.

These unit root tests and cointegration tests are all based on asymptotic properties of the tests, thus, sometimes, they could be misleading and one should do more than just tests of hypotheses to determine what is going on with data.

□

8.11 Quiz

Question 8.2

1 standard error at time 1 and no change at $t \neq 1$

□

Question 8.4

(a), (b), (c)

□

Question 8.6

will cause the system to change but will eventually return the system to its steady state. Find the article in `https://www.sciencedirect.com/science/article/abs/pii/S0921800906000395` □

Question 8.8

See the case study for Chapter 8.

□

Question 8.9

See Section 8.7.2 and Table 8.3

□