

Series 8

1. In this exercise, we revisit the deposits of Sheet 7, Exercise 1 (data `varve.dat`). We saw that an $ARIMA(1,1,1)$ model fitted the logarithmic data. Now we would like to use this model to predict the next 30 values.
 - a) Use the $ARIMA(1,1,1)$ model found in Sheet 7, Exercise 1 to compute the predictions. Plot the predicted value and the prediction intervals. What do you notice?
R Hint: `predict()`, `plot()`, `lines()`
 - b) Compute the first-order differences `y = diff(log(d.varve))`, and then fit an $ARMA(1,1)$ process to them. For this new series y , now, make and plot predictions based on the $ARMA(1,1)$ model. What do you notice when comparing these to part a)?
 - c) Which transformations are necessary if we want to obtain predictions for the original time series? Describe how to proceed (without actually performing the calculations you describe) when using the data from parts a) and b), respectively.

2. In this exercise we analyse three datasets using different tools: a dataset on **AT&T bonds**, one on **pine data** and the already known **airline data**. You can load the **AT&T bond data** as follows:

```
> t.url <- "http://stat.ethz.ch/Teaching/Datasets/WBL/attbond.dat"
> d.att <- ts(scan(t.url, skip=1))
```

The dataset contains the daily capitalised earnings value of an AT&T bond from april 1975 until december 1975.

To load the **pine data**, call:

```
> t.url <- "http://stat.ethz.ch/Teaching/Datasets/WBL/foehre.dat"
> d.foe <- ts(scan(t.url, skip=1), start=1107)
> d.foe <- window(d.foe, start=1201, end=1500)
```

This dataset contains the measurements of the year rings of an 858 year old Douglas Fir from 1107 until 1964. We use only the data from 1201 until 1500. The **airline data** are already available in R as a `ts` object:

```
> d.air <- AirPassengers
```

It contains monthly data on numbers of airpassengers.

- a) Fit an ARIMA/SARIMA Model for all of these datasets. Use transformations if suitable. Compute a prediction and plot it along with the prediction band. Try different models and compare the outcome. Also analyse the residuals using `acf()`/`pacf()` and `qqnorm()`.
R Hint: `predict(..., n.ahead=...)`
 For the time series plot, the `acf` and the `pacf` you can either plot them separately or do it simultaneously with the function `f.acf()` (which can be downloaded (or copied) from <http://stat.ethz.ch/WBL/Source-WBL-2/R/f.acf.R>).
- b) Now fit a polynomial model with a season if necessary. Again compute the prediction and plot it. Check the residuals.
- c) Now use exponential smoothing. Again make predictions, plot the predicted timeseries, including confidence bounds and check the residuals.
R Hint: `?HoltWinters`

Exercise hour: Monday, April 14.