

Chapter 3

Univariate Time Series: Linear Models

3.1 Introduction

Time series is a sequence of numerical data in which observations are measured at a particular instant of time. The frequency of observation can, for example, be annual, quarterly, monthly, daily, etc. The main goal of time series analysis is to study the dynamics of the data.

In this chapter we introduce basic time series models for estimation and forecasting of financial data. Further details about theory of time series analysis can be found in Hamilton (1994), Greene (2000), Enders (2004), Tsay (2002) and others.

3.2 Stationarity and Autocorrelations

3.2.1 Stationarity

A time series $\{Y_t\}$ is said to be *strictly stationary* if for all integers i, j and all possible integers k the multivariate distribution function of $(Y_i, Y_{i+1}, \dots, Y_{i+k-1})$ is identical to $(Y_j, Y_{j+1}, \dots, Y_{j+k-1})$. In practice we are very often interested in consequences of this assumption regarding moments of the distribution. If Y_i and Y_j have identical distribution this implies that their means are identical, thus $E[Y_t]$ does not depend on time and equal to some constant μ . Also, because the pairs (Y_i, Y_{i+s}) and (Y_j, Y_{j+s}) have identical bivariate distributions it follows that the autocovariances

$$\text{cov}(Y_t, Y_{t+s}) = E[(Y_t - \mu)(Y_{t+s} - \mu)] = \lambda_s$$

depend only on the time lag s . This implies also that Y_t have constant variance $\lambda_0 = \sigma^2$.

A stochastic process whose first and second order moments (means, variances, and covariances) do not change with time is said to be second order stationary. More precisely, a time series Y_t is called stationary if the following conditions are satisfied:

$$E[Y_t] = \mu, \quad E[(Y_t - \mu)^2] = \gamma_0, \quad E[(Y_t - \mu)(Y_{t-s} - \mu)] = \gamma_s \quad \text{for all } t$$

Here μ , γ_0 , and γ_k are finite-valued numbers that do not depend on time t .

3.2.2 Autocorrelation

The autocorrelations of a stationary process are defined by $\rho_s = \frac{\gamma_s}{\gamma_0}$. These correlations describe the short-run dynamic relations within the time series, in contrast with the trend, which corresponds to the long-run behaviour of the time series.

The simplest possible autocorrelations occur when a stationary process consists of uncorrelated random variables. In this case $\rho_0 = 1$, $\rho_s = 0$ for all $s > 0$. Such time series is called *white noise*.

It is important when modeling financial returns to appreciate that if $\{Y_t\}$ is white noise then Y_t and Y_{t+s} are not necessarily independent for $s > 0$.

The *partial autocorrelation* ϕ_s at lag s measures the correlation of Y_t values that are s periods apart after removing the correlation from the intervening lags. It equals the regression coefficient on Y_{t-s} when Y_t is regressed on a constant, Y_{t-1}, \dots, Y_{t-s} .

Time series prediction To describe the correlations, we imagine that our observed time series comes from a stationary process that existed before we started observing it. We denote the past of the stationary process Y_t by $\mathcal{Y}_{t-1} = \{Y_{t-1}, Y_{t-1}, \dots\}$, where the "dots" mean that there is no clear-cut beginning of this past. We call it also the information set available at time point $t - 1$. The least squares predictor of Y_t based on the past \mathcal{Y}_{t-1} is the function $f(\mathcal{Y}_{t-1})$ that minimizes $E[(Y_t - f(\mathcal{Y}_{t-1}))^2]$. This predictor is given by the conditional mean $f(\mathcal{Y}_{t-1}) = E[Y_t | \mathcal{Y}_{t-1}]$ with corresponding (one-step-ahead) prediction errors $e_t = Y_t - f(\mathcal{Y}_{t-1}) = Y_t - E[Y_t | \mathcal{Y}_{t-1}]$.

The process e_t is also called the innovation process, as it corresponds to the unpredictable movements in Y_t . If the observations are jointly normally distributed, then the conditional mean is a linear function of the past observations

$$E[Y_t | \mathcal{Y}_{t-1}] = a + p_1 Y_{t-1} + p_2 Y_{t-2} + \dots$$

Here a models the mean $E[Y_t] = \mu$ of the series. From the above equation we get $\mu = a + \sum p_k \mu$, so that $\mu = (1 - \sum p_k)^{-1} a$. As the process is assumed to be stationary, the coefficients p_k do not depend on time and the innovation process e_t is also stationary. It has the following properties:

- $E[e_t] = 0$ for all t
- $E[e_t^2] = \sigma^2$ for all t ;
- $E[e_s e_t] = 0$ for all $s \neq t$.

Here the variance σ^2 is constant over time.

3.2.3 Example: Variance Ratio Test

Very often a predictability of stock returns is linked to the presence of autocorrelation in the returns series. If stock returns form an iid process, then variances of holding period returns should increase in proportion to the length of the holding period. If the log return is constant, then under the rational expectation hypothesis stock prices follows a random walk

$$p_{t+h} = \mu + p_{t+h-1} + u_{t+h} = \mu + \mu + p_{t+h-2} + u_{t+h} + u_{t+h-1} = p_t + \mu h + \sum_{i=0}^{h-1} u_{t+i}.$$



360°
thinking.

Deloitte.

Discover the truth at www.deloitte.ca/careers

© Deloitte & Touche LLP and affiliated entities.

Variance of the returns forecasts

$$\text{var} [p_{t+h} - p_t] = \sum_{i=0}^h E [u_{t+i}^2] = h\sigma^2$$

due to the independence. Alternatively, if log returns are iid, then

$$\text{var} [r_{t,t+h}] = \text{var} [r_{t,t+1} + r_{t+1,t+2} + \dots + r_{t+h-1,t+h}] = h\text{var} [r_{t+1}]$$

The variance-ratio statistic is defined as

$$VR_h = \frac{1}{h} \frac{\text{var} [r_{t,t+h}]}{\text{var} [r_{t+1}]} = 1 + \frac{2}{h} \sum_{j=1}^{h-1} (h-j)\rho_j,$$

which should be unity if returns are iid and less than unity under mean reversion.

The variance ratio test is set up as $H_0: VR_h = 1$ and under the null

$$Z_h = \frac{VR_h - 1}{\sqrt{2(2h-1)(h-1)/3hT}} \sim N(0,1).$$

See Cuthbertson and Nitzsche (2004) for more details about the test. Let us consider as an example how to program the variance ratio test in EViews.

In this test uses overlapping h -period returns. As an input to the program, the workfile should contain a series of log prices p used to test for predictability. We start the program in a usual way.

```
smpl @all
```

```
!h=2
```

The variable `!h` denotes the horizon of the returns forecast. The next we create one period and h period returns.

```
smpl @first+1 @last
```

```
series r=p-p(-1)
```

In order to build the variance ratio statistics we need to have the actual number of observations (returns), mean and variance of returns series.

```
scalar T=@obs(p)
```

```
scalar mu=@mean(r)
```

```
scalar var1=@sumsq(r-mu)/(T-1)
```

```
smpl @first+!h @last
```

```
series rh=p-p(-!h)
```

```
scalar varh=@sumsq(rq-!h*mu)/(T-!h+1)
```

We can now compute the variance ratio statistic

```
scalar VRh=varh/(!h*var1)
```

```
scalar Zh=(VRh-1)/@sqrt((2*(2*!q-1)*(!q-1))/(3*!q*T))
```

We need a p-value in order to test the hypothesis. Two-sided significance level (p-value) can be calculated as follows

```
scalar Zh_level=2*(1-@cnorm(@abs(Zh)))
```

Finally, we create a table to report the results. We declare a new table VRTEST object with 2 rows and 5 columns, set the width of each column and write the context of each cell down.

```
table(2,5) VRTEST
Setcolwidth(VRTEST,1,15)
Setcolwidth(VRTEST,2,15)
Setcolwidth(VRTEST,3,10)
Setcolwidth(VRTEST,4,10)
Setcolwidth(VRTEST,5,13)
Setcell(VRTEST,1,1,"Nr of obs")
Setcell(VRTEST,1,2,"Horizon h")
Setcell(VRTEST,1,3,"VRh")
Setcell(VRTEST,1,4,"test stat Zh")
Setcell(VRTEST,1,5,"p-value")
Setcell(VRTEST,2,1,T,0)
Setcell(VRTEST,2,2,!h,0)
Setcell(VRTEST,2,3,VRh,4)
Setcell(VRTEST,2,4,Zh,4)
Setcell(VRTEST,2,5,Zh_level,5)
delete r mu rh T var1 varh Zh Zh_level
next
```

3.3 ARMA processes

A zero mean white noise process $\{\varepsilon_t\}$ can be used to construct new processes. We describe two commonly used examples first and afterwards their generalization – autoregressive-moving average (ARMA) model.

3.3.1 Autoregressive process

A simple way to model dependence between consecutive observations is

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \varepsilon_t,$$

where ε_t is white noise. Such process is called a *first-order autoregressive process* or *AR(1)* process. It is stationary if the coefficient $|\alpha_1| < 1$.

Since $E[\varepsilon_t] = 0$ it follows that under the stationarity condition the mean of the process $E[Y_t] = \frac{\alpha_0}{1-\alpha_1}$ and variance $\text{var}[Y_t] = \frac{\sigma_\varepsilon^2}{1-\alpha_1^2}$ where $\sigma_\varepsilon^2 = \text{var}[\varepsilon_t]$. An *AR(1)* process has autocorrelations $\rho_s = \alpha_1^s$ for $s > 1$.

A more general representation of the autoregressive process is

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \dots + \alpha_p Y_{t-p} + \varepsilon_t$$

and called an autoregressive process of order p , or in short, *AR(p)*.

3.3.2 Moving average process

Consider the process $\{Y_t\}$ defined by

$$Y_t = \alpha_0 + \varepsilon_t + \beta_1 \varepsilon_{t-1}$$

so Y_t is a linear function of the present and immediately preceding innovations. This process is called a moving average process of order 1 and denoted by $MA(1)$.

A $MA(1)$ process is always stationary with mean α_0 and variance $(1 + \beta_1^2) \sigma_\varepsilon^2$. Its autocorrelations are $\rho_1 = \frac{\beta_1}{1 + \beta_1^2}$ and $\rho_s = 0$ for $s > 1$.

Comparing two time series we see that a shock ε_t in $MA(1)$ process affects Y_t in two periods (only two positive autocorrelation coefficients), while a shock in the $AR(1)$ process affects all future observations with a decreasing effect.

The $MA(1)$ process may be inverted to give ε_t as an infinite series in Y_t, Y_{t-1}, \dots , namely

$$\varepsilon_t = Y_t + \beta_1 Y_{t-1} + \beta_1^2 Y_{t-2} + \dots$$

that is

$$Y_t = -\beta_1 Y_{t-1} - \beta_1^2 Y_{t-2} - \dots + \varepsilon_t.$$

Thus, $MA(1)$ time series can be represented as $AR(\infty)$ process. It is possible to invert $MA(1)$ process into a stationary AR process only if $|\beta_1| < 1$. This condition is known as *invertibility* condition.

A more general representation of a moving average process is

$$Y_t = \alpha_0 + \varepsilon_t + \beta_1 \varepsilon_{t-1} + \dots + \beta_q \varepsilon_{t-q}$$

and called a moving average process of order q , or in short, $MA(q)$.

3.3.3 ARMA process

It is possible to combine the autoregressive and moving average specification into $ARMA(p, q)$ model

$$Y_t = \alpha_1 Y_{t-1} + \dots + \alpha_p Y_{t-p} + \varepsilon_t + \beta_1 \varepsilon_{t-1} + \dots + \beta_q \varepsilon_{t-q}. \quad (3.3.1)$$

An $ARMA(p, q)$ time series can be represented in a shorter form using the notion of lag operator.

The lag operator L , is defined as $LY_t = Y_{t-1}$, the operator which gives the previous value of the series. This operator can also be used to represent the lags of the second or higher orders in the following way:

$$L^2(Y_t) = L(L(Y_t)) = L(Y_{t-1}) = Y_{t-2}.$$

In general $ARMA(p, q)$ process is

$$A(L)Y_t = B(L)\varepsilon_t,$$

where

$$A(L) = 1 - \alpha_1 L - \alpha_2 L^2 - \dots - \alpha_p L^p$$

$$B(L) = 1 - \beta_1 L - \beta_2 L^2 - \dots - \beta_q L^q$$

Stationarity requires the roots of $A(L)$ to lie outside the unit circle, and invertibility places the same condition on the roots of $B(L)$.

Table 3.1: Correlation patterns

Time series	acf	pacf
$AR(p)$	Infinite: decays towards zero	Finite: disappears after lag p
$MA(q)$	Finite: disappears after lag q	Infinite: decays towards zero
$ARMA(p, q)$	Infinite: damps out	Infinite: decays towards zero

3.3.4 Estimation of ARMA processes

$ARMA(p, q)$ models are generally estimated using the technique of maximum likelihood.

An often ignored aspect of the maximum likelihood estimation of $ARMA(p, q)$ models is the treatment of initial values. These initial values are the first p values of Y_t and q values of ε_t in (3.3.1). The exact likelihood utilizes the stationary distribution of the initial values in the construction of the likelihood. The conditional likelihood treats the p initial values of Y_t as fixed and often sets the q initial values of ε_t to zero. The exact maximum likelihood estimates (MLE) maximize the exact log-likelihood, and the conditional MLE maximize the conditional log-likelihood. The exact and conditional MLEs are asymptotically equivalent but can differ substantially in small samples, especially for models that are close to being non-stationary or non-invertible.

For pure AR models, the conditional MLEs are equivalent to the least squares estimates

Model Selection Criteria Before an $ARMA(p, q)$ may be estimated for a time series Y_t , the AR and MA orders p and q must be determined by visually inspecting the autocorrelation and partial autocorrelation functions for Y_t . If the autocorrelation function decays smoothly and the partial autocorrelations are zero after one lag, then a first-order autoregressive model is appropriate. Alternatively,

if the autocorrelations were zero after one lag and the partial autocorrelations decay slowly towards zero, a first-order moving average process would seem appropriate.

Alternatively, statistical model selection criteria may be used. The idea is to fit all $ARMA(p, q)$ models with orders p and q and choose the values of p and q which minimizes model selection criteria:

$$AIC(p, q) = \ln(\tilde{\sigma}^2(p, q)) + \frac{2}{T}(p + q)$$

$$BIC(p, q) = \ln(\tilde{\sigma}^2(p, q)) + \frac{\ln(T)}{T}(p + q)$$

where $\tilde{\sigma}^2(p, q)$ is the MLE of $\text{var}[\varepsilon_t] = \sigma^2$ without a degrees of freedom correction from the $ARMA(p, q)$ model.

3.3.5 Example: ARMA in EViews

We start our example from the simulation of ARMA process and then we take a look at its estimation. In order to illustrate the statements in Table 3.1, let us simulate $AR(3)$, $MA(2)$ and $ARMA(3, 2)$ processes and compute their autocorrelation and partial autocorrelation functions.

SIMPLY CLEVER

ŠKODA


We will turn your CV into
an opportunity of a lifetime



Do you like cars? Would you like to be a part of a successful brand?
 We will appreciate and reward both your enthusiasm and talent.
 Send us your CV. You will be surprised where it can take you.

Send us your CV on
www.employerforlife.com

In particular, we simulate

$$\begin{aligned} Y_t &= 0.8Y_{t-1} + 0.15Y_{t-2} - 0.1Y_{t-3} + u_t \\ Y_t &= u_t - 0.95u_{t-1} + 0.3u_{t-2} \\ Y_t &= 0.8Y_{t-1} + 0.15Y_{t-2} - 0.1Y_{t-3} + u_t - 0.95u_{t-1} + 0.3u_{t-2} \end{aligned} \quad (3.3.2)$$

To start with, we generate a series of uncorrelated normally distributed residuals (remember, command `nrnd` generates standard normally distributed random number)

```
series u=0.5*nrnd
```

Also, we have to generate initial values for the series. Since the highest order of the series is 3, let us generate first three values. This can be done by setting sample to only first three observations and assign zero values to all of three series.

```
smpl @first @first+2
```

```
series y1=0
```

```
series y2=0
```

```
series y3=0
```

Now, we set the sample for the rest of observations and generate series according to formulae (3.3.2)

```
smpl @first+3 @last
```

```
y1=0.8*y1(-1)+0.15*y1(-2)-0.1*y1(-3)+u
```

```
y2=u-0.95*u(-1)+0.3*u(-2)
```

```
y3=0.8*y1(-1)+0.15*y1(-2)-0.1*y1(-3)+u-0.95*u(-1)+0.3*u(-2)
```

Now, we are ready to build and inspect their correlograms. Remind, that in order to build a correlogram, one should click on the icon if the time series being investigated and choose **View/Correlogram...** option. The correlograms of three time series is given on Figures ??-??.

As we have expected, the autocorrelation function for the first series ($AR(3)$) damps out slowly towards zero while its partial autocorrelation function has spikes at first three lags. The autocorrelation function of the second series ($MA(2)$) has spikes at two first lags and disappears afterwards (becomes insignificant) while the partial autocorrelation function decays oscillating towards zero. Both autocorrelation and partial autocorrelation functions of the third series ($ARMA(3,2)$) decay slowly towards zero without any clear spikes.

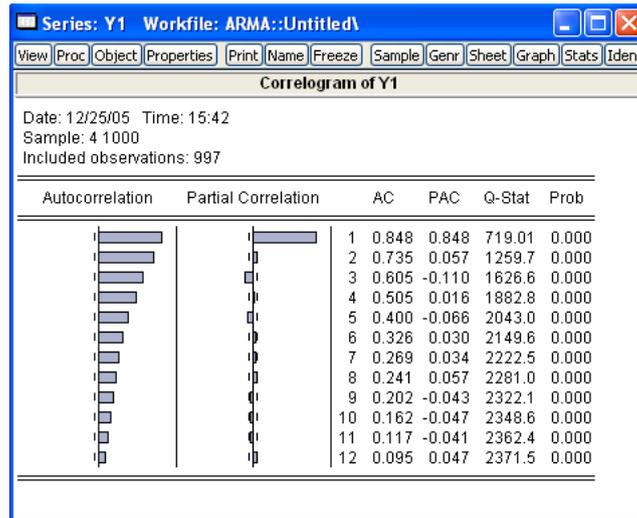


Figure 3.1: Correlogram of an $AR(3)$ process

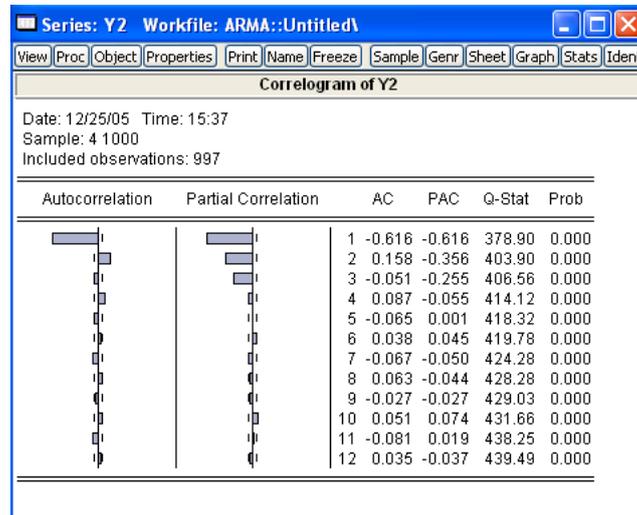


Figure 3.2: Correlogram of a $MA(2)$ process

Estimation An estimation of the ARMA processes is performed in EViews in the same way as OLS estimation of a linear regression. The only difference is in specifying autoregressive and moving average terms in the model. If the series has got autoregressive components, we should include terms $ar(1)$, $ar(2)$, etc, as regressors up to the required order. For example, to estimate the first series, type

$$y1\ c\ ar(1)\ ar(2)\ ar(3)$$

in the estimation equation box. EViews produces an output given in Figure ??

All coefficients are significant as expected and are very close to the true values.

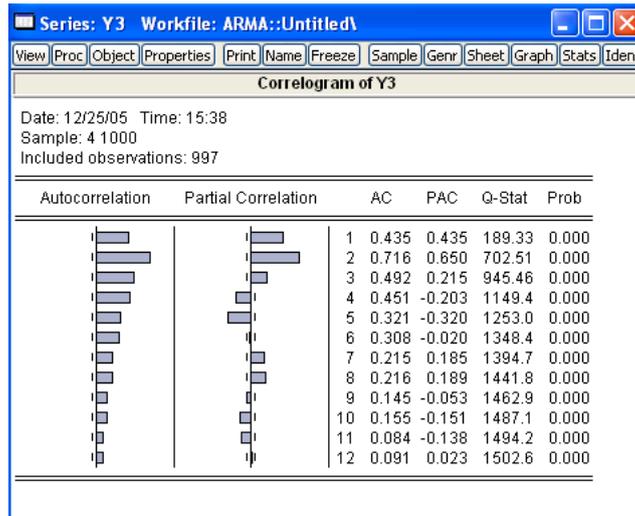


Figure 3.3: Correlogram of an $ARMA(3, 2)$ process

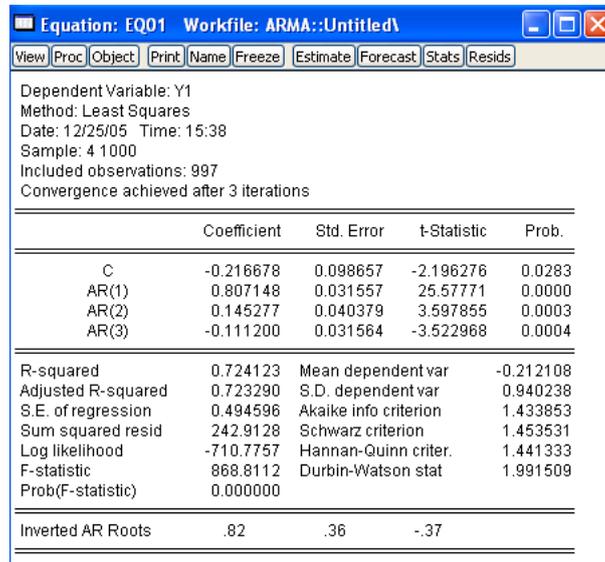


Figure 3.4: Estimation output of $ARMA$ process

Inference and tests can be performed in the same way as it was done for the OLS regression.

If one needs to estimate the model containing moving average components, $ma(1)$, $ma(2)$, etc terms should be included into the model specification. For example, to estimate the second time series, we write

$$y_2 = c + ma(1) + ma(2)$$

Autoregressive and moving average terms can be combined to estimate ARMA

model. Thus, specification of the third series looks like

```
y3 c ar(1) ar(2) ar(3) ma(1) ma(2)
```

After having estimated an ARMA model, one can check whether the estimated coefficients satisfy the stationarity assumptions. This can be done through **View/ARMA structure** of the **Equation** object. For the third series we obtain

AR Root(s)	Modulus	Cycle
0.815744	0.815744	
-0.101284 ± 0.082792i	0.130817	2.557972

No root lies outside the unit circle.
ARMA model is stationary.

MA Root(s)	Modulus	Cycle
0.370754 ± 0.922562i	0.994273	5.285871

No root lies outside the unit circle.
ARMA model is invertible.

Figure 3.5: Table of the roots of the estimated *ARMA* process

It says that our ARMA series is both stationary and invertible.

3.3.6 Programming example

If we had not known the order of the ARMA series, we would need to apply one of the information criteria to select the most appropriate order of the series. The following program illustrates how this can be done using the Akaike criterion.

First we need to define the maximal orders for autoregressive and moving average parts and store them into variables `pmax` and `qmax`. Also we need to declare a matrix object `aic` where the values of the Akaike statistic will be written for each specification of the ARMA process.

```
smpl @all
scalar pmax=3
scalar qmax=3
matrix(pmax+1,qmax+1) aic
```

Next, we define nested loops which will run through all possible ARMA specification with orders within the maximal values.

```
for !p=0 to pmax
```

```
  for !q=0 to qmax
```

As the number of lags included in the model increases we add a new AR term in the model. For this purpose we create a new string variable `textsf%order` containing the model specification.

```
    if !p=0 then %order=""
```

```
    else
```

```
      for !i=1 to !p
```

```
        %order=%order+" ar("+@str(!i)+")"
```

```
      next
```

```
    endif
```

I joined MITAS because
I wanted **real responsibility**

The Graduate Programme
for Engineers and Geoscientists
www.discovermitas.com



Month 16

I was a construction supervisor in the North Sea advising and helping foremen solve problems

Real work
International opportunities
Three work placements





We perform the same procedure with the MA term specification.

```

if !q=0 then %order=%order+" "
else
for !i=1 to !q
%order=%order+" ma("+@str(!i)+" )"
next
endif

```

Once the model specification is determined and written in the variable `%order` we can use a substitution to estimate the corresponding model.

```

equation e.ls y3 c %order
%order=""

```

The last command nullify the variable `%order` for the use in the next step of the loops. Now we can write the value of the Akaike criterion for the current in the table.

```

aic(!p+1,!q+1)=e.@aic
next
next
delete e

```

After the program run, the values of the Akaike criterion are stored in the table `aic`. Now we can choose that specification of the ARMA model which produces the smallest AIC value.