

JBS Master in Finance Econometrics Module

Lent 2010

Thilo Klein

<http://thiloklein.de>

**Computer Lab Sessions 2&3
ARIMA, ARCH and GARCH Models**

Contents

Exercise 1. Estimation of a quarterly ARMA model of the US Producer Price Index (PPI).....	2
Exercise 2. ARMA model selection.....	3
Exercise 3. Seasonality and ARIMA modeling	4
Exercise 4. GARCH model of US PPI.....	4
Exercise 5. GARCH-M	5
Exercise 6. GARCH and TARCH models	6
Exercise 7. GARCH models.....	6

Exercise 1. Estimation of a quarterly ARMA model of the US Producer Price Index (PPI)

Use `quarterly.csv`.

- Plot the PPI in level and first difference. Are these sequences stationary?
- Take logs and first difference of PPI (the resulting variable is a measure of what?). Plot again and compute ACF and PACF and do ADF test. Why is it important to take logs? Has the variable a unit root?
- What kind of model suggests the ACF and PACF? What happens with the PAC and PACF at lag 4? What could be the explanation?
- Estimate an ARMA(p,q), for the following cases (p,q)=(1,0), (2,0), (1,1), (1,(1,4)) ((1,4) means a MA term at lags 1 and 4), (2,1), (1,(0,4)). For each of these models report SSR, AIC and Q(5), Q(10) and Q(15). Which is the best model? Does this model pass all the diagnostic tests (Normality, no autocorrelation, no ARCH terms)?
- Estimate the ARMA models (p,q)=(1,1) and (1,(1,4)) over the period 1960.1-1989.3. Obtain the one-step-ahead forecast (static forecast in R) and the one-step-ahead forecast error from each. Evaluate the forecast performance of these models.

Solution: p.87-93 Enders (2004).

- Generate first difference of $\log(\text{ppi})$ and plot the series
 - `LPPI <- ts(log(PPI), start=c(1960, 1), freq=4)`
 - `LPPI.d1 <- diff(LPPI, lag=1, differences=1)`
 - `par(mfrow=c(2,1)); plot(LPPI); plot(LPPI.d1)`
 Answer: The first is for sure not stationary, the second one maybe.
- The autocorrelation and partial autocorrelation decay quickly: it seems to be a stationary series:
 - `acf(LPPI.d1); pacf(LPPI.d1)`
 ADF test (without trend!!) lets us reject the null hypothesis that LPPI.d1 has a unit root. The series is stationary.
 - `library(tseries)`
 - `adf.test.1(LPPI.d1, kind=2, k=2)`
 - `adf.test.2(LPPI.d1, L=2, int=T, trend=F)`
 It is important to take logs because the resulting variable is a measure of inflation.
- It could be an AR process, since the PACF decay more quickly than the ACF. It could also be an ARMA process, since up to lag 4 the PACs are significant (and then do not decay so quickly).

Note that the autocorrelation of order 4 is measuring the correlation between observations for same quarters of different years. This is reflecting some kind of seasonality. Many economic time series tend to exhibit seasonality. For example, people spend more money in Christmas than in any other time of the year, and normally the Christmas expenditure is correlated with how much we spent last year (the expenditure in the same quarter of previous year). And of course the same happens with prices.
-

Model	SSR	AIC	Q(5), Q(10), Q(15)
AR(1)	0.021057	-1026.62	13.17, 24.34, 26.80
AR(2)	0.020138	-1031.1	8.43, 17.93, 20.99
ARMA(1,1)	0.013351	-1036.24	4.99, 14.50, 17.40

ARMA(2,1)	0.019601	-1033.63	4.93, 14.45, 17.38
ARMA(1,(1,4))	0.012312	-1038.29	2.67, 9.30, 12.25
ARMA(1,(0,4))	0.019691	-1033.07	4.84, 13.03, 15.96

The ARMA(1,(1,4)) seems to fit our data best, it has the lowest SSR, AIC and Ljung-Box statistics. Let us now run some diagnostic tests on the model residuals.

Autocorrelation. Ljung-Box statistic: Is there any evidence of autocorrelation up to lag 5, 10 and 15? No, there isn't!

- `library(ccgarch)`
- `nNA <- is.na(arma1d$res)==F`
- `ljung.box.test(arma1d$res[nNA])`

Normality. Jarque-Bera test for the null of normality: There is evidence of non-normality. We reject the null hypothesis of normal distributed residuals. We should improve our model!

- `jarque.bera.test(arma1d$res[nNA])`

ARCH terms. ARCH LM test: We reject the hypothesis of no autocorrelation of the squared residuals, i.e. there is evidence of ARCH behaviour (heteroskedasticity, the conditional variance is not constant). This could also explain the non-normality result. We will learn about how to model this phenomenon.

- `library(FinTS)`
- `ArchTest(c(arma1d$res), lags=2)`

e) Select the training sample over the period 1960.1-1989.3.

- `LPPI.d1.w <- window(LPPI.d1, start=c(1960,1), end=c(1989,3))`

Use generic function `predict()` to obtain the one-step-ahead forecast and the one-step-ahead forecast error.

- `predict(arma1e.11, n.ahead=1)`

We find that the forecast error is lower for the ARMA(1,(1,4)) model and the point forecast of this model is also closer to the actual observation:

- `window(LPPI.d1, c(1989,4), c(1989,4))`

Exercise 2. ARMA model selection

Use `arima.csv`. Choose the best model for each variable (y_1, \dots, y_7) applying the modified Box-Jenkins methodology. Do it step by step reporting the results at each step.

Solution:

For	The best model is:	Original process:
y_1	AR(1)	$y_t = 0.9 \cdot y_{t-1} + \varepsilon_t$
y_2	ARMA(1)	$y_t = -0.9 \cdot y_{t-1} + \varepsilon_t$
y_3	AR(2)	$y_t = 0.9 \cdot y_{t-1} - 0.2 \cdot y_{t-2} + \varepsilon_t$
y_4	White noise	$y_t = \varepsilon_t$
y_5	MA(1)	$y_t = \varepsilon_t + 0.9 \cdot \varepsilon_{t-1}$
y_6	MA(2)	$y_t = \varepsilon_t + 0.9 \cdot \varepsilon_{t-1} + 0.8 \cdot \varepsilon_{t-2}$
y_7	ARMA(1,1)	$y_t = 0.9 \cdot y_{t-1} + \varepsilon_t + 0.8 \cdot \varepsilon_{t-1}$

Exercise 3. Seasonality and ARIMA modeling

Use quarterly.csv, we want to find the best seasonal ARIMA model for money (money defined as M1) (M1NSA series).

- a) Plot the M1NSA in level and first difference. Are these sequences stationary?
- b) Take logs and first difference of M1NSA (the resulting variable is a measure of money growth). Plot again and compute ACF and PACF. What happens at lags 4, 8, 12, etc?
- c) Take a seasonal difference (over the first differenced series) and compute again ACF and PACF, what can you observe now?
- d) What kind of model suggests the ACF and PACF?
- e) Estimate a SARIMA(1,1,0)(0,1,1) and a SARIMA(0,1,1)(0,1,1). What is the best model?

(Important: sometimes it is not necessary the regular difference, it will be enough with a seasonal difference; so as rule of thumb first take the seasonal difference and then first regular difference only if needed.)

Solution: p.95-99 Enders (2004).

(D4DLM1NSA is the seasonal difference of the first difference of log(M1NSA))

Exercise 4. GARCH model of US PPI

In exercise 1, we saw that there is evidence of conditional heteroscedasticity after fitting an ARIMA model to this data. Let's try now to model the variance of the process (use quarterly.csv).

- a) Formally test for ARCH errors (using the residuals from the ARIMA model).
- b) Use the ACF and PACF of the squared residuals as an indication of the order of the GARCH process. How many ARCH terms seem to be needed? Estimate the model.
- c) Test again for remaining ARCH terms (and compute again ACF and PACF). What can you conclude? Observe carefully the estimated coefficients, what problems do you identify?
- d) Estimate now a GARCH(1,1), do you still have the same problems? Tabulate ACF and PACF and test for autocorrelation up to lag 4 in squared residuals.
- e) Produce one-step-ahead forecast with this model.

Solution: See also Enders (2004) p. 123-126.

- a) First estimate an ARMA(1,(1,4)) for the mean (from the lab session we know that this is the best model for this time series).
 - `arma4a <- arma(LPPI.d1, lag=list(ar=1,ma=c(1,4))); summary(arma1d)`
 - Plot the residuals of this estimation:
 - `plot(arma4a$res)`
 - It seems that the residuals have a non-constant variance.
 - Formally test it with an ARCH LM test with null hypothesis H0: No autocorrelation in squared residuals up to order 4.
 - `library(FinTS)`
 - `ArchTest(c(arma4a$res), lags=4)`

The test statistic $n \cdot R^2$ is Chi-squared with 4 degrees of freedom. We reject the null. Then proceed to estimate an ARCH model for the variance.

- b) As for the mean, the ACF and PACF will give you information about what kind of GARCH(p,q) could be the best model, i.e. they could indicate the p and the q. (remember GARCH models are in some sense ARMA models for the variance). Let us try an ARCH(4) model for the variance. To do this, we first need to install the `rgarch` package (follow the description on my website). Then load the package, specify mean and variance equation, and run the joint mean and variance model.

➤ `library(rgarch)`

➤ `spec <- ugarchspec(variance.model = list(model = "fGARCH", submodel = "GARCH", garchOrder = c(4,0)), mean.model = list(armaOrder = c(1,4), include.mean = F), fixed.pars = list(ma2 = 0 ,ma3 = 0))`

➤ `sgarch.fit <- ugarchfit(data=c(LPPI.d1), spec = spec); sgarch.fit`

- c) MA(4) term is not significant and ARCH effects still remain.

- d) Estimate GARCH(1,1) model with mean equation ARMA(1,(1,4))

➤ `spec <- ugarchspec(variance.model = list(model = "fGARCH", submodel = "GARCH", garchOrder = c(1,1)), mean.model = list(armaOrder = c(1,4), include.mean = T), fixed.pars = list(ma2 = 0,ma3 = 0))`

Verify that the correlogram of squared residuals is clean. Do ARCH LM test again.

➤ `ArchTest(sgarch.fit@fit$resid, lags=4)`

It could be a problem at lag 4. But we can't reject the null of no autocorrelation in the squared residuals up to lag 4. We do not proceed further here. You should improve the model to clean the arch effects further.

- e) A one-step-ahead forecast can be had by setting `n.ahead` to 1. Let us see what happens with 200 steps ahead.

➤ `ugarchforecast(sgarch.fit, n.ahead=200)`

As time goes on, the forecast tends to the long run mean and variance of the process.

Exercise 5. GARCH-M

Use [arch.csv](#).

- Estimate an ARIMA model for the series `ym` (return on a portfolio) following Box-Jenkins methodology. Why might someone conclude that the residuals appear to be white noise?
- Perform the LM test for ARCH errors.
- Estimate an ARCH-M process (using the standard deviation in the mean equation), with an ARCH(1) for the variance.
- Check the ACF and the PACF. Do they appear to be satisfactory? Try other formulations for the ARCH-M process. Interpret the coefficient of the std. dev. in the mean equation.

Solution

- a) You should estimate a MA(3,6).

➤ `arma5a <- arima(LPPI.d1.w, order=c(0,0,6), fixed=c(0, 0, NA, 0, 0, NA, NA))`

One might conclude that the residuals appear to be white noise because there is no autocorrelation left on residuals. But the residuals are not normal and they present conditional heteroskedasticity.

- b) Heteroskedasticity Test: ARCH

➤ `ArchTest(arma5a$resid, lags=4)`

We reject the null of no autocorrelation in squared residuals.

- c) Now estimate an ARCH(1) variance model and choose `garchInMean=T` and `inMeanType=1` to obtain an ARCH-M with the standard deviation in the mean equation.
- `spec <- ugarchspec(variance.model = list(model = "fGARCH", submodel = "GARCH", garchOrder = c(1,0)), mean.model = list(armaOrder = c(0,6), garchInMean = T, inMeanType = 1), fixed.pars = list(ma1 = 0, ma2 = 0, ma4 = 0, ma5 = 0))`
- An ARCH test indicates that we still have significant correlation in squared residuals.
- d) Try to improve the model estimating now a GARCH(1,1) for the variance, and eliminating the term MA(3). Then we have a better model with low correlations in standardized residuals and squared standardized residuals.

Exercise 6. GARCH and TARARCH models

The file NYSE.xls contain the daily values of the New York Stock Exchange Composite Index. Reproduce the results of section 10, chapter 3 of Enders (2004) (ignore the IGARCH estimation subsection).

Exercise 7. GARCH models

Use `garch.csv`. The variables, `Ret`, `Inf`, `dtbill`, are the S&P Composite index return, the US inflation rate, and the first difference of the three-month Treasury bill rate. The mean equation has the following form $Ret=c+b_1Ret(-1)+b_2Inf(-1)+b_3dtbill(-1)+\varepsilon$

- a) Test for ARCH terms in the squared residuals from this equation (and plot the ACF and PACF of the squared residual).
- b) Try to model the heteroskedasticity using: GARCH and TARARCH models.
- c) Which model seems to be the best one? Justify your answer.

Solution

- a) Test for ARCH terms in the squared residuals from this equation (and plot the ACF and PACF of the squared residual)
 - `garch.ts <- ts.union(ret=ts(garch$ret),inf=ts(garch$inf),dtbill=ts(garch$dtbill))`
 - `lm7 <- dynlm(ret ~ L(ret,1) + L(inf,1) + L(dtbill,1), data=garch.ts)`

The ARCH test results strongly suggest the presence of ARCH in the residuals. We reject the null of no serial correlation of order one.

 - `ArchTest(lm7$resid, lags=4)`

The correlogram shows clearly that there is correlation of order one of residuals squared.

 - `par(mfrow=c(2,1)); acf(lm7$resid); pacf(lm7$resid)`
- b) GARCH(1,1) probably is enough to take into account all the autocorrelation on squared residuals. First generate a matrix of external regressors.
 - `n <- dim(garch)[1]`
 - `ret_1 <- c(NA, garch$ret[1:(n-1)])`
 - `inf_1 <- c(NA, garch$inf[1:(n-1)])`
 - `dtbill_1 <- c(NA, garch$dtbill[1:(n-1)])`
 - `ex.reg <- data.frame(ret_1, inf_1, dtbill_1)[2:n,]`

Then specify and fit the model

- `spec <- ugarchspec(variance.model = list(model = "fGARCH", submodel = "GARCH", garchOrder = c(1,1)), mean.model = list(armaOrder = c(0,0), external.regressors = ex.reg))`
- `sgarch.fit <- ugarchfit(data=garch$ret[2:n], spec = spec); sgarch.fit`

Note that the parameters verify the conditions for stationarity of the conditional variance (they are both >0 and the sum is <1).

A correlogram of the residuals squared shows that now the correlogram is clean. You can also test formally for additional ARCH terms. In fact an ARCH(1) is enough to generate a clean correlogram. Therefore this could be as well an acceptable model.

TARCH. Even though the previous model seems to be good, we can try to see if there are some asymmetric responses to negative and positive shocks. The command is the same as for the GARCH, except that we would now choose:

- `submodel="TGARCH"`

Because the term γ_1 , which is $\varepsilon_{t-1}^2 \cdot I(\varepsilon_{t-1} < 0)$, is significant it seems that there are indeed asymmetric responses of the conditional variance to positive and negative shocks.

- c) The second model can imply a negative conditional variance if the shock was positive and large enough in period $t-1$ (note that ε_{t-1}^2 has a positive coefficient). And of course a variance can only be positive. Asymmetric effects seem to be in place, therefore GARCH(1,1) seems satisfactory. We could proceed further and try to fit another asymmetric GARCH model, for example an EGARCH.

Source: Exercises 4 to 6 are modified versions of Enders' (2004) exercises (chapter 3).