



Avd. Matematisk statistik

KTH Teknikvetenskap

Project work in SF2943 Time series analysis, spring semester 2012

Jan Grandell, Timo Koski

1 Introduction and general instructions

The purpose of this set of project assignments is to make the student familiar with the practical handwork and theoretical effort required in time series analysis of real life data.

The project work is to be submitted in the form of two written reports, called Report A and Report B. **It is recommended that you work in groups of three to produce the piece of studies required for the reports.**

1.1 Structure of the reports

Report A consists of parts A.1 and A.2.

- A.1 'Choose your own time series'
- A.2 'Classical Time Series Modeling'
 - A.2.1 Sea level in Stockholm 1825–1984
 - A.2.2 Simulation of ARMA processes

Report B consists of parts B.1 and B.2.

- B.1 'Analysis of your own time series'. Here the group does an analysis of the time series chosen in A.1. by three different methods.
- B.2 'Parameter estimation, ARCH and GARCH processes'
 - B.2.1 Statistical estimation for time series
 - B.2.2 ARCH and GARCH processes as models for foreign exchange rates

Detailed descriptions of each of these assignments are given in Section 2 and Section 3 below.

1.2 Software for Time Series Analysis SF2943

We recommend that the project is done using MATLAB. The necessary course related m-files, not generally available in MATLAB, can be downloaded from

<http://www.math.kth.se/matstat/gru/sf2943/projects2012.html>.

The data files are also accessed via this page.

MATLAB has, in fact, an extensive toolbox, *System Identification Toolbox*, which can be used for analysis and modeling of time series. If you have access to this toolbox, you may use it. There are also several m-functions in MATLAB *Signal Processing Toolbox* (statistical signal processing) that may turn out to be useful. *Do report explicitly which software (or toolboxes) you have used!*

Every legitimate copy of the course textbook supplies a CD-ROM with a time series package ITSM2000. This contains, e.g., routines for model choice. The course syllabus does not allot time for any training in using ITSM2000, but you are allowed use it on your own in these projects.

1.3 How do you submit your two reports?

- Each group (of three) writes and submits two joint reports. The names (first name, family name) and Swedish *personnummer* of each member of the group are to be given on the cover page of both reports.
- Report A and Report B are submitted as two different documents. Files in pdf-format are preferred. The reports are to have the structure given above, e.g., Report A contains A.1, A.2: A.2.1, A.2.2.
- The reports are to be sent by email to the address `tryd.math.kth.se`. The subject of the email is to be of the form: `reportAsf2943xyz` and `reportBsf2943xyz` respectively, where x, y, and z are the family names of the students in the group.

1.4 Deadlines for submission of reports

Report A to be submitted by 24.00, Wednesday May 2, 2012

Report B to be submitted by 24.00, Monday May 28, 2012

1.5 Warning for websupported plagiarism

Please do not try this.

2 Report A

2.1 Report A.1

Choose a time series according to your own interests. If you have to write the data by hand into a file it is enough to use around 100 values. You may get the data from any source except from the course textbook by Brockwell and Davis, *Introduction to Time Series and Forecasting* (henceforth referred to as B&D). Further it must be real data and *not* randomly generated data. You might for instance find data on the web. The web page

<http://www.robjhyndman.com/TSDL/>

of the Time Series Data Library by Rob J. Hyndman from Monash University is a good source. The page

<http://www.math.kth.se/matstat/links.htm#Data>

has links to other repositories of data sources that can be used for A.1.

Recommendation: Financial and economic time series, especially those dealing with stock prices and business cycles (konjunkturcykler), are often difficult to analyze. Hence, think twice before choosing this kind of data.

Your report on Part A.1 must contain a description of the data (including where you found it) and a plot of the raw data set.



Your task (to be reported in Report B.1) will be to, during the course and invoking what you have learned, do three different analyses of the data. You may very well do these analyses at the same time when you do the other parts of your project. Find more on this in Section 3.1 below.

2.2 Report A.2

2.2.1 Report A.2.1: Sea level in Stockholm, 1825 to 1984

This project consists of making a *classical decomposition* of data of sea level in Stockholm, 1825 to 1984.

The sea level in Stockholm has been measured since 1774, the longest sea level measurement series in the world. From 1825 there exists a complete series of data from every month. This series will be used here. At the beginning the measurements were done at Slussen; now the instrument is placed on Skeppsholmen. The first figure gives the total series from 1825 to 1984, the second one is a ten years series. The trend in the series is an effect of the land elevation.



The m-files `acf.m`, `acvf.m`, `diffd.m`, `ljungbox.m`, `ranktest.m`, `smoothma.m`, `smoothpf.m` and `seascomp.m` can be used. The data can be found in the MATLAB data-file `sealevel.mat` or in the textfiles `sealevel.dat` and `sldate.dat`. Save them in an appropriate library, where MATLAB can find them.

The textfile `sldate.dat` also contains the months and years of the data, as first row and column. `sealevel.dat` only contains the data and is easier to read into MATLAB (use `load sealevel.dat`). Even easier to use the mat-file directly.

For the classical decomposition, see B&D, Section 1.5.

1. The total data vector `s1` consists of 1920 observations. Select 600 of them, representing 50 consecutive years. Begin at a random time. Store the selected data in appropriate vector. Use then the m-file `acf.m` to compute and plot the sample autocorrelation function of the selected data. The command `y=acf(x)` will compute the sample autocorrelation function of the time series `x`. Remember that the index n represents the correlation $\rho(n+1)$ as the indices of vectors in MATLAB begin with 1 and not 0.

The command `y=acf(x,plott)` where `plott` is an arbitrary number also draws a plot, with lines $\pm 1.96/\sqrt{n}$; see B&D, Example 1.4.6 for an explanation. Comment on your plot. What conclusions can you draw? Does the plot say something about a period?

2. For estimating trend and season factor, use first Method S1, described in B&D, p. 31. What do you think is an appropriate period p ?
3. Estimate the seasonal component s_t and the deseasonalized data vector d_t .

The MATLAB-command `[d,s]=seascomp(x,p)` does that, where `x` is the original data vector and `p` the period. Plot s_t and d_t .

4. Test the deseasonalized data d for remaining trend, by using the *rank test* described in B&D, p. 37. Use the m-file `ranktest` which computes, in B&D's terminology,

$$\frac{|P - \mu_P|}{\sigma_P}$$

What is the result of the test? Shall the hypothesis "no trend" be rejected?

5. Estimate the trend in the deseasonalized data vector d by using the m-file `smoothpf`.

The command `[c,m,z] = smoothpf(x,grad)` estimates the trend in data vector $\mathbf{x}$ by an appropriate polynomial of degree `grad`. The coefficients of the polynomial in `coeff` will appear in `c`; `m` will be the trend vector and `z` the estimated residuals $x - m$. Try a linear and a quadratic trend. Plot the residuals. Does a quadratic trend seem to be a better model than a linear one?

6. Test the the residual series `z` in 5 for independence by using the Ljung-Box test (B&D, p. 36). Use $h = 40$. The test statistics can be computed by the m-file `ljungbox`. What is the result? Use the significance level 1%.
7. Make plots of the original data, the deseasonalized data d_t and the estimated seasonal component s_t .

2.2.2 Report A.2.2: Simulation of ARMA processes

This part consists of simulation and analysis of some ARMA processes. We shall first provide some help with the pertinent files and then give the assignments.

Description of MATLAB files

You can use the following m-files.

- **simarma** simulates an $\text{ARMA}(p, q)$ process. The m-file first simulates normally distributed innovations, and from them an ARMA process is constructed. The syntax is

```
y = simarma(fi,theta,n,s2,seed)
```

where **fi** is the vector (in B&D's notation, e.g. p. 83) $\phi = (\phi_1, \phi_2, \dots, \phi_p)$ and **theta** is $\theta = (\theta_1, \theta_2, \dots, \theta_q)$. Furthermore **n** is the number of observations to be generated and **s2** is the white noise variance σ^2 . The integer **seed** is a seed for the random generation and must be set for every simulation. This makes it possible to regenerate the same series. In all simulations you have to report the seed. The output **y** will be a vector with all simulated observations. To generate an AR process you just have to set **theta** to 0, and if **fi** is set to 0 an MA process is simulated.

- **causal** tells you if the specified ARMA process is causal. We want to simulate causal processes, admitting stationary solutions. The syntax is **causal(fi)**, where **fi** is as above.
- **predarma** gives prediction of an ARMA process. The syntax is

```
[y,se2] = predarma(x,fi,theta,s2,h)
```

where **x** is the vector of known observations, **fi**, **theta** and **s2** are as above, and **h** is the forward time step, i.e. the time ahead that the prediction is made. The output **y** is the predicted value and **se2** the mean square error of the prediction. See B&D, Section 3.3. In B&D's notation, $x = (X_1, X_2, \dots, X_n)$ and $y = P_n X_{n+h}$.

- **innov** is used in **predarma**.
- **armaacvf** gives the autocovariance function of an ARMA process. Do not confuse this with the m-file **acvf**, which gives the sample autocovariance function, based on an observed time series. The syntax is

```
y = armaacvf(fi,theta,N,s2)
```

with `fi`, `theta` and `s2` as above and `N` is the maximal lag; that is, the covariance is computed up to $h = N$. The values of the autocovariance function is found in the vector `y`. For the autocovariance function of an ARMA process, see B&D, Section 3.2. The autocorrelation function is, as you know, $\rho(h) = \gamma(h)/\gamma(0)$, so dividing `y` above with its first element gives the autocorrelation function.

- `psi` is used in `armaacvf`. This function computes the ψ -parameters in the linear representation of the process, see B&D, p. 51.
- `roots2ar` computes the AR parameters from the roots in the generating polynomial.
- `arroots` is the inverse, i.e. it gives the roots of the generating function.

AR(2) process and ARMA(2,2) process

The defining equation is

$$X_t - \phi_1 X_{t-1} - \phi_2 X_{t-2} = Z_t$$

Read Example 3.2.4 (B&D, p. 91). One case is missing in that example, namely $\xi_1 = \xi_2$, $|\xi| = |\xi_1| = |\xi_2| > 1$. In that case the process is causal and the covariance function is

$$\sigma^2 \xi^{|h|} \left(\frac{|h|}{\xi^2 - 1} + \frac{\xi^2(3\xi^2 - 1)}{(\xi^2 - 1)^2} \right)$$

1. Choose two complex conjugate roots of the generating function, so that the resulting process is causal. Choose the roots so that their modulus (absolute values) are near 1. Compute ϕ_1 and ϕ_2 according to Example 3.2.4 and simulate 100 observations of the process and draw a plot. The simplest way to compute the coefficients from the roots is to use the Matlab function `roots2ar`. You can calculate with complex numbers in Matlab as usual, just write e.g. `3+4i`. You can also use polar coordinates, e.g. `3*exp(-2/3 i)`.

Caution! If you chose two conjugate complex roots and compute ϕ_1 and ϕ_2 according to Example 3.2.4 they can be complex with small imaginary part, due to numerical imprecision. In that case take the real parts.

Plot the autocorrelation function of the process, and also compute the sample autocorrelation function based on the simulations. Plot this sample correlation function and compare the plots.

2. Choose two different real roots of the characteristic function, and make plots of simulations, and autocorrelation functions as in 1.
3. Compare the autocorrelation plots in 1 and 2. What is the main qualitative difference?
4. Consider case 1, two conjugate complex roots. The autocovariance function is given in (3.2.12) on p. 91. The expression for the autocorrelation function is quite simple. Motivate that it is

$$r^{-|h|} \frac{\sin(|h|\theta + \psi)}{\sin(\psi)},$$

where ψ as in (3.2.13) and the notations as in the example.

If one wants to model a process with strong positive dependence, that is with autocorrelation function near 1 for a range of h , how should r and θ in (3.2.12) be chosen? Choose a θ and try with two or three different r , so that you have processes with various strong positive dependence. Plot simulations of the processes (100 values) and plot the autocorrelation functions and the sample autocorrelations functions.

5. Simulate 10 observations $x_1, x_2, \dots, x_{10}$ of a ARMA(2,2) process with parameters $\phi_1 = 0.4$, $\phi_2 = 0.5$, $\theta_1 = 1$, $\theta_2 = 0.8$. Check that the process is causal. Make predictions up to time 20 from these 10 observations. Plot the predictions and the root mean square errors of the predictions, $\sqrt{E((X_{10+h} - P_{10}X_{10+h})^2)}$, $h = 1, 2, \dots, 10$.
6. The process in 5 is a normal process. What is the expected value of $X_{11} - \hat{X}_{11}$? Show that the mean square error $E((X_{11} - \hat{X}_{11})^2)$ is equal to the variance $V(X_{11} - \hat{X}_{11})$. See B&D, p. 65. Note that the mean square error is computed in 6. Show that

$$P(\hat{X}_{11} - \lambda_{\alpha/2} r_{11} \leq X_{11} \leq \hat{X}_{11} + \lambda_{\alpha/2} r_{11}) = 1 - \alpha,$$

where r_{11} is the root mean square error $\sqrt{E(X_{11} - \hat{X}_{11})^2}$ of $\hat{X}_{11}$.

The interval

$$(\hat{X}_{11} - \lambda_{\alpha/2} r_{11}, \hat{X}_{11} + \lambda_{\alpha/2} r_{11})$$

is a *prediction interval* for X_{11} with confidence level $1 - \alpha$. Compute this interval.

3 Report B

3.1 Report B.1

Do three different analyses of the data you chose and reported in A.1. Here the field is open for own insights and skill in analysis. **Note, however, that three different methods of elimination of trend and seasonal components do not constitute an acceptable study.**

The idea behind this project is to get an understanding of the multitude of problems which may—and will—occur when you are working with real data not chosen in the classroom mainly to fit the methods. These kinds of difficulties are rarely discussed in the textbooks, which, as is proper as such, focus on the mathematical theory and algorithms. **This is the reason why you must choose your time series at an early stage of the course. This is mandatory: You may have happened to select a difficult set of data, but, nevertheless, You are NOT allowed to change the time series you earlier reported in A.1, when you get to do the work required for B.1!**

3.2 Report B.2

3.2.1 Report B.2.1: Estimation of parameters of AR and ARMA processes

The data you shall use are in the files `temp.dat`, `timech.dat` and `el.dat`.

The data in `el.dat` is the monthly production of Australia's production of electricity in GWh Jan 1956–August 1995. The data in `temp.dat` is the monthly mean temperature in New York City, Jan 1946–Dec 1959. The data in `timech.dat` is the annual change in the earth's rotation 1821–1970 (10^{-5} s). The new m-files are `burg.m`, `pacf.m`, `yuwaest.m`, `pergram.m`, `specdens.m`, `specarma.m`, `boxcox.m` and `boxcoxf.m`.

1. Consider first monthly temperatures in the variable `temp`. Plot the temperatures. Plot the periodogram (see B&D, p. 121). Make one plot with no smoothing (no weights), and one with the weights $[3 \ 3 \ 2 \ 1]$ ($[W(0) \ W(1) \ W(2) \ W(3)] = [3 \ 3 \ 2 \ 1]/15$), see B&D, p. 125). Use the m-file `pergram`. The command `pergram(x,w,plott)` gives the periodogram. With two arguments, the the plot is drawn with no weight function if `w` is a scalar (number). If `w` is a vector, this is the weights $w(0), w(1), \dots$. The weights are symmetric and automatically normalized (you need not to enter weights with sum 1).

You shall produce the two plots and explain why the period seems to be 12 in the plots.

2. Compute the autocovariance function of the `temp` series. Use the m-file `acvf`. Then compute the partial autocovariance function (see B&D, p. 94) using the m-file `pacf`. The call `pacf(g)`, with `g` the ACVF, computes the PACF up to lags equal to the number of observations; `pacf(g,n)` computes the PACF up to lag n . The call `pacf(g,n,plott)`, where `plott` is an arbitrary number, also draws a plot. If you put the second argument `n=[]` then the PACF is computed and plotted up to lag equal to number of observations. Finally, `pacf(g,n,plott,c1)`, where `c1` is an arbitrary number, also draws 95% confidence bounds (see B&D, e.g. p. 97). The PACF plot suggests an AR model. Fit an appropriate order p (see B&D, Examples 3.2.6 and 3.2.9 on p. 95 and p. 99).

You shall produce the PACF plot with probability lines, give the order of the AR process you think is appropriate and motivate your choice of order.

3. Use `yuwaest` to estimate the parameters in the model in 2. The syntax is `[fi,s2,C] = yuwaest(y,p)`, where `y` is the time series, `fi` is the estimated autoregressive model parameter, `s2` is the estimated WN variance and `C` is the estimated covariance matrix of the estimated parameter vector `fi`. The diagonal of C thus contains the estimated variances of the estimated parameters.

Give the estimates and approximate 95% confidence limits of the parameter ϕ_1 .

You shall give the estimates of the ϕ parameters and the estimated WN variance σ^2 . Give also the confidence interval above.

4. Compute and plot the spectral density for the $AR(p)$ process fitted in 3. Use `specarma`. Type `help specarma` for the syntax. Find the dominating period from the maximum of the spectral density.

You shall produce the plots and again explain the period.

5. Plot the data in `timech`. There seems to be a linear trend, which suggests that an ARIMA model could be appropriate (B&D, Chapter 6). We shall study the difference series $Y_t = (1 - B)X_t$. Use the m-file `diffd` to construct the series y_t . Compute the autocovariance function of y_t and after that the PACF. Model this series as a suitable AR process. Estimate the parameters, and construct 95% confidence limits for one of the parameters.

You shall produce a plot of the PACF, explain your choice of model, and give the estimated parameters together with one confidence interval. You shall also express the ARIMA model of the original time series

$\{X_t\}$ that is fitted (see B&D, Section 6.1), that is in the form

$$X_t = a_1 X_{t-1} + a_2 X_{t-2} + \cdots + b_1 Z_t + \cdots + b_q Z_{t-q+1}$$

The parameters $a_1, a_2, \dots$ shall be written as numbers and observe that X_t only appears on the left-hand side.

6. The data file `e1.mat` (or `e1.dat`) contains Australia's monthly production of electricity in GWh Jan 1956–Aug 1995. Plot the time series. You see that the production fluctuates more and more and thus seems to be non-stationary. One can use a transformation of the data to make the series more compatible with a stationary series, see B&D, p. 188. A Box-Cox transformation is given by the logarithm of the data or by a power of the data. In most common cases the variance or the standard deviation of the data seems to be proportional to the mean. To stabilize the fluctuations one may in these cases use the square root or the logarithm of the data as the transformation; in B&D's notation, λ is 0 or $1/2$.

The matlab file `boxcox` is a script that gives a plot in which you can enter the data and the λ value. You can, by using the slide, see how the fluctuations stabilize when entering a new λ , which can be done manually or by a slide. The script `boxcox` uses the m-file `boxcoxf`. Make an appropriate transformation and then by some method eliminate the trend and the seasonality. Plot the estimated dual (see project 1, part 5.)

3.2.2 Report B.2.2: ARCH and GARCH models for modelling financial time series

The purpose of this project is to understand the ARCH and GARCH models for modeling financial time series. The data you shall use is in the file `logret_DEM_USD.mat`. The data in this file are daily log-returns of foreign exchange rates (FX rates) for the German mark (DEM) quoted against the US dollar (USD). If S_t is the FX rate between the German mark and the US dollar at day t then the log-return X_t over day t is $X_t = \log(S_{t+1}/S_t)$ ($\log x$ denotes the natural logarithm with base e).

To load the data into MATLAB, use the command `load logret_DEM_USD`. Then you will get a vector named `dem` containing the log-returns for the DEM/USD exchange rate. Note that you have to run MATLAB from the same directory where you have the file `logret_DEM_USD.mat`.

The construction of ARCH and GARCH time series makes it difficult to explicitly compute interesting quantities that a bank or financial institution might be interested in, for instance the *Value-at-Risk*. Therefore simulation of a fitted model is needed to compute such quantities.

1. Simulate the ARCH(1) process with different choices of the parameters $\alpha_0 > 0$ and $\alpha_1 > 0$. This can be done using the m-file `archsim`. Simulate samples of different length and different starting positions y_0 . The command is `[y,sigma]=archsim(a0,a1,y0,n)`, where `a0` and `a1` are the parameters in the ARCH-model, `y0` is the initial value and `n` is the sample size. The vector `y` is the simulated ARCH process and `sigma` is the volatility process, $\sigma_t = (\alpha_0 + \alpha_1 X_{t-1}^2)^{1/2}$.

Your report should include the following:

- (a) One plot of an ARCH process $\{Y_t\}$ of length $n = 500$ with parameters α_0 and $\alpha_1 < 1$, starting at $y_0 = 0$. You should also include a plot with the corresponding volatility process σ_t .
- (b) One plot of an ARCH process $\{Y_t\}$ of length $n = 500$ with parameters α_0 and $1 < \alpha_1 < 2e^\gamma$, starting at $y_0 = 0$. Here γ is Euler's constant, $\gamma \approx 0.57721$. See the additional lecture notes by Jan Grandell.
- (c) One plot of an ARCH process $\{Y_t\}$ of length $n = 500$ with parameters α_0 and $2e^\gamma < \alpha_1$, starting at $y_0 = 0$.

Can you see any qualitative differences between these plots?

In (a) and (b) the distribution of the ARCH process will eventually converge to a stationary distribution with (a) finite variance and (b) infinite variance. In (c) it will not converge to a stationary distribution.

2. Simulate the GARCH(1,1) process with different choices of the parameters $\alpha_0 > 0$, $\alpha_1 > 0$ and $\beta_1 > 0$. This can be done using the m-file `garchsim`. Simulate samples of different length and different starting positions y_0 . The command is

```
[y,sigma] = garchsim(a_0,a_1,b_1,y_0,sigma0,n,seed)
```

where `a0`, `a1` and `b1` are the parameters of the GARCH model, `y_0` is the initial value, `sigma0` the initial volatility, `n` is the sample size and `seed` is the random number seed. For every simulation you do, give the seed. The vector `y` is the simulated GARCH process and `sigma` is the volatility process, $\sigma_t = (\alpha_0 + \alpha_1 X_{t-1}^2 + \beta_1 \sigma_{t-1}^2)^{1/2}$.

Your report should include the following:

- (a) One plot of a GARCH process $\{Y_t\}$ of length $n = 500$ with parameters α_0 , α_1 and β_1 , starting at $y_0 = 0$, $\sigma_0 = \alpha_0$, such that $\alpha_1 + \beta_1 < 1$. You should also include a plot with the corresponding volatility process σ_t .

- (b) One plot of a GARCH process $\{Y_t\}$ of length $n = 500$ with parameters α_0 , α_1 and β_1 , starting at $y_0 = 0$, $\sigma_0 = \alpha_0$, such that $\alpha_1 + \beta_1 > 1$.

Do you see any qualitative difference of these plots?

In (a) the distribution of the GARCH process will eventually converge to a stationary distribution whereas in (b) it will not converge to a stationary distribution.

3. Plot the FX log-returns in `logret_DEM_USD.mat`. This plot should be included in your report.
4. We shall fit a GARCH(1,1) model to this data set. This can be done using the GARCH parameter estimation functions in MATLAB's Financial Toolbox. Unfortunately it is unavailable for the students (or?). Therefore we give the estimates $\alpha_0 = 8.440 \cdot 10^{-7}$, $\alpha_1 = 0.06434$, $\beta_1 = 0.9195$.
5. By assuming that the estimated parameters are correct and that the observed FX log-returns $\{Y_t\}$ comes from the estimated GARCH(1,1) process we can, if we assume some starting position, compute the volatility process for the observed FX log-returns. Let $y_0 = 0$ and $\sigma_0 = 0.005$. Compute the volatility process $\{\sigma_t\}$ for the observed FX log-returns by:

$$\sigma_t = (\alpha_0 + \alpha_1 Y_{t-1}^2 + \beta_1 \sigma_{t-1}^2)^{1/2}, \quad Y_0 = 0, \sigma_0 = 0.005.$$

Plot $\{\sigma_t\}$ and compare with the FX log-returns. Note how the volatility process increases as the variation of the FX data becomes large. Include the plot in your report.

6. An important issue for many financial institutions is to compute the *Value-at-Risk* for a cash flow over some period of time. The *Value-at-Risk* is defined as follows. If X is a random variable interpreted as some random amount we have earned (or lost if $X < 0$) at some time T , then $\text{VaR}_p(X)$ is the number x_p such that $\mathbb{P}(X \leq -x_p) = p$. For instance, if $p = 0.01$ then $\text{VaR}_{0.01}(X)$ is the amount $x_{0.01}$ such that the probability of losing more than $x_{0.01}$ is 0.01. If X has normal distribution with mean 0 and variance 1 then $-\text{VaR}_p(X)$ is the p th quantile of the standard normal distribution. Of particular interest to financial institutions is to compute the *Value-at-Risk* over a 10 day period. This is the objective of this exercise.

Assume that we have observed the FX log-returns in the given data set and that the last date in that series is today, day 500. We want to compute $\text{VaR}_{0.01}(W_{510} - W_{500})$ where W_{510} is the amount held in the

German mark day 510 given that $W_{500} = 1000000$ DEM. This means that we want to compute the amount $z_{0.01}$ such that the probability of losing more than $z_{0.01}$ over the next 10 days is 0.01.

Let X be the 10 day log-return from today: $X = Y_{501} + Y_{502} + \dots + Y_{510}$. The profit between day 500 and day 510 is

$$\begin{aligned} W_{510} - W_{500} &= \frac{W_{510} - W_{500}}{W_{500}} \cdot W_{500} = \left(\frac{W_{510}}{W_{500}} - 1 \right) \cdot W_{500} \\ &= (e^X - 1) \cdot W_{500} = (e^X - 1) \cdot 10^6. \end{aligned}$$

The value $z_{0.01}$ can be estimated by simulating a large number, $N = 10000$ say, trajectories of the GARCH process $y_{501}, y_{502}, \dots, y_{510}$ and determining $z_{0.01}$ as the amount such that $0.01 \cdot N = 100$ of the trajectories have $w_{510} - w_{500} = 10^6 \cdot (\exp\{y_{501} + y_{502} + \dots + y_{510}\} - 1) \leq z_{0.01}$. A useful function to determine $z_{0.01}$ is `prctile` in the m-file `prctile.m`.

In your report on B.2.2 part 6. you should include:

- A histogram over the outcomes of the $N = 10000$ samples of $W_{510} - W_{500}$. This can be done with the MATLAB function `hist`.
- An estimate for $z_{0.01}$ based on your simulations of the $N = 10000$ paths of the GARCH process.
- The Matlab code for the simulations and calculation of Value-at-Risk.