Applied Time Series Analysis
FS 2012 – Week 08

Marcel Dettling

Institute for Data Analysis and Process Design

Zurich University of Applied Sciences

marcel.dettling@zhaw.ch

http://stat.ethz.ch/~dettling

ETH Zürich, April 16, 2012

Applied Time Series Analysis
FS 2012 – Week 08

Non-Stationary Models: ARIMA and SARIMA

Why?

We have seen that many time series we encounter in practice show trends and/or seasonality. While we could decompose them and model the stationary part, it might also be attractive to directly model a non-stationary series.

How does it work?

There is a mechanism, "the integration" or "the seasonal integration" which takes care of the deterministic features, while the remainder is modeled using an ARMA(p,q).

There are some peculiarities!

→ **see blackboard!**

Applied Time Series Analysis
FS 2012 – Week 08

Example: Monthly Oil Prices

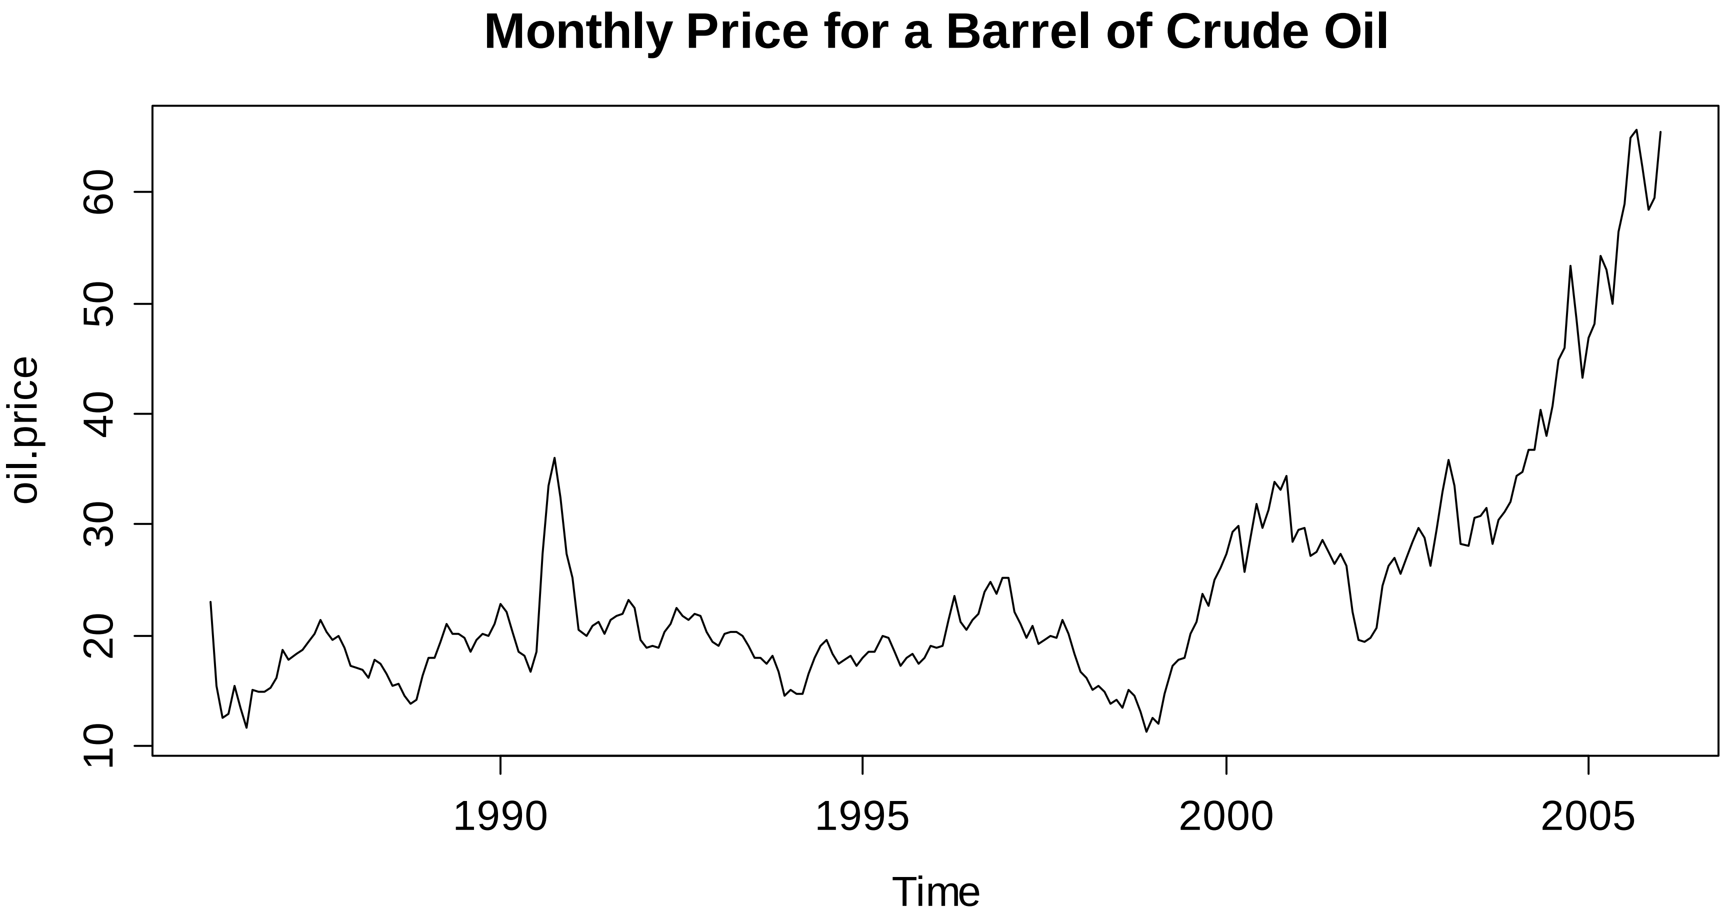

Taking the Logarithm is Key

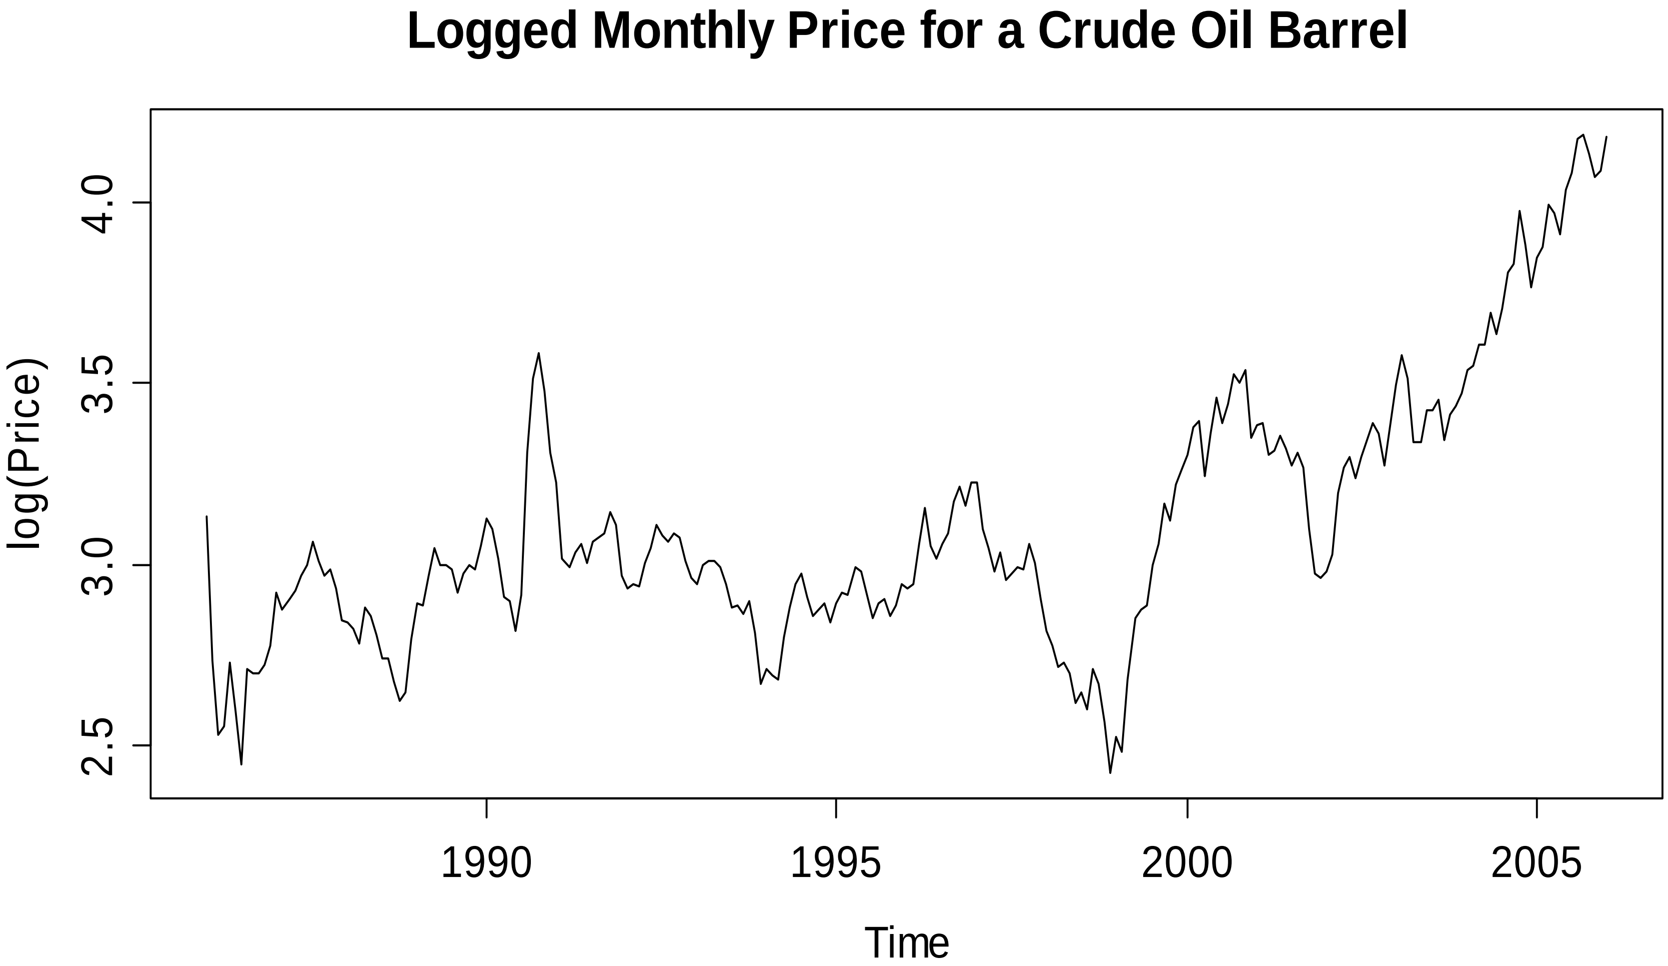

Applied Time Series Analysis
FS 2012 – Week 08

Differencing Yields a Stationary Series

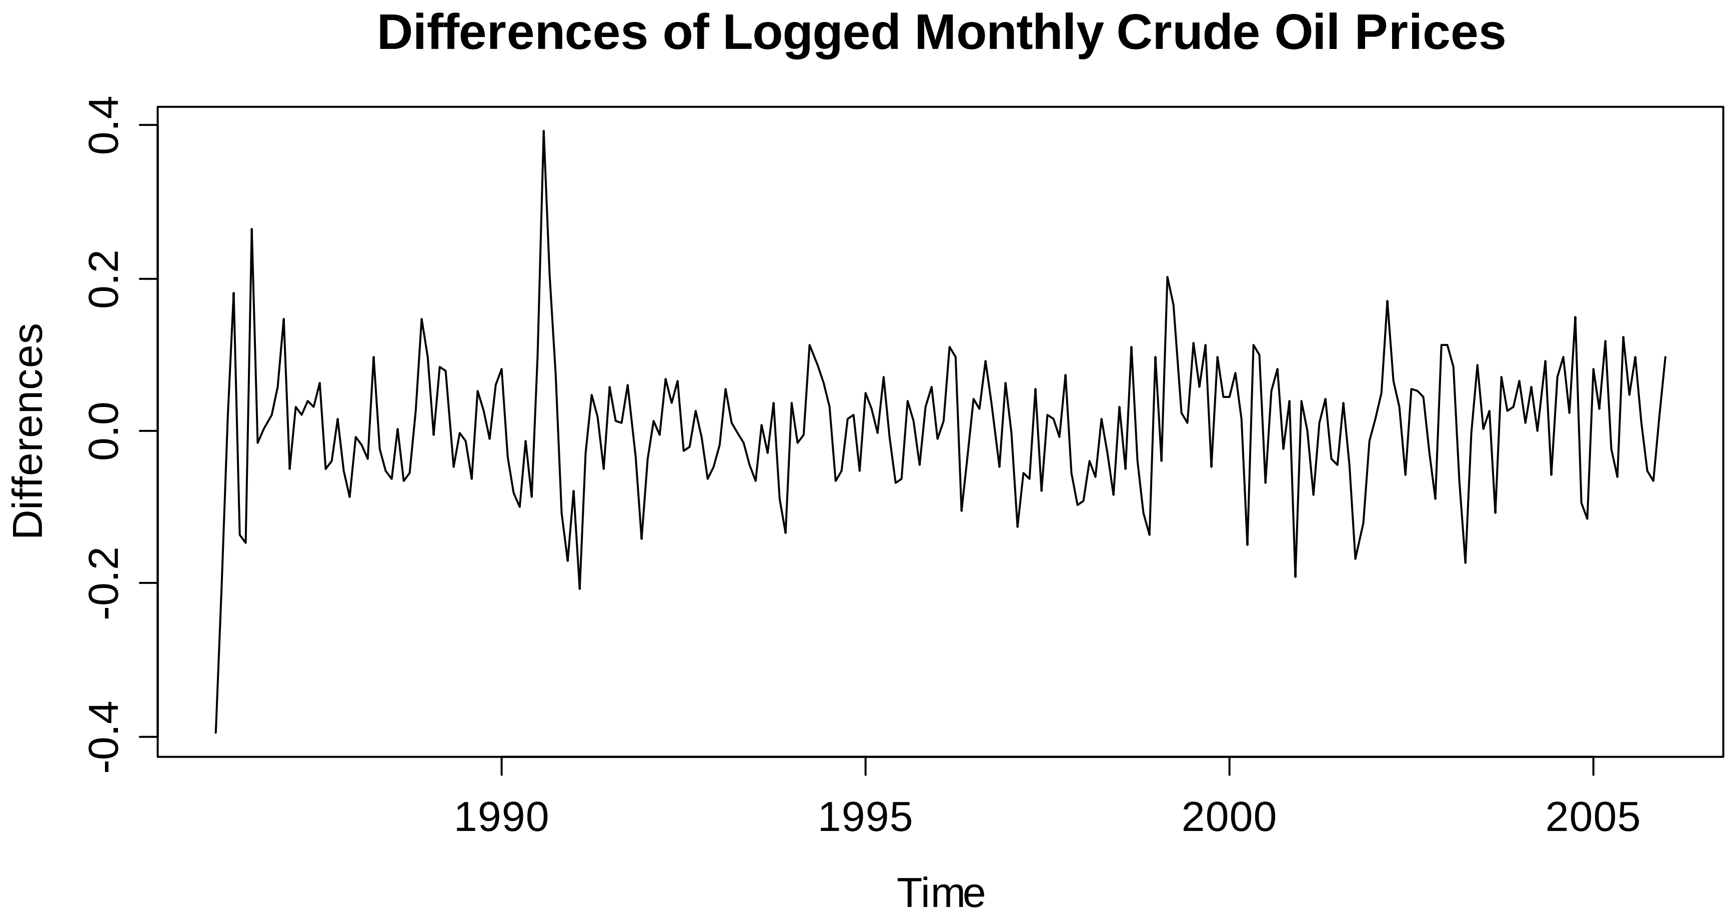

ARIMA(p,d,q)-Models

Idea: Fit an ARMA(p,q) to a time series where the d^{th} order difference with lag 1 was taken before.

Example: If $Y_t = X_t - X_{t-1} = (1-B)X_t \sim ARMA(p,q)$,
then $X_t \sim ARIMA(p,1,q)$

Notation: With backshift-operator B()

$$\Phi(B)(1-B)^d X_t = \Theta(B)E_t$$

Stationarity: ARIMA-models are usually non-stationary!

Advantage: it‘s easier to forecast in R!

Applied Time Series Analysis
FS 2012 – Week 08

ACF/PACF of the Differenced Series

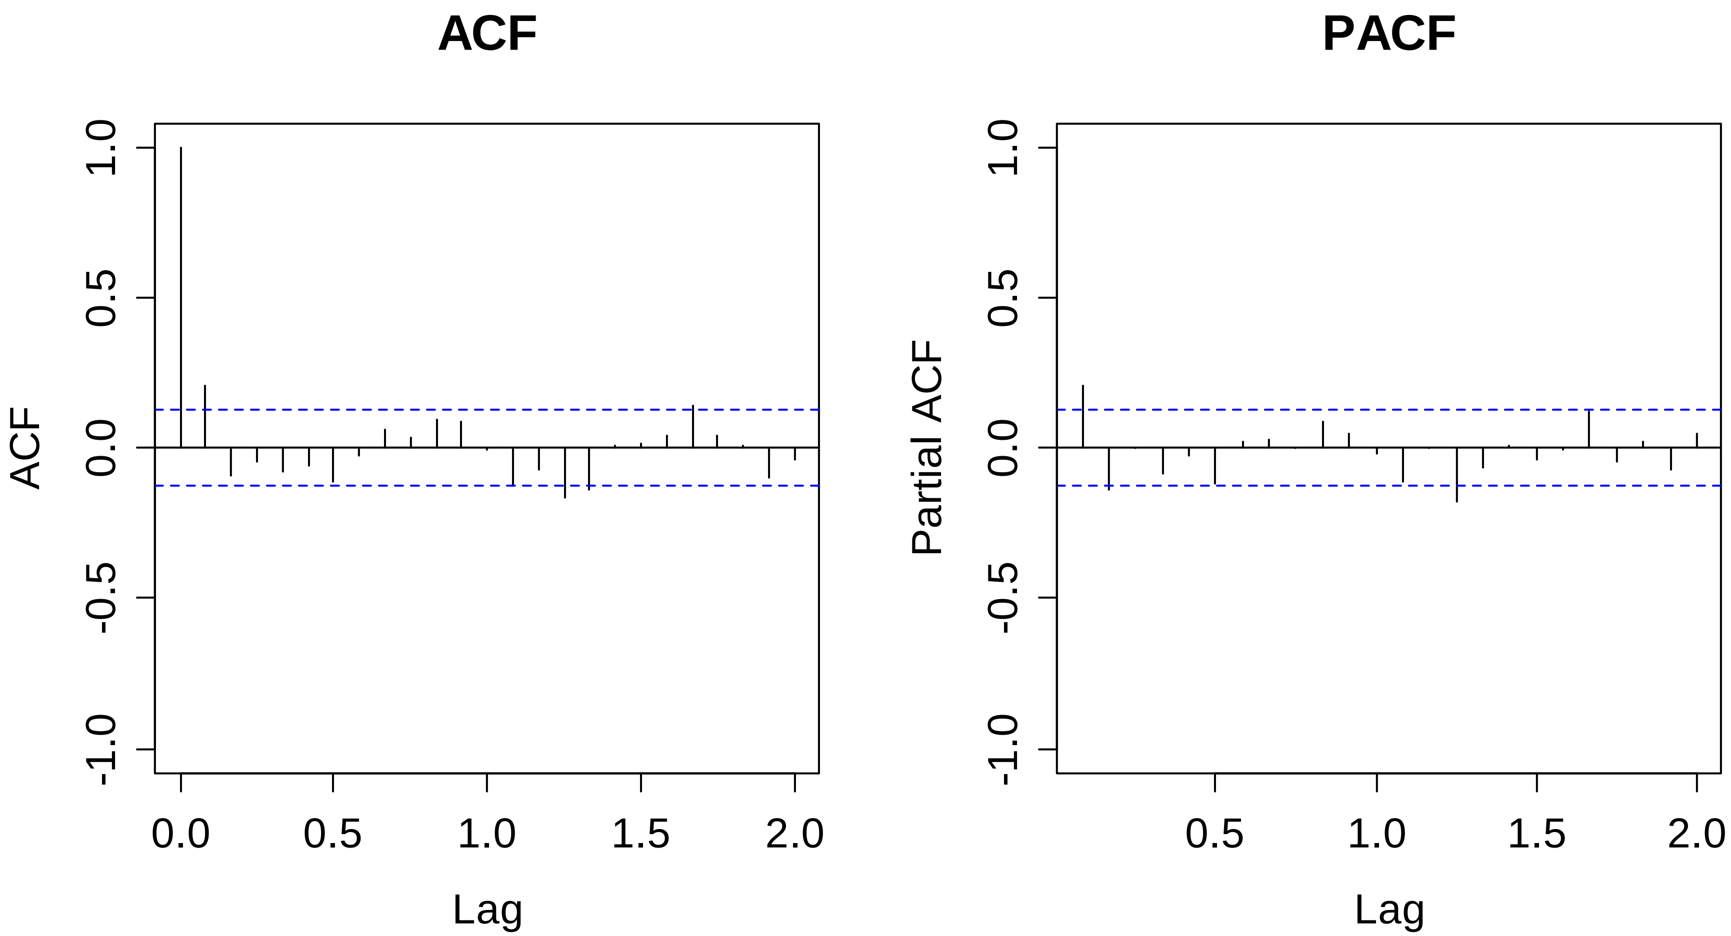

Fitting an ARIMA in R

We start by fitting an ARIMA(1,1,2) to the oil series:

```
> arima(lop, order=c(1,1,2))

Call:
arima(x = lop, order = c(1, 1, 2))

Coefficients:
         ar1      ma1      ma2
      0.8429  -0.5730  -0.3104
s.e.  0.1548   0.1594   0.0675

sigma^2 = 0.0066:  ll = 261.88,  aic = -515.75
```

Alternative Fitting

Instead of fitting an ARIMA(1,1,2) to the logged oil series, we can also take the differenced log-oil series and fit an ARMA(1,2) to it.

IMPORTANT:

In this case, we have to do fitting without including an intercept (why?), thus:

```
> arima(diff(log(oil.price)), order=c(1,0,2),
        include.mean=FALSE)
```

Meaning of the Model / Recipe

We can rewrite the ARIMA(1,1,2) model as an ARMA(2,2), **see blackboard...**

Some guidelines on how to fit ARIMA models to observed time series can also be found **on the blackboard...**

Applied Time Series Analysis
FS 2012 – Week 08

Residual Analysis of the ARIMA(1,1,2)

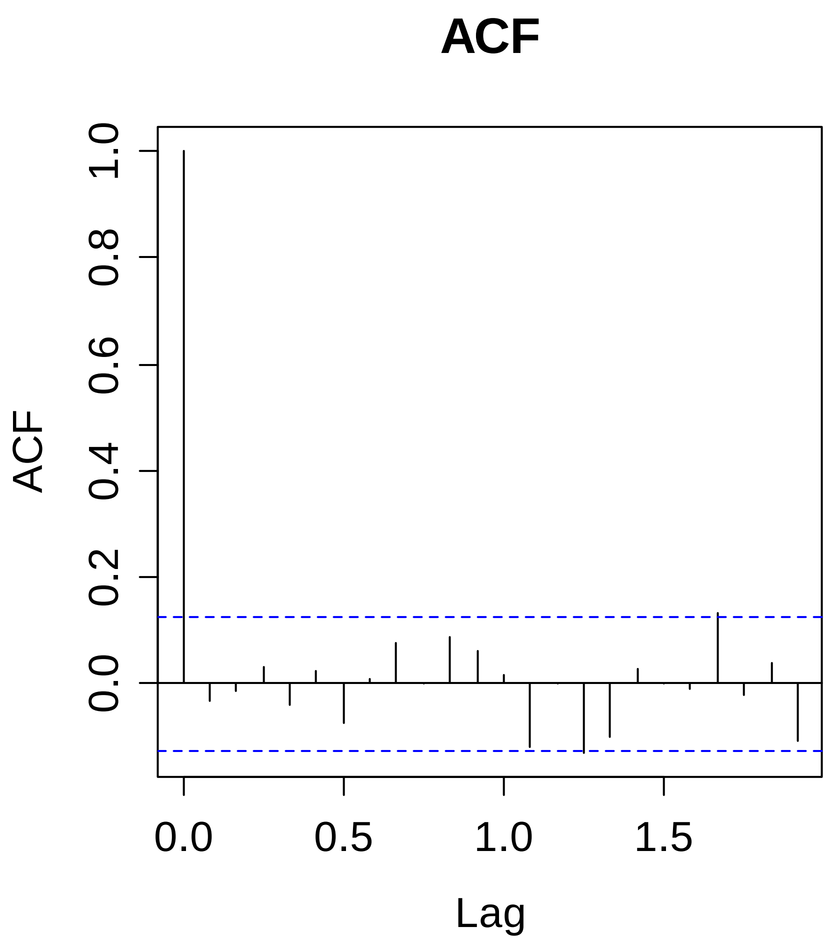

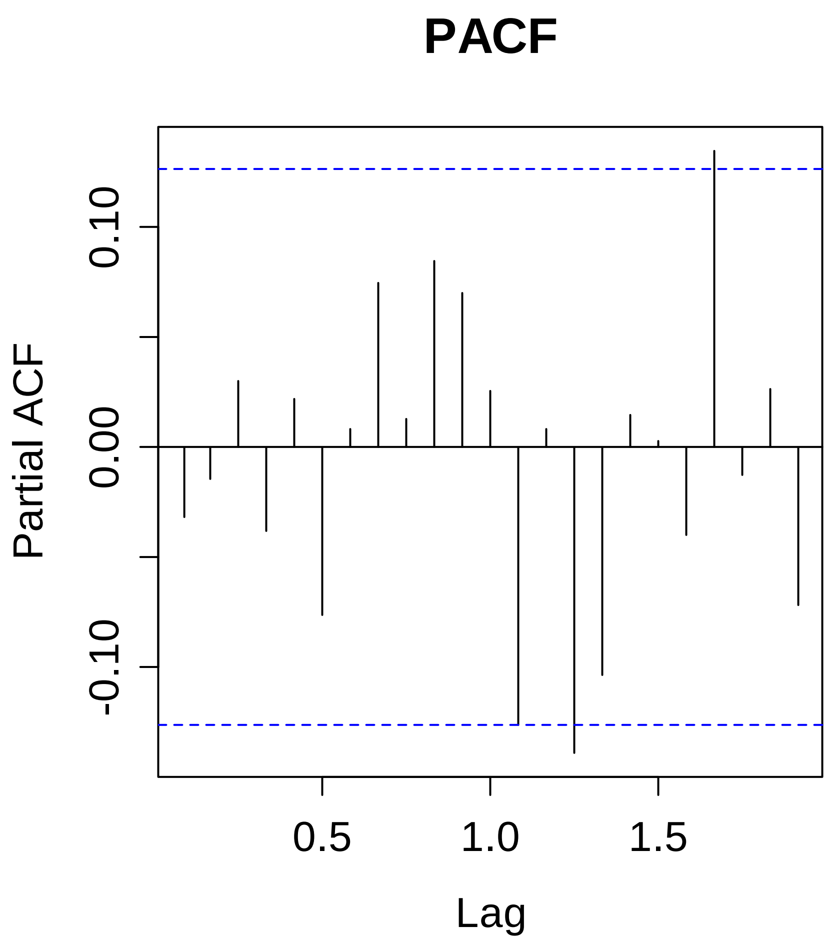

SARIMA(p,d,q)(P,D,Q)^s

= a.k.a. Airline Model. We are looking at the log-trsf. airline data

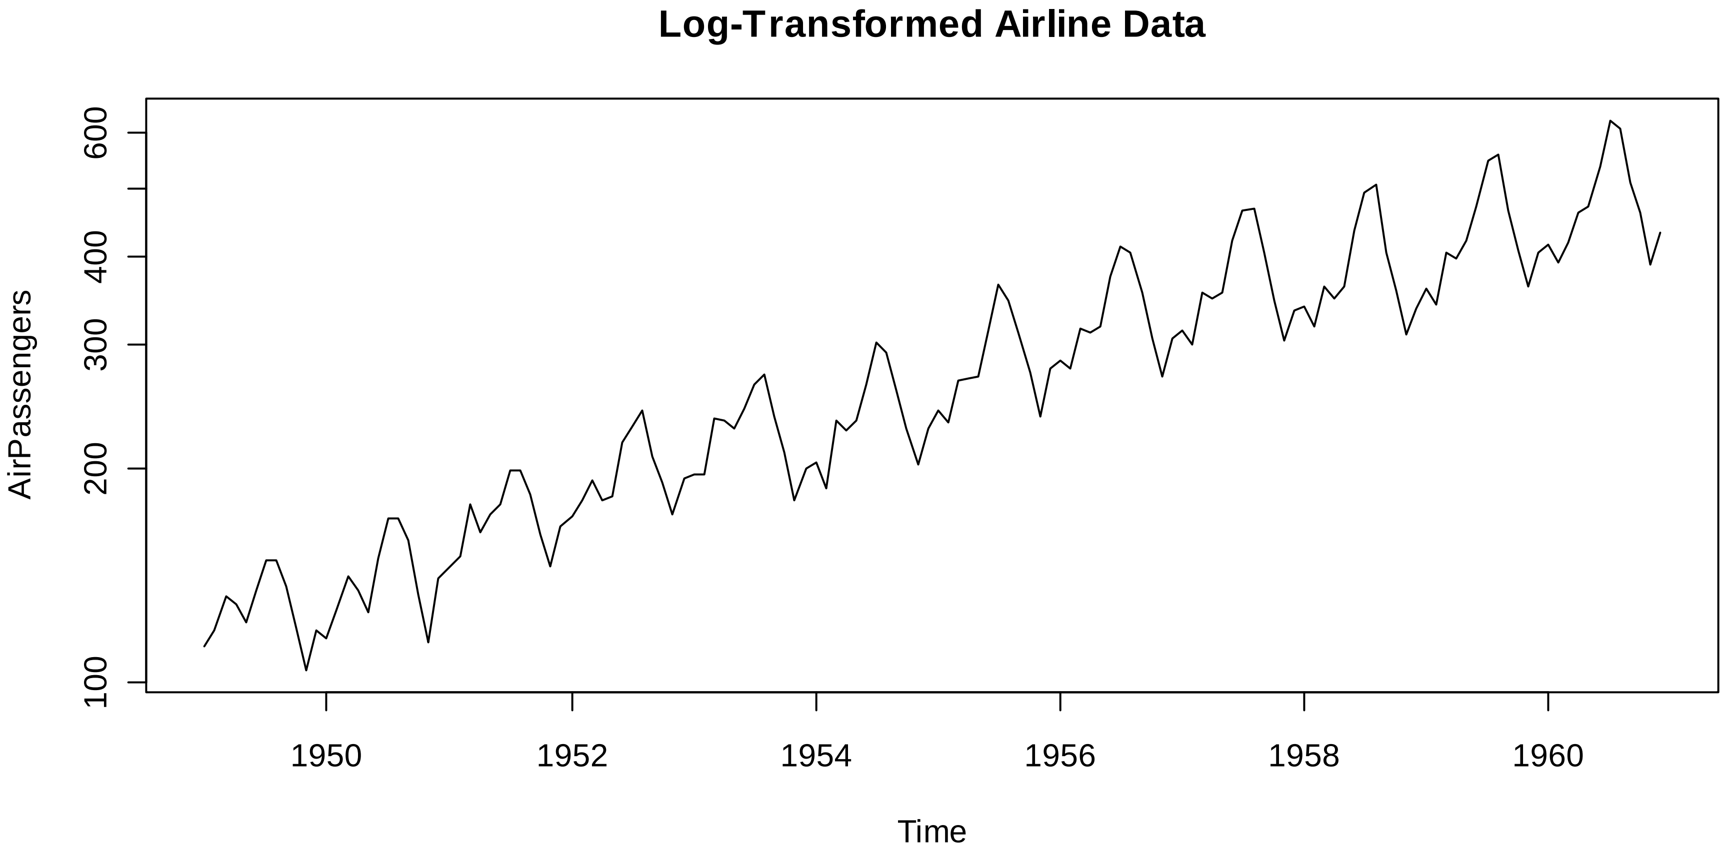

Seasonal Differencing Helps…

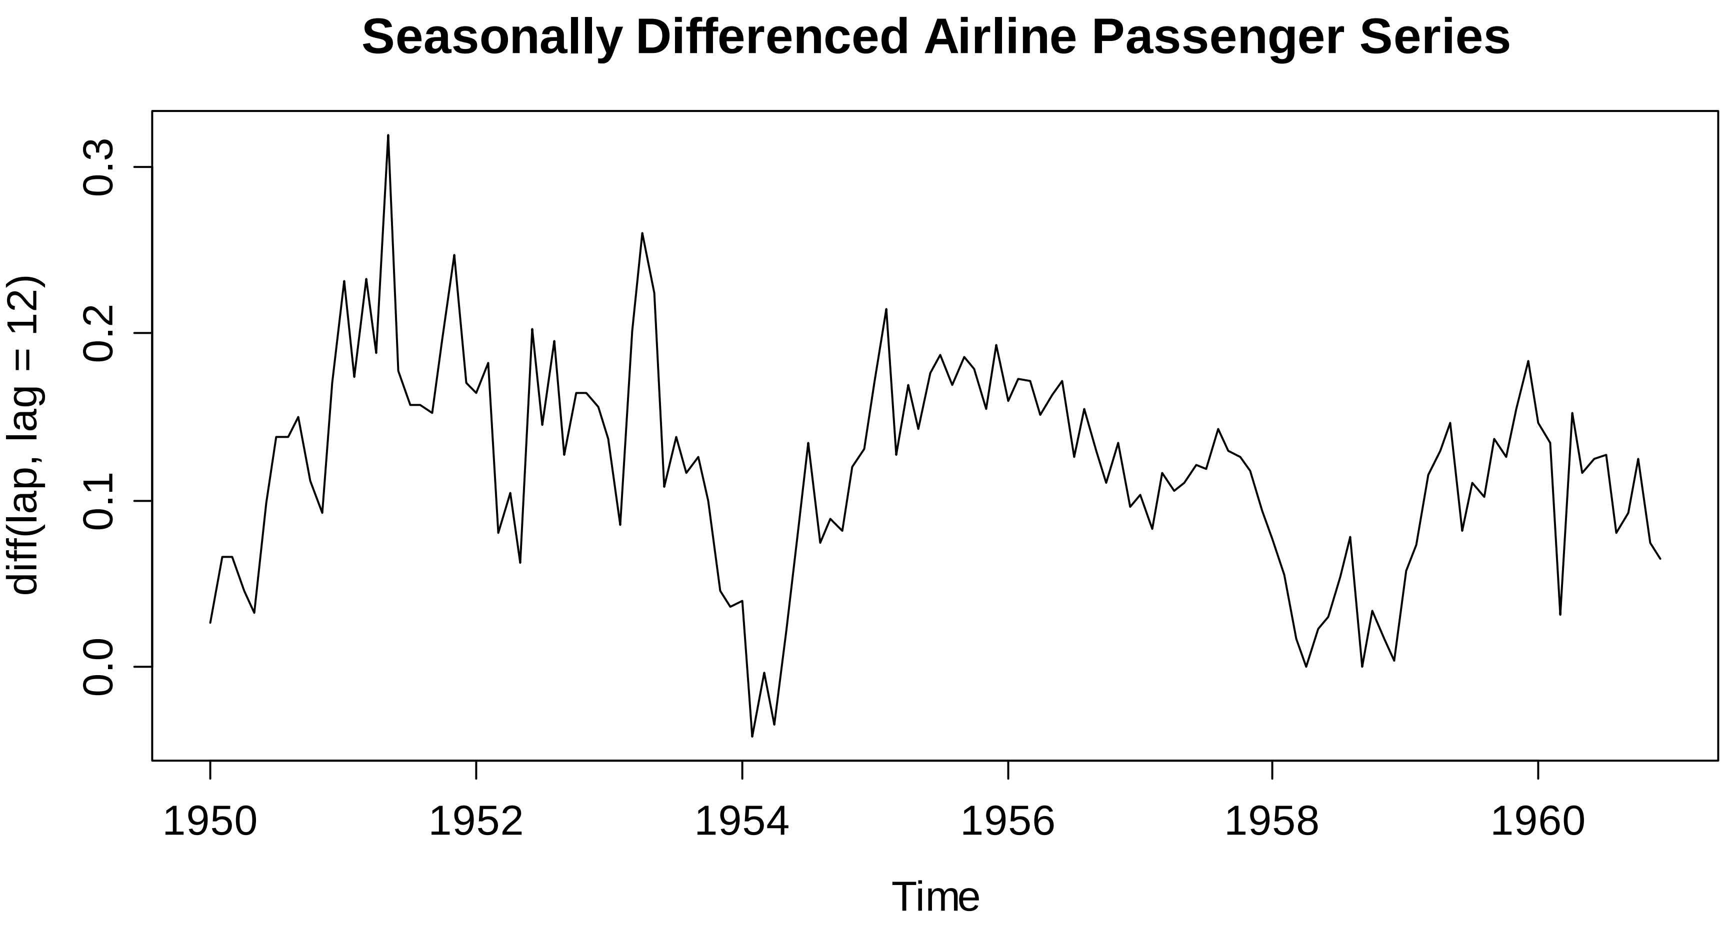

Applied Time Series Analysis
FS 2012 – Week 08

… But More Is Needed!

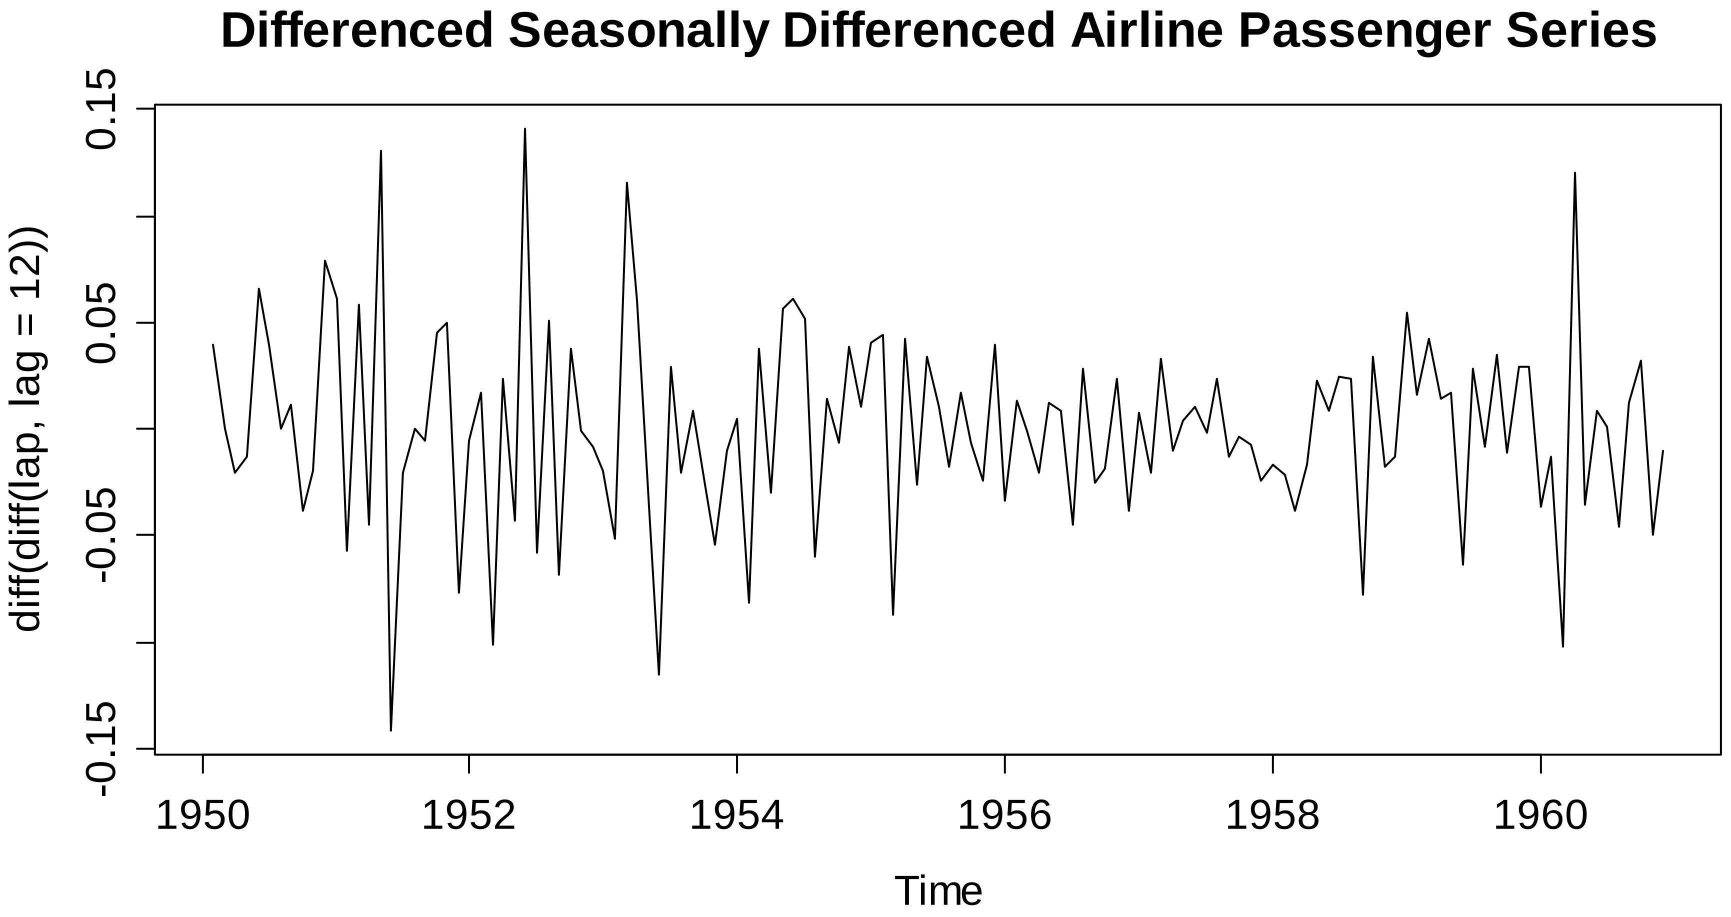

Applied Time Series Analysis
FS 2012 – Week 08

SARIMA(p,d,q)(P,D,Q)^s

We perform some differencing… (→ **see blackboard**)

ACF/PACF of SARIMA(p,d,q)(P,D,Q)^s

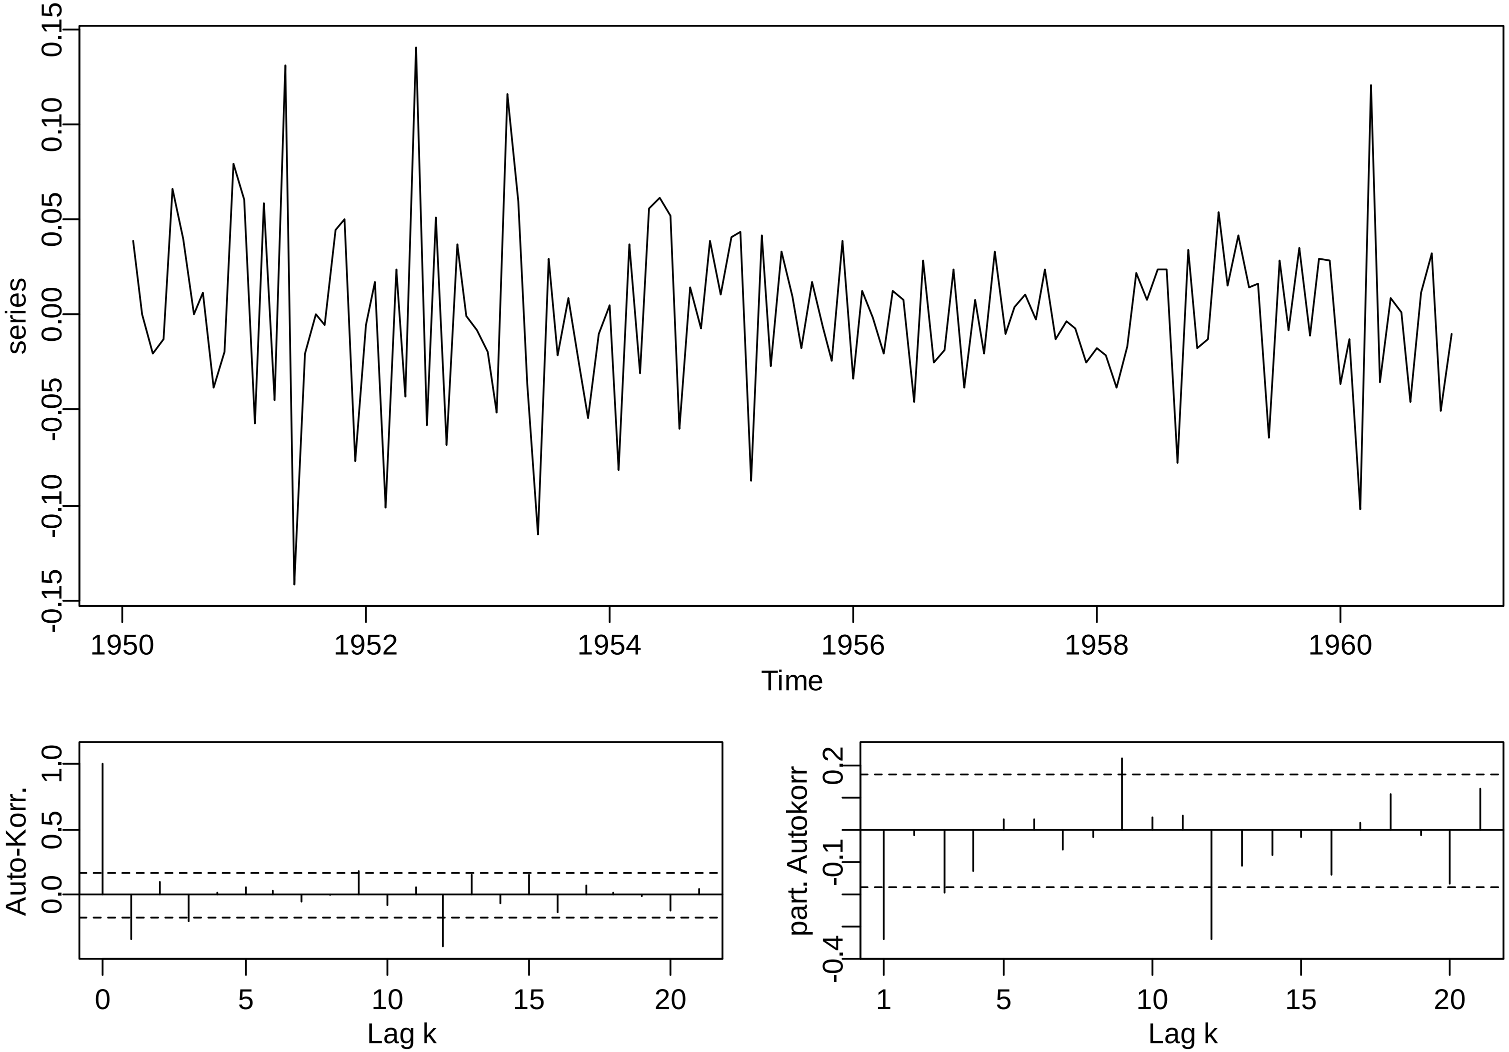

Modeling the Airline Data

Since there are "big gaps" in ACF/PACF:

$$Z_t = (1+\beta_1 B)(1+\gamma_1 B^{12})E_t$$
$$= E_t + \beta_1 E_{t-1} + \gamma_1 E_{t-12} + \beta_1 \gamma_1 E_{t-13}$$

This is an MA(13)-model with many coefficients equal to 0, or equivalently, a SARIMA$(0,1,1)(0,1,1)^{12}$.

Note: **Every SARIMA$(p,d,q)(P,D,Q)^s$ can be written as an ARMA(p+sP,q+sQ), where many coefficients will be equal to 0.**

*SARIMA(p,d,q)(P,D,Q)^s^*

The general notation is:

$$Z_t = (1-B)^d (1-B^s)^D X_t$$

$$\Phi(B)\Phi_s(B^s)Z_t = \Theta(B)\Theta_s(B^s)E_t$$

Interpretation:

- one typically chooses d=D=1
- s = periodicity in the data (season)
- P,Q describe the dependency on multiples of the period

→ **see blackboard...**

Forecasting Airline Data

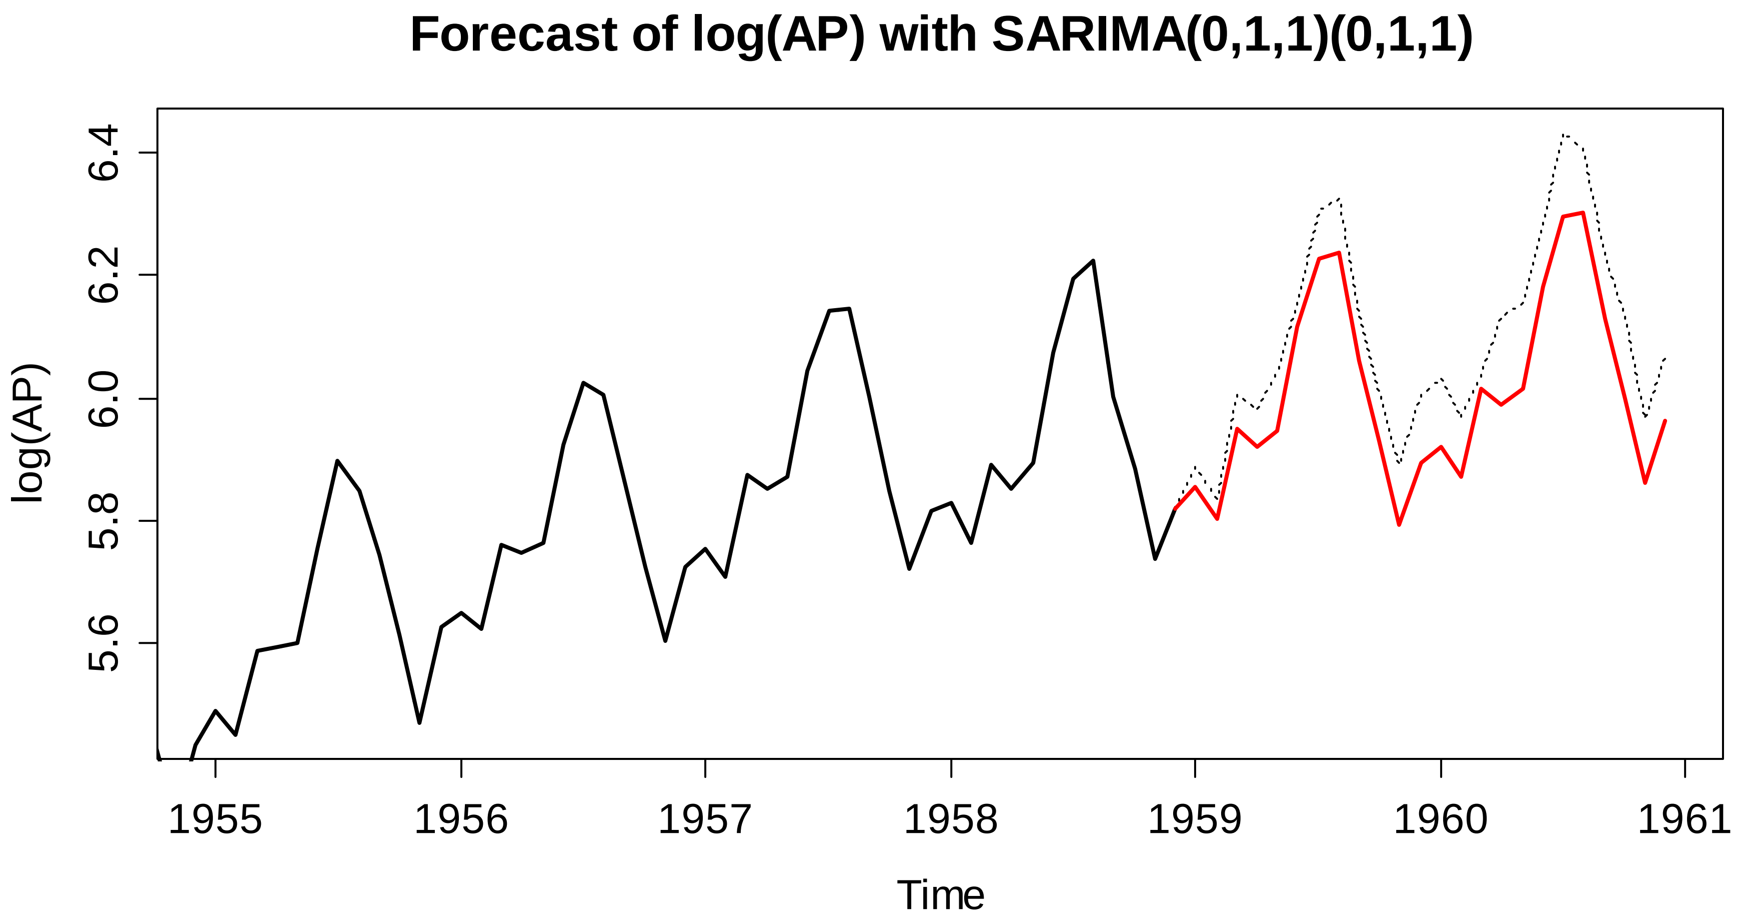

Residual Analysis of SARIMA(0,1,1)(0,1,1)

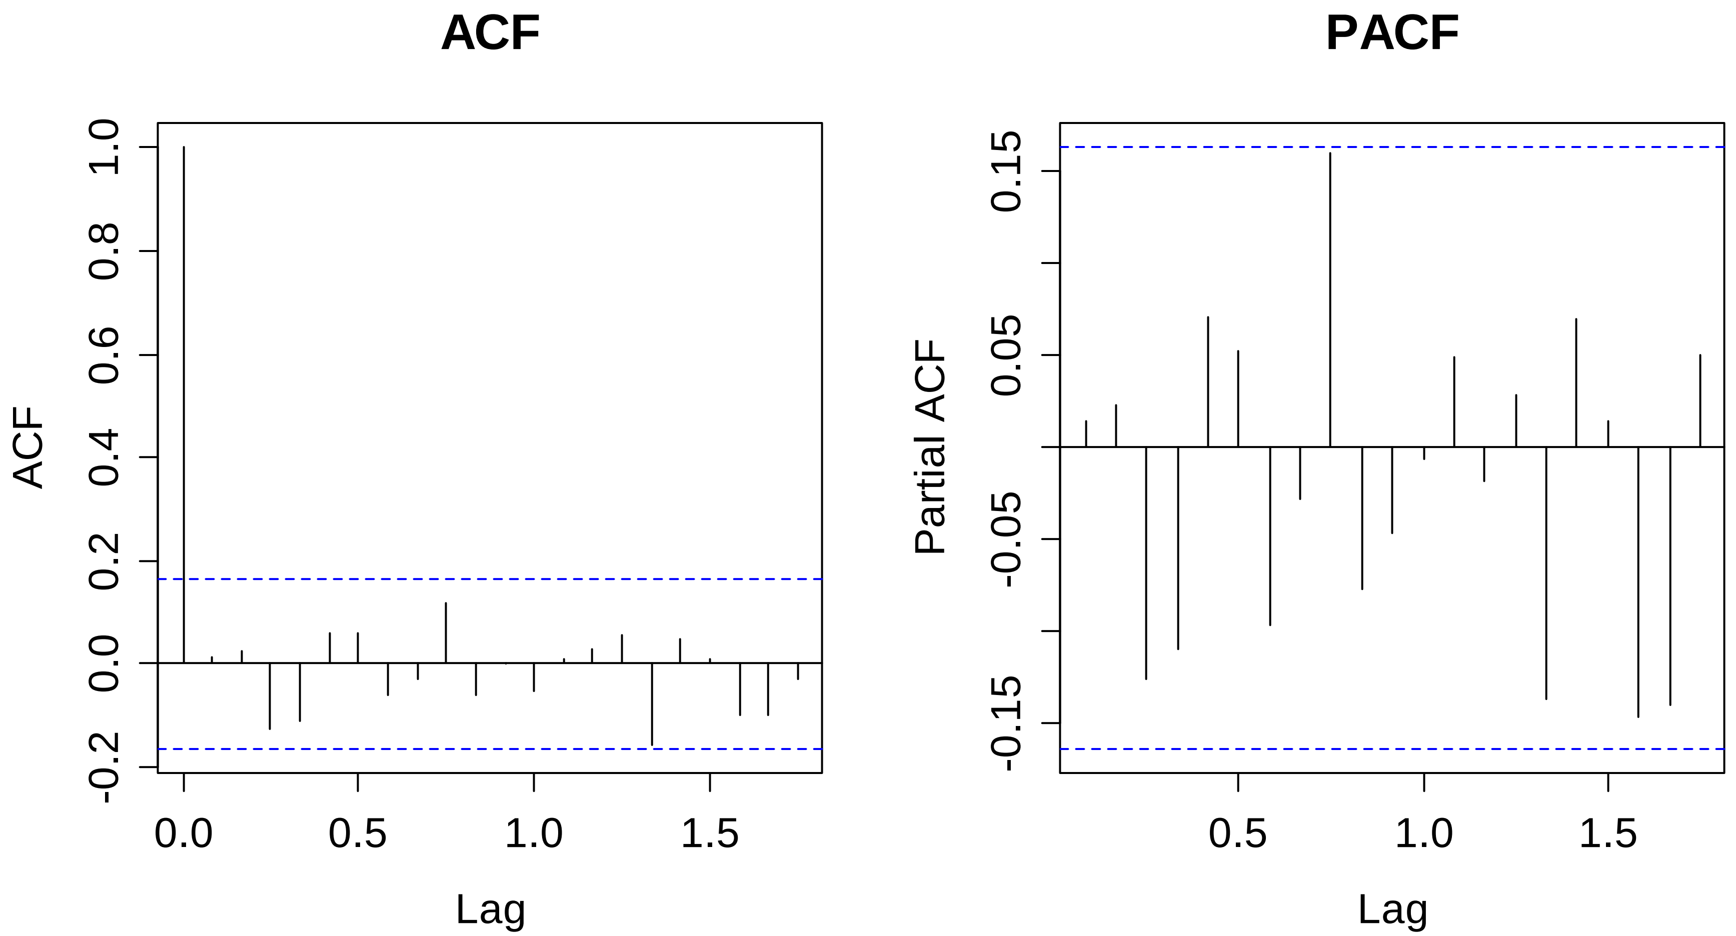

Outlook to Non-Linear Models

What are linear models?

Models which can be written as a linear combination of X_t
i.e. all AR-, MA- and ARMA-models

What are non-linear models?

Everything else, e.g. non-linear combinations of X_t,
terms like X_t^2 in the linear combination, and much more!

Motivation for non-linear models?

- modeling cyclic behavior with quicker increase then decrease
- non-constant variance, even after transforming the series

Applied Time Series Analysis
FS 2012 – Week 08

SMI Log-Returns

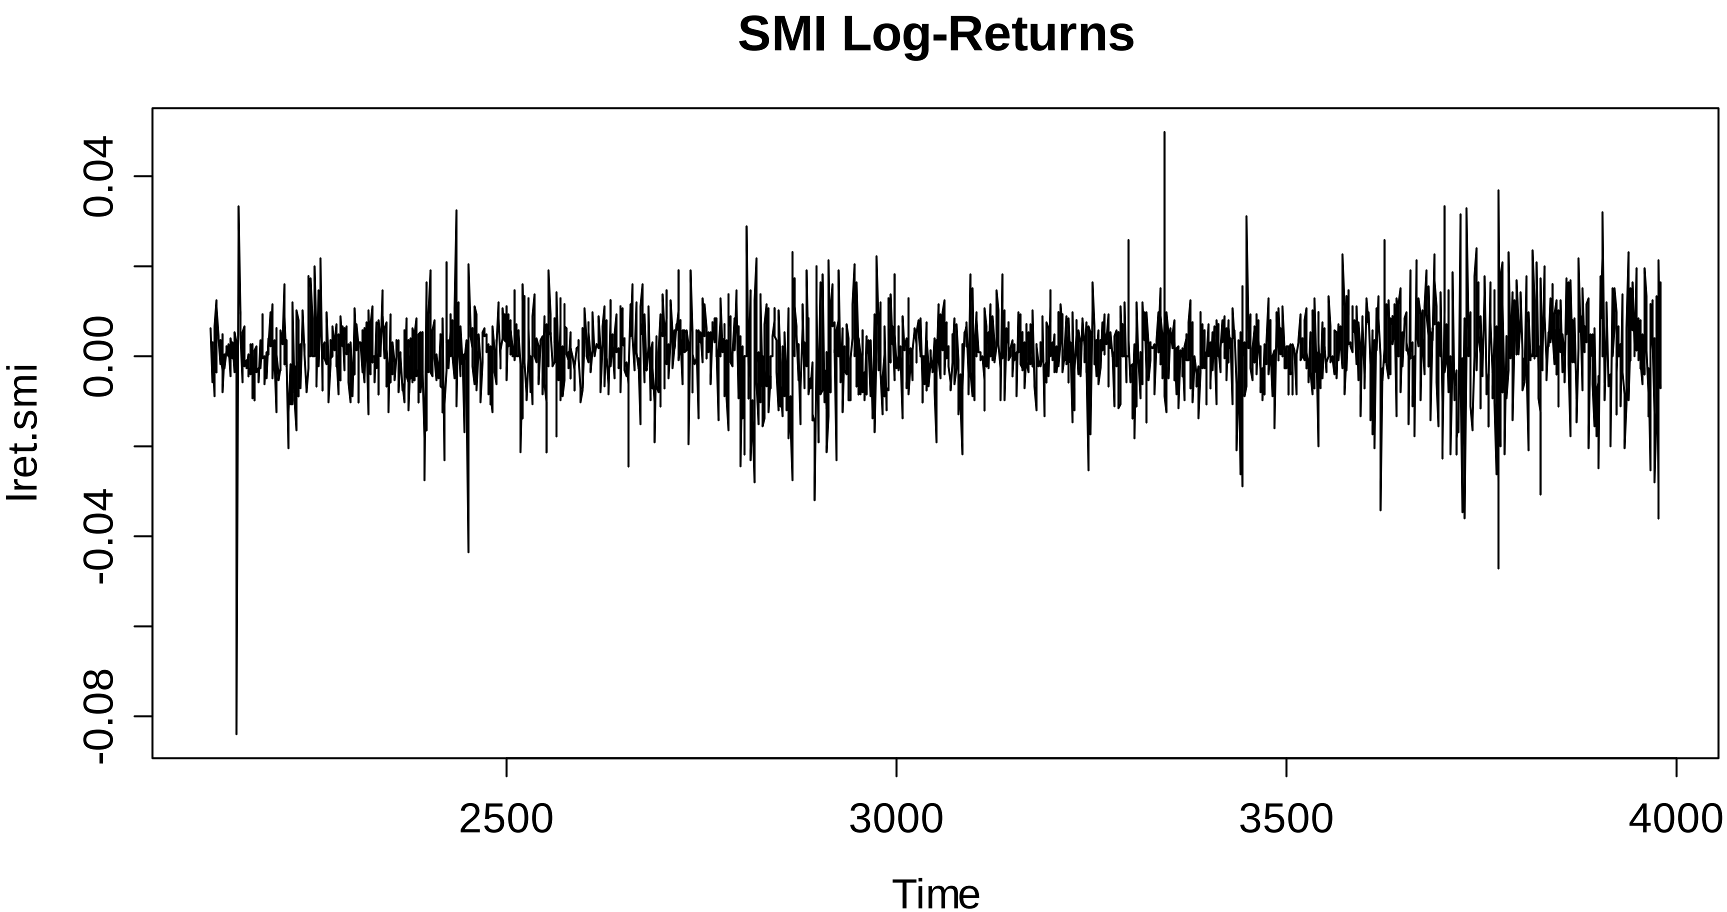

Applied Time Series Analysis
FS 2012 – Week 08

Normal Plot of SMI Log-Returns

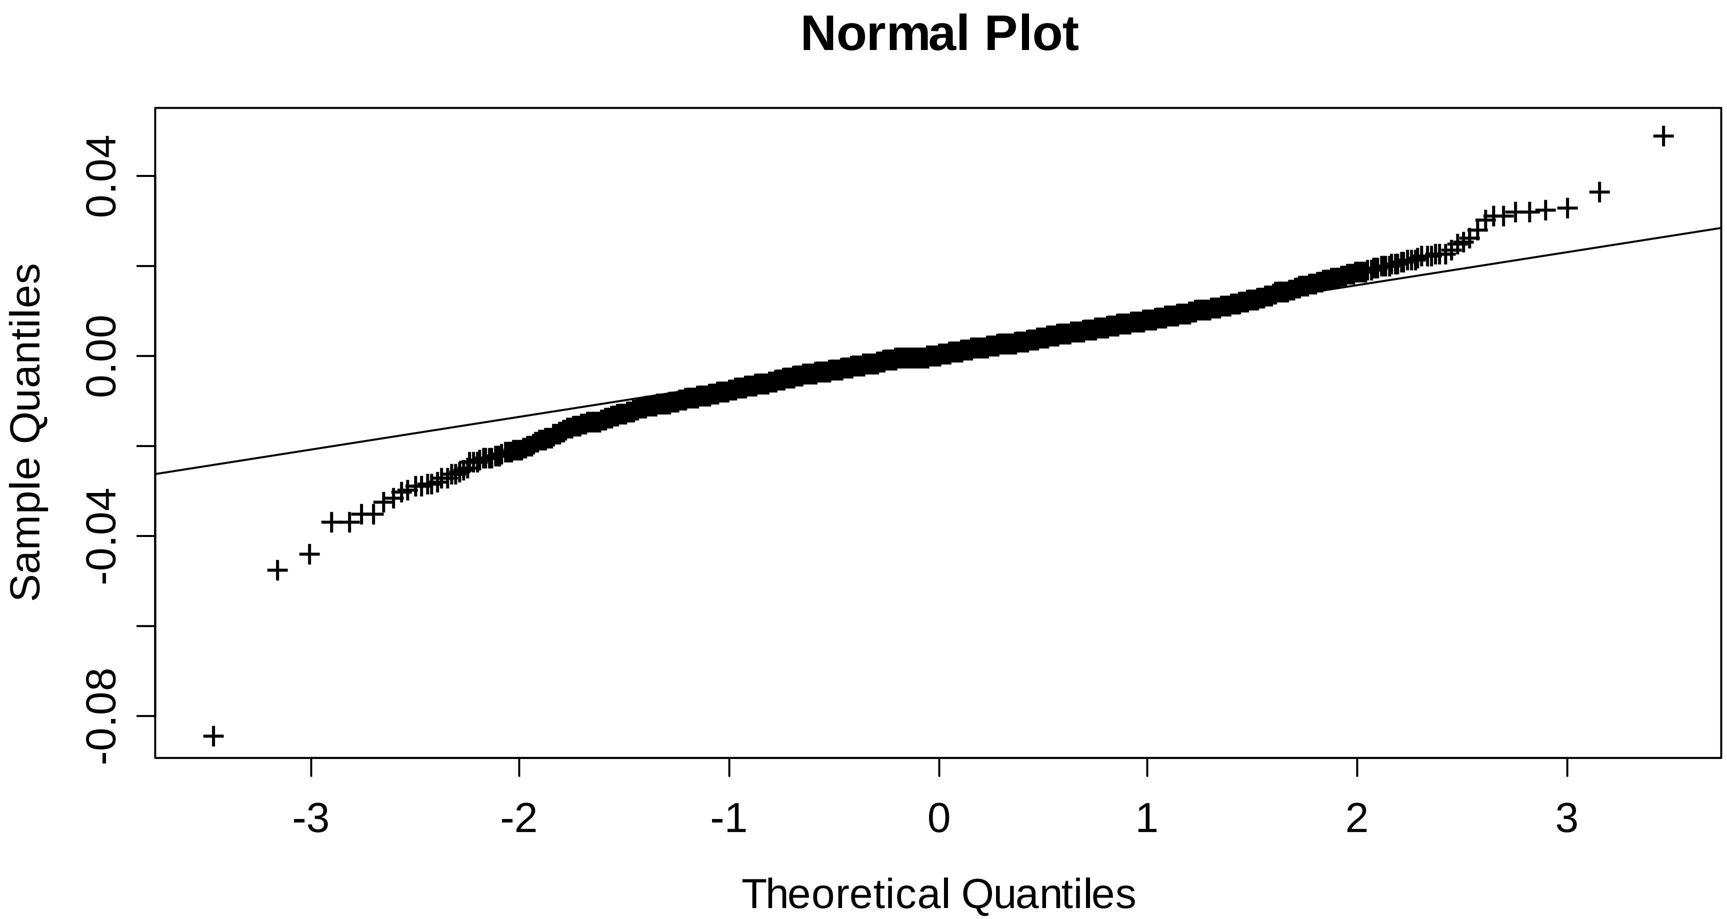

ACF of SMI Log-Returns

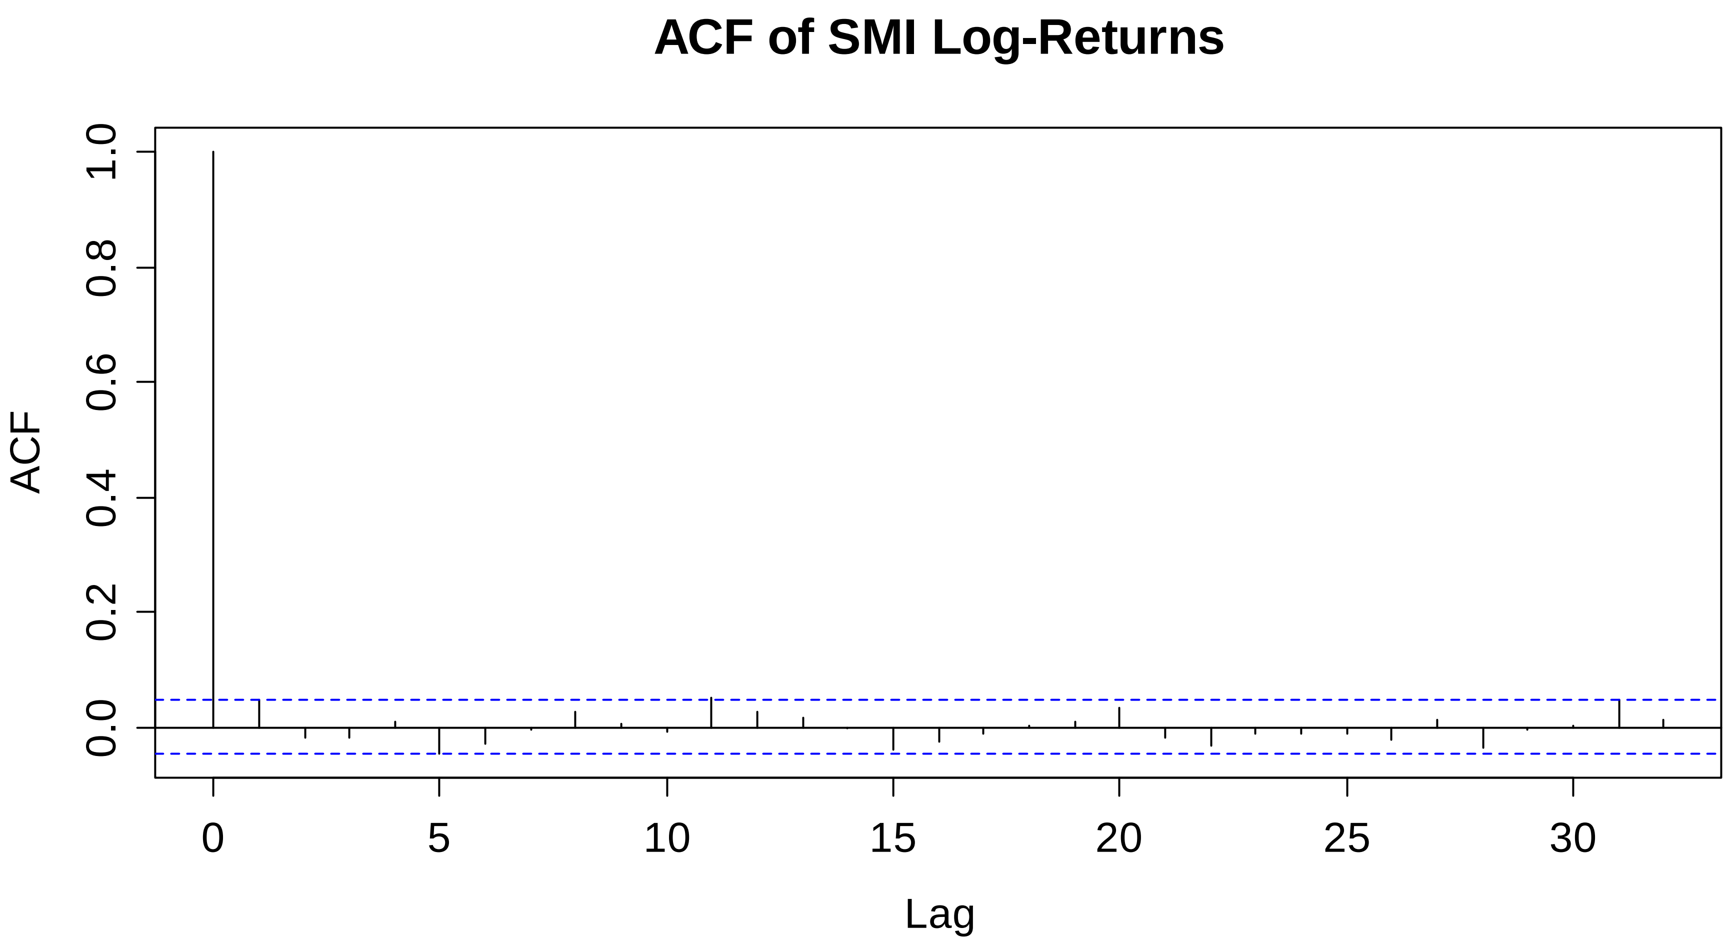

ACF of of Squared SMI Log-Returns

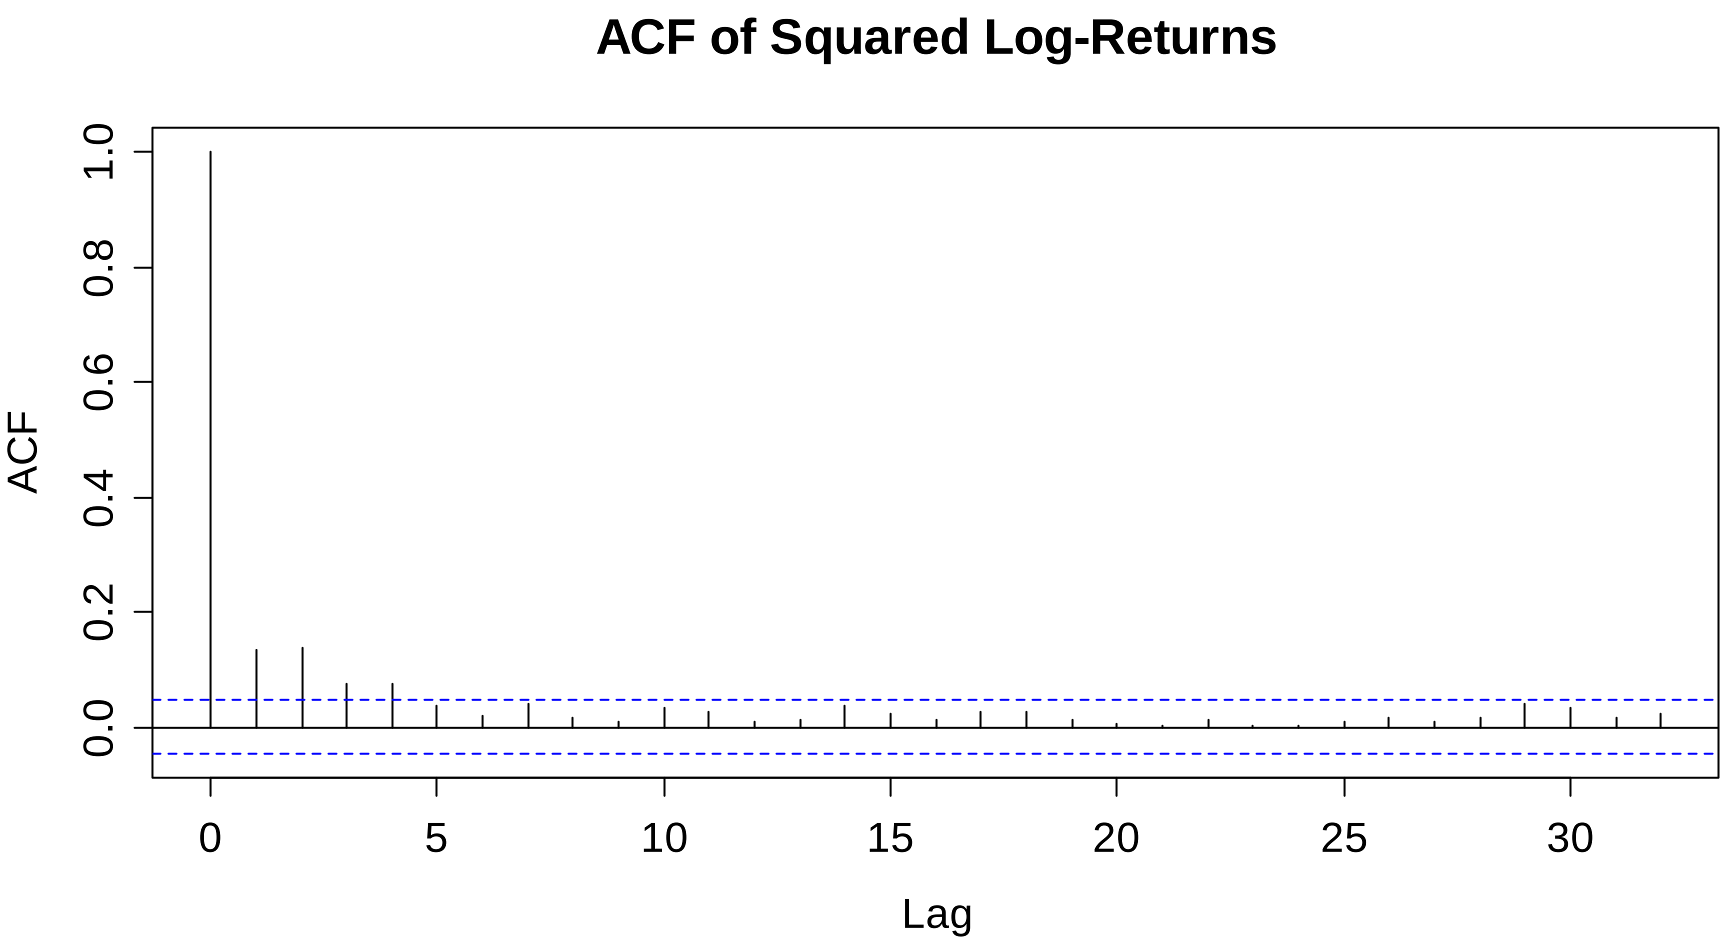

Applied Time Series Analysis
FS 2012 – Week 08

The ARCH / GARCH Model

→ **See blackboard...**

Applied Time Series Analysis
FS 2012 – Week 08

Model Choice

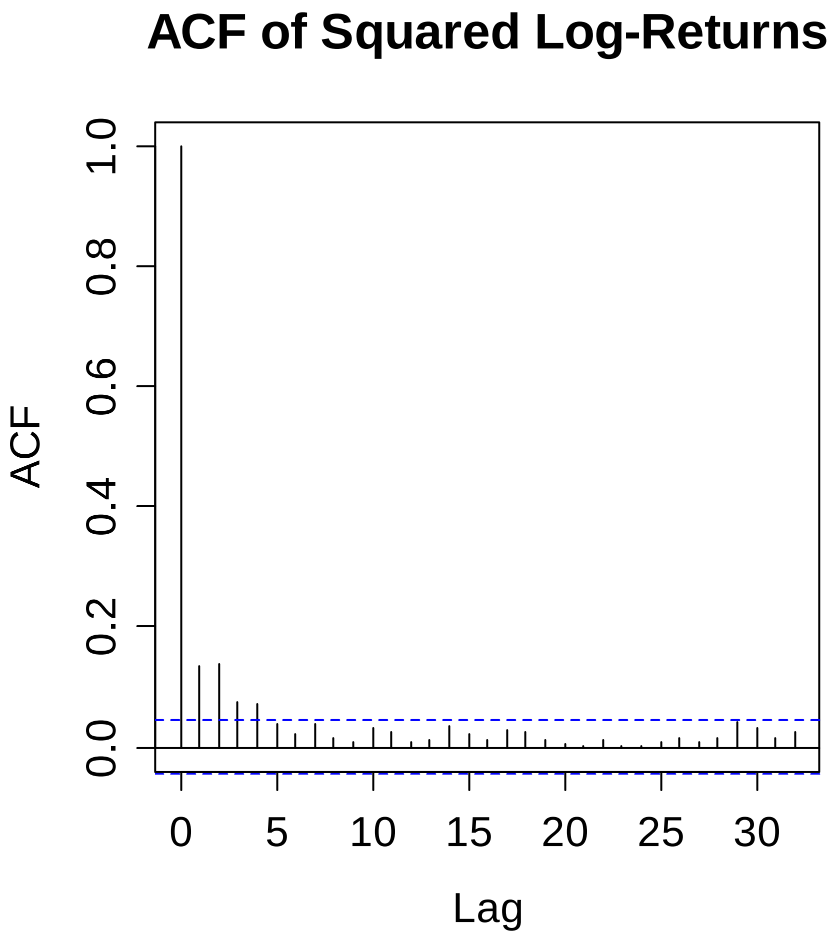

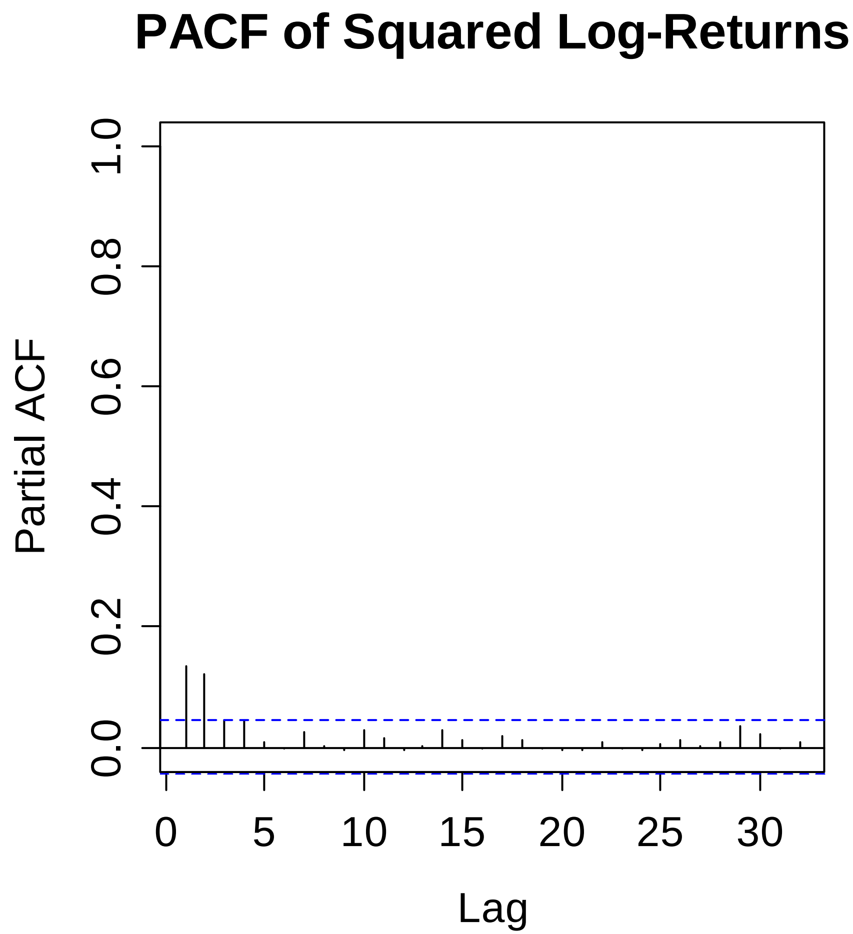

Fitting an ARCH(2) Model

R allows for convenient fitting...

```
> fit <- garch(lret.smi, order = c(0,2))
> fit

Call: garch(x = lret.smi, order = c(0, 2))

Coefficient(s):
       a0          a1          a2
6.568e-05  1.309e-01  1.074e-01
```