


## Solution to Series 7


1. a) In some areas the variance is much smaller than in others. The “peak” in the middle indicates that a logarithmic transformation must first be applied to the data. If we look at the correlogram, we notice that the ordinary autocorrelations decay far too slowly. Even for large lags, they still lie outside the confidence band:

```
> plot(d.varve)
> acf(d.varve)
> pacf(d.varve)
```




Even the time series of logarithmic data cannot yet be regarded as stationary, since it exhibits clear trends (first increasing, then decreasing), which can however be eliminated by taking first differences:

```
> plot(log(d.varve))
> acf(log(d.varve))
> pacf(log(d.varve))
```



The plot of first differences for the transformed series shows that stationarity can now be assumed:

```
> plot(diff(log(d.varve)))
> acf(diff(log(d.varve)))
> pacf(diff(log(d.varve)))
```



- b) The correlogram plotted in part a) indicates an ARIMA(1,1,1) process (or perhaps an ARIMA(0,1,1) process). Fitting these two models, we see that the ARIMA(1,1,1) model is very good at describing the logarithmic data. In both fitted models, the algorithm converges; of the two models, ARIMA(1,1,1) has a smaller AIC.

The estimated coefficients are  $\hat{\beta}_1 = -0.84$  for the fitted ARIMA(0,1,1) model and  $\hat{\alpha}_1 = 0.25$ ,  $\hat{\beta}_1 = -0.91$  for the fitted ARIMA(1,1,1) model. For both models, the estimated mean is  $\hat{\mu} = -0.00127$ , which leads us to assume the data do not need correcting by their mean. Furthermore, the estimated error variances are 0.224 (for the ARIMA(0,1,1) model) and 0.2138 (for the ARIMA(1,1,1) model).

Thus the ARIMA(0,1,1) model looks as follows:

$$Y_t = X_t - X_{t-1}$$


$$Y_t = E_t - 0.84 \cdot E_{t-1}; \quad \sigma_{E_t}^2 = 0.224$$

For the ARIMA(1,1,1) model, we similarly have

$$Y_t = X_t - X_{t-1}$$


$$Y_t = 0.25 \cdot Y_{t-1} + E_t - 0.91 \cdot E_{t-1}; \quad \sigma_{E_t}^2 = 0.214$$

**Residuals of the fitted ARIMA(0,1,1) process:**



The first ordinary (=first partial) autocorrelation clearly lies outside the confidence band. Thus the residuals cannot be considered independent.

### Residuals of the fitted ARIMA(1,1,1) process:



These residuals no longer exhibit any undesired structure.

#### R commands:

##### ARIMA(0,1,1) model:

```
> mean(diff(log(d.varve))) # -0.001271813
> r.varve.m1 <- arima(log(d.varve), order=c(0,1,1))
> r.varve.m1$code # Code = 0, i.e. convergence
> r.varve.m1
```

Result:

```
Call:
arima(x = log(d.varve), order = c(0, 1, 1))
```

Coefficients:

```
      ma1
    -0.8421
s.e.    0.0411
```

```
sigma^2 estimated as 0.224: log likelihood = -234.77, aic = 473.53
```

```
> plot(resid(r.varve.m1))
> acf(resid(r.varve.m1))
> pacf(resid(r.varve.m1))
```

##### ARIMA(1,1,1) model:

```
> r.varve.m2 <- arima(log(d.varve), order=c(1,1,1)) # Code = 0, i.e. convergence
> r.varve.m2$code
> r.varve.m2
```

Result:

```
Call:
arima(x = log(d.varve), order = c(1, 1, 1))
```


Coefficients:

```
      ar1      ma1
    0.2461  -0.9140
s.e.    0.0590   0.0234
```

```
sigma^2 estimated as 0.2138: log likelihood = -226.65, aic = 459.3
```

```
> f.acf(resid(r.varve.m2)) > acf(resid(r.varve.m2))
> pacf(resid(r.varve.m2))
```


- c) The correlation structure of the residuals has already been examined in part b). The residuals of the ARIMA(1,1,1) process do look normally distributed:





**R commands:**

```
> qqnorm(r.varve.m2$resid)
> qqline(r.varve.m2$resid)
```

2. a) `> library(tseries)`  
`> t.x <- read.table("http://stat.ethz.ch/Teaching/Datasets/WBL/GARCH.dat")`  
`> t.x <- ts(t.x[,1])`  
`> plot(t.x)`



```
> par(mfrow=c(2,1))
> acf(t.x)
> acf(t.x^2)
```



The timeseries is not stationary, since the variance does not seem to be constant.

The ACF-plot shows no dependencies, the ACF-plot of the squared timeseries shows some kind of exponential decay.

```
b) > mAIC <- matrix(rep(NA, 9), nrow=3)
> colnames(mAIC) <- c("ARCH0", "ARCH1", "ARCH2")
> rownames(mAIC) <- c("GARCHO", "GARCH1", "GARCH2")
> for (i in 0:2){ for (j in 0:2){
  if(i!=0 | j!=0){
    fit <- garch(t.x, order=c(i, j), trace=F)
    mAIC[i+1, j+1] <- AIC(fit)
  }}
> mAIC
```

	ARCH0	ARCH1	ARCH2
GARCHO	NA	-1749.609	-1746.861
GARCH1	-1721.815	-1783.530	-1771.216
GARCH2	-1715.435	-1781.700	-1776.897

```
> min(mAIC, na.rm=T)
[1] -1783.53
```

The AIC suggests a GARCH(1,1).

```
c) > fit <- garch(t.x, order=c(1,1), trace=F)
> round(coef(fit), 5)
      a0      a1      b1
0.00006 0.10222 0.86595
```

3.
  - Timeseries 1: ARMA(2,1) with  $\alpha_1 = 0.8$ ,  $\alpha_2 = -0.2$ ,  $\beta_1 = 0.7$ .
  - Timeseries 2: ARMA(1,3) with  $\alpha_1 = -0.6$ ,  $\beta_1 = -0.7$ ,  $\beta_2 = 0.4$ ,  $\beta_3 = 0.6$ .
  - Timeseries 3: ARIMA(2,1,0) with  $\alpha_1 = 0.4$ ,  $\alpha_2 = -0.5$ .

The data has been simulated with the R-Code:


```
> set.seed(7)
> ts1 <- arima.sim(n=200, model=list(order=c(2, 0, 1), ar=c(0.8, -0.2),
  ma=c(0.7)))
> ts2 <- arima.sim(n=200, model=list(order=c(1, 0, 3), ar=-0.6,
  ma=c(-0.7, 0.4, 0.6)))
> ts3 <- arima.sim(n=199, model=list(order=c(2, 1, 0), ar=c(0.4, -0.5)))
```

In order to determine the appropriate order of the models one should check `acf` and `pacf`. Then one should fit several models, that the `acf` and `pacf` suggest. The residuals and the AIC should be used to decide, which model is best and the coefficients for this model estimated.


4. a) **AR(2) with  $\alpha_1 = 0.9$  and  $\alpha_2 = -0.5$ :**

The ordinary autocorrelations describe a dampened sine curve, and the partial autocorrelations are cut off at lag  $k = 2$ . For the simulated model (`set.seed(79)`), the estimate of partial autocorrelation at lag 2 is  $\hat{\rho}_{part}(2) = -0.45$ , which is a reasonably close to its theoretical counterpart,  $\rho_{part}(2) = \alpha_2 = -0.5$ .

Plot of theoretical ordinary and partial autocorrelations:



Plots for a simulated time series of length  $n = 200$ :



AR(2) with `ar=c(0.9,-0.5)`

**Complete R code:**

```
## Plotting theoretical ordinary autocorrelations
> plot(0:30, ARMAacf(ar=c(0.9,-0.5), lag.max=30), type="h", ylab="ACF")


# Plotting theoretical partial autocorrelations
> plot(1:30, ARMAacf(ar=c(0.9,-0.5), lag.max=30, pacf=T), type="h",
      ylab="PACF")

## Simulation
> set.seed(79)
> r.sim1 <- arima.sim(n=200, model=list(ar=c(0.9,-0.5)))
## Plotting
> plot(r.sim1)
> acf(r.sim1, lag=30)
> acf(r.sim1, type="partial", lag=30)
> str(acf(r.sim1, type="partial"))
> acf(r.sim1, type="partial")$acf[2]
[1] -0.44574
## Estimate of 2nd partial autocorrelation: -0.44574
```


b) **MA(3) with  $\beta_1 = 0.8$ ,  $\beta_2 = -0.5$  and  $\beta_3 = -0.4$ :**

The ordinary autocorrelations are cut off at lag  $k = 3$ .

Plot of theoretical ordinary and partial autocorrelations:



Plots for a simulated time series of length  $n = 200$ :



MA(3) with  $ma=c(0.8,-0.5,-0.4)$


**Complete R code:**

```
> plot(0:30, ARMAacf(ma=c(0.8,-0.5,-0.4), lag.max=30), type="h",
      ylab="ACF")
> plot(1:30, ARMAacf(ma=c(0.8,-0.5,-0.4), lag.max=30, pacf=T),
      type="h", ylab="PACF")
> set.seed(79)
> r.sim2 <- arima.sim(n=200, model=list(ma=c(0.8,-0.5,-0.4)))
> plot(r.sim2)
> acf(r.sim2, lag=30)
> pacf(r.sim2, lag=30)
```


c) **ARMA(1,2) with  $\alpha_1 = -0.75$ ,  $\beta_1 = -1$  and  $\beta_2 = 0.25$ :**

The ordinary autocorrelations do not have a real cut-off point. Because of this, the MA(2) part of the model is difficult to identify in the correlogram of ordinary autocorrelations in this example. However, the autocorrelations do decay exponentially (in absolute terms – their signs alternate), which indicates that an AR component is present. The partial autocorrelations of an ARMA(1,2) process should behave like MA(2) for lags  $k > 1$ , and its ordinary autocorrelations like AR(1) for lags  $k > 2$ . The partial autocorrelation present here might also allow for an AR(3) model to describe the ARMA(1,2) model. (This goes into the topic of approximating ARMA by AR( $\infty$ ) models.

Plot of theoretical ordinary and partial autocorrelations:



Plots for a time series of length  $n = 200$ :



ARMA(1,2) with  $ar=c(-0.75)$ ,  $ma=c(1,-0.25)$

(The R code is along the same lines as that of the previous parts.)