

Lecture 8

7 Seasonal Models

Some times series exhibits certain cyclical/periodical behavior.

Example: Quarterly earning per share of a company of Johnson and Johnson for the period 1st quarter 1960 to the last quarter 1980, see Figure 2.13. This data possess some specific characteristics:

- the earnings grew exponentially and had a strong seasonality;
- the volatility of the earnings increased over time.

Note: the cyclical pattern repeats itself every year, so that then periodicity of the series is 4 (length of the cycle is 4 quarters). If monthly data are considered, for example monthly sales of Wal-Mart stores, then the periodicity is 12.

Seasonal models are useful dealing with environmental times series, e.g. pricing weather derivatives and energy future prices.

- In applications seasonality can be removed from the data, leading to a seasonally adjusted time series (with no seasonal effect), that can be used for inference.
- The procedure of removal of seasonality is called seasonal adjustment.
- Most published economic data are seasonally adjusted, for example growth rate of gross domestic product (GD), unemployment rate.

However, in some applications such as forecasting, seasonality is important and the data must be handled taking into account seasonality. Forecasting is a major objective of financial time series analysis. Below we discuss some econometric models, useful in modeling seasonal time series.

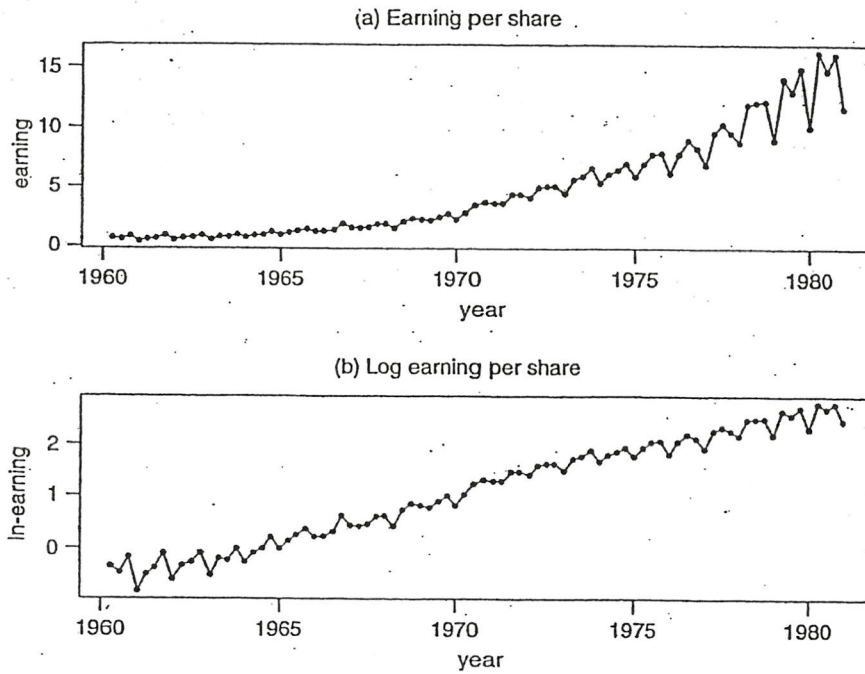


Figure 2.13. Time plots of quarterly earning per share of Johnson and Johnson from 1960 to 1980: (a) observed earning and (b) log earning.

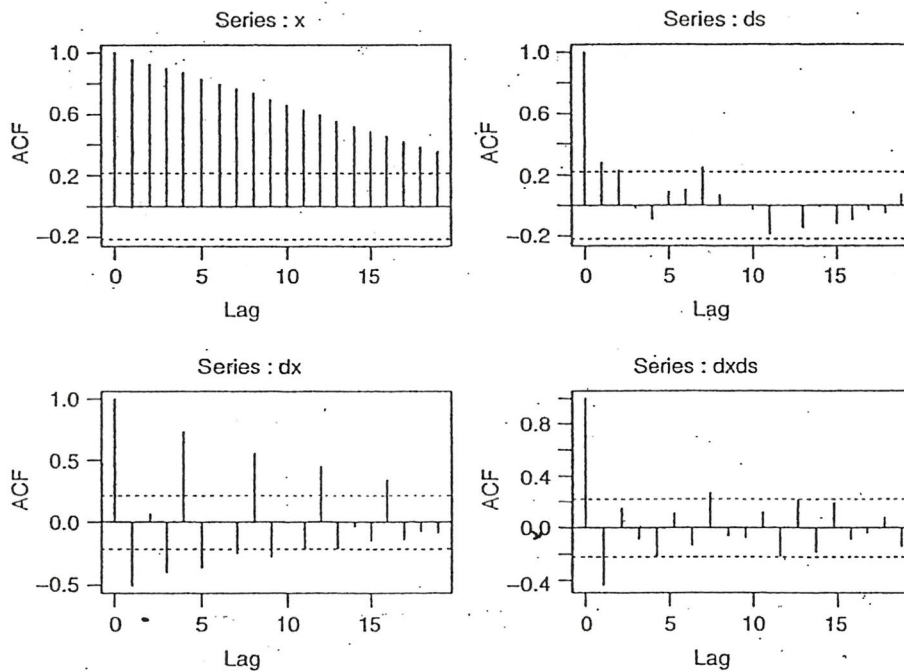


Figure 2.14. Sample ACF of the log series of quarterly earning per share of Johnson and Johnson from 1960 to 1980, where x_t is the log earning, dx is the first differenced series, ds is the seasonally differenced series, and $dx ds$ denotes series with regular and seasonal differencing.

In the models we discussed before dependence and ACF between current observation X_t and the past observation X_{t-k} was decaying fast as k was increasing.

For some time series, dependence and the decay of ACF shows a repeating cyclical behavior.

Examples:

- inventory level on a Monday will be similar to that on previous Mondays
- Sales of ice cream in June of one year will be strongly correlated with sales in June in previous years.

Cyclic patterns or seasonal patterns can be very effectively used to improve the forecast performance.

ARIMA model we discussed allow for modeling seasonal and non-seasonal dependence.

7.1 Example: Airline Passenger Data

Figure 5.1 shows the plot of the monthly number X_t of international airline passengers from January 1949 to December 1960.

Table 5.1 shows data in a two-way tables:

- for different years (columns)
- within months (horizontal lines) in the same year

Figure 5.2 shows seasonality and trend in more close-up format

- monthly seasonal pattern for each year
- change during entire 12 -ear period.

Such change seems to have increasing slope over time.

TABLE 5.1 Monthly Passenger Totals (Measured in Thousands) in International Air Travel — BJR Series G

Year	January	February	March	April	May	June	July	August	September	October	November	December
1949	112	118	132	129	121	135	148	148	136	119	104	118
1950	115	126	141	135	125	149	170	170	158	133	114	140
1951	145	150	178	163	172	178	199	199	184	162	146	166
1952	171	180	193	181	183	218	230	242	209	191	172	194
1953	196	196	236	235	229	243	264	272	237	211	180	201
1954	204	188	235	227	234	264	302	293	259	229	203	229
1955	242	233	267	269	270	315	364	347	312	274	237	278
1956	284	277	317	313	318	374	413	405	355	306	271	306
1957	315	301	356	348	355	422	465	467	404	347	305	336
1958	340	318	362	348	363	435	491	505	404	359	310	337
1959	360	342	406	396	420	472	548	559	463	407	362	405
1960	417	391	419	461	472	535	622	606	508	461	390	432

5.1 SEASONAL DATA 113




Figure 5.1 Time series plot of the monthly international airline passenger data—BJR series G.




Figure 5.2 Interaction plot of the monthly international airline passenger data—BJR series G.

Observations:

- The number of passengers is seasonal with a peak in July and lowest number in January and November.
- consistent patterns of trend and seasonality allow to make good forecasts.

Data transformation

Notice:

- Differences between peaks and valleys and variability seem to increase over time when the number of passengers gets larger.

An important step in data analysis is to use a data transformation to make variability more stable.

At this point we ignore seasonality and use log transformation: we replace X_t by the log data $z_t = \log(X_t)$

Figures 5.3 and 5.4 show the log transformed data. We see:

- trend is more linear
- variability more or less constant
- the lines connecting months are almost parallel (that means that year effect is almost additive)




Figure 5.3 Time series plot of the natural logarithm of the number of airline passengers.




Figure 5.4 Interaction plot of the natural logarithm of the number of airline passengers.

7.2 Seasonal ARIMA models

Important feature of seasonal time series:

- if the season is s period long, then observations that are s time intervals apart are alike.

In our case $s = 12$. So holiday month December is similar to previous December and correlated with the data from December of the following year.

We have two relationships going simultaneously:

- (i) between monthly observations within a year.
- (ii) between observations in the same month for subsequent years.

Therefore we need to build two time series models:

- one for relationship between successive observations within a year.
- one that links the months in successive years.

Having this in mind, model building is similar to that we used for non-seasonal models.

- First, if data is nonstationary we use differencing to make data stationary so that correlation dies out quickly.
- For seasonal data, we may need to use not only regular differencing $\nabla z_t = z_t - z_{t-1}$, but also seasonal differencing $\nabla_s z_t = z_t - z_{t-s}$.

For log transformed airline data with $s = 12$ months seasonality, we can proceed as follows:

- we try first difference $\nabla z_t = z_t - z_{t-1}$,
- we try seasonal difference $\nabla_{12} z_t = z_t - z_{t-12}$
- we try the two combined:

$$\nabla \nabla_{12} z_t = \nabla(z_t - z_{t-12}) = z_t - z_{t-1} - z_{t-12} + z_{t-13}.$$

Figure 5.6 shows the plots of the log transformed airline data z_t , the first difference ∇z_t and the combined difference $\nabla\nabla_{12}z_t$.

Figure 5.7 shows the corresponding ACF's.

- Figure 5.7b shows the seasonal autocorrelation spikes at 12, 24, 36 which is indication of seasonality with $s = 12$.
- Figures 5.6c and 5.7c show that after seasonal differencing data $\nabla_{12}z_t$ remains non-stationary, however after differencing twice, $\nabla\nabla_{12}z_t$ looks as stationary time series. ACF spikes at 12, 24, 36 disappear.
- We conclude that we need one regular and one seasonal differencing.
- For further modeling we need to use the time series $w_t = \nabla\nabla_{12}z_t$.

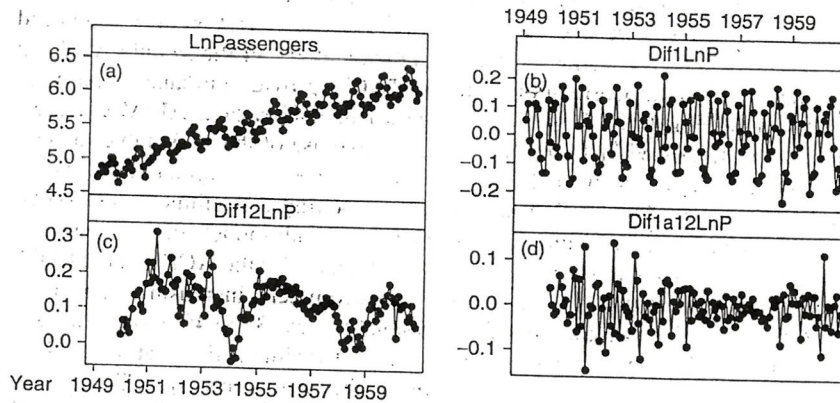


Figure 5.6 Time series plots of (a) the log transformed series z_t , (b) the first difference ∇z_t , (c) the seasonal difference $\nabla_{12}z_t$, and (d) the combined first difference and seasonal difference $\nabla\nabla_{12}z_t$.

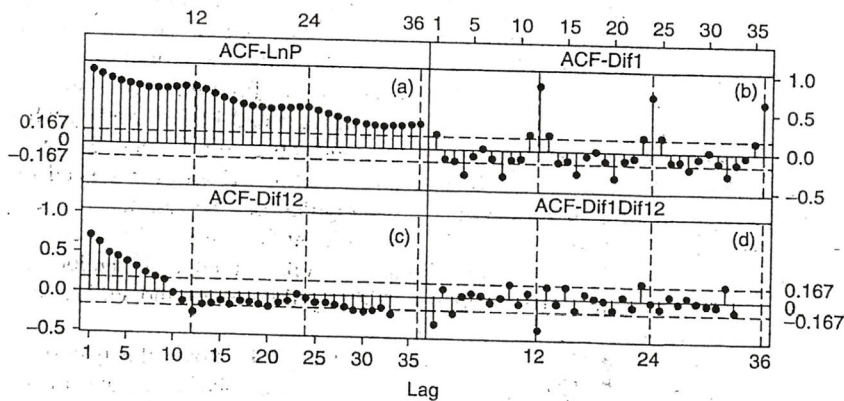


Figure 5.7 ACFs for (a) the log transformed series z_t , (b) the first difference ∇z_t , (c) the seasonal difference $\nabla_{12}z_t$, and (d) the combined first difference and seasonal difference $\nabla\nabla_{12}z_t$. The horizontal lines at ± 0.167 are approximate $2 \times$ standard error limits for the sample autocorrelations $\pm 2SE \approx \pm 2/\sqrt{n}$ where n is the sample size.

Model identification

The methodology is similar to ARMA models where ACF and PACF provide the guidance.

Differences:

- differently from regular time series, we need to look at the order of two models: one seasonal of each 12th month, and one regular: month-to-month
- In non-seasonal models it is sufficient to consider ACF and PACF up to 20-25 lags.
- In seasonal models we should consider ACF and PACF up to 3 or 4 multiples of seasonality, i.e. 36 or 48 in our example.

Modeling seasonal pattern: we look at similarities that are 12 lags apart.

- In Figure 5.7d of ACF of w_t we see negative spike at the lag 12, but later seasonal correlation cuts off (no correlation at lag 24 and 36).
- This suggests we should use the seasonal MA(1) model to model the 12th month seasonal pattern:

$$w_t = b_t - \Theta b_{t-12}$$

where b_t is not necessary a white noise process.

Modeling regular pattern: Next, we turn to the regular time series model, looking for patterns between successive months.

In Figure 5.7d we see significant correlation at lag 1, which cuts off at lag 2,3,4, ... This indicates that for modeling the month-to month data we need an additional MA(1) model:

$$b_t = \varepsilon_t - \theta\varepsilon_{t-1}$$

where ε_t is a white noise sequence.

We get

$$\begin{aligned} w_t &= b_t - \Theta b_{t-12} \\ &= \varepsilon_t - \theta\varepsilon_{t-1} - \Theta(\varepsilon_{t-12} - \theta\varepsilon_{t-13}) \\ &= \varepsilon_t - \theta\varepsilon_{t-1} - \Theta\varepsilon_{t-12} + \Theta\theta\varepsilon_{t-13}. \end{aligned}$$

Final model: Recall that by w_t we denote

$$w_t = \nabla \nabla_{12} z_t = z_t - z_{t-1} - z_{t-12} + z_{t-13}.$$

Putting all together, we get the model for the airline passenger number:

$$z_t - z_{t-1} - z_{t-12} + z_{t-13} = \varepsilon_t - \theta \varepsilon_{t-1} - \Theta \varepsilon_{t-12} + \Theta \theta \varepsilon_{t-13}.$$

Alternative form: we can write the above model in a simpler form using "backshift operator".

Backshift operator L :

$$\begin{aligned} LX_t &= X_{t-1} \text{ shifts time back by 1;} \\ L^2 X_t &= X_{t-2} \text{ shifts time back by 2;} \\ L^{12} X_t &= X_{t-12} \text{ shifts time back by 12.} \end{aligned}$$

Then the regular differencing:

$$\nabla z_t = (1 - L)z_t = z_t - Lz_t = z_t - z_{t-1}$$

The seasonal differencing:

$$\nabla_{12} z_t = (1 - L^{12})z_t = z_t - L^{12}z_t = z_t - z_{t-12}$$

The above model can be written as

$$\begin{aligned} w_t = \nabla \nabla_{12} z_t &= (1 - L)(1 - L^{12})z_t = z_t - z_{t-1} - z_{t-12} + z_{t-13}, \\ (1 - \theta L)(1 - \Theta L^{12})\varepsilon_t &= \varepsilon_t - \theta \varepsilon_{t-1} - \Theta \varepsilon_{t-12} + \Theta \theta \varepsilon_{t-13}; \end{aligned}$$

The above model can be written as

$$(1 - L)(1 - L^{12})z_t = (1 - \theta L)(1 - \Theta L^{12})\varepsilon_t.$$

This model is called multiplicative airline seasonal model.

Model fitting and check

We have a data set of 144 observations.

To demonstrated the use of the model to make short term forecasts, we set aside the last year 12 observations, to be able to compare the forecast with actual data. We use 132 observation for the model fitting.

Check of fit:

- Table 5.2 shows that the MA(1) coefficients of the seasonal part and regular part are significant
- Ljung Box test shows no significant correlation in residuals
- Figure 5.10 shows no significant ACF and PACF which confirms absence of correlation in residuals.

We conclude that the model we obtained for the log of the number of international passengers fits well to the data:

$$z_t = z_{t-1} + z_{t-12} - z_{t-13} + \varepsilon_t - 0.34\varepsilon_{t-1} - 0.63\varepsilon_{t-12} + 0.214\varepsilon_{t-13}.$$

TABLE 5.2 ARIMA (0, 1, 1) × (0, 1, 1)₁₂ Model Summary for the Airline Passenger Data

Model term	Coefficient	Standard error	<i>t</i>	<i>p</i>
MA 1	0.3407	0.0868	3.93	0.000
SMA 12	0.6299	0.0766	8.23	0.000
Differencing: 1 regular, 1 seasonal of order 12.				
Number of observations: Original series 132, after differencing 119.				
Residuals: SS = 0.151421; MS = 0.001294; <i>df</i> = 117.				
Modified Box–Pierce (Ljung–Box) chi-square statistic:				
Lag	12	24	36	48
Chi square	7.5	19.6	30.5	38.7
<i>df</i>	10	22	34	46
<i>p</i> -value	0.679	0.607	0.638	0.770




Figure 5.10 Plots of (a) ACF and (b) PACF of the residuals after fitting a seasonal ARIMA $(0, 1, 1) \times (0, 1, 1)_{12}$ model to the log airline data from January 1949 to December 1959.

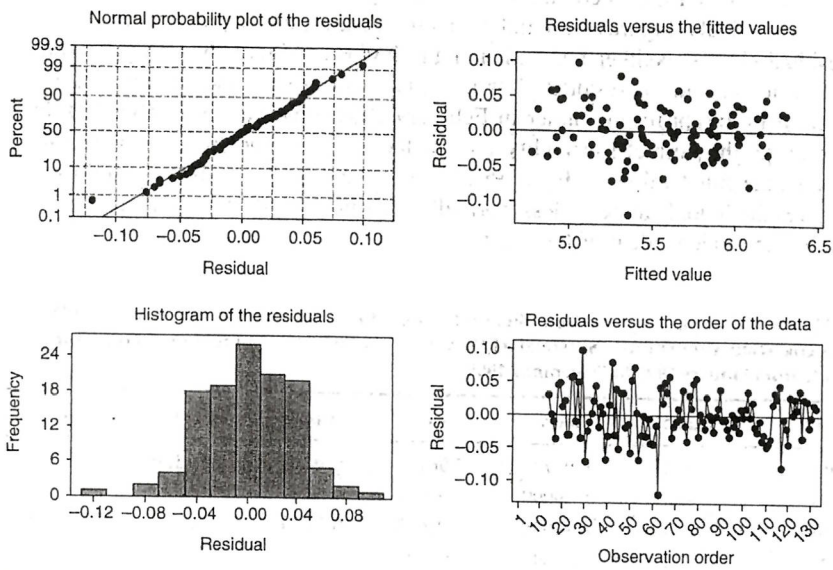


Figure 5.11 Residual checks after fitting a seasonal ARIMA $(0, 1, 1) \times (0, 1, 1)_{12}$ model to the log airline data from January 1949 to December 1959.

Forecasting of a seasonal model

To develop a forecast equation we follow the same procedure as for regular time series.

k -step ahead forecast $\hat{z}_t(k)$ of z_{t+k} . We assume $1 \leq k \leq 12$.

We start with the model for the fitted time series:

$$z_t = z_{t-1} + z_{t-12} - z_{t-13} + \varepsilon_t - \theta\varepsilon_{t-1} - \Theta\varepsilon_{t-12} + \theta\Theta\varepsilon_{t-13}.$$

We rewrite the model as

$$z_{t+k} = z_{t+k-1} + z_{t+k-12} - z_{t+k-13} + \varepsilon_{t+k} - \theta\varepsilon_{t+k-1} - \Theta\varepsilon_{t+k-12} + \theta\Theta\varepsilon_{t+k-13}.$$

The forecast we obtain using our previous rules:

$$\begin{aligned}\hat{z}_t(k) &= [z_{t+k-1} + z_{t+k-12} - z_{t+k-13} + \varepsilon_{t+k} - \theta\varepsilon_{t+k-1} - \Theta\varepsilon_{t+k-12} + \theta\Theta\varepsilon_{t+k-13}] \\ &= [z_{t+k-1}] + [z_{t+k-12}] - [z_{t+k-13}] + [\varepsilon_{t+k}] - \theta[\varepsilon_{t+k-1}] - \Theta[\varepsilon_{t+k-12}] + \theta\Theta[\varepsilon_{t+k-13}] \\ &= \hat{z}_t(k-1) + z_{t+k-12} - z_{t+k-13} - \theta[\varepsilon_{t+k-1}] - \Theta\varepsilon_{t+k-12} + \theta\Theta\varepsilon_{t+k-13}.\end{aligned}$$

For example, 1-months ahead forecast is

$$\hat{z}_t(1) = z_t + z_{t-11} - z_{t-12} - \theta\varepsilon_t - \Theta\varepsilon_{t-11} + \theta\Theta\varepsilon_{t-12}.$$

Using estimates $\theta = 0.34$ and $\Theta = 0.63$ we can compute our forecast using EViews.

To demonstrate how it works, suppose using the data up to the end of the year 1959 we wish to forecast passengers numbers for the year 1960. The forecasts include 95% confidence intervals.

- Table 5.3 shows 1-step ahead forecasting results for z_t . They are plotted in Figure 5.14
- Figure 5.15 shows forecasts for the actual passengers numbers $y_t = \exp(z_t)$.

TABLE 5.3 Forecasts with 95% Prediction Intervals for 1960 for the Airline Passenger Data on Log Scale After Fitting a Seasonal ARIMA (0, 1, 1) × (0, 1, 1)₁₂ Model to the Log Airline Data from January 1949 to December 1959

Time	Number	Forecast	Lower	Upper	Actual	Difference
1960-1	133	6.03771	5.96718	6.10823	6.03309	-0.00462
1960-2	134	5.99099	5.90652	6.07546	5.96871	-0.02228
1960-3	135	6.14666	6.05023	6.24308	6.03787	-0.10879
1960-4	136	6.12046	6.01341	6.22751	6.13340	0.01293
1960-5	137	6.15698	6.04026	6.27369	6.15698	-0.00000
1960-6	138	6.30256	6.17692	6.42819	6.28227	-0.02029
1960-7	139	6.42828	6.29432	6.56224	6.43294	0.00466
1960-8	140	6.43857	6.29677	6.58037	6.40688	-0.03169
1960-9	141	6.26527	6.11604	6.41450	6.23048	-0.03479
1960-10	142	6.13438	5.97807	6.29069	6.13340	-0.00098
1960-11	143	6.00539	5.84231	6.16846	5.96615	-0.03924
1960-12	144	6.11358	5.94401	6.28316	6.06843	-0.04515




Figure 5.14 Forecasts together with the actual observations and 95% prediction intervals for 1960 for the airline passenger data on the log scale after fitting a seasonal ARIMA $(0, 1, 1) \times (0, 1, 1)_{12}$ model to the log airline data from January 1949 to December 1959.




Figure 5.15 Forecasts together with the actual observations for 1960 for the airline passenger data in actual units.

Example. We fit to the quarterly earnings per share of Johnson and Johnson the airline model. Parameters were estimated using the exact maximum likelihood. The fitted model is:

$$(1 - L)(1 - L^4)x_t = (1 - 0.678L)(1 - 0.314L^4)\varepsilon_t, \quad \hat{\sigma}_\varepsilon = 0.089.$$

The standard errors for parameters θ and Θ are 0.080 and 0.101.

To check if the model adequate, i.e. if residuals $\hat{\varepsilon}_t$ of the estimated model are uncorrelated, we can use the Ljung-Box statistic. The Ljung-Box statistic of residuals show $Q(12) = 10.0$ with p -value 0.44. The model appear to be adequate.

Forecasting. To illustrate forecasting performance we fix 76 observations and reserve the last eight data point for forecasting evaluation.

We compute 1-step to 8-step forecast using fitted model at the forecast origin $t = 76$. Since the model is for the log earnings, we take the antilog transformation, to obtain forecast per share.

Figure 2.15 shows the forecast performance of the model:

- a) the observed data is shown by solid line
- b) forecast is shown by dots
- c) the dashed lines show the 95 percent confidence band.




Figure 2.15. Out-of-sample point and interval forecasts for the quarterly earning of Johnson and Johnson. The forecast origin is the fourth quarter of 1978. In the plot, the solid line shows the actual observations, dots represent point forecasts, and dashed lines show 95% interval forecasts.

Example (Deterministic seasonal behaviour).

This example demonstrates deterministic seasonal behaviour. We consider monthly simple returns of the CRSP Decile 1 index from January 1960 to December 2003, 528 observations.

Figure 2.16 shows the time plot of simple returns r_t which does not contain any clear pattern of seasonality. The sample ACF of the return series contains significant lags at 12, 24, 36 and lag 1.

This indicates that the seasonal behaviour might be deterministic.

To check if the seasonal component is deterministic and eliminate it we can use dummy variables. We define the dummy variable Jan_t for January, setting

$$Jan_t = 1 \text{ if } t \text{ is January, } Jan_t = 0 \text{ if } t \text{ is not January,}$$

and employ the simple regression model

$$r_t = \beta_0 + \beta_1 Jan_t + e_t.$$

The right panel of Figure 2.16 shows the time plot and sample ACF of the residuals series of simple regression model. We see that

- there is not significant correlation at any multiple of 12.
- ACF suggests that the seasonal pattern was successfully removed by the January dummy variable.
- we conclude that the seasonal behaviour in the monthly simple returns is due to January effect.

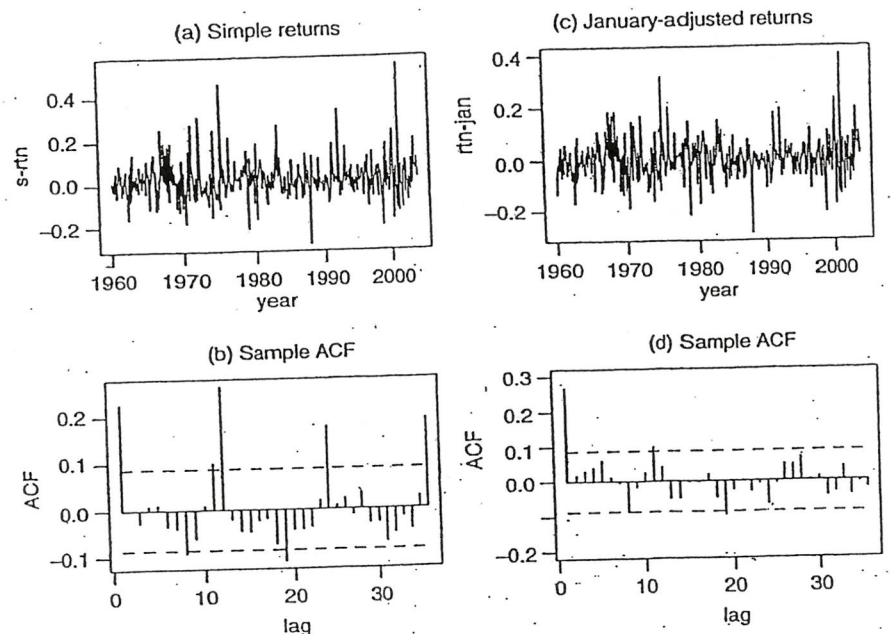


Figure 2.16. Monthly simple returns of CRSP Decile 1 index from January 1960 to December 2003: (a) time plot of the simple returns, (b) sample ACF of the simple returns, (c) time plot of the simple

SCA Demonstration
Output edited.

tsm ml. model (1)(12).dec1=c1+(12)noise.
--
estim ml. hold resi(r1)

SUMMARY FOR UNIVARIATE TIME SERIES MODEL -- M1

VAR.		ORIGINAL		DIFFERENCING				
VARIABLE OR CENTERED								
DEC1	RANDOM	ORIGINAL	NONE					
PAR.	VAR.	NUM./	FACTOR	ORDER	CONS-	VALUE	STD	T
LABEL	NAME	DENOM.			TRAIINT		ERROR	VALUE
1	C1	CNST	1	0	NONE	.0004	.0003	1.11
2	DEC1	MA	1	12	NONE	.9213	.0205	44.90
3	DEC1	AR	1	1	NONE	.2496	.0419	5.95
4	DEC1	AR	2	12	NONE	.9943	.0094	105.71

EFFECTIVE NUMBER OF OBSERVATIONS 515
R-SQUARE 0.207
RESIDUAL STANDARD ERROR. 0.705662E-01

estim ml. method exact. hold resi(r1)

SUMMARY FOR UNIVARIATE TIME SERIES MODEL -- M1

VAR.		ORIGINAL		DIFFERENCING				
VAR.		OR CENTERED						
DEC1	RANDOM	ORIGINAL	NONE					
PAR.	VARI.	NUM./	FACTOR	ORDER	CONS-	VALUE	STD	T
LABEL	NAME	DENOM.			TRAIINT		ERROR	VALUE
1	C1	CNST	1	0	NONE	.0002	.0002	.67
2	DEC1	MA	1	12	NONE	.9989	.0156	63.99
3	DEC1	AR	1	1	NONE	.2638	.0424	6.23
4	DEC1	AR	2	12	NONE	.9963	.0058	170.55

EFFECTIVE NUMBER OF OBSERVATIONS 515
R-SQUARE 0.283
RESIDUAL STANDARD ERROR. 0.670734E-01

8 Multivariate time series

Markets under economic globalization are interacting. Knowing how the markets are interrelated is of great importance in finance. In this section we discuss jointly multiple time series.

Multivariate time series observations at time t consists of k components:

$$\mathbf{r}_t = (r_{1t}, \dots, r_{kt})$$

observed in time t .

We use boldface to indicate vectors and matrices.

Example: Investor holds stocks of IBM, Microsoft, Exxon Mobil, General Motors and Wal-Mart. They form 5-dimensional daily return series:

- r_{1t} denotes daily log return of IBM stock;
- r_{2t} denotes daily log return of Microsoft, and so on.

Another example: investor interested in global investment, may consider

- return series of S&P index of the US,
- FTSE 100 index of UK,
- Nikkei 225 index of Japan.

The series is 3-dimensional.

We shall study now econometric model for analyzing the multivariate time series \mathbf{r}_t .

8.1 Stationarity

We say that k - dimensional time series

$$\mathbf{r}_t = (r_{1t}, \dots, r_{kt})'$$

weakly stationary if its first and second moments are time-invariant.

The mean vector consists of expectations of components:

$$\boldsymbol{\mu} = (\mu_1, \dots, \mu_k)' \equiv (E[r_{1t}], \dots, E[r_{kt}])'$$

The covariance matrix $\boldsymbol{\Gamma}_0 = [\Gamma_{ij}(0)]$ is $k \times k$ matrix, which (i, j) element is the covariance between components r_{it} and r_{jt} :

$$\Gamma_{ij}(0) = Cov(r_{it}, r_{jt}).$$

It can be written as

$$\boldsymbol{\Gamma}_0 = E[(\mathbf{r}_t - \boldsymbol{\mu})(\mathbf{r}_t - \boldsymbol{\mu})'].$$

8.1.1 Cross-Correlation Matrix

The cross-correlation matrix is an analog of correlation matrix. Denote by \mathbf{D} a $k \times k$ diagonal matrix consisting of standard deviations of r_{it} :

$$\mathbf{D} = \text{diag}(\sqrt{\Gamma_{11}(0)}, \dots, \sqrt{\Gamma_{kk}(0)}).$$

The Cross-correlation matrix of \mathbf{r}_t is defined as

$$\boldsymbol{\rho}_0 \equiv [\rho_{ij}(0)] = \mathbf{D}^{-1}\boldsymbol{\Gamma}_0\mathbf{D}^{-1}.$$

The (i, j) -th element of $\boldsymbol{\rho}_0$ is

$$\rho_{ij}(0) = \frac{\Gamma_{ij}(0)}{\sqrt{\Gamma_{ii}(0)\Gamma_{jj}(0)}} = \frac{Cov(r_{it}, r_{jt})}{sd(r_{it})sd(r_{jt})}.$$

Such correlation coefficient is called contemporaneous, because it is the correlation of the two series at time t .

It is easy to see that

- $\rho_{ii}(0) = 1$, $-1 \leq \rho_{ij}(0) \leq 1$, $\rho_{ij}(0) = \rho_{ji}(0)$,
- Matrix ρ_0 is symmetric.
- Cross-correlation measures the strength of linear dependence between time series.

The lag-k cross-covariance is defined as

$$\Gamma_k \equiv [\Gamma_{ij}(k)] = E[(\mathbf{r}_t - \boldsymbol{\mu})(\mathbf{r}_{t-k} - \boldsymbol{\mu})']$$

where

$$\Gamma_{ij}(k) = Cov(r_{it}, r_{j,t-k}).$$

The lag-k cross correlation is defined as

$$\boldsymbol{\rho}_k \equiv [\rho_{ij}(k)] = \mathbf{D}^{-1} \Gamma_k \mathbf{D}^{-1}$$

where

$$\rho_{ij}(k) = \frac{\Gamma_{ij}(k)}{\sqrt{\Gamma_{ii}(0)\Gamma_{jj}(0)}} = \frac{Cov(r_{it}, r_{j,t-k})}{srd(r_{it})srd(r_{jt})}.$$

For $k > 0$, the correlation coefficient $\rho_{ij}(k)$ measures the linear dependence between r_{it} and $r_{j,t-k}$ which occurred prior time t .

Note:

- Matrices Γ_k and $\boldsymbol{\rho}_k$ are symmetric
- Under weak stationarity Γ_k and $\boldsymbol{\rho}_k$ do not depend on time t , only on the lag k .

8.2 Sample Cross-Correlation Matrices

Question: How to estimate cross-covariance matrix Γ_k and cross-correlation matrix ρ_k from the data

$$\mathbf{r}_t, \quad t = 1, \dots, N?$$

Γ_k can be estimated by

$$\hat{\Gamma}_k = \frac{1}{N} \sum_{t=k+1}^N (\mathbf{r}_t - \bar{\mathbf{r}})(\mathbf{r}_{t-k} - \bar{\mathbf{r}})'$$

where

$$\bar{\mathbf{r}} = \frac{1}{N} \sum_{t=1}^N \mathbf{r}_t$$

is the vector of sample mean, which is a consistent estimator for $\boldsymbol{\mu}$.

ρ_k can be estimated by

$$\hat{\rho}_k = \hat{\mathbf{D}}^{-1} \hat{\Gamma}_k \hat{\mathbf{D}}^{-1}, \quad k \geq 0,$$

where $\hat{\mathbf{D}}$ is the $k \times k$ diagonal matrix of the sample standard deviations of the components.

Example. Consider the monthly log returns of IBM stock and the S&P 500 index from Jan 1926 to Dec 1999 with 888 observations. Returns include dividend payments and are in percentages. We denote these two series by r_{1t} and r_{2t} . It has the form of a bivariate series

$$\mathbf{r}_t = (r_{1t}, r_{2t})'$$

Figure 8.1 shows time plots of components of \mathbf{r}_t . Figure 8.2 shows some scatterplots of the two series. They show that series are correlated. Indeed:

a) the correlation coefficient at lag 0 between series is 0.64, which are statistically significant at 5% level.

b) however, the cross correlations at lag 1 are weak.

Table 8.1 contains summary statistics and cross correlation matrices. For bivariate series, cross correlation matrices are 2×2 with four correlations.

Table 8.1 (c) simplifies CCM results using Tiao-Box notations:

- "=" means that correlation coefficient is greater or equal to $2/\sqrt{N}$ (significant at 5% level)
- "-" means that correlation coefficient is less or equal to $-2/\sqrt{N}$ (significant at 5% level)
- "." means that correlation coefficient is between $-2/\sqrt{N}$ and $2/\sqrt{N}$ (not significant at 5% level)

Table 8.1 (c) shows that significant correlations appear mostly at lag 1 and 3.

Conclusions: Analysis of CCM (cross-correlation matrices) shows that

a) S&P 500 index returns has some marginal autocorrelation at lag 1 and 3;

b) IBM stock returns depend weakly on the previous returns of the S&P 500 index, since cross-correlations at lag 1 and 3 are significant.

Figure 8.3 contains the sample autocorrelations and the cross-correlation of the two series. Dynamic relationship between two series is weak, but their contemporaneous correlation at lag 0 is statistically significant.

Table 8.1. Summary Statistics and Cross-Correlation Matrices of Monthly Log Returns of IBM Stock and the S&P 500 Index: January 1926 to December 1999

<i>(a) Summary Statistics</i>						
Ticker	Mean	Standard Error	Skewness	Excess Kurtosis	Minimum	Maximum
IBM	1.240	6.729	-0.237	1.917	-30.37	30.10
S&P 500	0.537	5.645	-0.521	8.117	-35.58	35.22

<i>(b) Cross-Correlation Matrices</i>									
Lag 1		Lag 2		Lag 3		Lag 4		Lag 5	
0.08	0.10	0.02	-0.06	-0.02	-0.07	-0.02	-0.03	0.00	0.07
0.04	0.08	0.02	-0.02	-0.07	-0.11	0.04	0.02	0.00	0.08

<i>(c) Simplified Notation</i>				
$\begin{bmatrix} + & + \\ \cdot & + \end{bmatrix}$	$\begin{bmatrix} \cdot & \cdot \\ \cdot & \cdot \end{bmatrix}$	$\begin{bmatrix} \cdot & - \\ - & - \end{bmatrix}$	$\begin{bmatrix} \cdot & \cdot \\ \cdot & \cdot \end{bmatrix}$	$\begin{bmatrix} \cdot & \cdot \\ \cdot & + \end{bmatrix}$

Table 8.2. Sample Cross-Correlation Matrices of Monthly Simple Returns of Five Indexes of U.S. Government Bonds: January 1942 to December 1999

Lag 1					Lag 2				
<i>Cross-Correlations</i>									
0.10	0.08	0.11	0.12	0.16	-0.01	0.00	0.00	-0.03	0.03
0.10	0.08	0.12	0.14	0.17	-0.01	0.00	0.00	-0.04	0.02
0.09	0.08	0.09	0.13	0.18	0.01	0.01	0.01	-0.02	0.07
0.14	0.12	0.15	0.14	0.22	-0.02	-0.01	0.00	-0.04	0.07
0.17	0.15	0.21	0.22	0.40	-0.02	0.00	0.02	0.02	0.22

<i>Simplified Cross-Correlation Matrices</i>									
$\begin{bmatrix} + & + & + & + & + \\ + & + & + & + & + \\ + & + & + & + & + \\ + & + & + & + & + \\ + & + & + & + & + \end{bmatrix}$	$\begin{bmatrix} \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & + \end{bmatrix}$								




Figure 8.1. Time plots of (a) monthly log returns in percentages for IBM stock and (b) the S&P 500 index from January 1926 to December 1999.

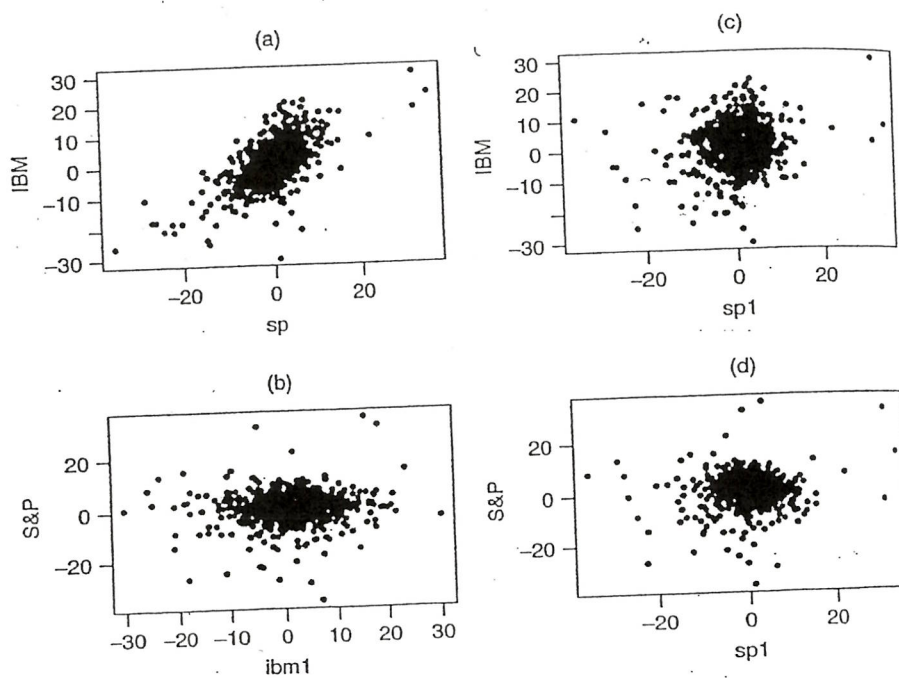


Figure 8.2. Some scatterplots for monthly log returns of IBM stock and the S&P 500 index: (a) concurrent plot of IBM versus S&P 500 (b) S&P 500 versus lag-1 IBM, (c) IBM versus lag-1 S&P 500, and (d) S&P 500 versus lag-1 S&P 500.

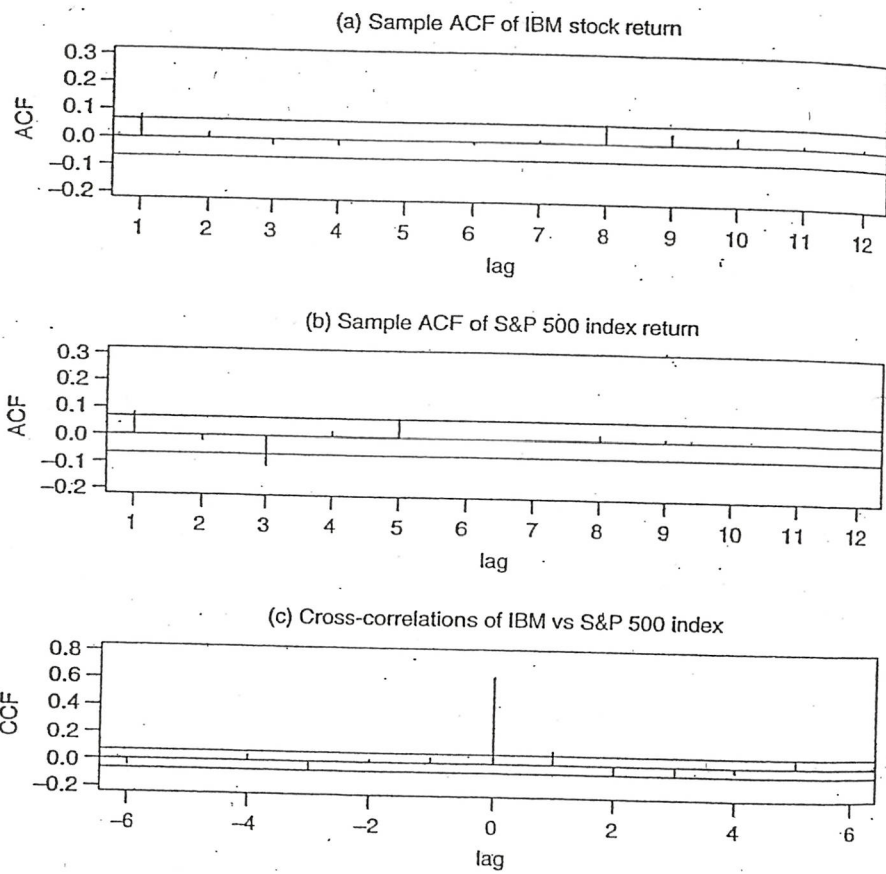


Figure 8.3. Sample auto- and cross-correlation functions of two monthly log returns: (a) sample ACF of IBM stock returns, (b) sample ACF of S&P 500 index returns, and (c) cross-correlations between IBM stock return and lagged S&P 500 index returns.

Example. Consider the simple returns of monthly indexes of US government bonds with maturities in 30, 20, 10, 5 and 1 year, from Jan 1942 to Dec 1999, 696 observations. Let $\mathbf{r}_t = (r_{1t}, \dots, r_{5t})$ be the returns series with decreasing time to maturity.

Figure 8.4 shows time plots of these time series. We see that the variability of 1-year bond is much smaller than that of returns with longer maturity.

The sample mean is

$$\hat{\boldsymbol{\mu}} = 10^{-2}(0.43, 0.45, 0.45, 0.46, 0.44)'$$

The sample standard deviations are

$$\hat{\boldsymbol{\sigma}} = 10^{-2}(2.53, 2.43, 1.97, 1.39, 0.53)'$$

The correlation matrix $\hat{\boldsymbol{\rho}}_0$ at lag 0 below shows that

- a) the series have high correlations,
- b) correlations between long-term bonds are higher than those between short term bonds.

Table 8.2 gives the lag-1 and lag-2 cross-correlation matrices of \mathbf{r}_t , and their simplified version. We see that

- 1) most significant correlations are at lag-1
- 2) at lag 1 the five returns series appear to be uncorrelated.
- 3) at lag 1 and lag-2, sample ACF of one-year bond returns are substantially higher than those of other series with longer maturities.

$$\hat{\boldsymbol{\rho}}_0 = \begin{bmatrix} 1.00 & 0.98 & 0.92 & 0.85 & 0.63 \\ 0.98 & 1.00 & 0.91 & 0.86 & 0.64 \\ 0.92 & 0.91 & 1.00 & 0.90 & 0.68 \\ 0.85 & 0.86 & 0.90 & 1.00 & 0.82 \\ 0.63 & 0.64 & 0.68 & 0.82 & 1.00 \end{bmatrix}$$

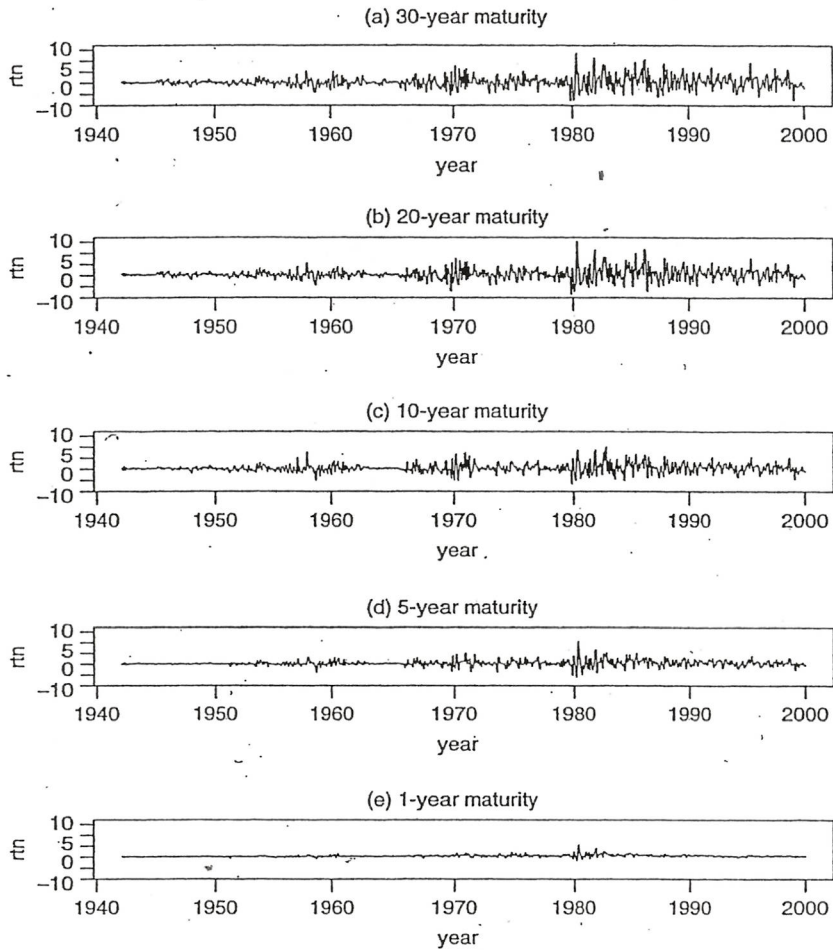


Figure 8.4. Time plots of monthly simple returns of five indexes of U.S. government bonds with maturities in (a) 30 years, (b) 20 years, (c) 10 years, (d) 5 years, and (e) 1 year. The sample period is from January 1942 to December 1999.

8.3 Multivariate Portmanteau test

This test generalizes univariate Ljung-Box statistics $Q(m)$ to multivariate case. The test is used to test that there are no auto-and-cross- correlations in the vector series \mathbf{r}_t .

The null hypothesis is

$$H_0: \boldsymbol{\rho}_1 = \dots = \boldsymbol{\rho}_m = 0.$$

The alternative is

$$H_a: \boldsymbol{\rho}_i \neq 0 \text{ for some } i = 1, \dots, m.$$

Decision is based on the test statistic $Q_k(m)$. We reject H_0 at 5% significance level if the p -value is less than 0.05.

Example. Applying the $Q_k(m)$ statistics to the bivariate monthly log returns of IBM stock and S&P 500 index of the previous example, we have

$$Q_2(1) = 9.81, \quad Q_2(5) = 47.06, \quad Q_2(10) = 71.64.$$

The p -values are 0.044, 0.0001, and 0.002, respectively. The test confirms the existence of serial correlation in the bivariate series at 5% significance level.

For 5-dimensional monthly simple returns of bond index in the second example above, we have $Q_5(5) = 1065.63$ which is highly significant.

The $Q_k(m)$ statistic is a joint test for checking the first m cross-correlation matrices of \mathbf{r}_t . If it rejects the null hypothesis, then we build a multivariate model for the time series to study the relationships between the component series.

8.4 Vector autoregressive models (VAR)

A simple vector model is the vector autoregressive model (VAR).

Var(1) model.

A multivariate time series \mathbf{r}_t follows a VAR model of order 1, if

$$\mathbf{r}_t = \boldsymbol{\phi}_0 + \boldsymbol{\Phi}\mathbf{r}_{t-1} + \mathbf{a}_t$$

where $\boldsymbol{\phi}_0$ is a k -dimensional vector, $\boldsymbol{\Phi}$ is a $k \times k$ matrix, and \mathbf{a}_t is a sequence of serially uncorrelated random vectors with zero mean and covariance matrix $\boldsymbol{\Sigma}$. In applications it is often assumed that \mathbf{a}_t is multivariate normal.

Consider the bivariate case: $k = 2$, $\mathbf{r}_t = (r_{1t}, r_{2t})'$, $\mathbf{a}_t = (a_{1t}, a_{2t})'$. The VAR(1) model consists of two equations:

$$r_{1t} = \phi_{10} + \Phi_{11}r_{1,t-1} + \Phi_{12}r_{2,t-1} + a_{1t},$$

$$r_{2t} = \phi_{20} + \Phi_{21}r_{1,t-1} + \Phi_{22}r_{2,t-1} + a_{2t},$$

where Φ_{ij} is the (i, j) th element of $\boldsymbol{\Phi}$, and ϕ_{i0} is the i th element of $\boldsymbol{\phi}_0$.

Comments:

- In first equation, Φ_{12} denotes the linear relationship between r_{1t} and $r_{2,t-1}$ in the presence of $r_{1,t-1}$.
- If $\Phi_{12} = 0$ then r_{1t} does not depend on $r_{2,t-1}$, and depends only on its own past.
- If $\Phi_{21} = 0$ then r_{2t} does not depend on $r_{1,t-1}$, when $r_{2,t-1}$ is given