Series 10

1. Let E_t be white noise with mean 0 and variance σ^2. Regard the following three processes:

i) $X_t = t + E_t$

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

iii) $Z_t = X_t - t$

a) Which of these three processes are stationary, and which are not? Why?

b) Compute the theoretical autocorrelation of the processes Y_t and Z_t and the cross-correlation between the two.
Hint: Write these processes using white noise E_t.

c) Simulate both Y_t and Z_t. To this end, assume that E_t follows a standard normal distribution $\mathcal{N}(0, 1)$. Simulate time series of length $n = 200$, and compare your empirical results to the theoretical ones of Part b).

2. In this exercise, we continue the investigation of the connection between advertising expenditure and sales started in Series 6, Exercise 2. We again examine the dataset `advert.dat` which contains the annual advertising expenditure (`ADVERT`) and the sales revenue (`SALES`) for a particular brand of vegetable stock (Lydia Pinkham's Vegetable Compound, 1907-1960).

Load and log-transform the data as follows:

```
> d.advert <- read.table("ftp://stat.ethz.ch/Teaching/Datasets/WBL/advert.dat",
                         header = TRUE)
> ts.advert <- ts(log(d.advert[, "ADVERT"]), start = 1907)
> ts.sales <- ts(log(d.advert[, "SALES"]), start = 1907)
```

In the following problems, we will refer to `ts.advert` and `ts.sales` as $X_{1,t}$ and $X_{2,t}$, respectively.

a) First look at the time series plots. Are the series stationary?
R-Hint: You can either use the generic plot function (first line), or a specific time series plot function (second line):

```
> plot(ts.union(ts.sales, ts.advert), plot.type = "single")
> ts.plot(ts.sales, ts.advert)
```

(Unless the time series have been joined using `ts.union()`, the command `plot(ts.sales, ts.advert)` will plot both of them against each other. It is easy to see that they are correlated.)

b) Compute the first-order differences $Y_{1,t} := X_{1,t} - X_{1,t-1}$ (`ts.adv.d1`) and $Y_{2,t} := X_{2,t} - X_{2,t-1}$ (`ts.sal.d1`). Are these processes stationary?

c) Regard the transfer function model

$$Y_{2,t} = \sum_{j=0}^{\infty} \nu_j Y_{1,t-j} + E_t \,.$$

Describe this model in words. From a business point of view, what is an illustrative interpretation of this model?

d) Estimate the coefficients ν_j. Proceed as described below; see also the lecture slides. Make an interpretation of the model you estimate. What do you notice about this model?
Instructions:

- First identify the row $Y_{1,t}$. Which model should you fit? Estimate its coefficients using the R function `arima()`.

- Compute D_t and Z_t for the transformed series.
  R-Hint: Suppose your R object with the ARMA model you used to fit $Y_{1,t}$ is called `r.fit.adv`. D_t then simply contains the residuals:
  ```
  > ts.D <- resid(r.fit.adv)
  ```
  To calculate Z_t, you can extract the AR (or possibly ARMA) coefficients of the fitted model and then use the R function `filter()`. For example, if `alpha` is a vector of AR coefficients, use
  ```
  > ts.Z <- filter(ts.sal.d1, c(1, -alpha), sides = 1)
  ```
  The argument `sides = 1` causes `filter()` to only take the past into account; cf. the help page.
  Since you should only consider time points present in both time series, use
  ```
  > ts.trans <- ts.intersect(ts.Z, ts.D)
  ```
- Draw the cross-correlogram for D_t and Z_t. Give an interpretation of them.
  `R-Hint:`
  ```
  > acf(ts.trans, na.action = na.pass)
  ```
- Estimate the coefficients ν_j using the forumla given in the lecture.
  R-Hint:
  The variance of the errors D_t can be found in `r.fit.adv$sigma2`. The command
  ```
  > acf(..., type = "covariance", na.action = na.pass)$acf[, 1, 2]
  ```
  also returns the cross-covariances between Y_t and Z_t.

Exercise hour: Monday, May 05.