Time series - introduction

Beáta Stehlíková

Time Series Analysis

Faculty of Mathematics, Physics and Informatics, Comenius University

Time series analysis

- We have monthly data - number of airline passengers:

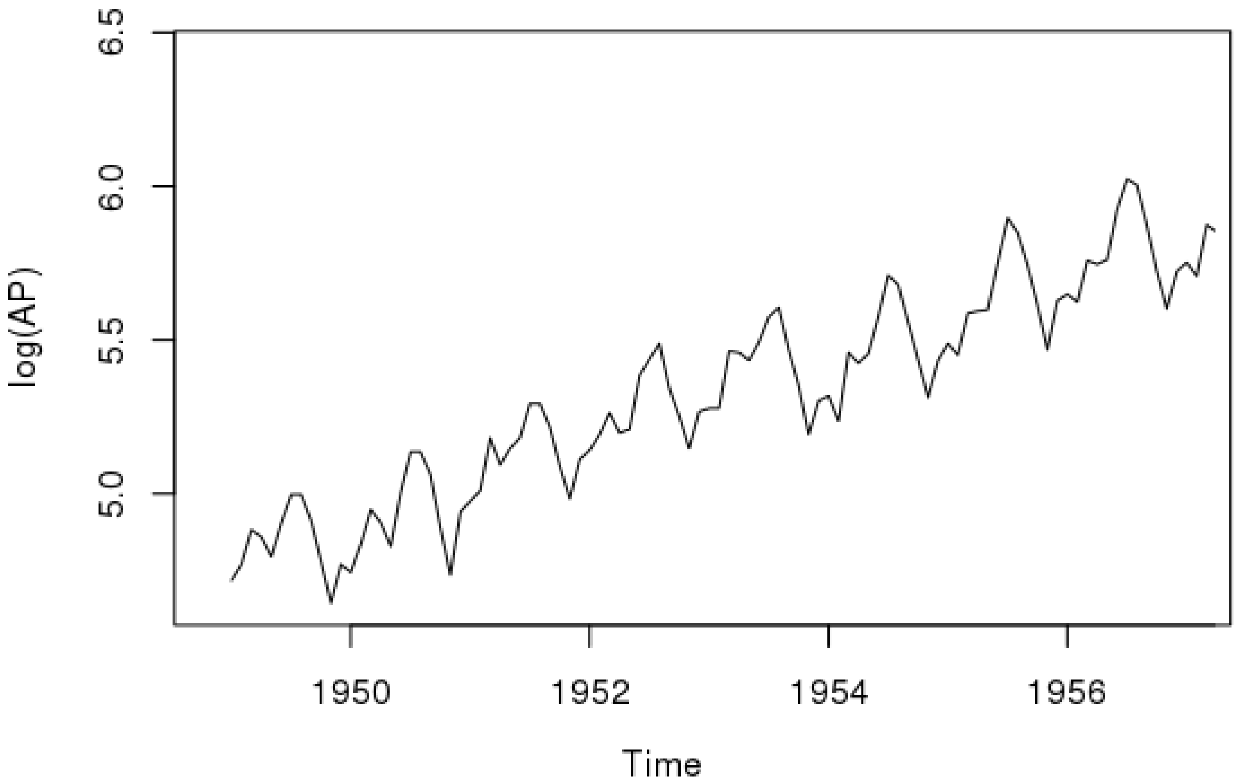

G. E. P. Box, G. M. Jenkins: **Time Series Analysis: Forecasting and Control**.

- Question: What the future number of passenger going to be?

Time series analysis

- We see an increasing trend and seasonality (months).
- Decomposition in R software:
  `plot(decompose(log(AP)))`

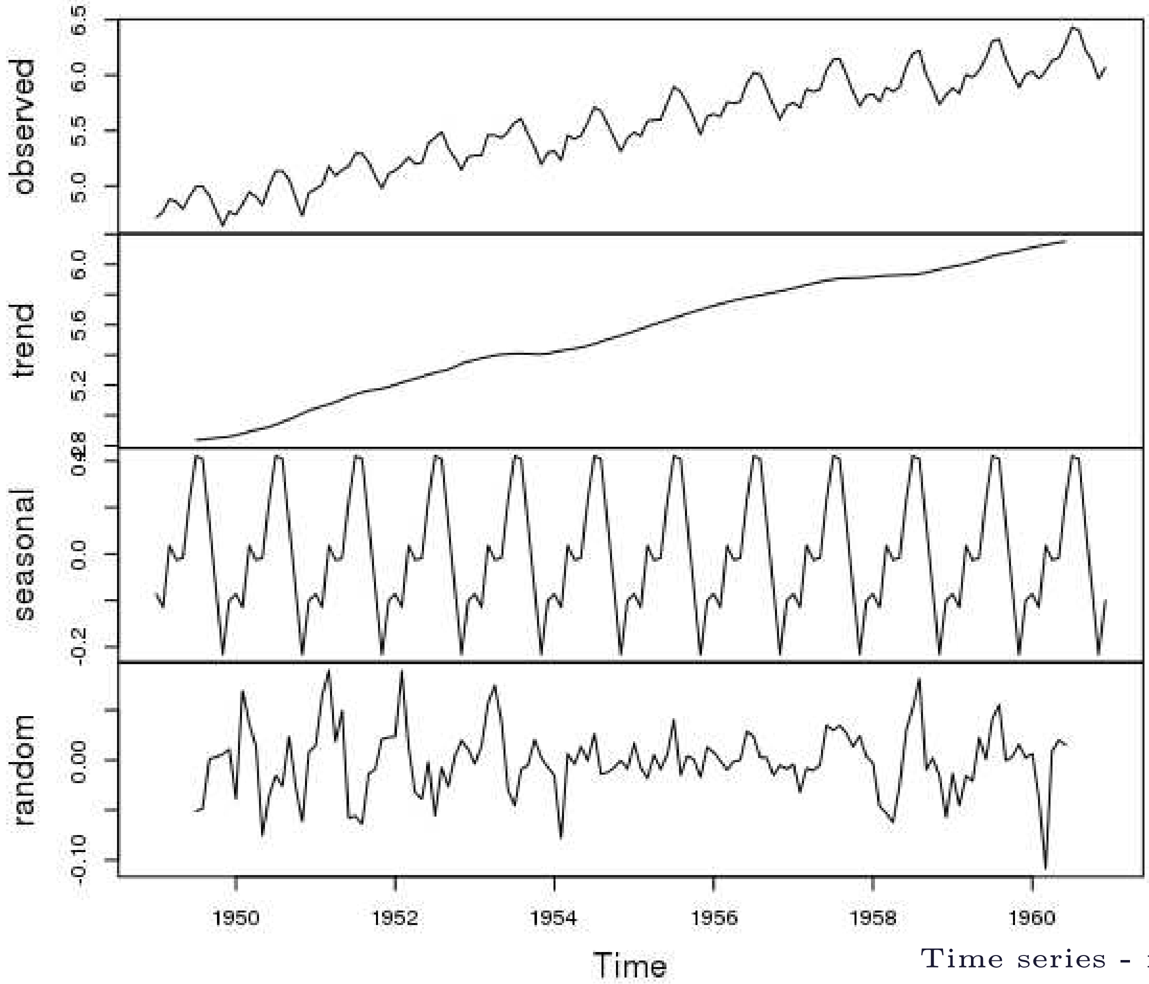

Time series analysis

- Intuition: without a shock, increasing trend and seasonality remain, so it is going to be something like:

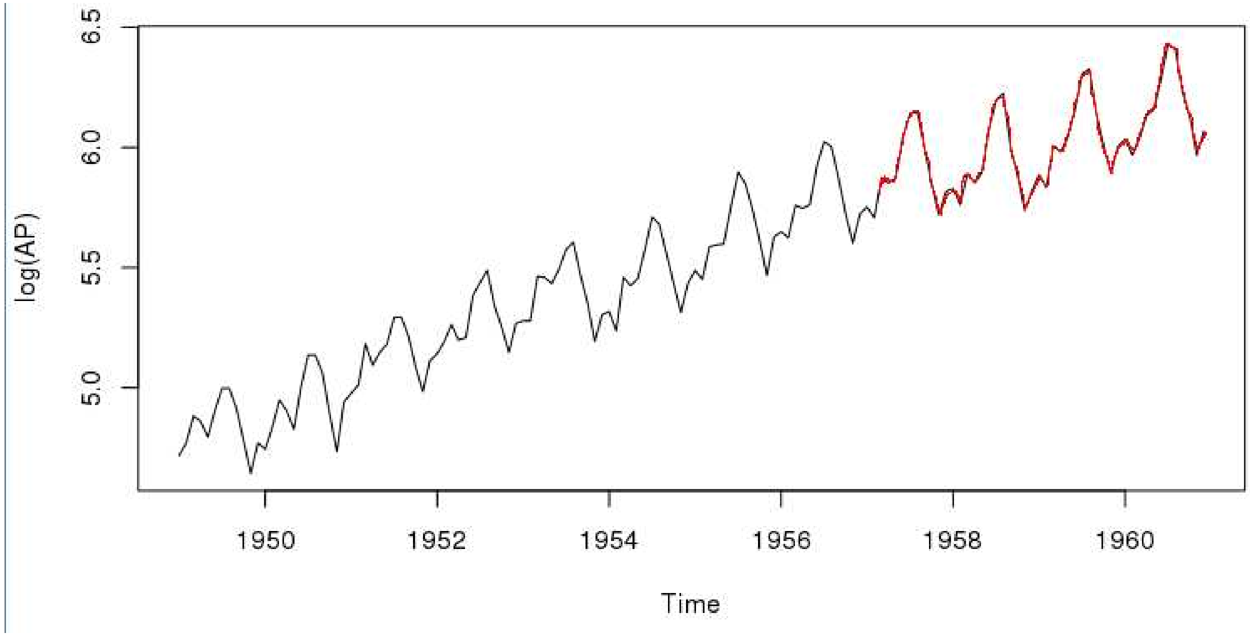

- Questions: How to express it quantitatively? How to determine the accuracy of the estimates, how to construct confidence intervals?

Box a Jenkins

A big part of the course will be based on the approach of Box and Jenkins

"The first paper you wrote with Jenkins has been considered as a ***breakthrough in statistics****. How do you become interested in time series?"*

Interview with G. E. P. Boxom after the celebration of his 80th birthday (1999):

http://halweb.uc3m.es/esp/Personal/personas/dpena/articles/boxIJFinter4.PDF

Modelling volatility

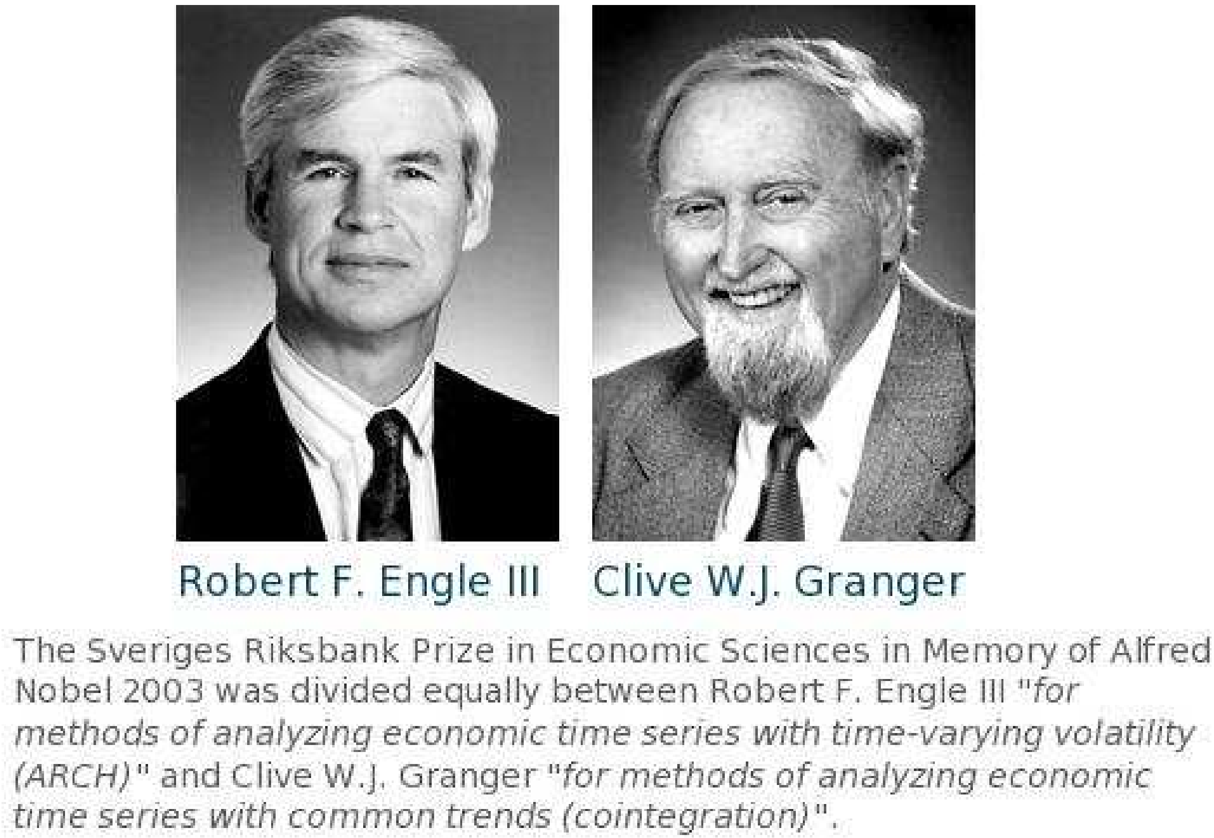

http://www.nobelprize.org

- ARCH model and its generalizations - also a part of this course

Curiosity

- The wife of R. Engle was born in Slovakia, in Prešov, in 2005 they visited Prešov together.

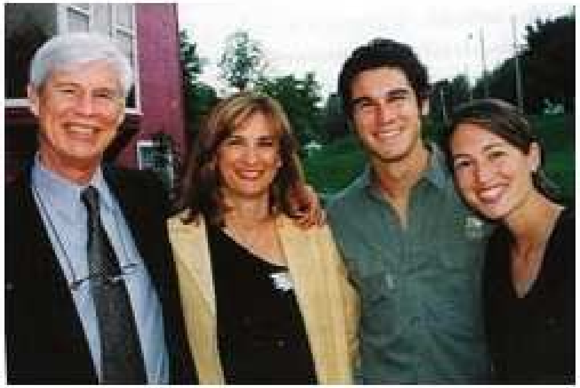

Rob, Marianne, Jordan and Lindsey in Williamstown 2002.

http://www.nobelprize.org/nobel_prizes/economics/laureates/2003/engle-autobio.html

http://www.presov.sk/portal/?c=12&id=3590 (in Slovak)

Basic concepts - outline of the lecture

- Time serie, moments
- Stationarity, ergodicity
- White noise
- Wold representation
- Autocorrelation function, tests about the autocorrelation function
- Lag operator

Time serie, moments

- Stochastic process $x_1, \ldots, x_T$ - it is fully determined by T-dimensional cumulative distribution function
- Usually we concentrate on the first two moments:
  - expected value $E[x_t]$
  - variance $Var[x_t]$
  - covariances $Cov[x_t, x_s]$, so called autocovariances

Stationarity, ergodicity

- Usually we have only one time series - one realization of the random process $\rightarrow$ in order to o some statistical inference we need additional assumptions
- For example: to estimate the expected value,... we need more than one realization of the random variable
- Ergodic process - sample moments computed from the time series with T observations converge as $T \rightarrow \infty$ to corresponding moments
- This concept makes sense only if we assume that $E[x_t] = \mu$, $Var[x_t] = \sigma^2, \ldots$ for $\forall t$

Stationarity, ergodicity

- Strong stationarity: joint CDF does not change when shifted in time
- Usually we work with a weaker assumption $\rightarrow$ weak stationarity:

$$
\begin{aligned}
(1) \qquad E[x_t] &= \mu \;\; \forall t \\
(2) \qquad Cov[x_t, x_s] &= \gamma(|t-s|) \;\; \forall t, s
\end{aligned}
$$

from (2) it follows: $Var[x_t] = const.$ for $\forall t$

- In what follows, under "stationarity" we will understand the weak stationarity

Stationarity - data

- Stationary time serie: data are reverting to some equilibrium value, around which they fluctuate
- Nonstationary time serie: for example with a trend

Stationarity - data

- EXAMPLE 1

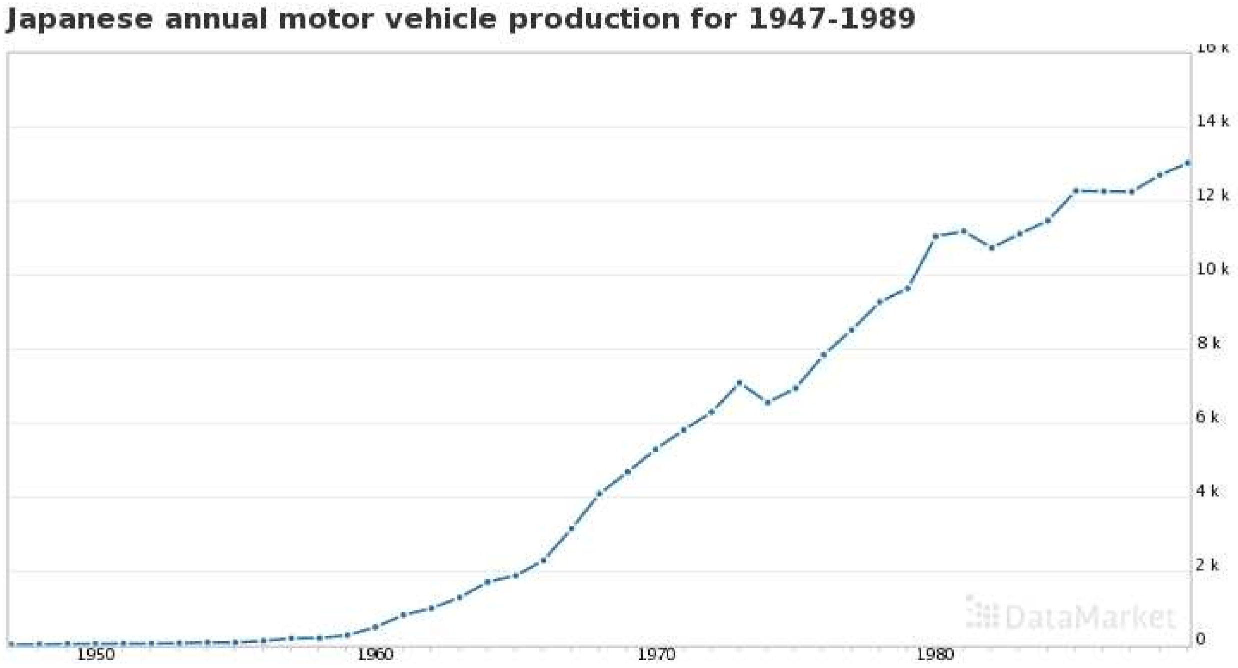

http://data.is/RVcjbL

We see:
increasing trend ⇒ expected value is not constant in time ⇒ time series is not stationary

Stationarity - data

- EXAMPLE 2
  Relative price: if two goods are on the same market, their relative price should fluctuate around some equilibrium value

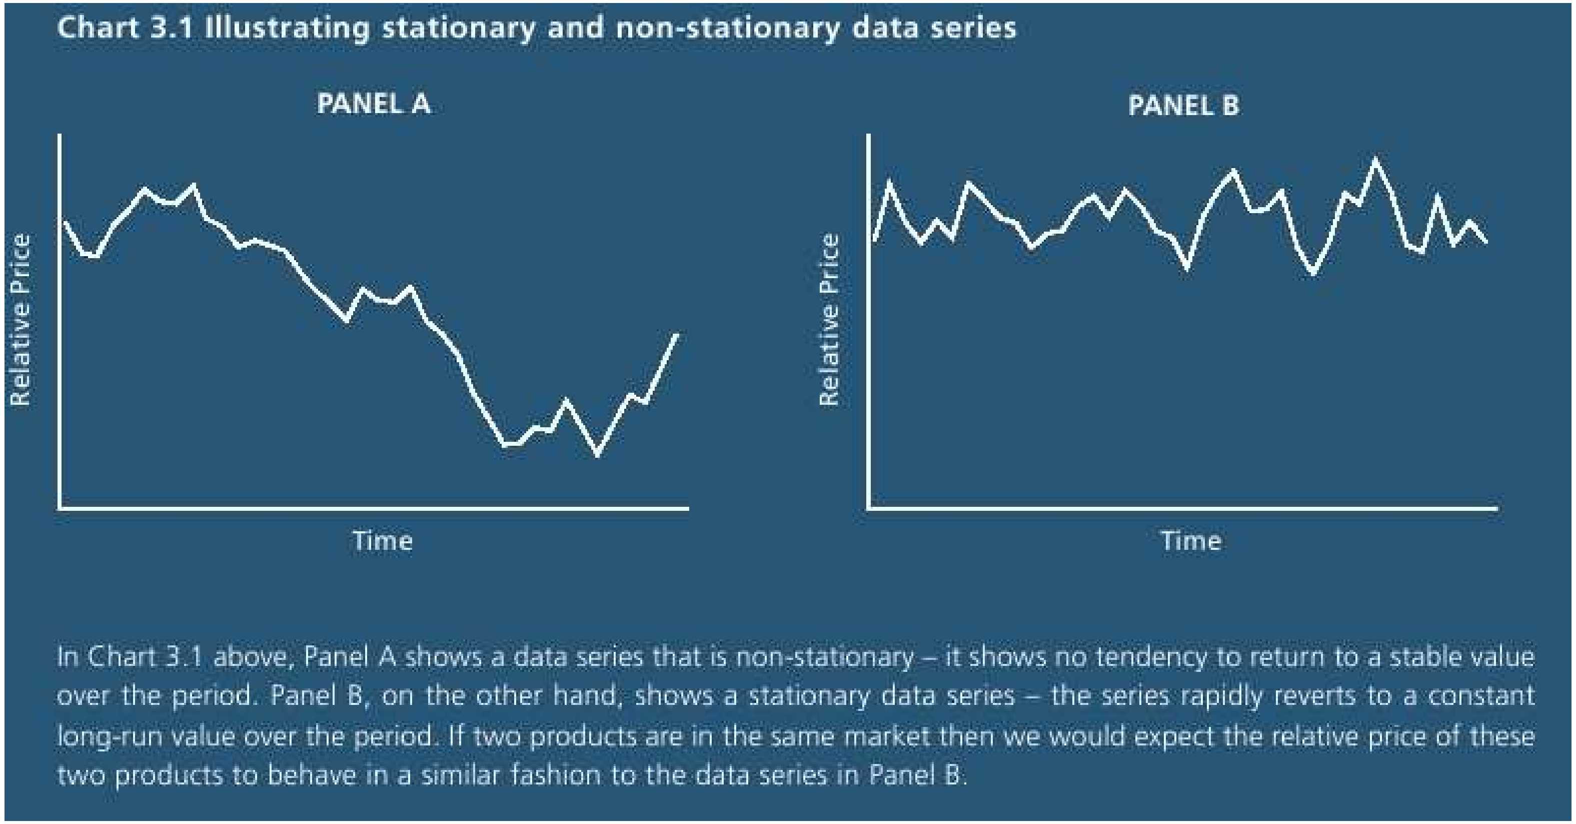

Chart 3.1 Illustrating stationary and non-stationary data series

In Chart 3.1 above, Panel A shows a data series that is non-stationary – it shows no tendency to return to a stable value over the period. Panel B, on the other hand, shows a stationary data series – the series rapidly reverts to a constant long-run value over the period. If two products are in the same market then we would expect the relative price of these two products to behave in a similar fashion to the data series in Panel B.

Stationarity - data

Chart 3.2: Price of Scottish Salmon relative to the price of Norwegian salmon in the UK

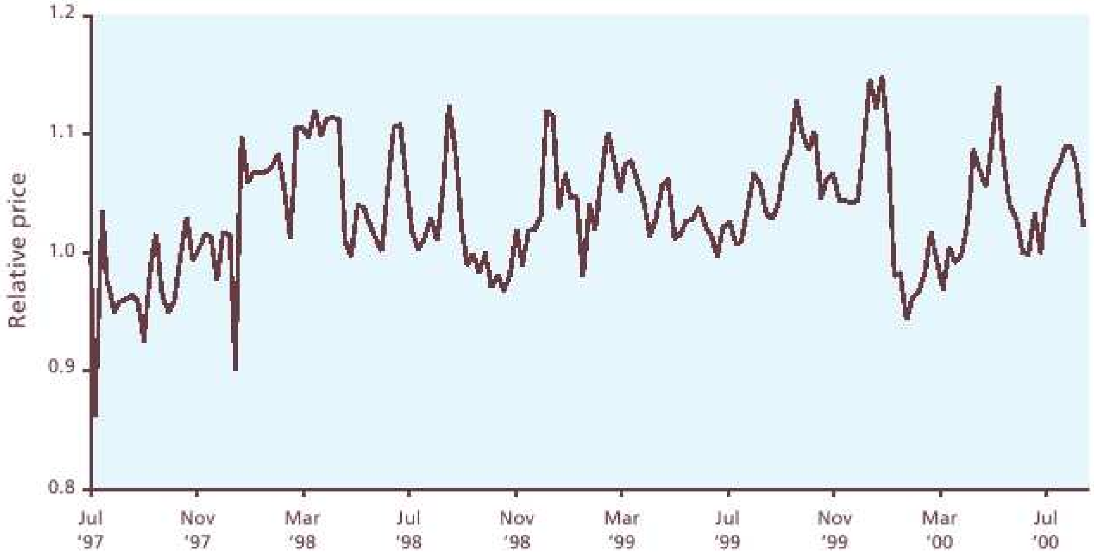

The above chart shows that the price of Scottish salmon relative to the price of Norwegian salmon in the UK appears to vary randomly around a constant long-run value, which suggests that the relative price is stationary. The econometric test for stationarity confirms that the relative price of Scottish salmon is stationary, which is what we would expect to observe if Scottish and Norwegian gutted salmon compete in the same product market in the UK.

An Introduction to Quantitative Techniques in Competition Analysis, Lexecon, 2005

http://www.crai.com/ecp/publications/full_list.htm.

White noise

- White noise $\{u_t\}$ - an important example of a stationary process

$$
\begin{aligned}
E[u_t] &= 0 && \forall t \\
Var[u_t] &= \sigma^2 && \forall t \\
Cov[u_t, u_s] &= 0 && \forall t \neq s
\end{aligned}
$$

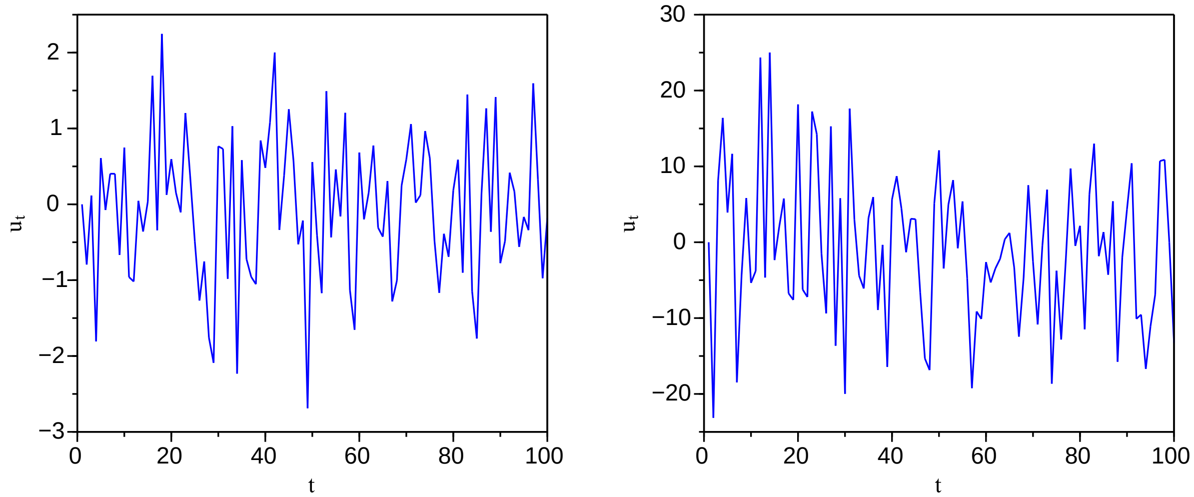

- Basic random process, we will use it to define others, also models for data

Stationarity - example 1

- Let u_t be a white noise, we define

$$x_t = u_t + u_{t-1}$$

- We compute:

$$E[x_t] = 0, \quad Var[x_t] = 2\sigma^2$$

$$Cov[x_t, x_s] = \begin{cases} \sigma^2 & \text{for } |t-s| = 1 \\ 0 & \text{for } |t-s| = 2, 3, \ldots \end{cases}$$

$\rightarrow$ process is stationary

Stationarity - example 2

- Let u_t be a white noise; we define

$$y_t = \begin{cases} u_1 & \text{for } t = 1 \\ y_{t-1} + u_t & \text{for } t = 2, 3, \ldots \end{cases}$$

- y_t can be written as $y_t = \sum_{i=1}^{t} u_i$
- We compute:

$$E[y_t] = 0, \quad Var[y_t] = t\,\sigma^2$$

$$Cov[y_t, y_s] = \sigma^2 \, \min(t, s)$$

$\rightarrow$ process is not stationary

Stationarity - examples 1, 2

- Comparison of trajectories of the processes from the previous examples:
  - ◇ left - stationary process (ex. 1)
  - ◇ right - nonstationary process (ex. 2), we can see the increasing dispersion

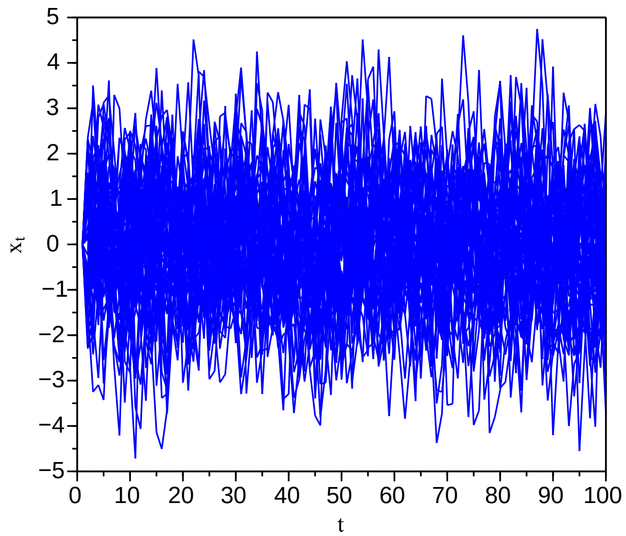

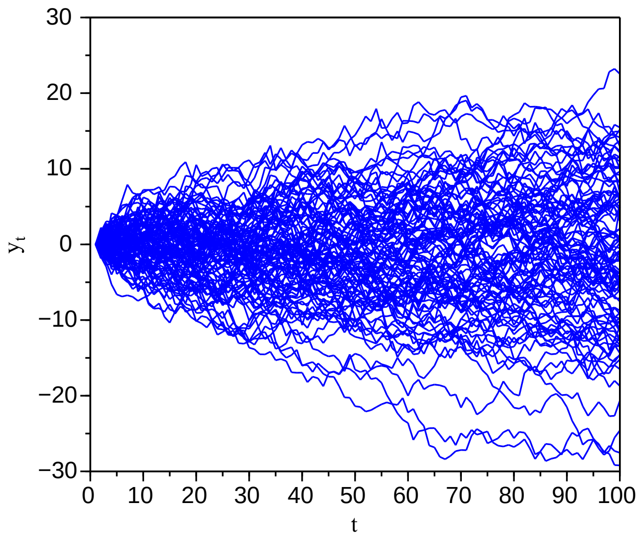

Stationarity - example 3

- Let $\{u_t\}_{t=-\infty}^{\infty}$ be a white noise; define

(3) $$x_t = \mu + \sum_{j=0}^{\infty} \psi_j u_{t-j},$$

where the coefficients ψ_j satisfy the condition $\sum_{j=0}^{\infty} \psi_j^2 < \infty, \psi_0 = 1$

- We compute:

$$E[x_t] = \mu, \;\; Var[x_t] = \sum_{j=0}^{\infty} \psi_j^2$$

$$Cov[x_t, x_{t+k}] = \sigma^2 \sum_{j=0}^{\infty} \psi_j \psi_{k+j}$$

$\rightarrow$ process is stationary

Woldo representation

- Previous example: process x_t, which has the form (3), is stationary
- It can be shown:
  Every stationary process x_t can be written in the form (3) , i. e.

  $$x_t = \mu + \sum_{j=0}^{\infty} \psi_j u_{t-j},$$

  where the coefficients ψ_j satisfy
  $\sum_{j=0}^{\infty} \psi_j^2 < \infty, \psi_0 = 1$ and u_t is a white noise
- This representation of a stochatic process is called the Wold representation (Wold, 1938)

Autocorrelation function - motivation

- Left: $y_t = u_t$, right: $y_t = 0.9y_{t-1} + u_t$
- Sample path and how the value y_t depends on y_{t-1}

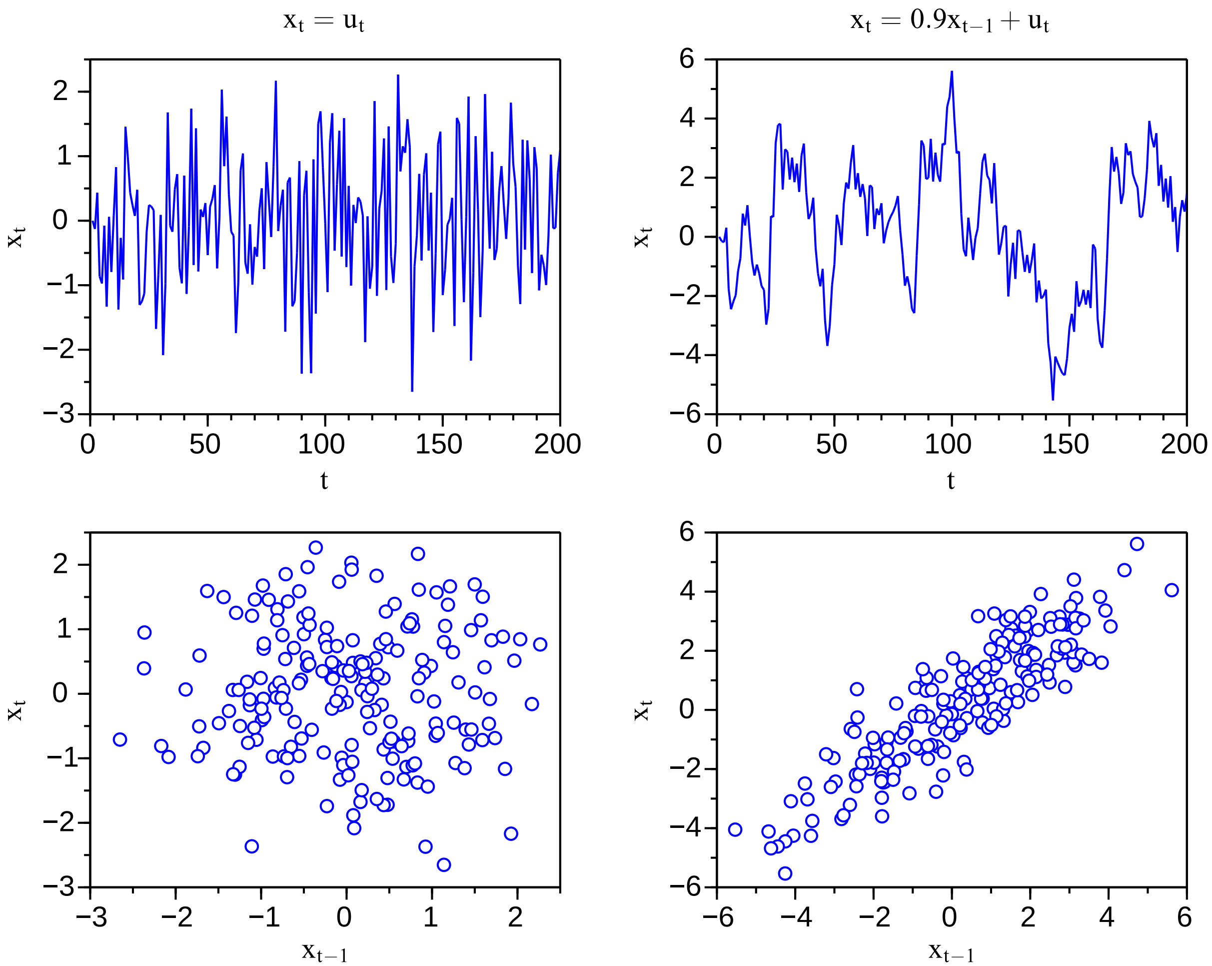

Autocorrelation function

- Autocorrelation function (ACF) of a stationary process:

$$\rho(\tau) = \frac{\gamma(\tau)}{\gamma(0)}$$

t. j.

$$\rho(\tau) = Cor(x_t, x_{t+\tau}) = \frac{Cov(x_t, x_{t+\tau})}{\sqrt{Var(x_t)\,Var(x_{t+\tau})}}$$

- Properties:

$$\rho(0) = 1, \;\; \rho(\tau) = \rho(-\tau)$$

$\rightarrow$ it suffices to compute $\rho(\tau)$ for $\tau = 1, 2, \ldots$

Autocorrelation function

- Ergodic process $\rightarrow$ expected value, variance, covariances can be consistently estimated from the data $x_1, \ldots, x_T$:

$$\hat{\mu} = \frac{1}{T}\sum_{t=1}^{T} x_t, \quad \hat{\gamma}(0) = \frac{1}{T}\sum_{t=1}^{T}(x_t - \hat{\mu})^2$$

$$\hat{\gamma}(\tau) = \frac{1}{T}\sum_{t=1}^{T-\tau}(x_t - \hat{\mu})(x_{t+\tau} - \hat{\mu})$$

$\rightarrow$ consistent estimate of the ACF:

$$\hat{\rho}(\tau) = \frac{\hat{\gamma}(\tau)}{\hat{\gamma}(0)}$$

- asymptotically unbiased

Autocorrelation function - tests

- Estimate of the ACF in the case of a white noise
  - asymptotically unbiased
  - variance $\approx 1/T$
  - approximate 95 % confidence interval: $\pm 2/\sqrt{T}$, it is often plotted together with estimated ACF
- In the case of a stochastic process, for which $\rho(\tau) = 0$ for $\tau > k$, for these τ we have

$$Var[\hat{\rho}(\tau)] \approx \frac{1}{T}\left(1 + 2\sum_{j=1}^{k}\hat{\rho}(j)^2\right)$$

Autocorrelation function - tests

- Testing, if a give time series is a white noise:
  1. confidence interval $\pm 2/\sqrt{T}$ for every autocovariance separately
  2. testing the zero value of $\rho(1), \ldots, \rho(m)$ together:
     - ◇ Box & Pierce, 1970: if H_0 holds, asymptotically

$$Q = T \sum_{j=1}^{m} \rho(j)^2 \sim \chi^2_m$$

     - ◇ Ljung & Box, 1978: modification with better properties in finite samples:

$$Q = T(T+2) \sum_{j=1}^{m} \frac{\rho(j)^2}{T-j} \sim \chi^2_m$$

     - ◇ *degress of freedom change when we are testing residuals from a model*

ACF of a time series - example 1

- We simulate 150 independent realizations of $N(0,1)$ distribution (so a white noise):

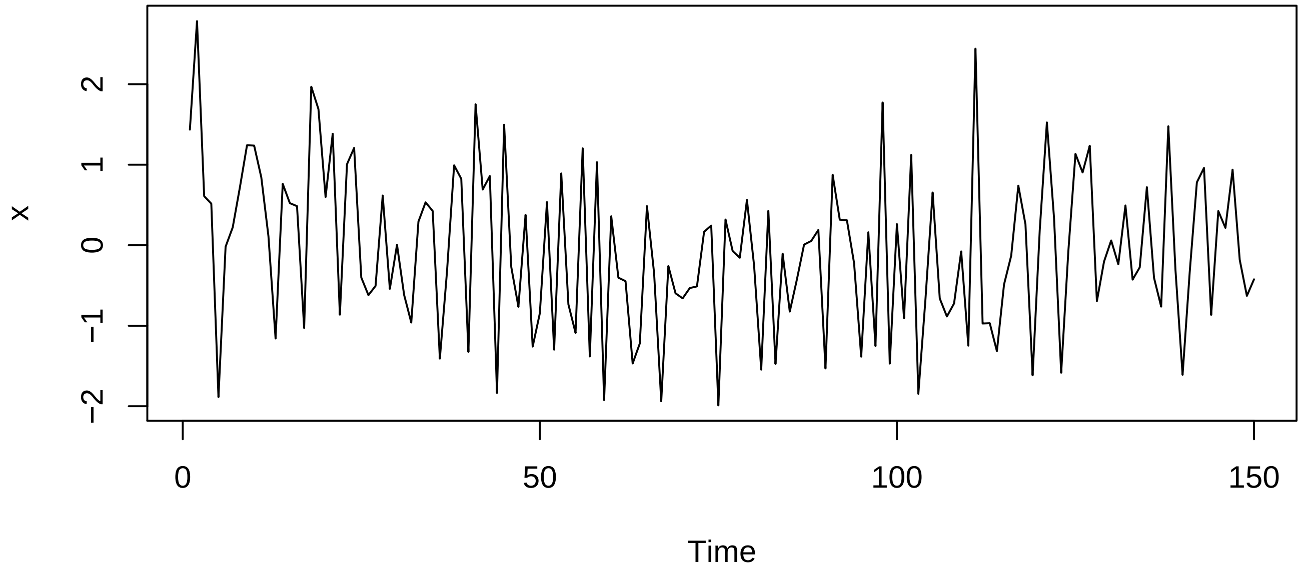

- ACF in R software:
  - ◇ function **`acf`**, for example. **`acf(x)`**
  - ◇ autocorrelation for lag 0 is always equal to 1, we can omit it, for example: **`plot(acf(x)[1:20])`**

ACF of a time series - example 1

- Output:

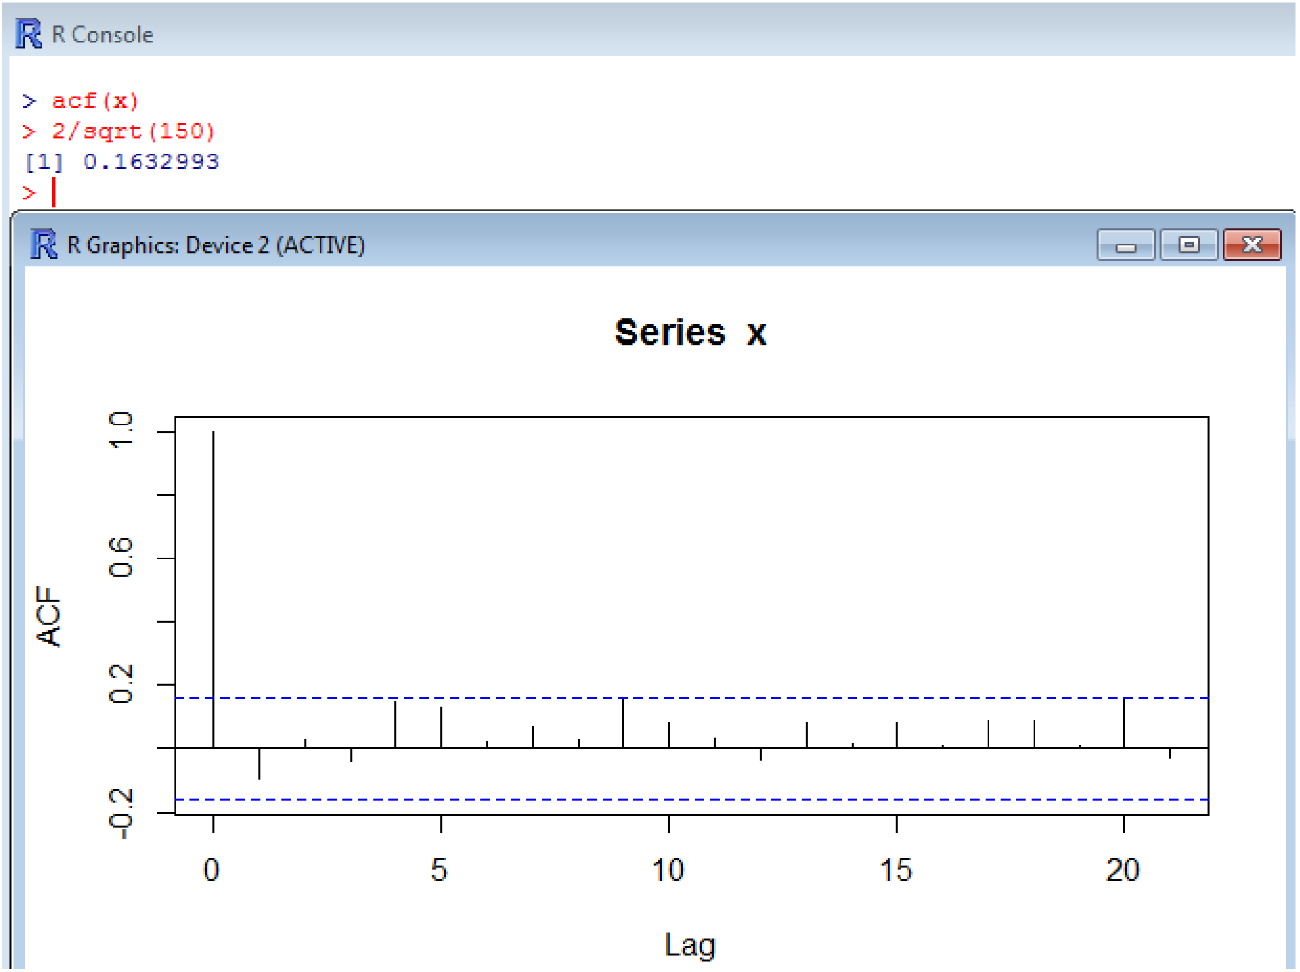

ACF of a time series - example 1

- Without lag 0:

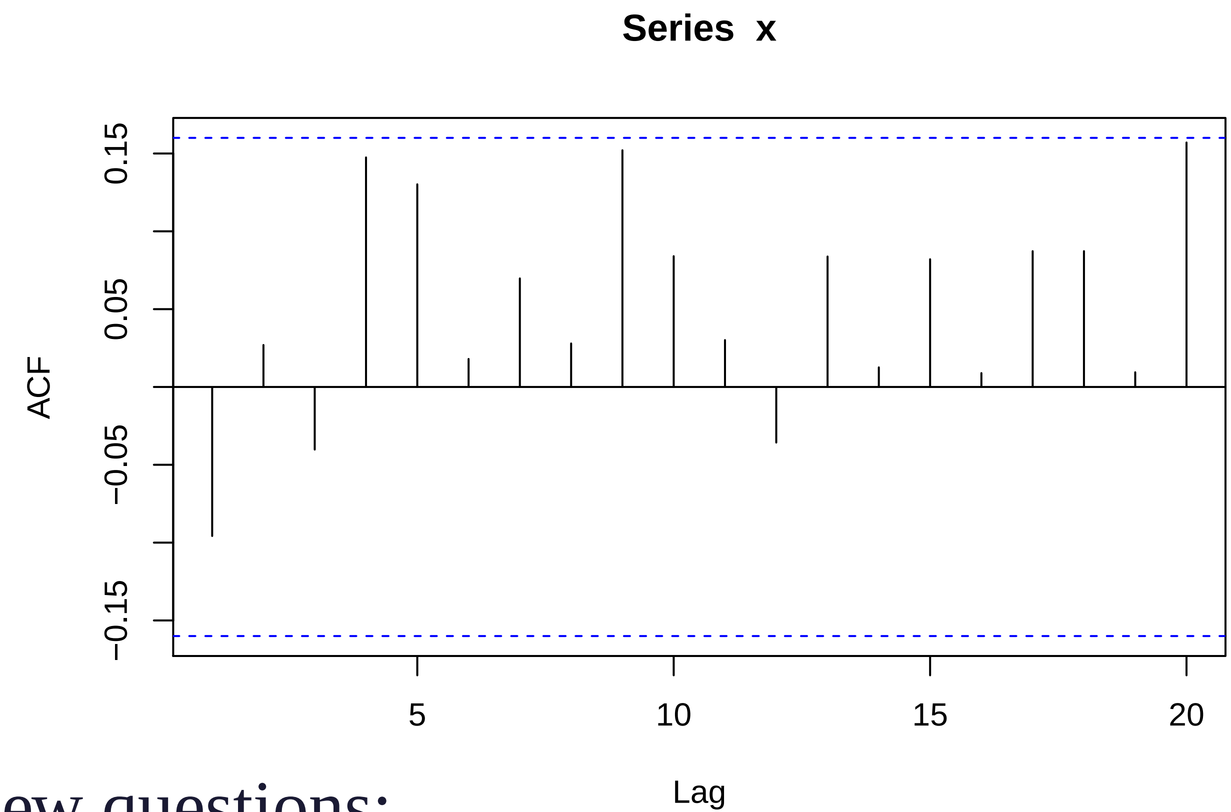

- Review questions:
  - What hypothesis can we test for every autocorrelation?
  - When is this hypothesis rejected?
  - What is the result in this concrete example?

ACF of a time series - example

- Ljung-Box statistics:
  - in R - function `Box.test`
  - for example, we want to test whether the first three autocorrelations are zero:

```
Box.test(x, lag=3, type="Ljung")
```

- Review questions:
  - How is the test statistics computed?
  - What is the probability distribution of the statistics under the null hypothesis?
  - When is the null hypothesis rejected?
  - How can we compute the p-value?

ACF of a time series - example

- We get:

```
> Box.test(x,lag=3,type="Ljung")

        Box-Ljung test

data:  x
X-squared = 1.7655, df = 3, p-value = 0.6225
```

- Extracting p-values:
  `Box.test(x,lag=3,type="Ljung")$p.value`
  $\rightarrow$ we write a loop and plot the p-hodnoty for all of the lags and compare them with 0.05 $\rightarrow$ what follows from this plot?

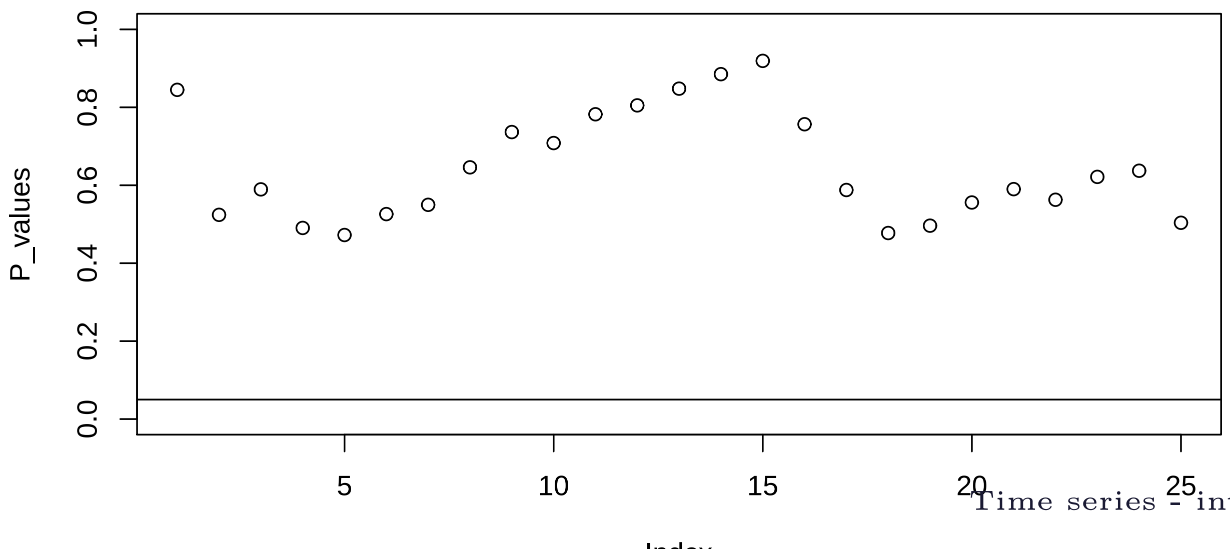

Lag operator

- Lag operator L - useful when studying time series
- Shifts the value of the process one period backwards:

$$Lx_t = x_{t-1}$$

- Properties:
  - ◇ powers: $L^2 x_t = L(Lx_t) = x_{t-2}$
  - ◇ $L^0 = 1$ is an identity: $(1 - L)x_t = x_t - x_{t-1}$
  - ◇ working with powers: $L^2(L^3) = L^5$
  - ◇ multiplication:
    $(1 - 0.5L)(1 - 0.3L) = 1 - 0.8L + 0.15L^2$
  - ◇ if c is a constant, then for example
    $(1 - 0.1L + 2L^2)c = (1 - 0.1 + 2)c = 2.9c$

Practice problems

Compute the expected value, variance and autocorrelation function for the following processes:

1. $x_t = u_t + u_{t-2}$, where u is a white noise
2. $x_t = u_t u_{t-1}$, where u_t are independent realizations of a random variable with zero expected value zero and a finite variance
3. $x_t = u_t$ for t odd and $x_t = \frac{\sqrt{2}}{2}(u_{t-1}^2 - 1)$ for t even, where u is a white noise with distribution $N(0, 1)$

Practice problems - remark

- It might be useful to check the computation by simulating the process and computing its sample ACF (and maybe running this for a couple of times).
- For example for the problem 2:

```
> set.seed(1234)          # for reproducibility
> n <- 500                # length of u
> u <- rnorm(n)           # we need concrete distribution
> x <- u[2:n]*u[1:(n-1)] # process x
> plot(acf(x)[1:20])     # sample ACF
```

Practice problems - remark

- Output from this code: