

Financial Risk

Lecture 1

Aims of the course

- Basic (market) risk management tools using extreme value theory (MS)
- Basic credit risk models (AH)
- Operational risk (ML)
- Reading project presentations to introduce your friends to different concepts of financial risks
- Not all risks may be handled by mathematical models

Financial Risks

- Credit risk
- Market risk
- Operational risk
- Insurance risk
- Liquidity risk
- Reputational risk
- Legal risk
- and so on ...



Market Risk

- Our main focus in the first four lectures will be on market risk
- To this end we will consider buying and holding a portfolio of assets for some period of time
- What we may lose during is time period is denoted L and is called the *loss variable*

Returns

- We may be interested in what we may lose from time t_1 to time t_2
- Letting p_{t_1} and p_{t_2} denote the value of the portfolio at time t_1 and t_2 , respectively, we define the *absolute return*

$$\bar{r}_{t_2} = p_{t_2} - p_{t_1}$$

- The *relative return*

$$\tilde{r}_{t_2} = \frac{p_{t_2} - p_{t_1}}{p_{t_1}}$$

- And the *log-return*

$$r_{t_2} = \ln\left(\frac{p_{t_2}}{p_{t_1}}\right)$$

A note on data used

- In most examples throughout the first four lectures two data sets, both found on the Facebook page and course web page, are used
- One in which stocks from the US market are divided into deciles based on market cap, i.e. the market value of all of a company's outstanding shares
- The other is daily Facebook stock quotes from 120930 to 140930

Loss variable

- So if we are interested in what we may lose in dollars, euros or SEK we define the loss variable L as the negative absolute return
- If we are interested in what we may lose in terms of a percentage of the portfolio value we define the loss variable L as the negative relative return
- Note that (Taylor expansion) log-returns and relative returns will be close for small relative returns

Value at Risk (VaR)

- One of the most common notions in financial risk management is that of Value at Risk (VaR)
- VaR may be used to determine the amount of regulatory capital to set aside for different types of risks
- Letting L be loss variable for our portfolio we may define VaR as

$$VaR_p = \inf\{x: P(L > x) \leq 1 - p\}$$

- So VaR gives us the smallest amount (or percentage) we may lose with a certain probability

Expected Shortfall (ES)

- Given that the VaR is exceeded, one may wonder how bad this can be
- We may quantify this in terms of the *expected shortfall*

$$ES_p = E[L | L > VaR_p]$$

Numerical approximation

- We may compute ES_p using the the average of $VaR_{p'}$'s, where $0 < p' < p$

$$ES_p = \frac{1}{1-p} \int_p^1 VaR_{p'} dp'$$

- For an arbitrary p we may approximate ES_p with e.g.,

$$\frac{1}{N} \sum_{i=1}^N VaR_{p+(1-p)(i-1)/N}$$

Value at Risk (VaR)

- Typically $p = 0.95$ or $p = 0.99$
- If we have that $L \sim N(\mu, \sigma)$ then (exercise)

$$VaR_p = \mu + \sigma z_p$$

and (exercise)

$$ES_p = \mu + \sigma \frac{\varphi(z_p)}{1 - p}$$

where φ is the standard normal density and z_p is such that $P(Z \leq z_p) = p$ for Z standard normal.

Value at Risk (VaR)

- So if we assume that all loss variables have normal distribution, there is no point in taking this course 😊
- Assuming normality of loss variables and returns facilitates calculations but (in most cases) also gives underestimated risks, especially in periods where the market is "nervous" ...

Extremal events

- Typically we are interested in hedging big losses
- Therefore, we are interested in models of extremal events, i.e., models of big losses
- It turns out that "standard" distributions, such as the normal distribution, are not sufficient
- Furthermore, the distribution of L is typically unknown...

Extreme Value Theory

- We know that we, even though the distribution of a (i.i.d.) sample may be unknown, can approximate the distribution of the sample mean by the normal distribution if the sample size is sufficiently large (CGS)
- Is there a "similar" scheme for the maximum or minimum of a (i.i.d.) sample of, say, log-returns (r_1, \dots, r_n) ?

$$M_n = \max\{r_1, \dots, r_n\} \text{ or } m_n = \min\{r_1, \dots, r_n\}$$

- In what follows we will typically use EVT distributions for maxima, but as we are interested in big losses we will use negated returns and the relation

$$M_n = \max\{-r_1, \dots, -r_n\} = -\min\{r_1, \dots, r_n\}$$

- When the theory for different distributions is presented we dispense with the minus signs but in applications we will use negative returns as we are interested in modeling losses, where we define a loss as a positive quantity

Extreme Value Theory

- If we assume that, e.g. log-returns are serially independent and have distribution function

$$F(x) = P(r_t \leq x)$$

it holds that

$$\begin{aligned} P(M_n \leq x) &= P(r_1 \leq x, \dots, r_n \leq x) = \prod_{i=1}^n P(r_i \leq x) \\ &= [F(x)]^n \end{aligned}$$

Degeneration

- But what happens if we let the number of observations increase, i.e. let $n \rightarrow \infty$?
- Then $[F(x)]^n \rightarrow 0$ or $[F(x)]^n \rightarrow 1$ depending on if $x < u$ or $x \geq u$ where u is the upper end point of r_t (typically $u = \infty$ for log-returns)
- So we need something more to get a non-trivial limit...

Appropriate sequences

- We need sequences $\{\alpha_n\}, \{\beta_n\}$ such that the distribution of

$$M_{n*} = \frac{M_n - \beta_n}{\alpha_n}$$

converges to a non-trivial limit

- We sometimes refer to $\{\alpha_n\}$ and $\{\beta_n\}$ and the scaling and location sequences, respectively

Limiting distributions

- It turns out that if the limit exists its distribution function will be (Generalized extreme value distribution)

$$F_*(x) = \exp\left[-(1 + \xi x)^{-1/\xi}\right]$$

for $x < -1/\xi$ if $\xi < 0$ and for $x > -1/\xi$ if $\xi > 0$

- The special case $\xi = 0$ gives

$$F_*(x) = \exp[-\exp(-x)]$$

for $-\infty < x < \infty$

Three types

- Type I, $\xi = 0$, the Gumbel distribution

$$F_*(x) = \exp[-\exp(-x)], -\infty < x < \infty$$

- Type II, $\xi > 0$, the Fréchet distribution

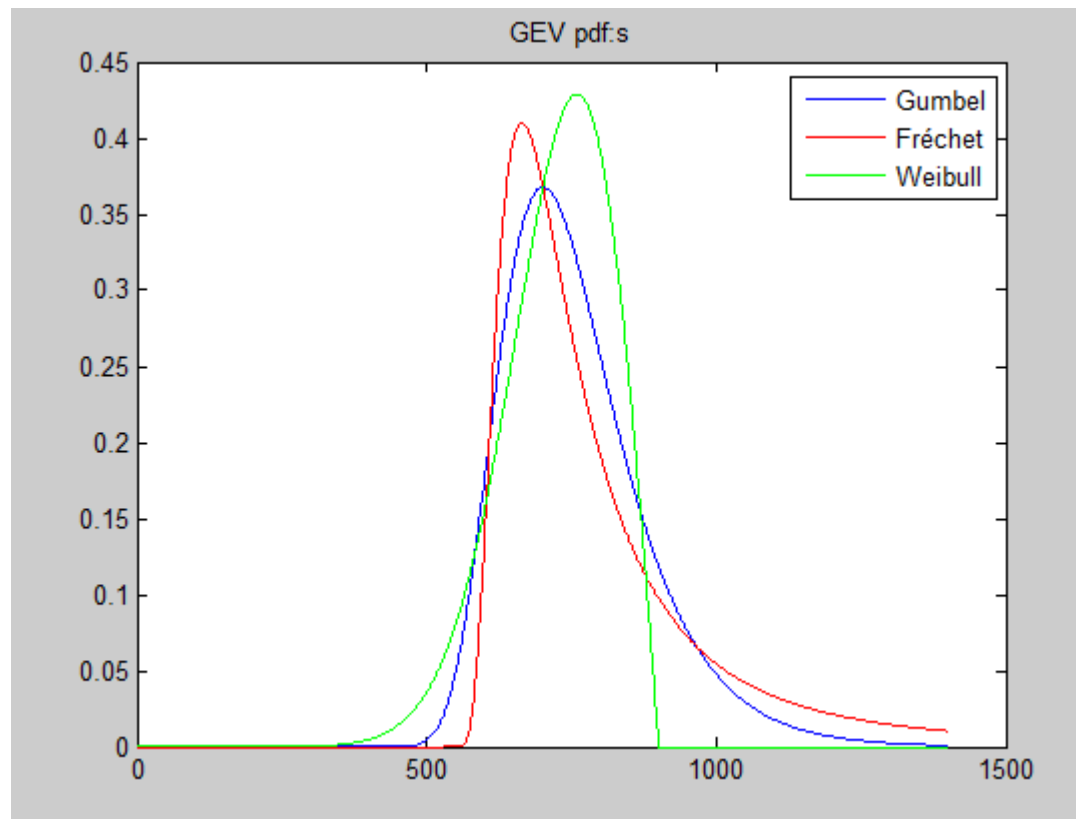
$$F_*(x) = \begin{cases} \exp[-(1 + \xi x)^{-1/\xi}], & x > -1/\xi \\ 0 & , \textit{otherwise} \end{cases}$$

- Type III, $\xi < 0$, the Weibull distribution

$$F_*(x) = \begin{cases} \exp[-(1 + \xi x)^{-1/\xi}], & x < -1/\xi \\ 1 & , \textit{otherwise} \end{cases}$$

Three types

- The pdf:s for $\xi = 0$, $\xi = 0.5$ and $\xi = -0.5$



Exercises

- Using $\alpha_n = 1/\lambda$ and $\beta_n = \frac{\ln n}{\lambda}$ show that the distribution of M_{n^*} for i.i.d. $Exp(\lambda)$ r.v.'s converges to the Gumbel
- Using $\alpha_n = 1/n$ and $\beta_n = 1 - 1/n$ show that the distribution of M_{n^*} for i.i.d. $Uni(0,1)$ r.v.'s converges to the Weibull

In practice

- In real world problems we typically do not have to worry about what

$$F(x) = P(r_t \leq x)$$

looks like in order to fit extreme value distributions

- What we do have to worry about however is that returns typically not are independent or stationary

In practice

- Stationarity means that the distribution (of returns) is constant over time, which is typically not the case in real-life situations
- Typically one finds correlations in second moments of returns and that volatility (standard deviation) is not constant over time but rather that there are clusters of high volatility.
- It turns out (lecture 4) that we may use the same distributions for maxima of dependent sequences if we can adjust for the dependence, at least if the dependence is not too strong

In practice

- Also, we do not really have to worry about what the sequences $\{\alpha_n\}, \{\beta_n\}$ look like
- It is so since if M_{n*} has limiting distribution $F_*(\cdot)$ we have for $y = \alpha_n x + \beta_n$ that

$$P(M_n \leq y) \approx F_*\left(\frac{y - \beta_n}{\alpha_n}\right)$$

- This means that the values of α_n and β_n may be involved in parameter estimations but since parameters are unknown anyway this is not a problem

Taking one step further

- So what to do when data is not stationary?
- It is an empirical fact that log-return data is typically not stationary, i.e. the distribution is not constant over time, and trying to estimate constant parameters of distributions from non-stationary data is not a good idea since non-stationarity implies that distribution parameters change over time...
- To remedy this we may first "devolatilize" or standardize loss data using some mean and volatility series $\{\hat{\mu}_t\}$ and $\{\hat{\sigma}_t\}$

$$z_t = \frac{-r_t - \hat{\mu}_t}{\hat{\sigma}_t}$$

Taking one step further

- Hopefully the series $\{z_t\}$ appears stationary and we fit some (maybe EVT) distribution F to this series
- We will then get

$$VaR_{p,t} = \hat{\mu}_t + \hat{\sigma}_t F^{-1}(p)$$

where $F^{-1}(p)$ denotes the p -quantile of F , i.e. the VaR_p for the stationary series $\{z_t\}$

- This is what we probably would use in real life to get a VaR-series that tracks returns and volatility clusters, instead of a constant VaR which would typically make our bank lose lots of money during volatility clusters and setting aside unnecessary regulatory capital in tranquil periods

Imposing stationarity

- There are several ways of “creating” the normalizing mean and volatility series $\{\hat{\mu}_t\}$ and $\{\hat{\sigma}_t\}$
- One simple, but often effective, way is to use windows of observations (negated returns), for which one computes means and standard deviations

Imposing stationarity

- So if we have a series of returns $\{r_t\}$ we get the mean and volatility series $\{\hat{\mu}_t\}$ and $\{\hat{\sigma}_t\}$ from

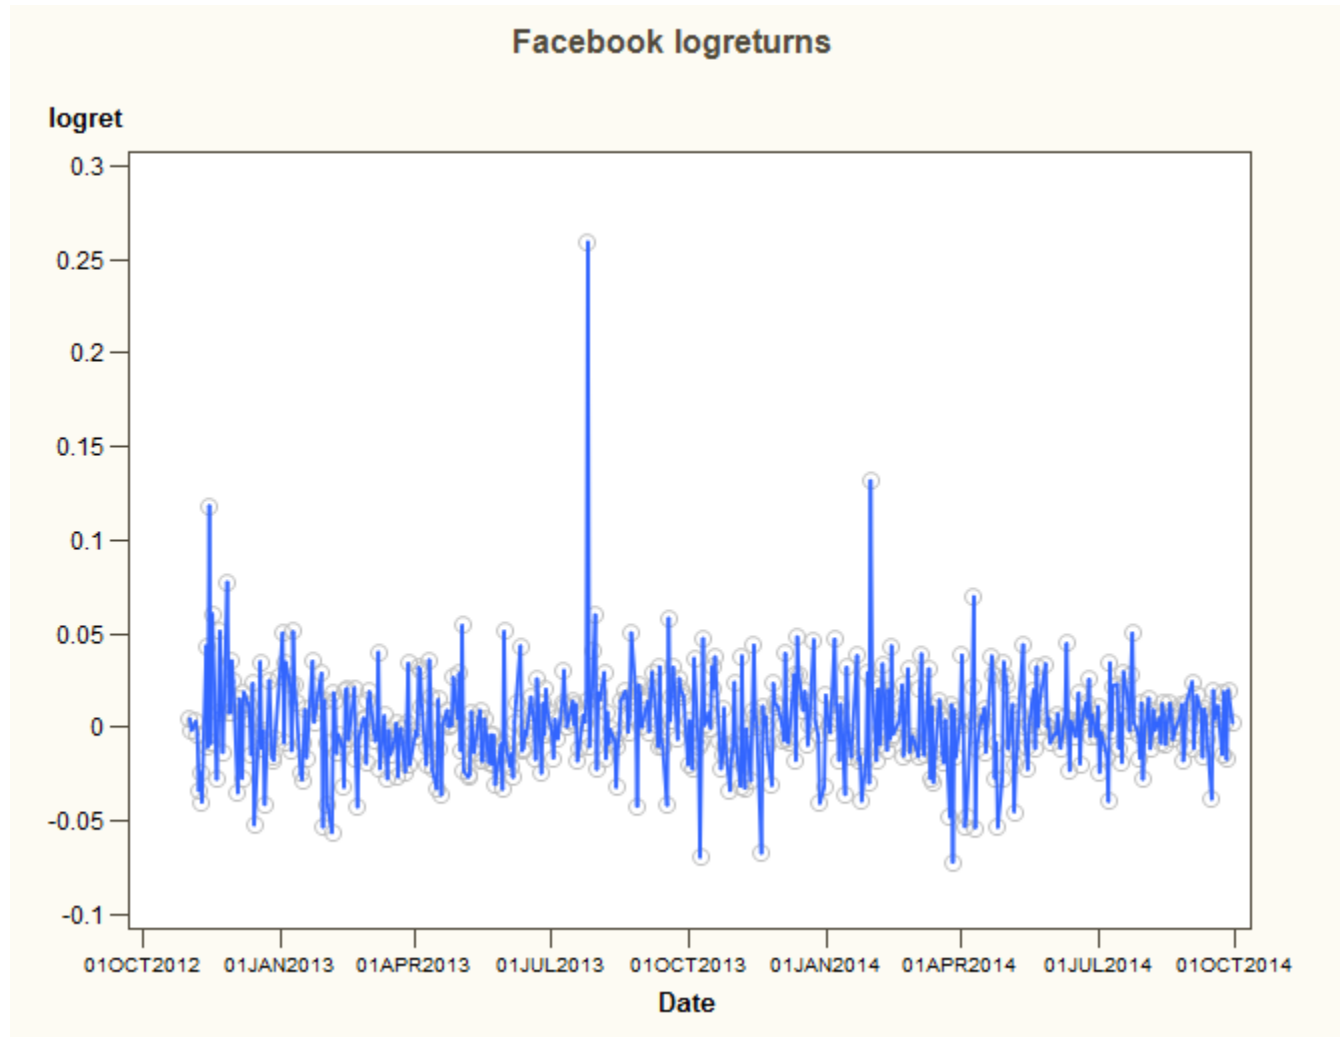
$$\hat{\mu}_t = -\frac{1}{\tau} \sum_{i=t-\tau}^{t-1} r_i$$

and

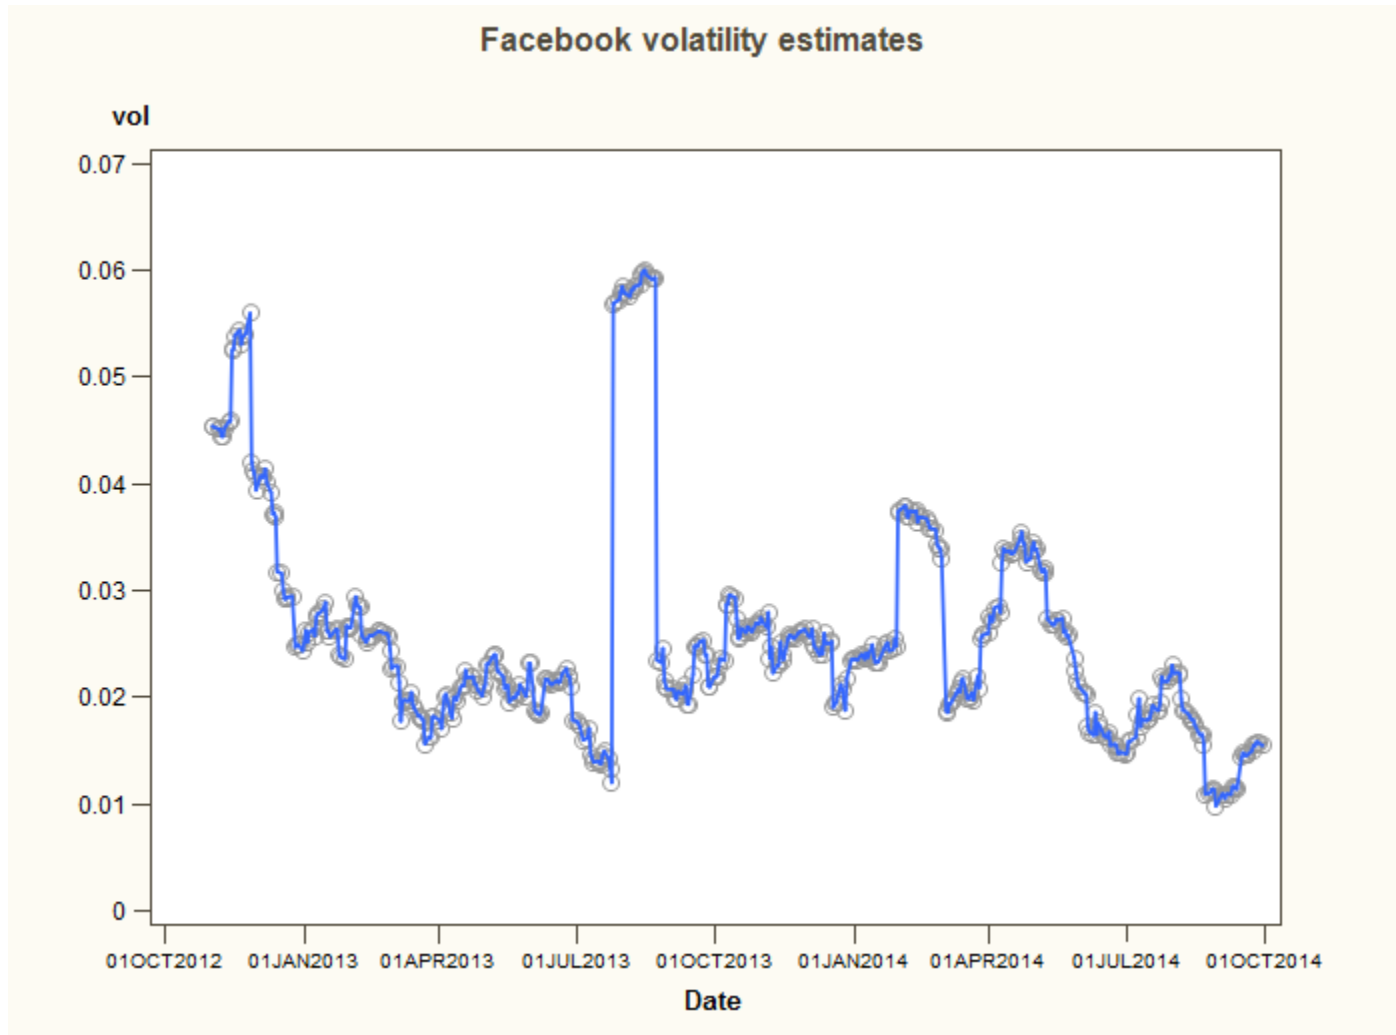
$$\hat{\sigma}_t = \sqrt{\frac{1}{\tau-1} \sum_{i=t-\tau}^{t-1} (r_i + \hat{\mu}_t)^2}$$

- For daily returns the window length, τ , is typically chosen to be 21 days, which amounts to an approximate trading month
- However adjustments of the window length may improve performance

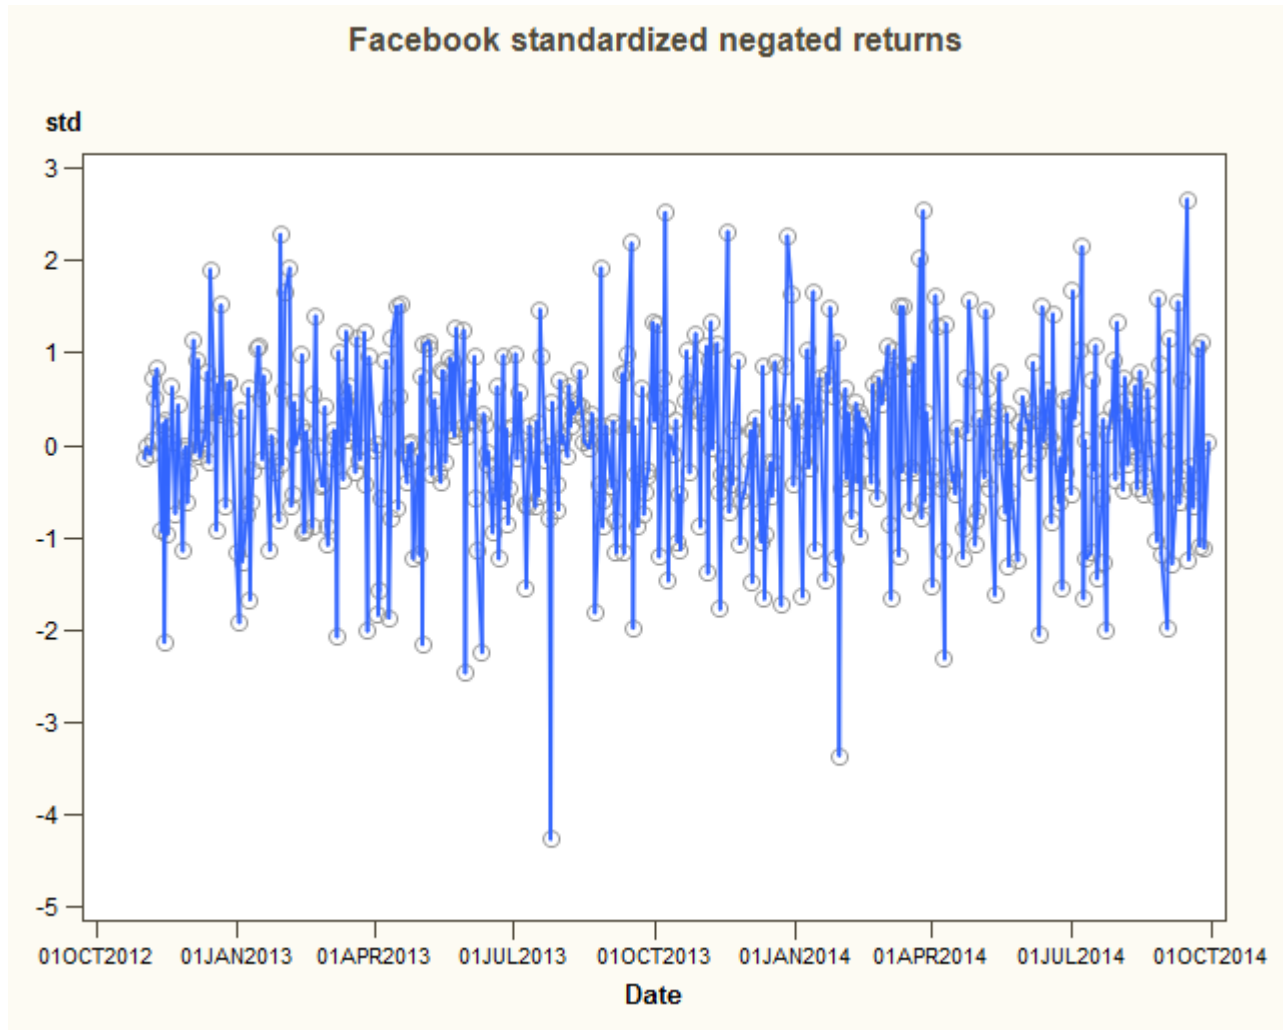
Imposing stationarity




Volatility estimates using 21-day windows




Standardized negated returns



Normality?



QQ-plot




VaR assuming normality

- Aside from the two very small observations it seems quite reasonable to assume that the standardized negated facebook logreturns follow a normal distribution
- This means that it may be reasonable (one day ahead) 95% VaR series is given by

$$VaR_{0.95,t} = -\hat{\mu}_t + 1.645\hat{\sigma}_t$$

- Note that "in real life" this is a prediction or forecast in the sense that we assume that the mean and volatility and normal distribution are the same at day t as at day $t - 1$ and since all components are estimated using information available at time $t - 1$

Plotting VaR and returns



Note that we have plotted negated VaR and (non-negated) returns in order to display losses.

Performance?

- There are 23 exceedances among the 480 observations, which amounts to roughly 4.8%
- This means that we are doing a pretty good job when estimating VaR
- However, it is common that one underestimates VaR when using normal distributions. The data at hand is "nice" ...

Summary and about tech proj 1

- We have seen that it is possible to compute and forecast VaR using very simple methods
- Next time we will look into fitting GEV distributions to data
- What has been covered today will be included in technical project 1 and in particular one must know how to create a dynamic VaR series under normal i.i.d. assumption starting from an arbitrary set of stock prices in order to pass the course
- Hence it is recommended that you download some data set, preferably two years of prices, from e.g. yahoo finance and starting playing around with it.
- Note that file downloaded from yahoo finance are csv-files that you may transfer to xlsx by using "text to columns" under "data" in Excel
- To get started check out the xlsx-file available on facebook and at the course web page