Series 1

1. We would like to illustrate various methods for descriptive decomposition and elimination of trends using the dataset `hstart`. This dataset contains monthly data on the start of residential construction in the USA within the time frame of January 1966 to January 1974. The data have undergone some transformation unknown to us (perhaps an index over some baseline value has been calculated, or perhaps the data are to be read as $x \cdot 10^?$ construction permits).

(Source: U. S. Bureau of the Census, Construction Reports.)

a) Read in the time series from a file; note that the function `ts()` expects a *vector* as an argument, whereas `read.table()` returns a *data frame*:

```
> hstart <- read.table("ftp://stat.ethz.ch/Teaching/Datasets/WBL/hstart.dat")
> hstart <- ts(hstart[, 1], start = 1966, frequency = 12)
```

Plot the time series. Is it stationary? If not, what kind of non-stationarity is evident? Into which components might this time series be decomposed sensibly?

b) **STL decomposition**

Decompose the time series into trend, seasonal component and remainder using the non-parametric STL method. Add the sum of trend and seasonal component to the plot of the time series you made in Part a).

R hint:

The decomposition is made using

```
> H.stl <- stl(hstart, s.window = "periodic")
```

The smoothing parameter for the seasonal effect is chosen by means of `s.window`. If `s.window = "periodic"`, the seasonal effect is estimated by averaging. It is also possible to specify a value for the smoothing parameter (an odd number). Try e.g. `H.stl.var <- stl(hstart, s.window = 15)`, and compare the result of this to `H.stl`. Incidentally, `summary()` can be used for displaying the values of `window`.

The trend estimation parameter can be set using `t.window`. Unlike `s.window`, this argument does have a default value (cf. the help file). Perhaps you could try to vary this parameter as well. The documentation for R and the help files give more details.

Trend, seasonal component and remainder of the STL decomposition are stored as

```
> H.stl$time.series[, "trend"]
> H.stl$time.series[, "seasonal"]
> H.stl$time.series[, "remainder"]
```

Have a look at the output of `str(H.stl)` for more details.

c) Plot the components of the STL decomposition and its estimates of monthly effects. Compare the monthly effects of the original decomposition, `H.stl`, to that of the decomposition with a chosen smoothing parameter, `H.stl.var` (cf. Part b)). Comment on the plots.

R hints:

```
> plot(H.stl)
> monthplot(H.stl$time.series[, "seasonal"])
```

d) The special filter

$$Y_t = \frac{1}{24}\Big(X_{t-6} + 2X_{t-5} + \ldots + 2X_t + \ldots + 2X_{t+5} + X_{t+6}\Big)$$

can be used for computing a trend estimate (cf. course notes). Calculate this filter and plot it together with the STL trend and the data in a single plot.

R hint: To calculate the filter, use

```
> H.filt <- filter(hstart, c(1, rep(2, 11), 1)/24)
```

e) The function `decompose()` performs a decomposition of a time series into trend, seasonal component and remainder similar to `stl()`, but using moving averages as in the filter of Part d) to estimate the trend. By using `decompose()`, we do not have to specify the filter by hand.
Perform a decomposition of the time series using `decompose()`, and add the resulting trend to the plot of task d).
R hint:

```
> H.decomp <- decompose(hstart, type = "additive")
```

Use `str(H.decomp)` or `?decompose` to get aware of the structure of the resulting object `H.decomp`.

f) Differencing
Try to remove the trend and seasonal effects by computing differences. After removing seasonal effects, choose some linear trend elimination method from the course notes and plot the outcome.
R hint:

```
> diff(hstart, lag = ...)
```

2. Have a look at the following set of time series. For every series, shortly answer the following questions (motivate your answers):

- Is the time series stationary?
- Is there a trend?
- Can one find some seasonal effect? If yes, what is the period?

a) Daily closing prices of the Swiss Market Index:

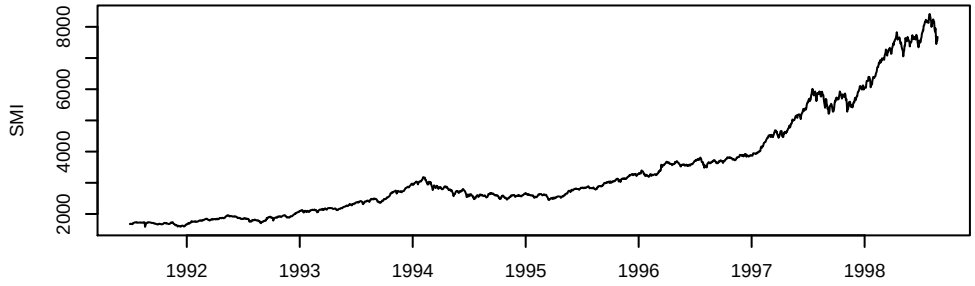

b) The underlying time series is the same as in Part a), the daily closing prices of the Swiss Market Index. Here, the data is plotted after taking differences to lag 1 of the logarithm of the time series.

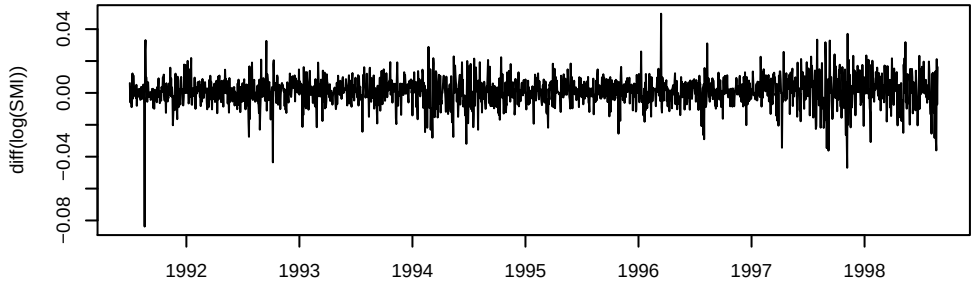

c) Quarterly earnings per share of Johnson & Johnson, an American manufacturer of pharmaceuticals and medical devices:

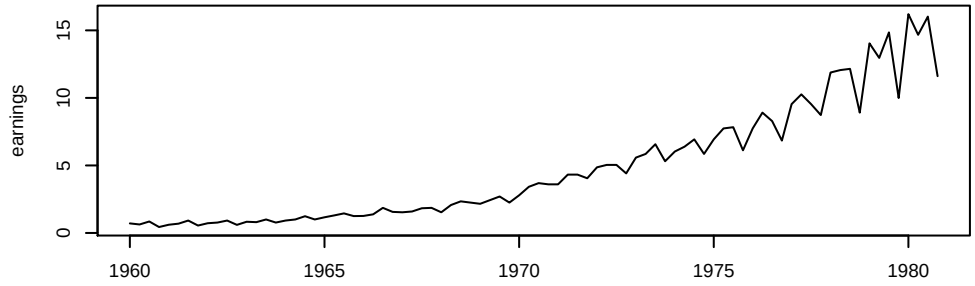

d) Monthly totals of international airline passengers departing from the USA, 1949 to 1960:

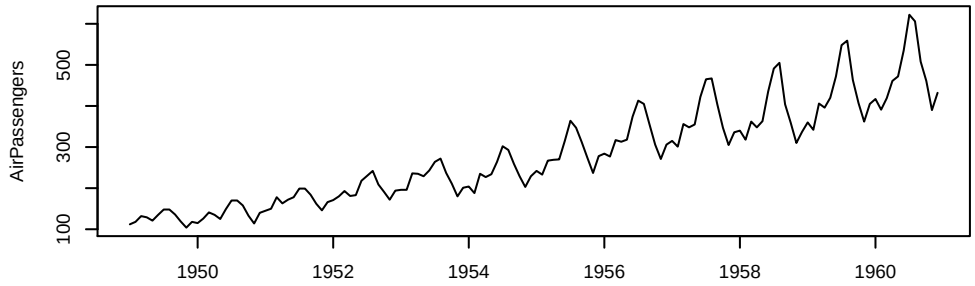

e) Again the "air passengers" time series of Part d), plotted after taking differences to lag 1 of the logarithm of the time series.

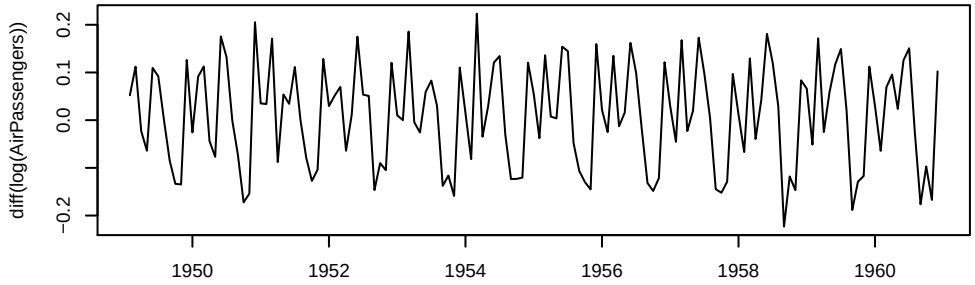

f) Yearly numbers of sunspots:

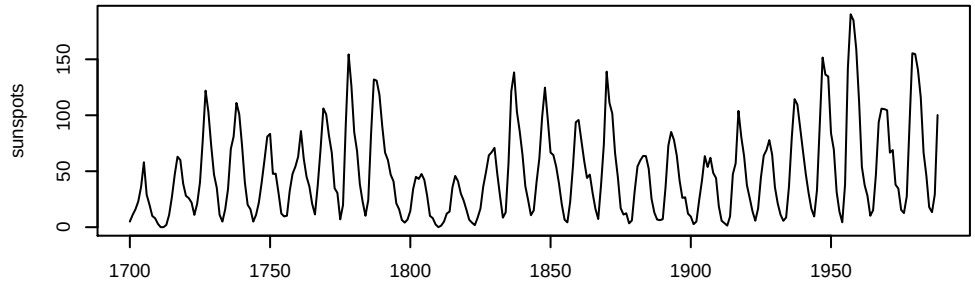

g) Widths of tree rings in the years from 1500 to 1979:

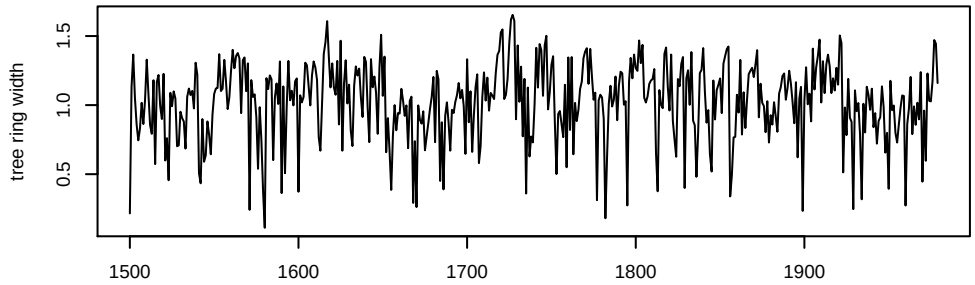

Preliminary discussion: Monday, February 28.

Deadline: Monday, March 07.