Introduction to Value at Risk

by

Olaf Menkens
School of Mathematical Sciences
Dublin City University (DCU)

St. Andrews, March 17th, 2009

A Quotation

The essence of risk management lies in maximizing the areas where we have some control over the outcome while minimizing the areas where we have absolutely no control over the outcome and linkage between effect and cause is hidden from us.

(See Bernstein [9], p. 197)

Value at Risk (VaR) – Definition

The concept of **Value at Risk** (VaR) measures the "risk" of a portfolio. More precisely, it is a statement of the following form:

With probability q the potential loss will not exceed the Value at Risk figure [$\rightarrow$ one sided confidence interval].

Speaking in mathematical terms, this is simply the $(1-q)$–quantile of the distribution of the d–day change of value for a given portfolio P. More specifically,

$$\mathbf{VaR}_{\boldsymbol{q,d}}(\boldsymbol{P}) = -\boldsymbol{F}_{\boldsymbol{P^d}}^{-1}(1-\boldsymbol{q}) \cdot \mathbf{PV}(\boldsymbol{P}) \ , \qquad (1)$$

where P^d is the change of value for a given portfolio over d days (the d–day return), F_{P^d} is the distribution function of P^d, and PV(P) is the present value of the portfolio P.

Value at Risk (VaR) – Remarks

- The above definition also known **absolute Value at Risk**. Correspondingly, the **relative Value at Risk** is defined as

$$\mathbf{E\!-\!VaR}_{\boldsymbol{q,d}}(\boldsymbol{P}) = -\left(\mathbb{E}\left[\boldsymbol{P}\right] + \boldsymbol{F}^{-1}_{\boldsymbol{P^d}}(\mathbf{1}-\boldsymbol{q}) \cdot \mathbf{PV}\left(\boldsymbol{P}\right)\right). \tag{2}$$

- The quantile function F^{-1} is a "generalized inverse" function

$$\begin{aligned} F^{-1}_{P^d}(q) &= \inf\left\{x : F_{P^d}(x) \geq q\right\} \qquad \text{for } 0 < q < 1 \\ &= \inf\left\{x : \mathbb{P}\left(P^d \leq x\right) \geq q\right\}. \end{aligned}$$

- It is often convenient to write the VaR in percent of a potential loss and not on a monetary base. In order to do so, remove the term $\mathrm{PV}(P)$ in equation (1) and (2) and replace $\mathbb{E}[P]$ by the mean of the underlying distribution F.

VaR – Further Remarks

- VaR has been developed by JP Morgan. More specific, they developed the so–called RiskMetrics and made it available publicly in 1994, which has been outsourced to a newly founded company, also called RiskMetrics (see also http://www.riskmetrics.com).
- The webpage http://www.gloriamundi.org contains a lot of information about VaR.

VaR – Pro and Contra

Pro:

- Easy to calculate (at least compared to other risk measures) and to understand (it is a monetary amount that focuses the mind).
- It is a common language of communication within the organisations as well as outside (e.g. regulators, auditors, shareholders).
- It is not really complicated, yet it is "messy" and "time–consuming".

Contra:

- It is not a coherent measure, more specifically it does not satisfy the sub–additivity axiom.
- It fails to recognize the concentration of risks.
- Most parametric approaches neglect the heavy tails and the skewness of the return series.

Calculating VaR

There are several approaches for calculating the Value at Risk–figure. The most popular are the

- **variance–covariance approach** (parametric) [$\rightarrow$ RiskMetrics],
- **historical simulation** (nonparametric),
- **Monte–Carlo simulation** (parametric), and
- **extreme value theory** (semiparametric).

Variance–Covariance Approach (VCA)

Assuming that the distribution of the observed returns are normally distributed, the VaR computation can be simplified considerably. This approach is a parametric one since it involves estimation of a parameter – the standard deviation.

With this assumption the d-day VaR to the q–quantile calculates to

$$\mathbf{VaR}_{\boldsymbol{q,d}}(\boldsymbol{X}) = -\sqrt{\boldsymbol{d}} \cdot \boldsymbol{N}_{\mathbf{0,1}}^{\mathbf{-1}}(\mathbf{1}-\boldsymbol{q}) \cdot \boldsymbol{\sigma} \cdot \mathbf{PV}(\boldsymbol{X})$$

for a **single asset** where

- $N_{0,1}^{-1}$ is the inverse of the standard normal distribution function,
- σ is the – estimated – daily standard deviation of the asset, and
- PV (X) is the present value invested in asset X.

Standard Normal Distribution and Quantiles

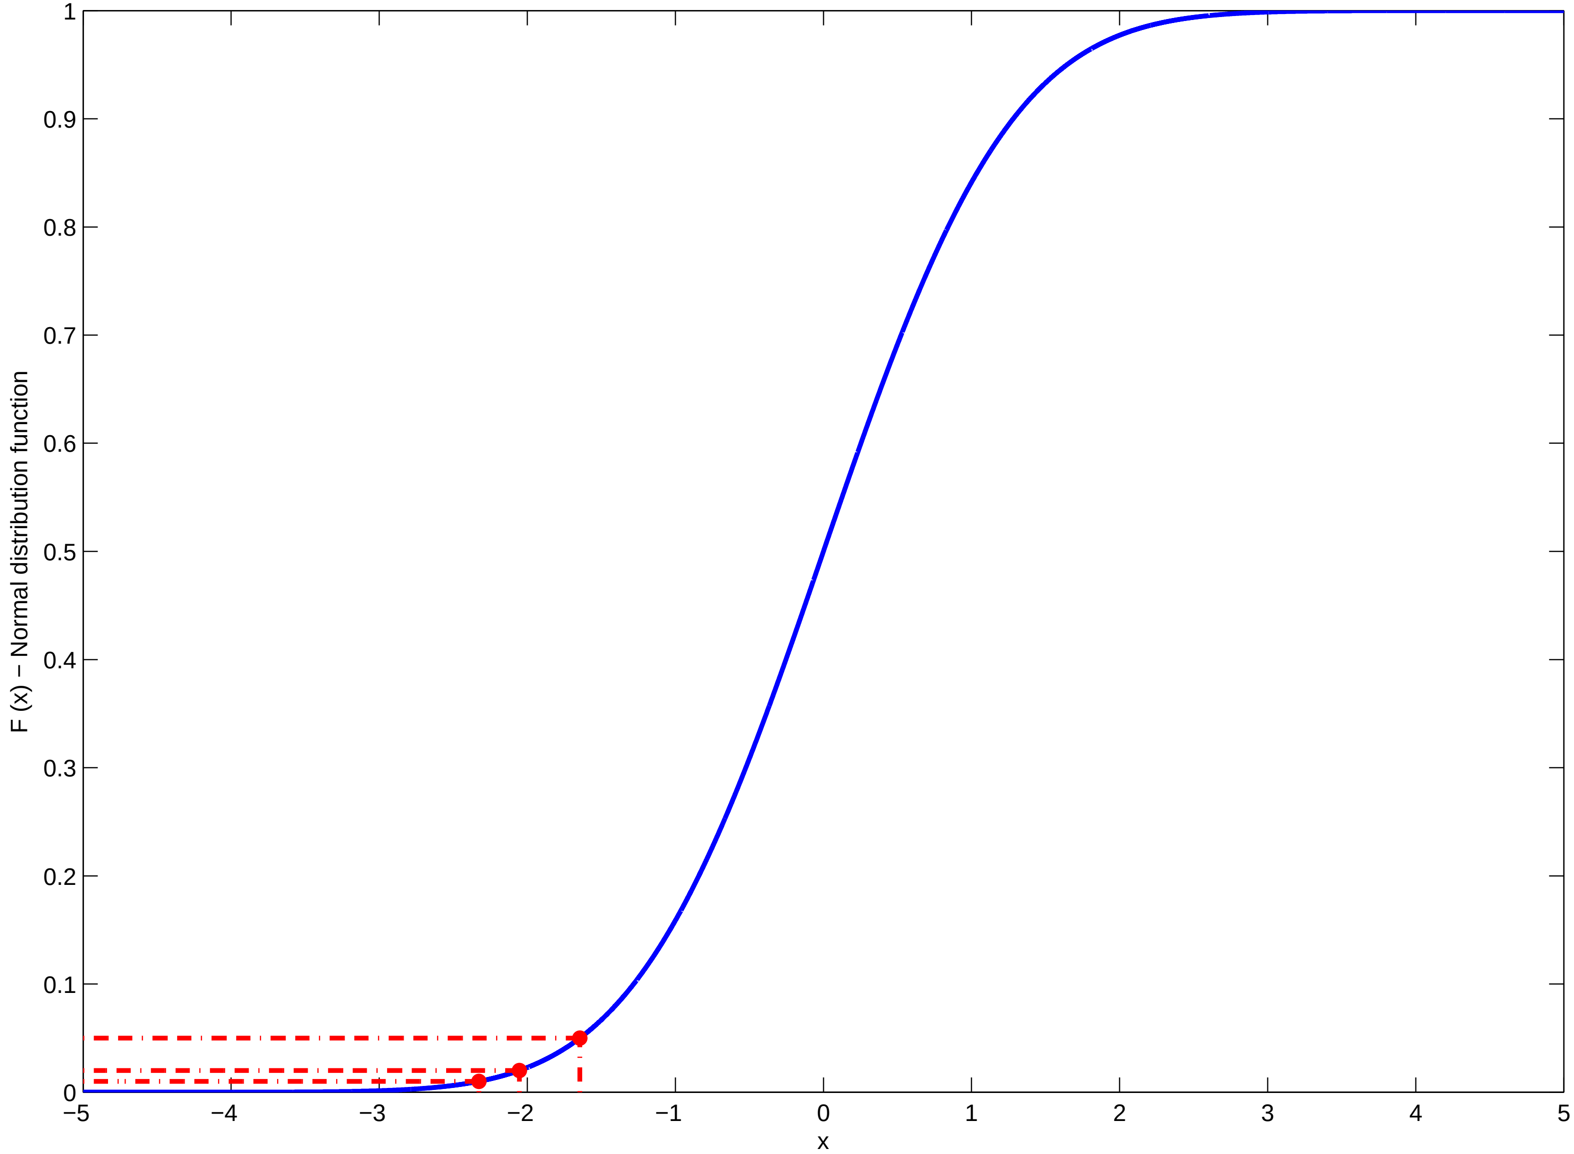

Standard Normal distribution. The dots mark the 0.01, 0.02, and 0.05 quantiles (or the 99, 98, and 95 percent Value at Risk, respectively), which have the values −2.3263, −2.0537, and −1.6449, respectively.

Inverse Standard Normal Distribution and Quantiles

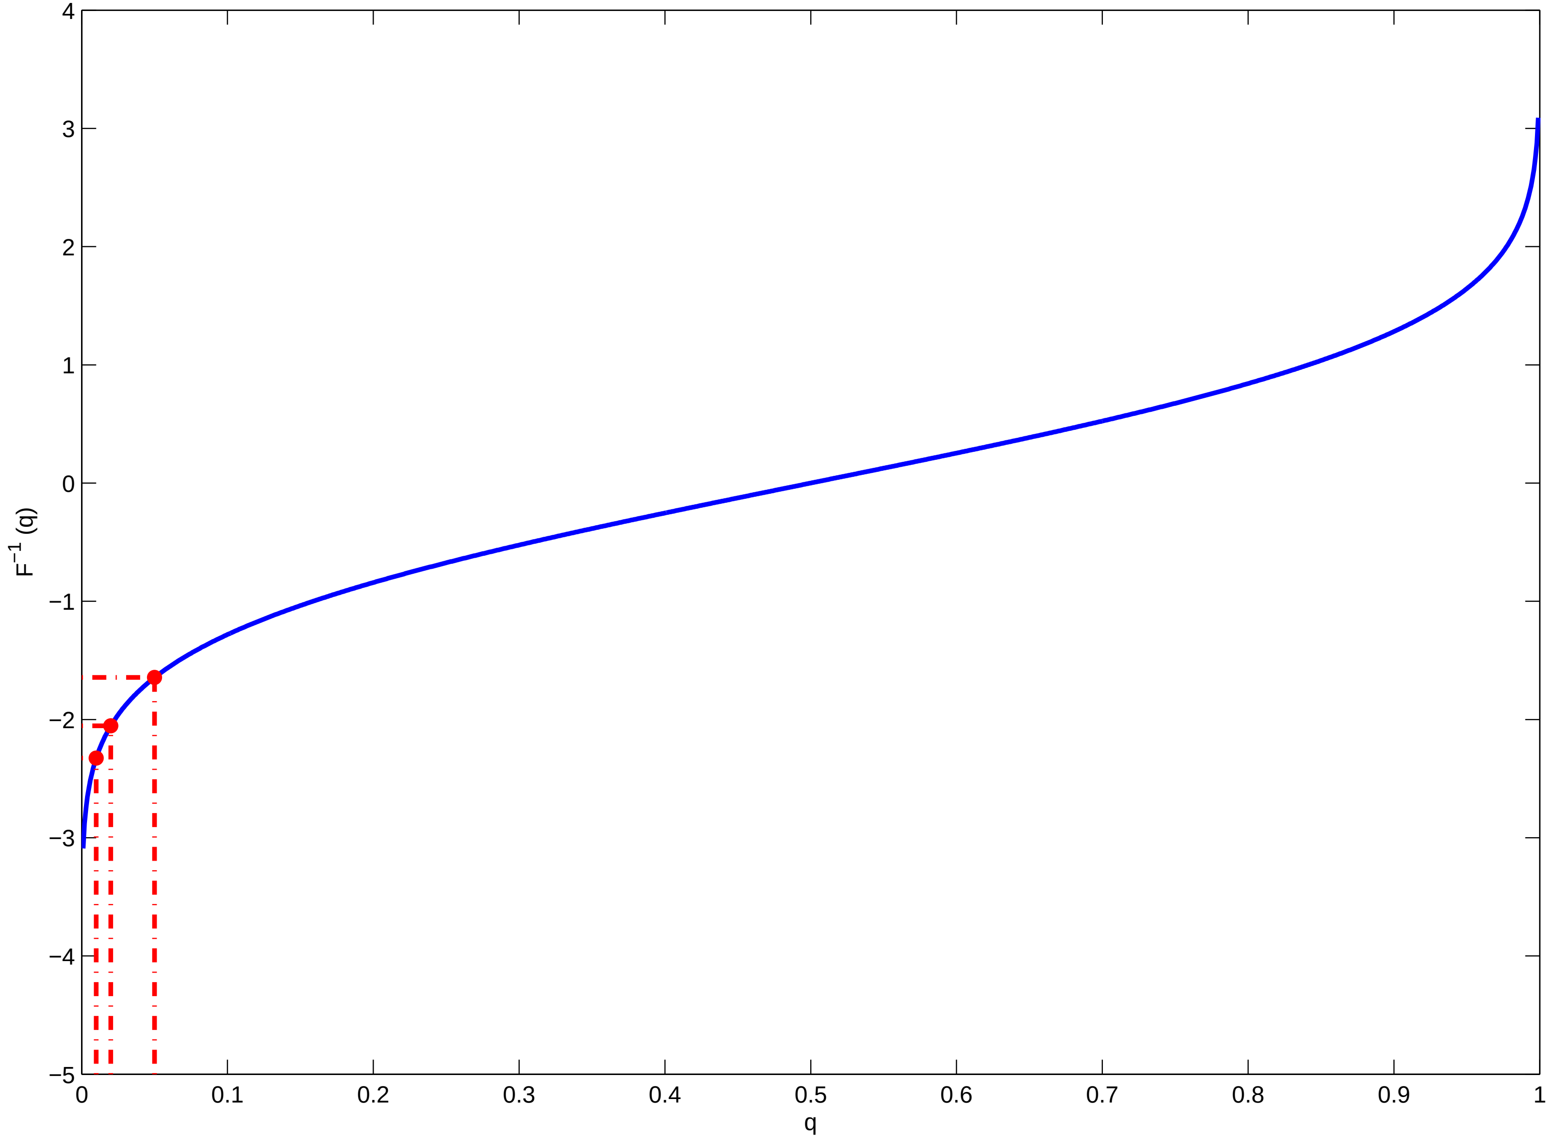

Inverse of the standard Normal distribution. The dots mark the 0.01, 0.02, and 0.05 quantiles (or the 99, 98, and 95 percent Value at Risk, respectively), which have the values −2.3263, −2.0537, and −1.6449, respectively.

Standard Normal Density and the 95% Quantile

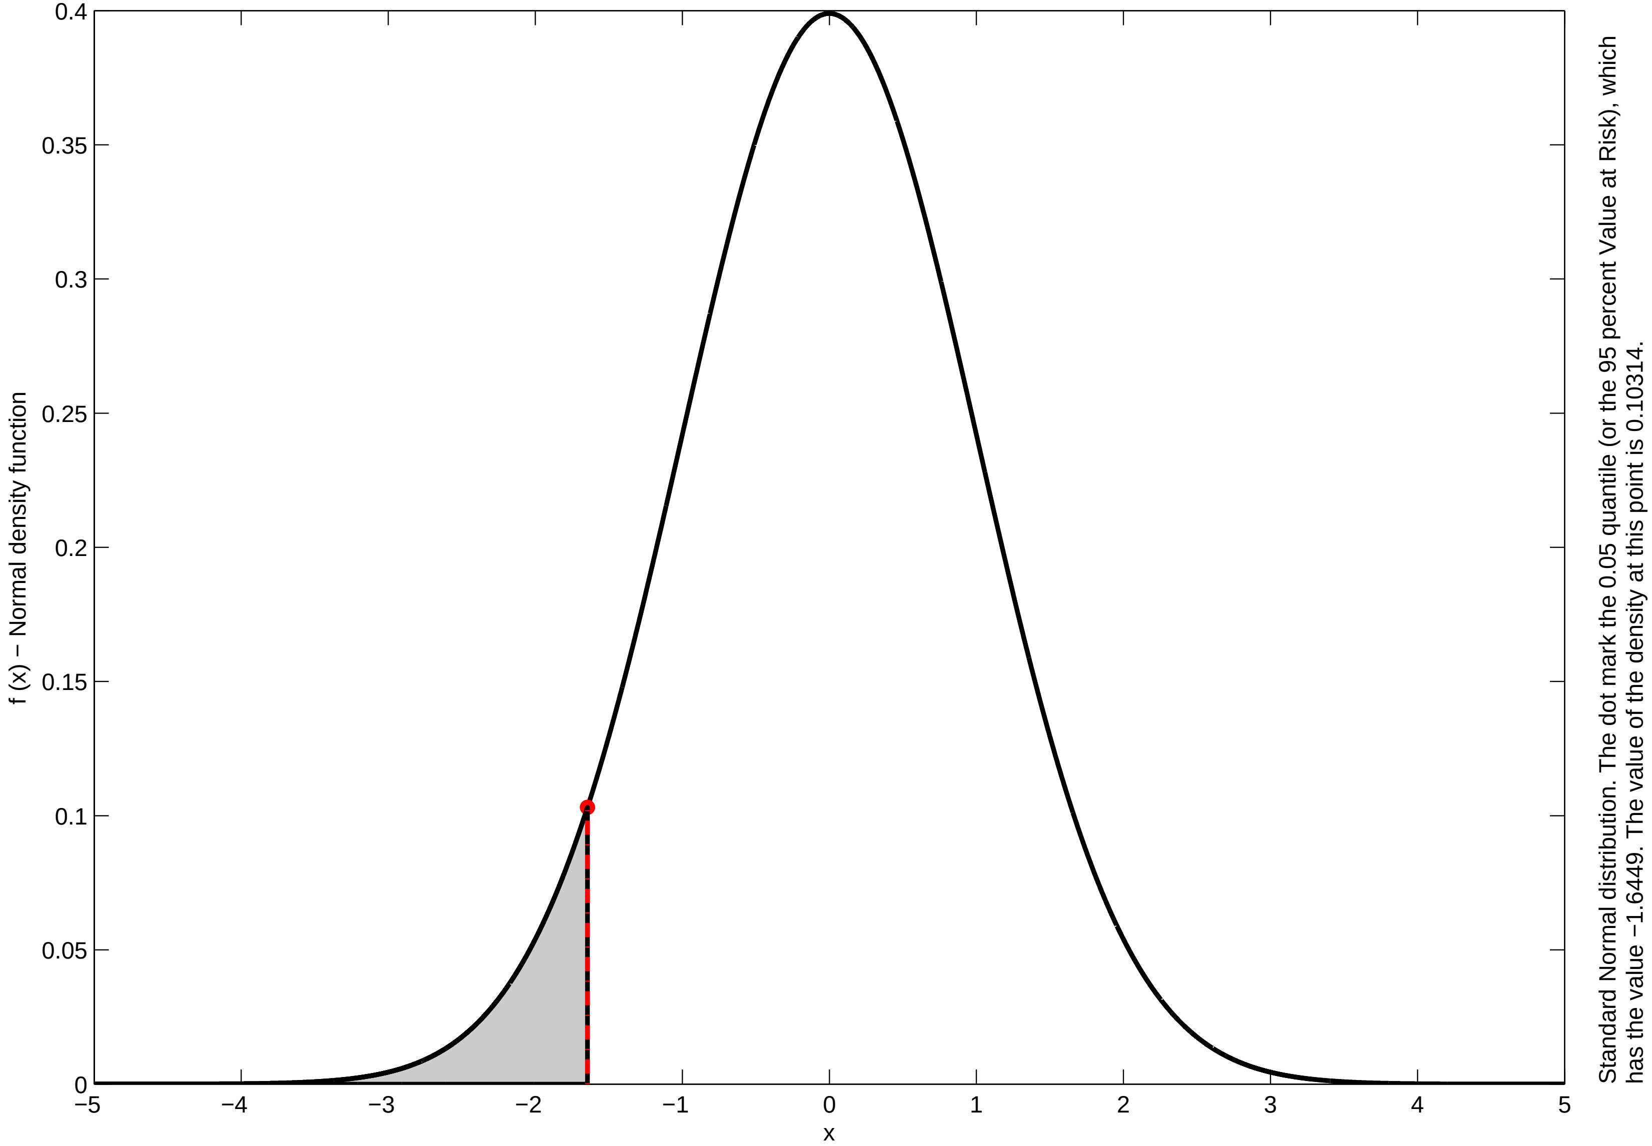

Standard Normal Density and the 98% Quantile

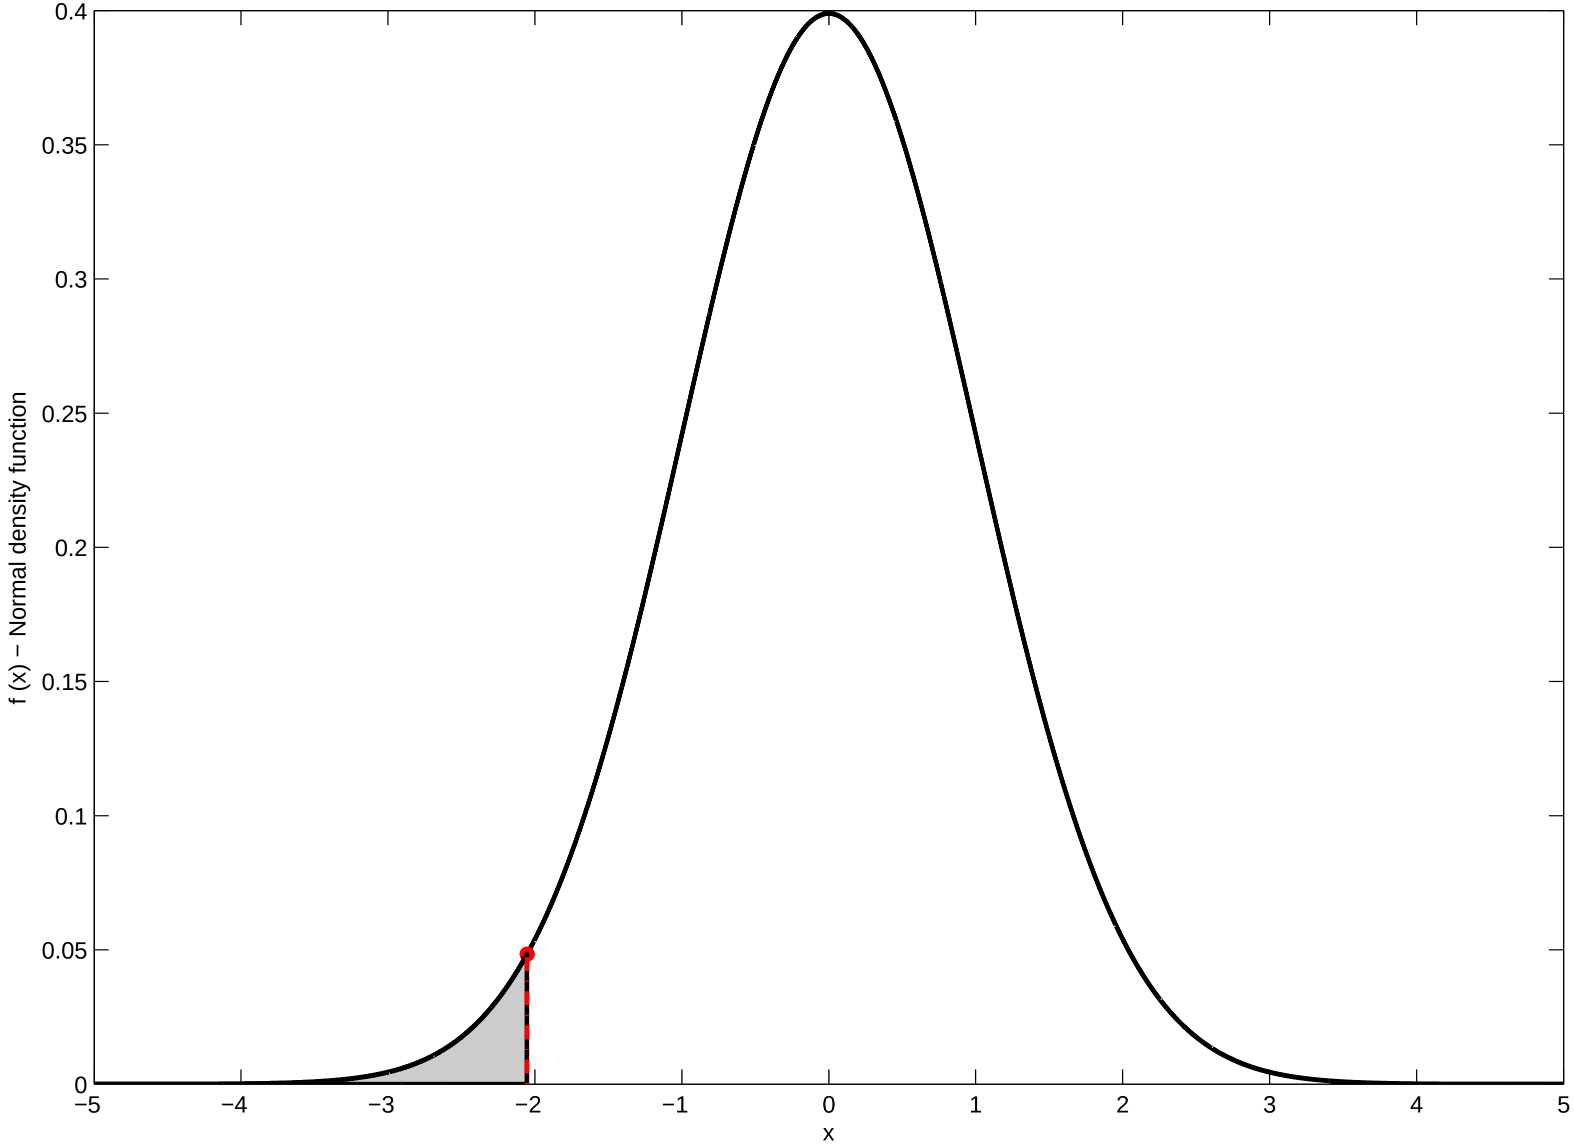

Standard Normal distribution. The dot mark the 0.02 quantile (or the 98 percent Value at Risk), which has the value −2.0537. The value of the density at this point is 0.048418.

Standard Normal Density and the 99% Quantile

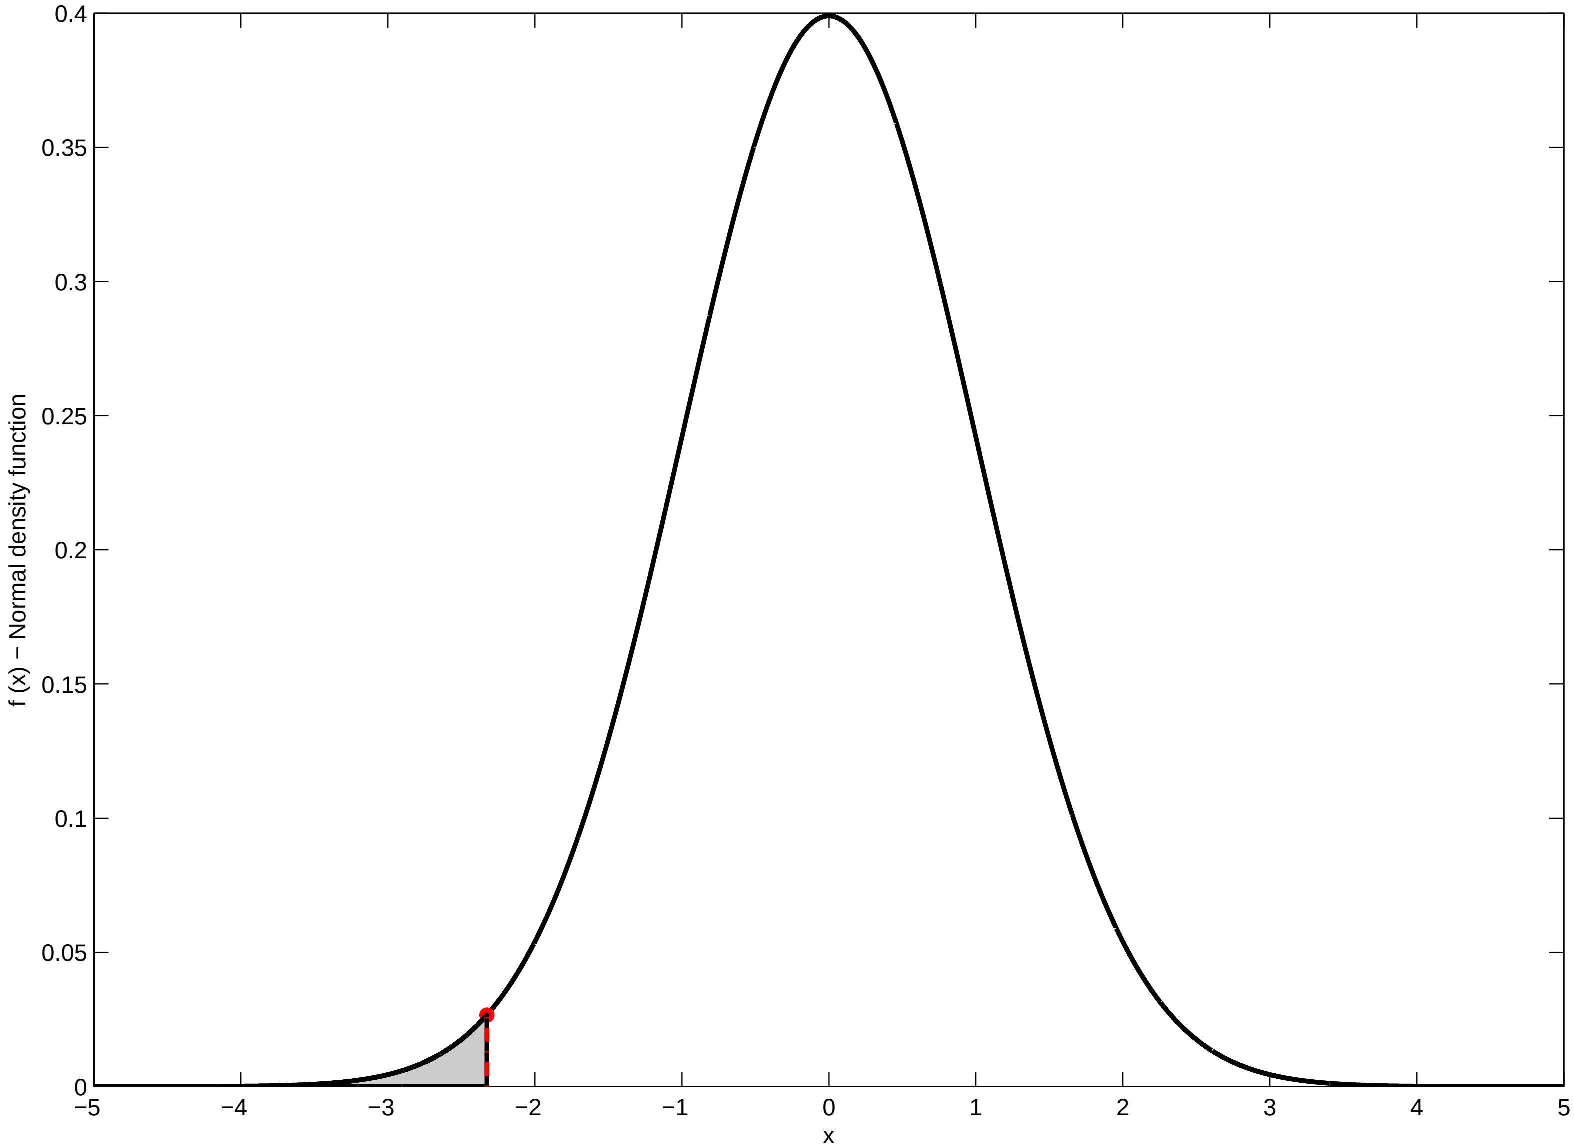

Standard Normal distribution. The dot mark the 0.01 quantile (or the 99 percent Value at Risk), which has the value −2.3263. The value of the density at this point is 0.026652.

VCA – Multiple Assets

The d-day VaR to the q–quantile is for a portfolio with n assets in this situation given by

$$\mathbf{VaR}_{\boldsymbol{q},\boldsymbol{d}}(\boldsymbol{X}) = -\sqrt{\boldsymbol{d}} \cdot \boldsymbol{N}_{\mathbf{0,1}}^{-\mathbf{1}}(\mathbf{1}-\boldsymbol{q}) \cdot \sqrt{\boldsymbol{Y}\boldsymbol{\Sigma}\boldsymbol{Y}^{\boldsymbol{t}}},$$

where

- $N_{0,1}^{-1}$ is the inverse of the standard normal distribution function (as above),
- Σ is the – estimated – $n \times n$ covariance matrix of the assets which are in the portfolio, and
- $Y = \mathrm{PV}(X)$ is a vector of length n and Y_i (for $i = 1, \ldots, n$) is the amount invested in asset i.

Problem: Some positions are **non–linear** in the underlying risk factors (such as options or the bond price–yield relationship).

VCA – Delta Approximation

Assumes that the non–linearity is sufficiently limited so that it is possible to get an accurate VaR estimation while ignoring the non–linearity.

The first–order Taylor series approximation of the change in the value of an option is given by

$$\Delta C \approx \delta \cdot \Delta X \,.$$

Thus, for very short holding periods, the VaR of an option can be approximated by

$$\begin{aligned} \mathrm{VaR}_{q,d}(C) &= \delta \cdot \mathrm{VaR}_{q,d}(X) \\ &= -\sqrt{d} \cdot N_{0,1}^{-1}(1-q) \cdot \sigma \cdot \mathrm{PV}(X) \cdot \delta \,, \end{aligned}$$

where C is a call option on the underlying asset X and δ is the delta of the option.

VCA – Delta–Gamma Approximation

Consider the second–order Taylor series approximation rather than the first–order one.

$$\Delta C \approx \delta \cdot \Delta X + \frac{\gamma}{2} (\Delta X)^2 \ .$$

Making the assumption that $\Delta U = (\Delta X)^2$ is normal distributed and independent of ΔX (which is normal distributed), this gives

$$\sigma = \sqrt{\delta^2 \sigma_X^2 + \left(\frac{\gamma}{2}\right)^2 \sigma_U^2} = \sqrt{\delta^2 \sigma_X^2 + \frac{\gamma^2}{4} \sigma_X^4} \, .$$

With this the Value at Risk calculates to

$$\mathbf{VaR}_{\boldsymbol{q,d}}(\boldsymbol{C}) = -\sqrt{\boldsymbol{d}} \cdot \boldsymbol{N}_{\mathbf{0,1}}^{\mathbf{-1}} \, (\mathbf{1} - \boldsymbol{q}) \cdot \boldsymbol{\sigma} \cdot \mathbf{PV} \, (\boldsymbol{X}) \cdot \sqrt{\boldsymbol{\delta}^{\mathbf{2}} + \frac{\boldsymbol{\gamma}^{\mathbf{2}}}{\mathbf{4}} \boldsymbol{\sigma}^{\mathbf{2}}} \, .$$

VCA – Résumé

Pro:

- Easy to calculate and to implement.

Contra:

- First–order approximation are reliable only if the portfolio is close to linearity.
- Second–order approximation assumes that $(\Delta X)^2$ is normally distributed and independent of ΔX while it is indeed χ^2–distributed and highly dependent of ΔX.
- There are many other delta–gamma approximations such as the ones of Wilson (1994, 1996), Feuerverger and Wong (2000), and Albanese et al. (2001).
- Numerical and simulation methods are improving rapidly (becoming both more sophisticated and faster). Thus the need for Greek–based VaR estimation diminish.

RiskMetrics

RiskMetrics is a well–known program which uses the variance–covariance approach. A detailed technical description of the method and the method for estimating the financial parameters can be found on the website of RiskMetrics.

The main difference is that the variance as well as the covariance are estimated by using an **exponential weighted moving average (EWMA)** of the square of price returns.

$$\boldsymbol{\sigma_{d,i} = \sqrt{\frac{1-\lambda}{d}\sum_{j=-\infty}^{i}\lambda^{i-j}R_j^2}}, \qquad \text{or}$$

$$\sigma_{d,i}^2 = \lambda\cdot\sigma_{d,i-1}^2 + (1-\lambda)\cdot\frac{R_i^2}{d}\,.$$

Analog for the covariance:

$$\sigma_{12_{d,i}}^2 = \lambda\cdot\sigma_{12_{d,i-1}}^2 + (1-\lambda)\cdot\frac{R_{1_i}^2\cdot R_{2_i}^2}{d}\,.$$

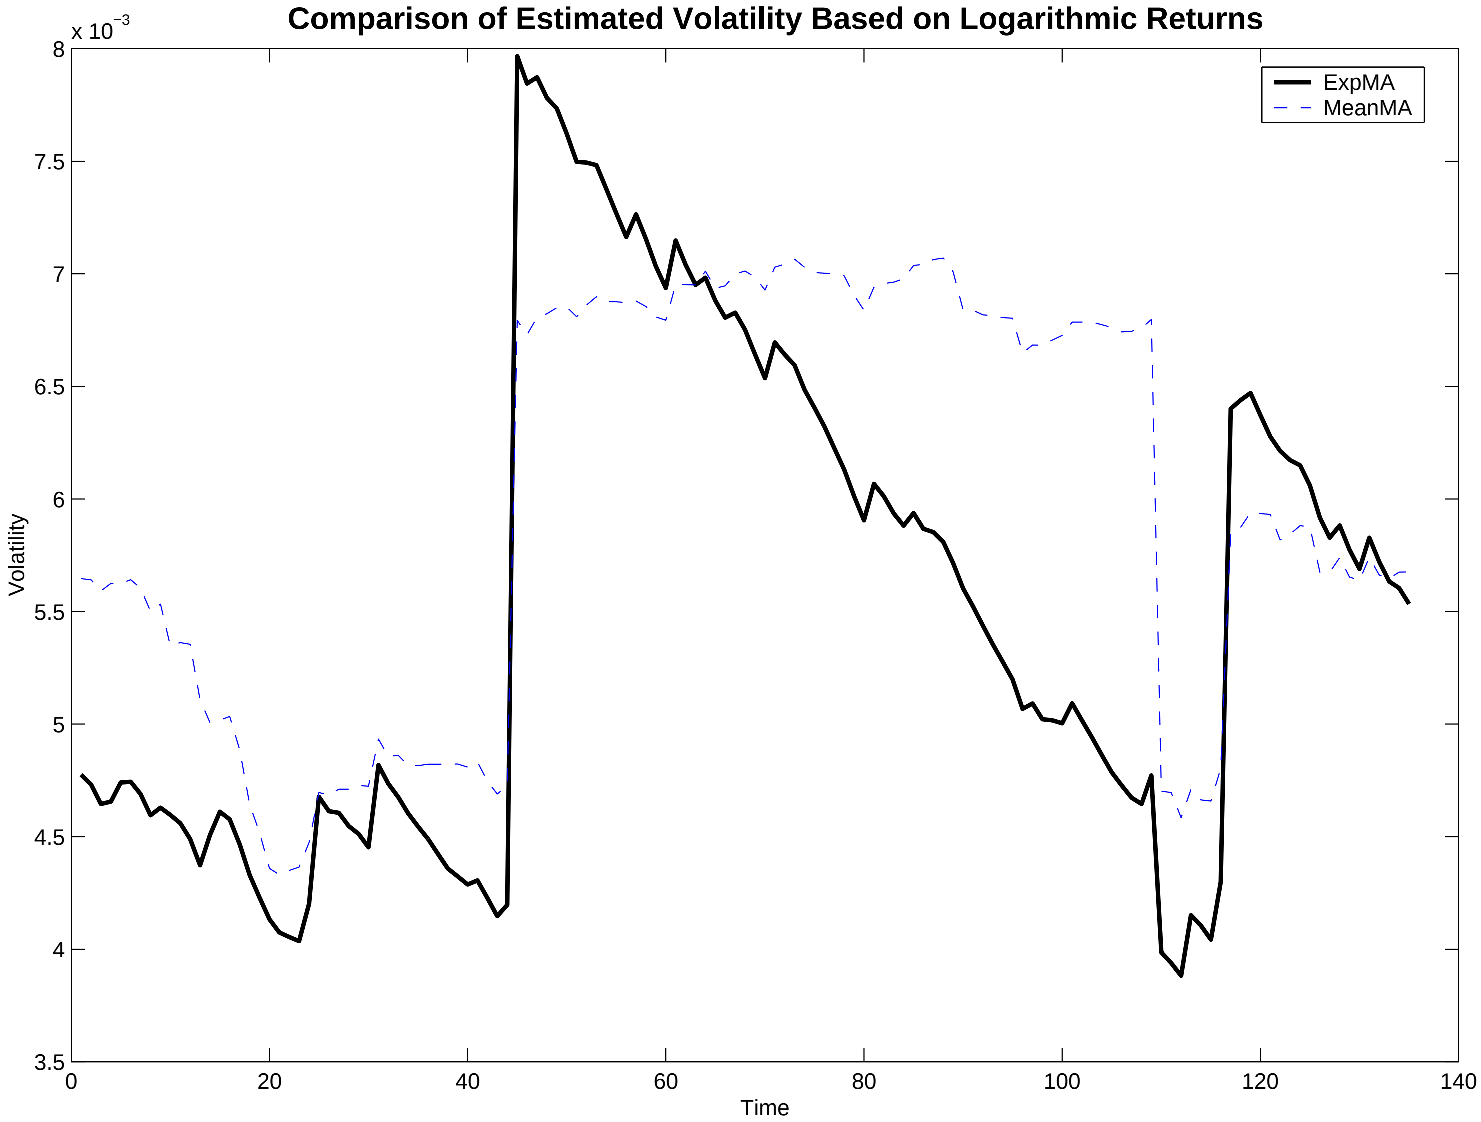

This volatility estimation is based on 65 observations. Moreover, the exponential moving average uses $\lambda = 0.97$. Calculated is the 10−day volatility.

Estimated Density and Fitted Normal Density for DJI

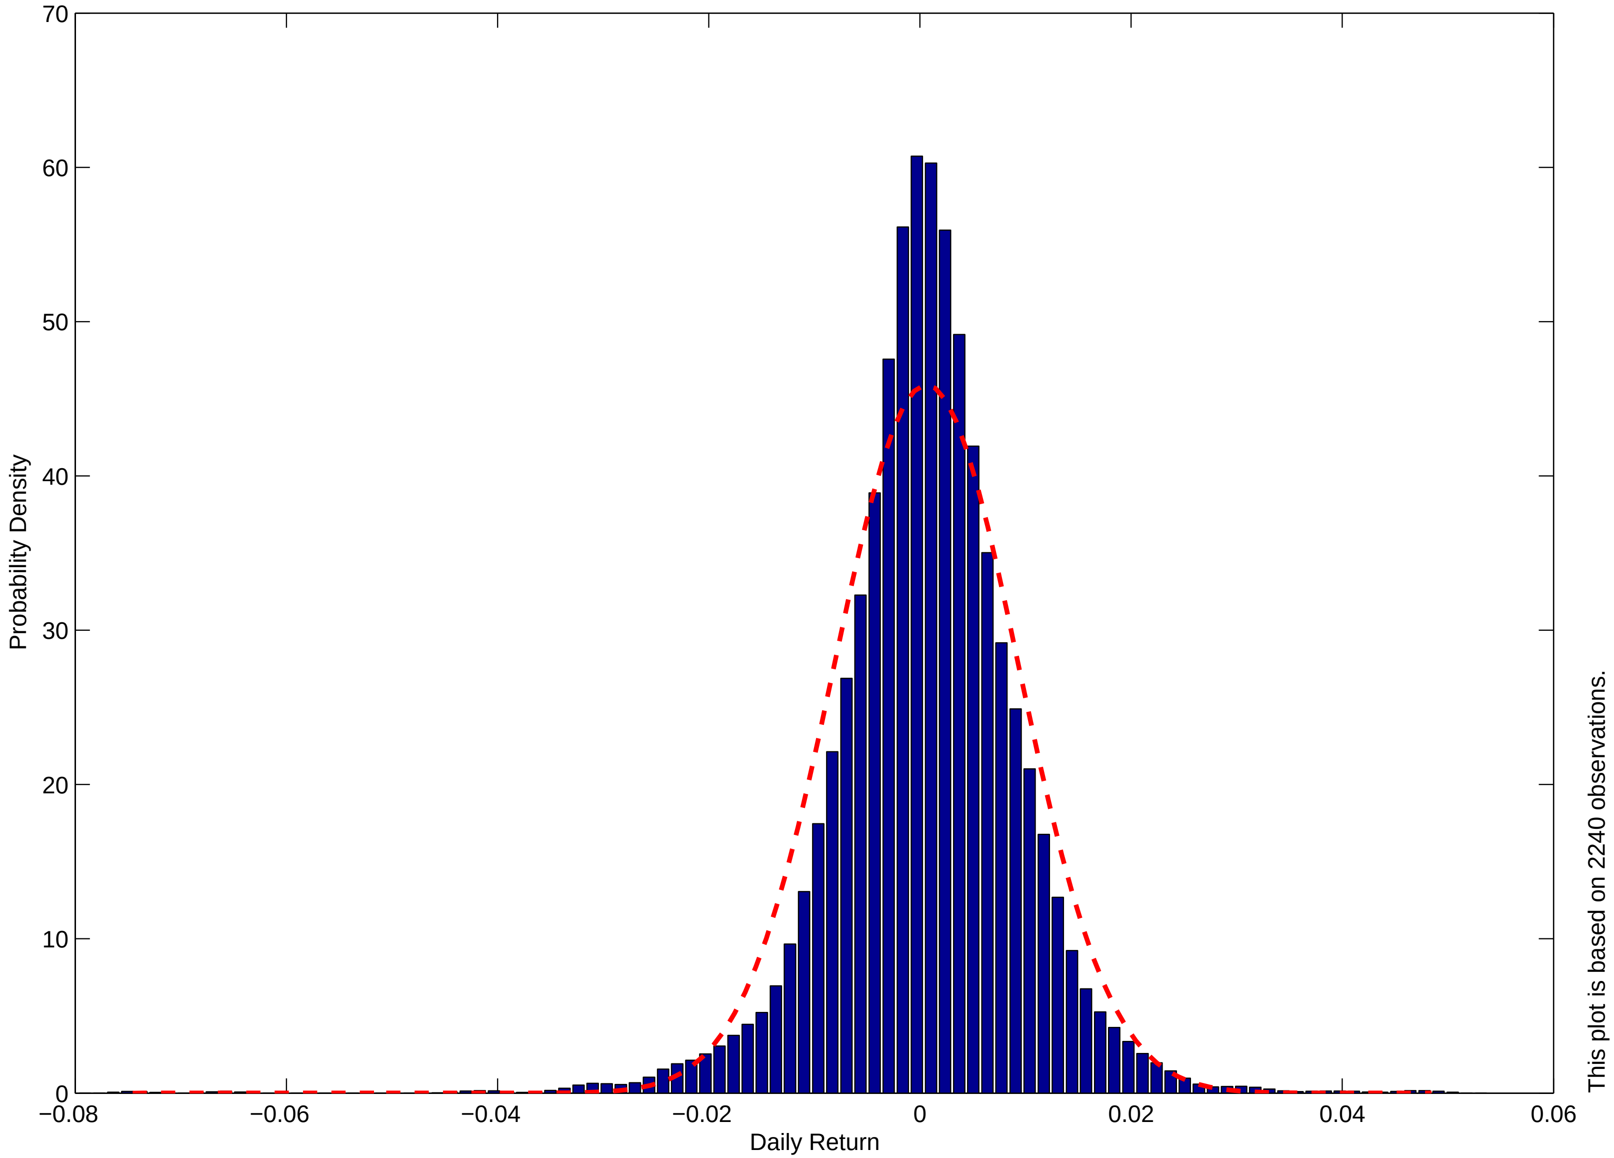

Left Tail of the Densities for DJI

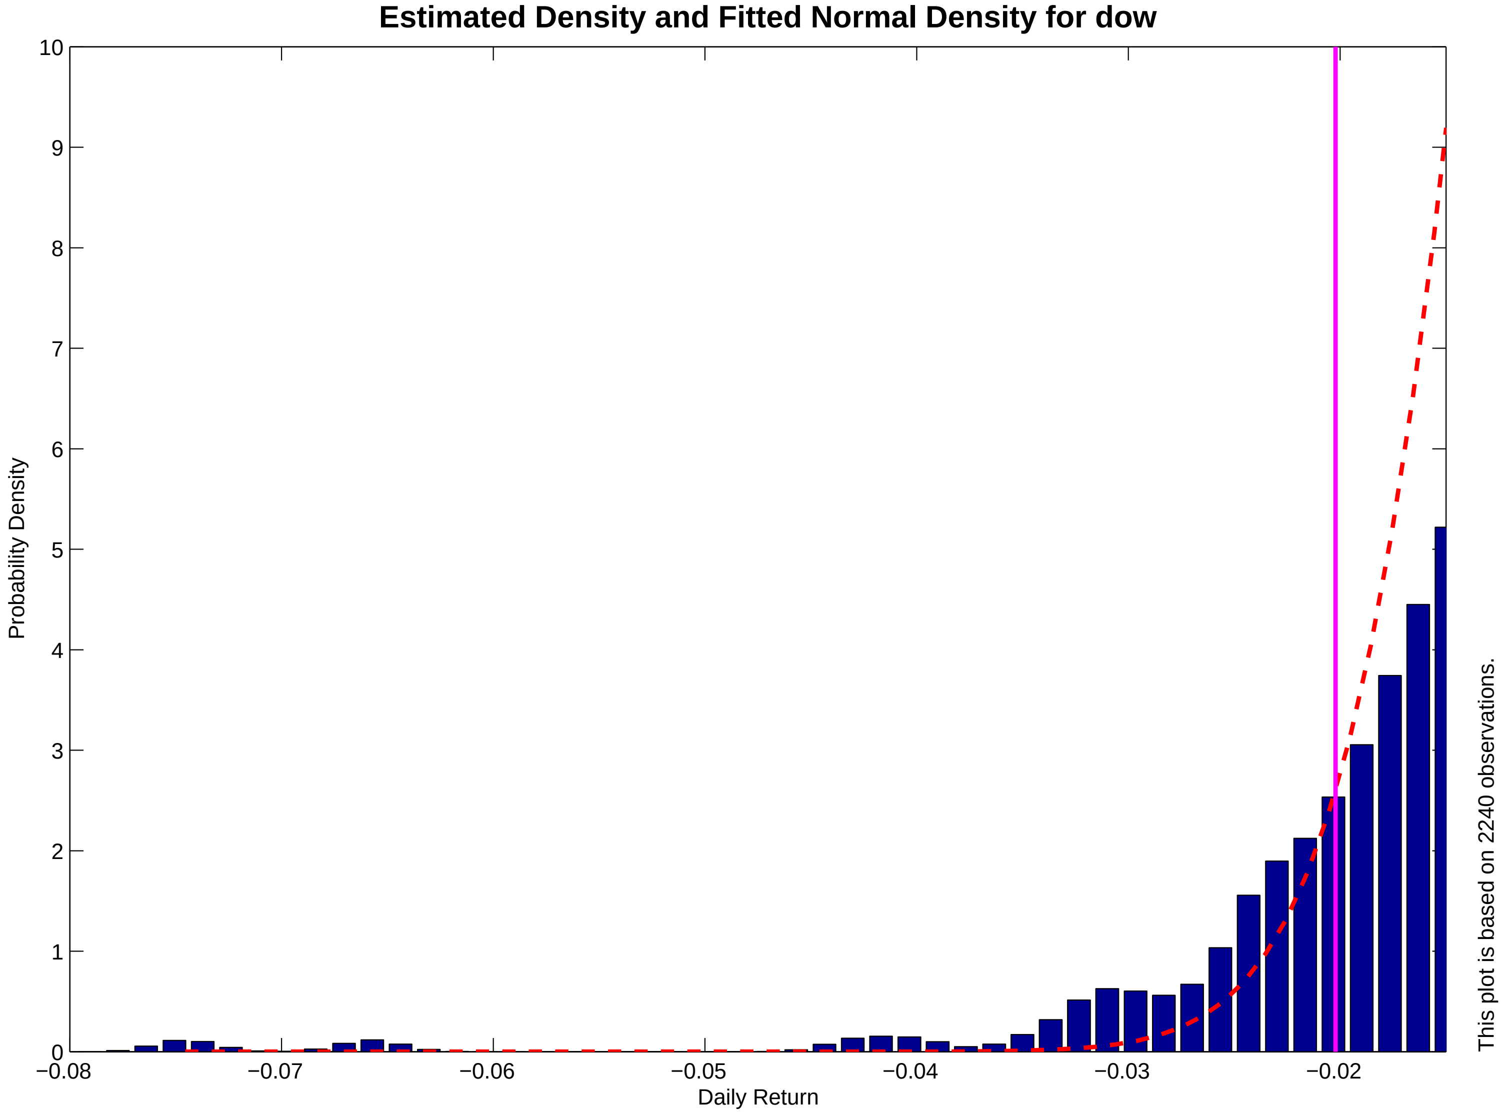

Estimated Density and Fitted Normal Density for DJI

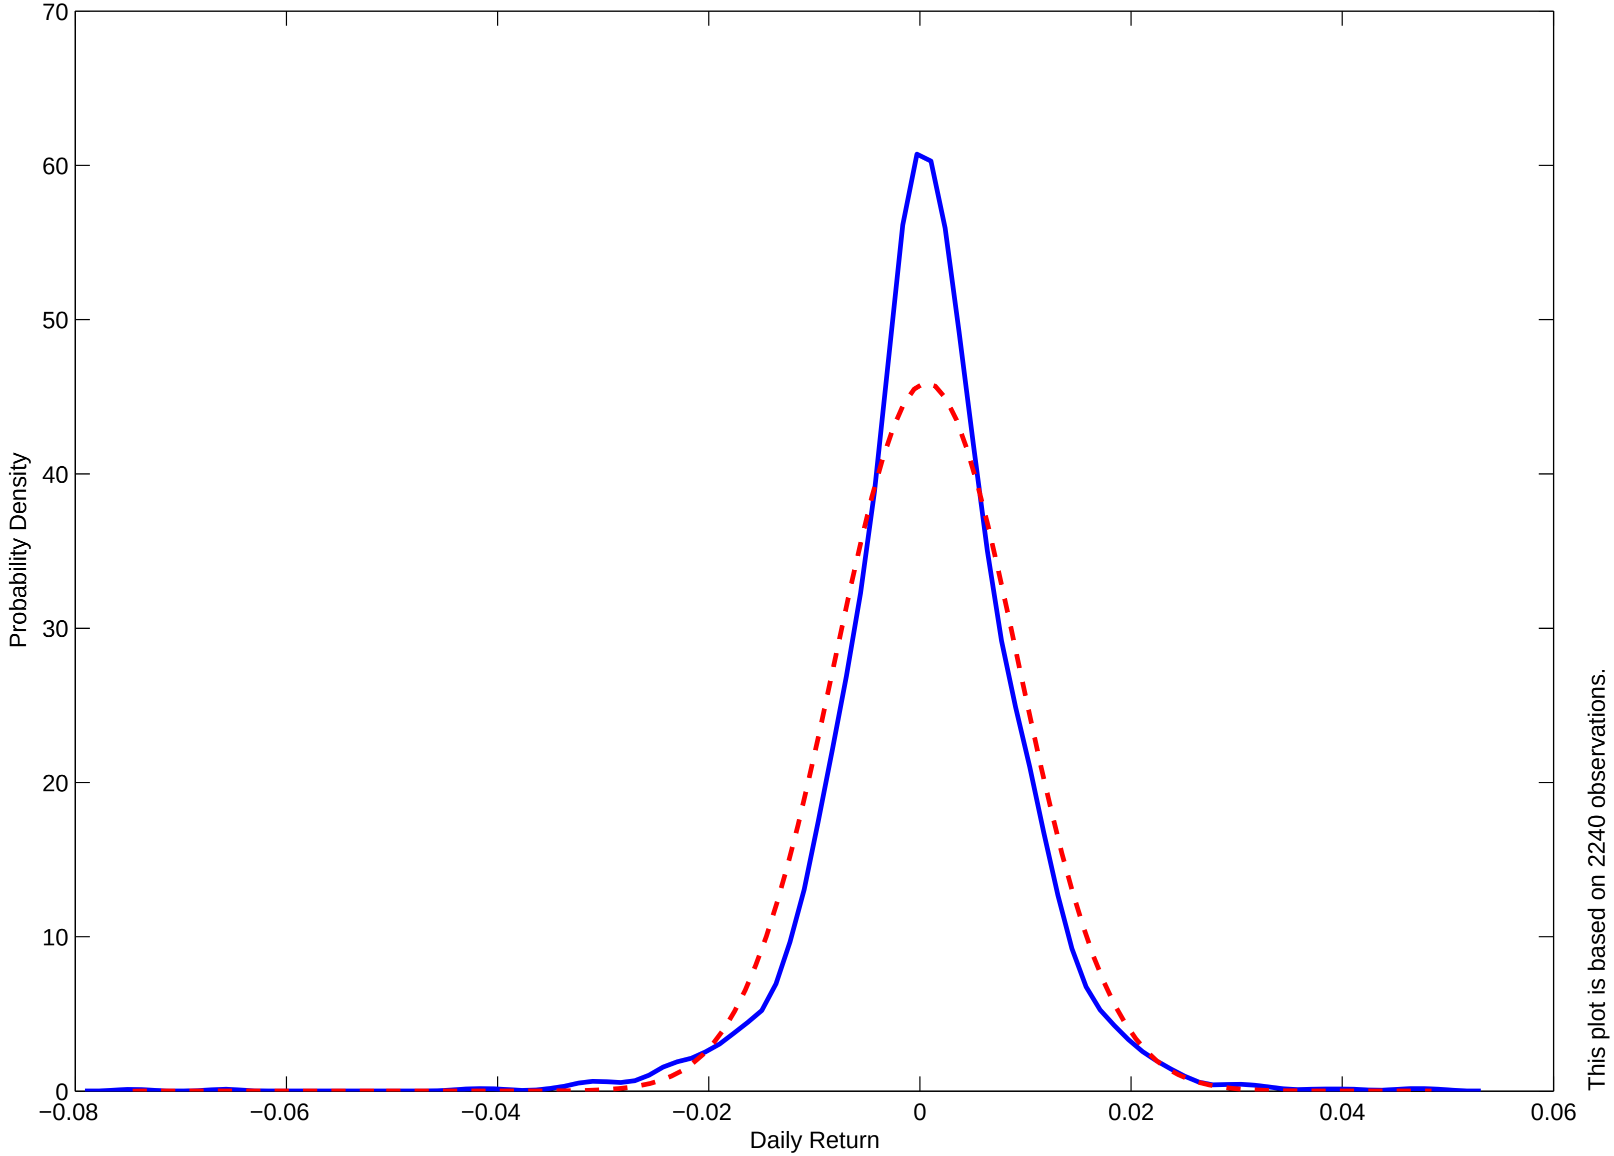

Estimated Density and Fitted Normal Density for eas

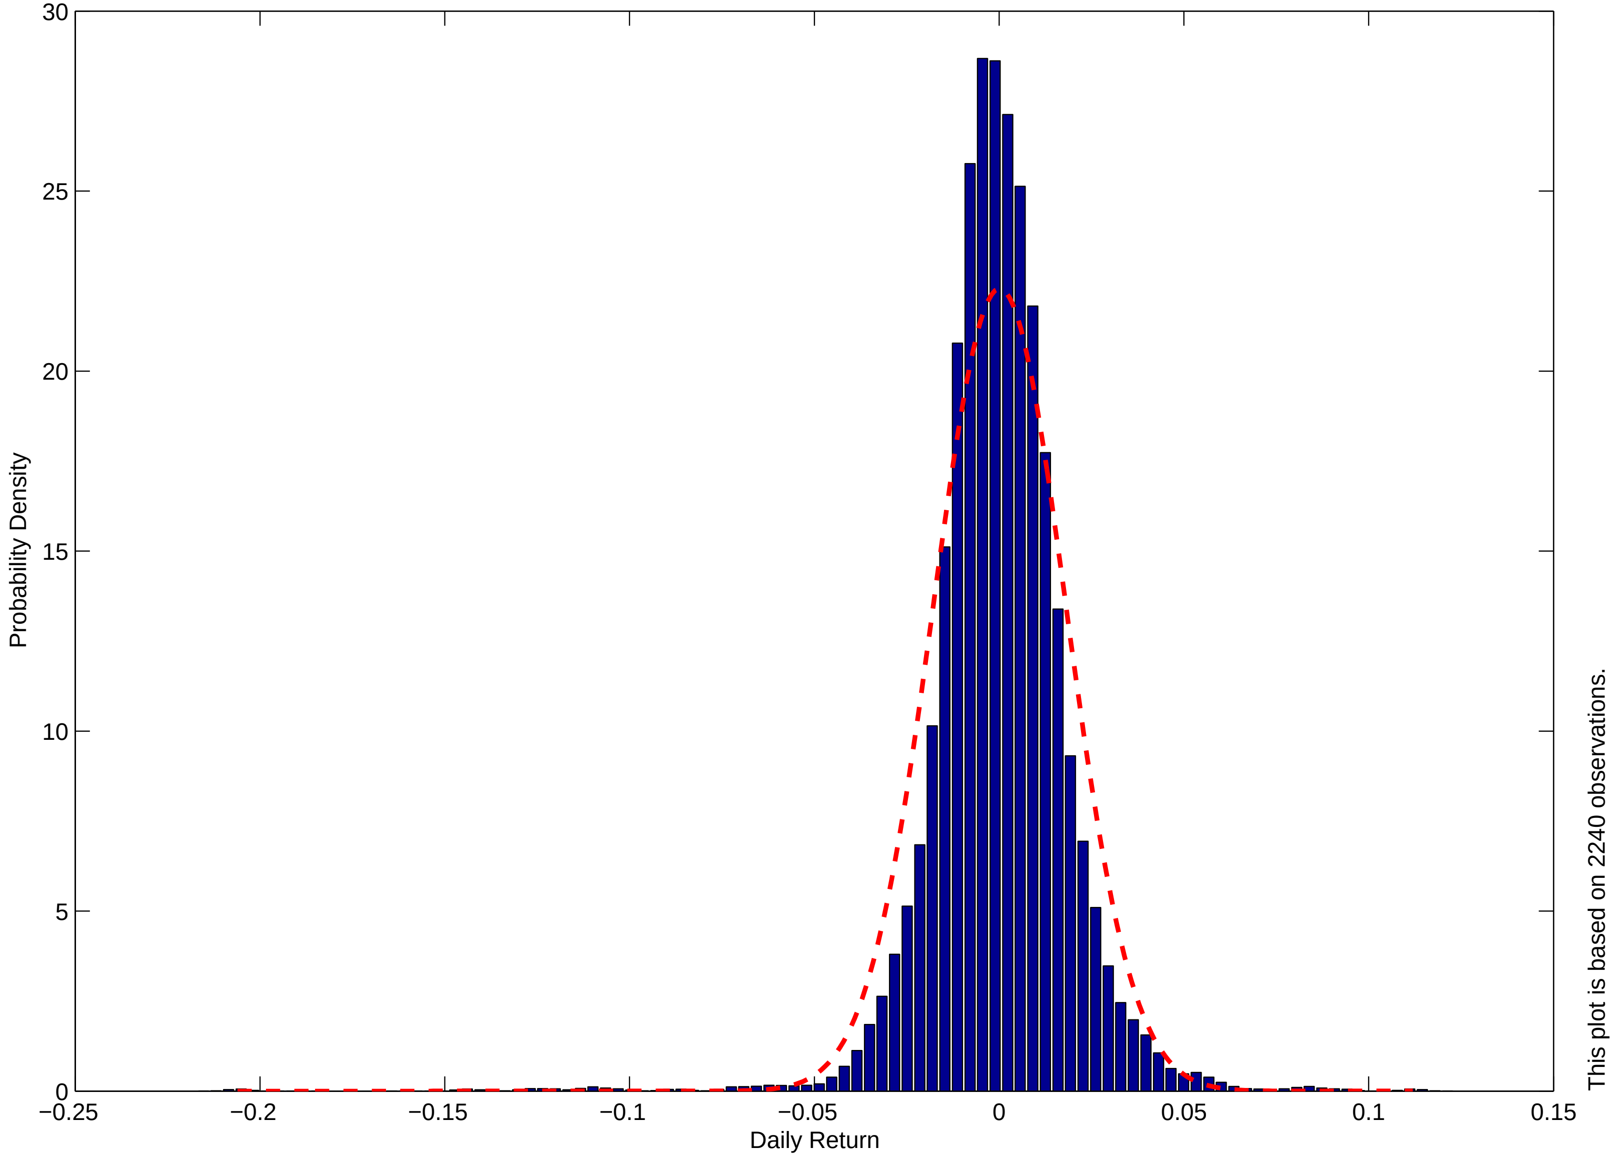

Estimated Density and Fitted Normal Density for eas

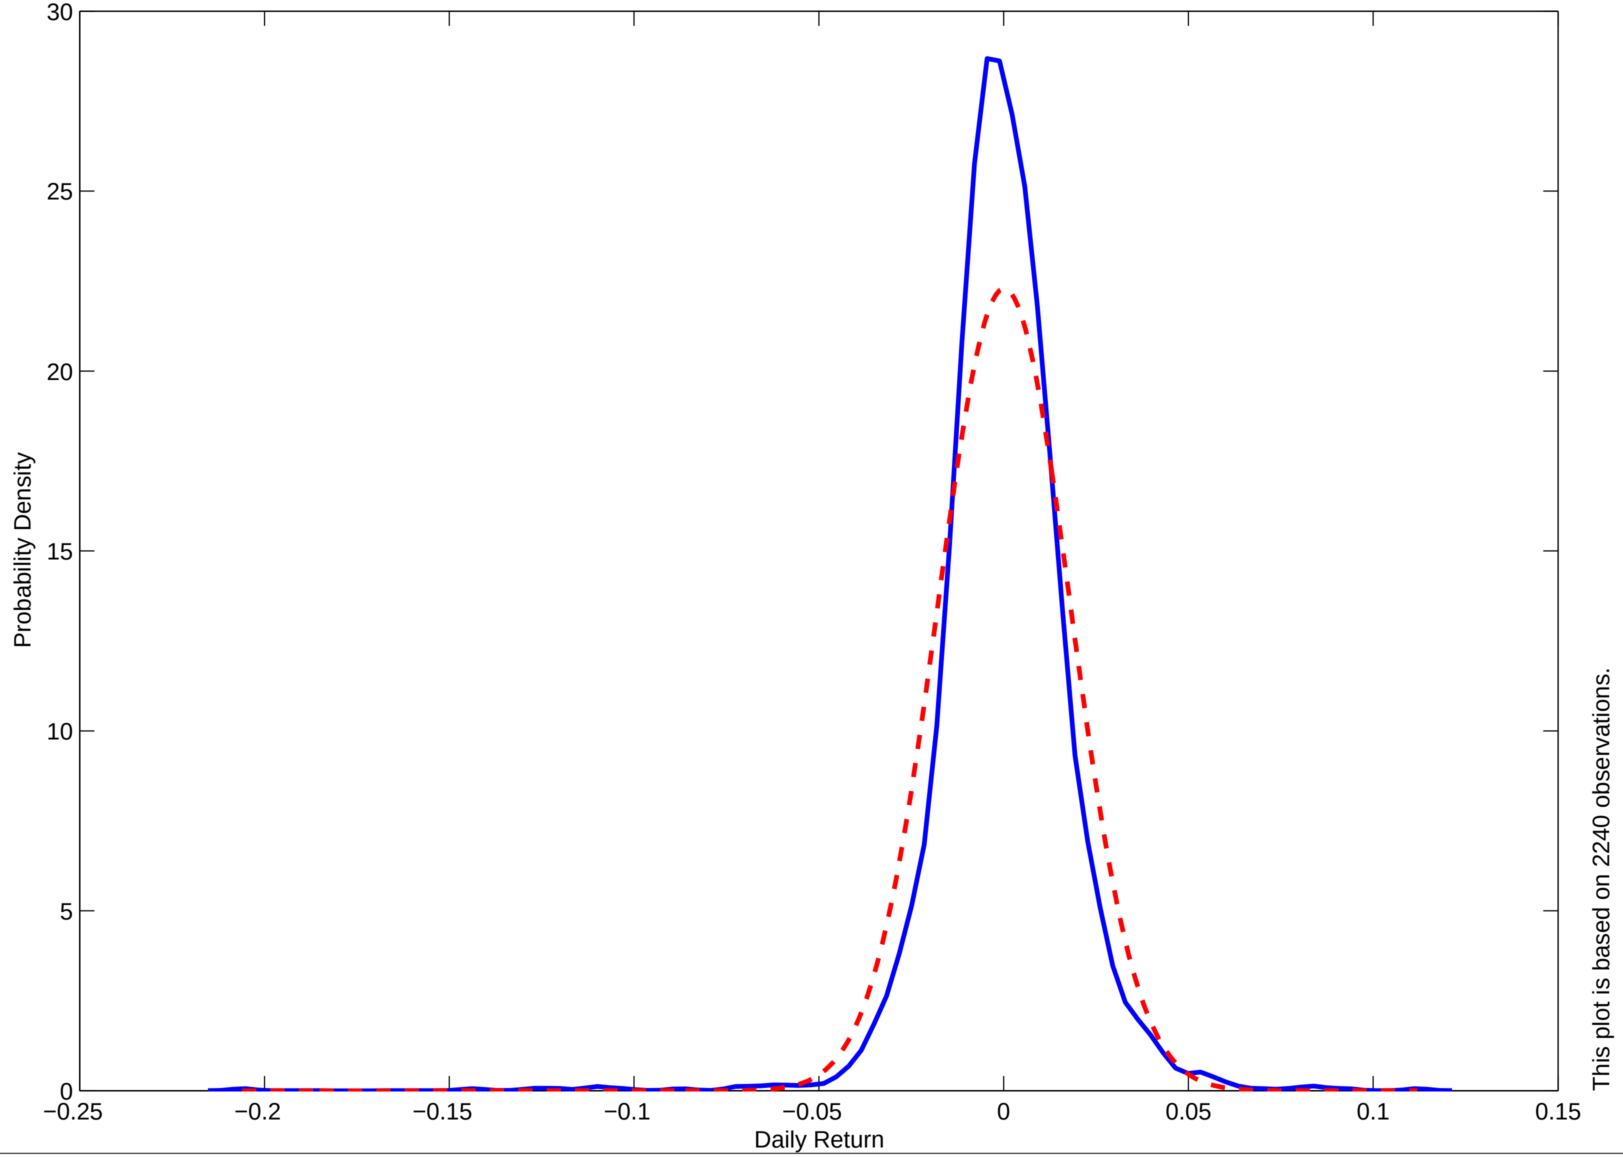

Historical Simulation

Historical simulation is based on order statistics. Given 100 observations the 99 percent quantile of the d–day returns is simply the lowest observation.

Let l be the number, which represents the q–th quantile of the order statistics with n observations. With this, x_l is the q–th quantile of an ordered time series X, which consists of n observations with $q = \frac{l}{n}$.

Estimating the q–quantile via order statistics is a generalisation of the median (which is the 50 percent quantile). While the median is in general a robust estimator, the robustness of the q–quantile depends on the quantile and the number of observations.

© Olaf Menkens

Observed Logarithmic Returns for DJI

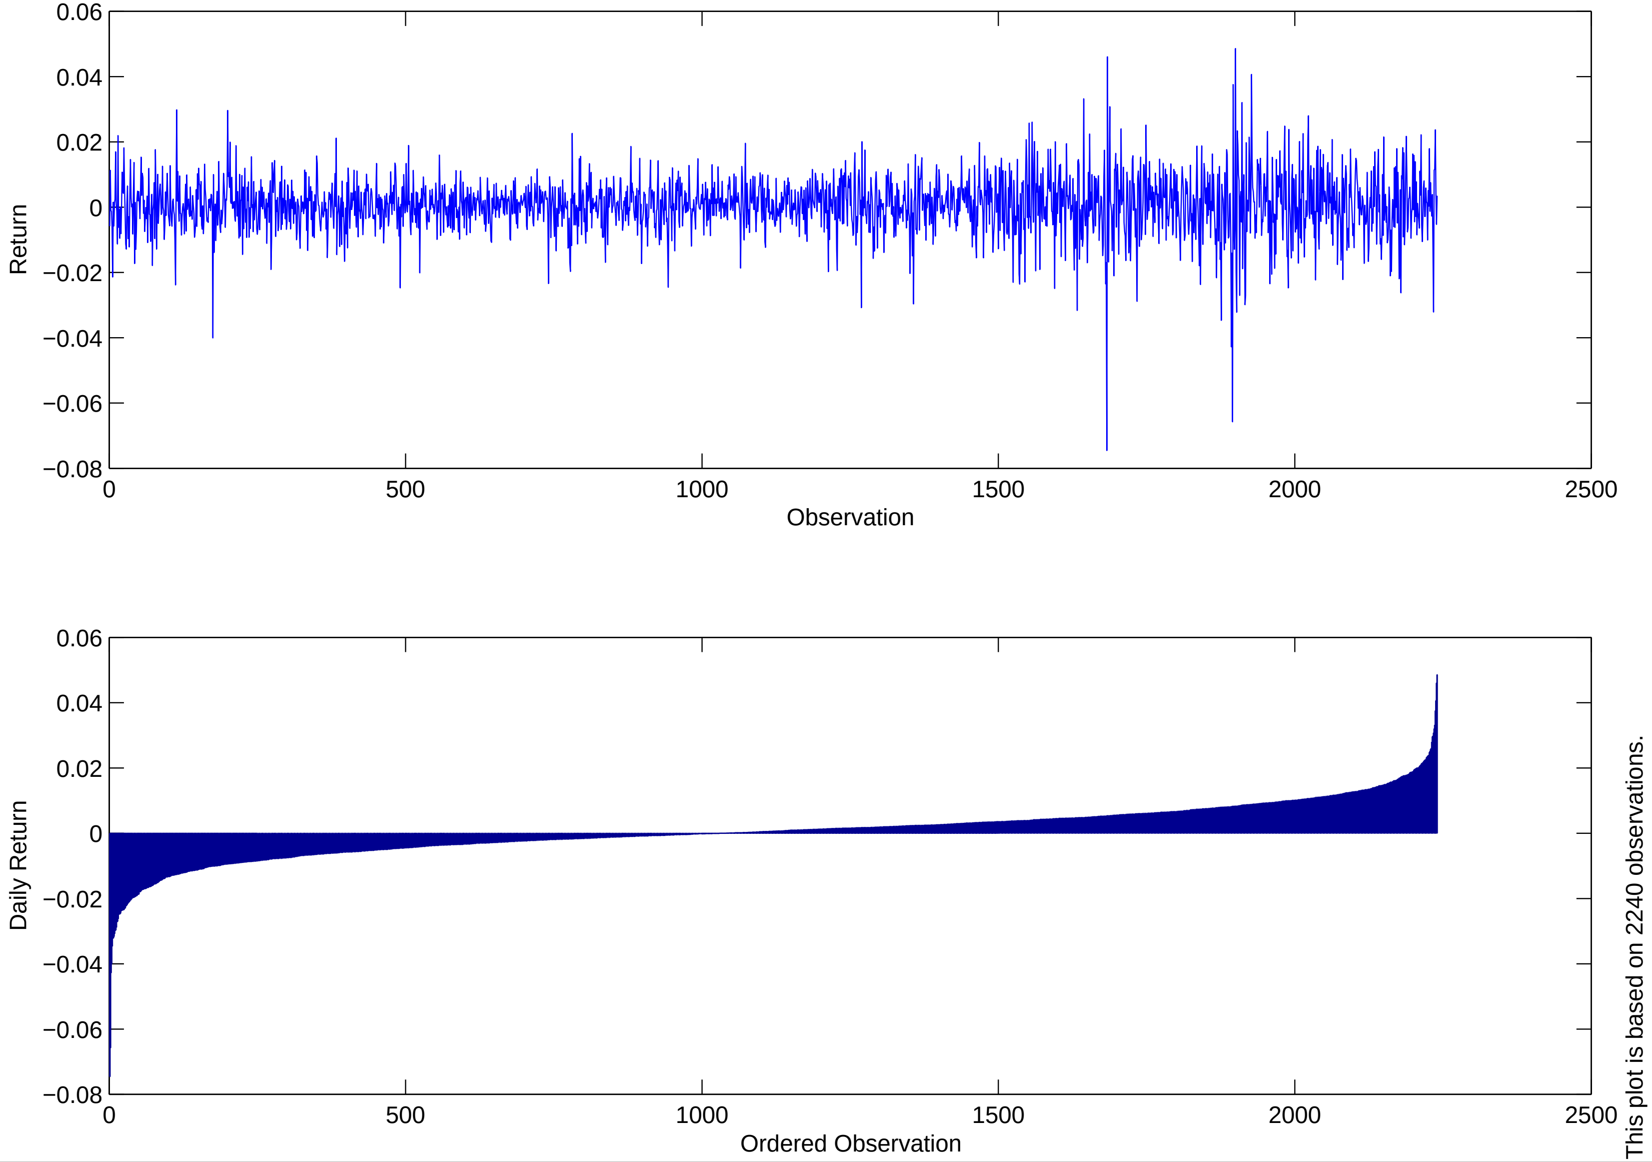

Observed Logarithmic Returns for eas

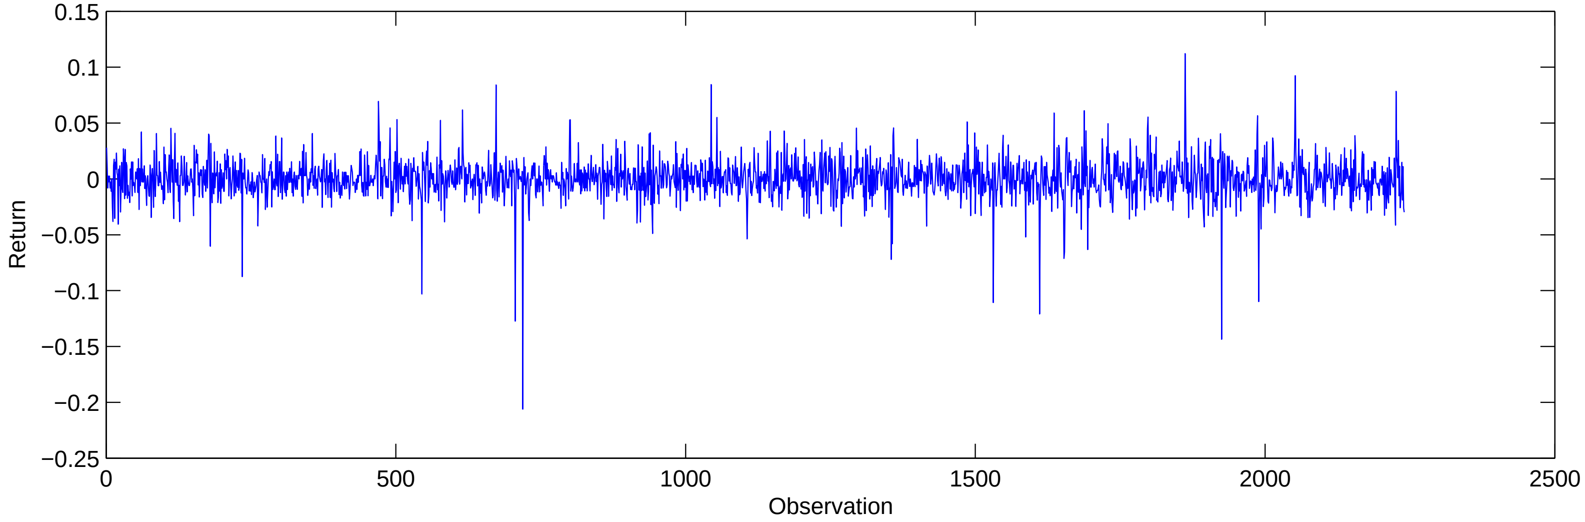

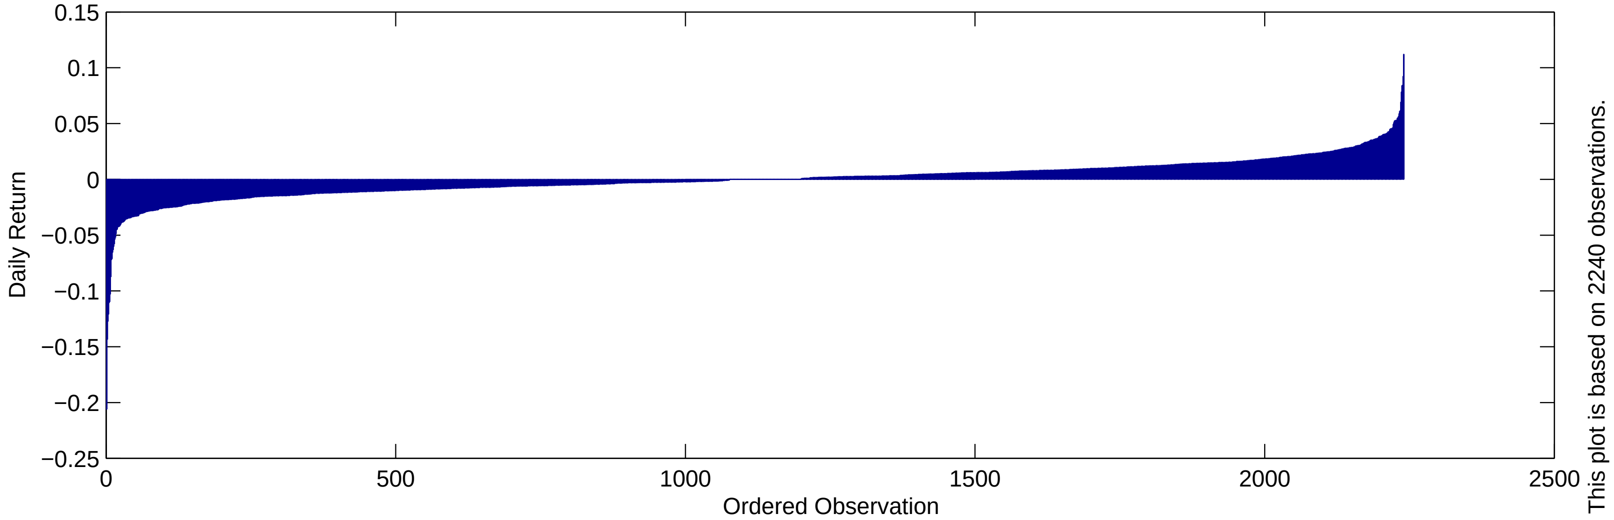

Historical Simulation – Error Estimation I

(i) $1-\alpha$ confidence interval

Hartung et al. [12] state that the $1-\alpha$ confidence interval for the q–th quantile of an order statistics, which is based on n points, is given approximately by $[x_r; x_s]$. Here r and s are the next higher natural numbers of

$$\begin{aligned} r^* &= n \cdot q - u_{1-\alpha/2}\sqrt{n \cdot q(1-q)} \qquad \text{and} \\ s^* &= n \cdot q + u_{1-\alpha/2}\sqrt{n \cdot q(1-q)}\,, \text{ respectively.} \end{aligned}$$

The notation u_α has been used for the α–quantile of the $N(0,1)$–distribution. This approximation can be used, if $q \cdot (1-q) \cdot n > 9$. Therefore this approximation can be used up to $q = 0.01$, if $n > 910$.

Obviously, these confidence intervals are not symmetric, meaning that the distribution of the error of the quantile estimation is not symmetric and therefore not normal distributed.

Historical Simulation – Error Estimation II

However, the error of the quantile estimation is asymptotically normal distributed (see e.g. Stuard and Ord [13]). Thus for large n the error is approximately normal distributed and it is possible to estimate the error as follows.

(ii) **Approximating the standard error**
Let X be a stochastic process with a differentiable density function $f > 0$. Then Stuart and Ord showed, that the variance of x_l is

$$\boldsymbol{\sigma_{x_l}^2 = \frac{q \cdot (1-q)}{n \cdot (f(x_l))^2}},$$

where f is the density function of X and f must be strict greater than zero.

1. Example: **The normal distribution**

For example, if $X \sim \mathcal{N}(0, \sigma^2)$ the the error calculates to

$$\begin{aligned}\sigma_x &= \frac{q \cdot (1-q)}{n} \cdot \frac{2\pi \cdot \sigma^2}{\exp\left(-\frac{x^2}{\sigma^2}\right)} \\ &= \frac{q \cdot (1-q)}{n} \cdot \frac{2\pi \cdot \sigma^2}{\exp\left(-y^2\right)},\end{aligned}$$

where the substitution $\sigma \cdot y = x$ has been used. This shows, that the error is **not** independent of the variance of the underlying process, if this underlying process is normal distributed.

Some Typical Paths of the Brownian Motion

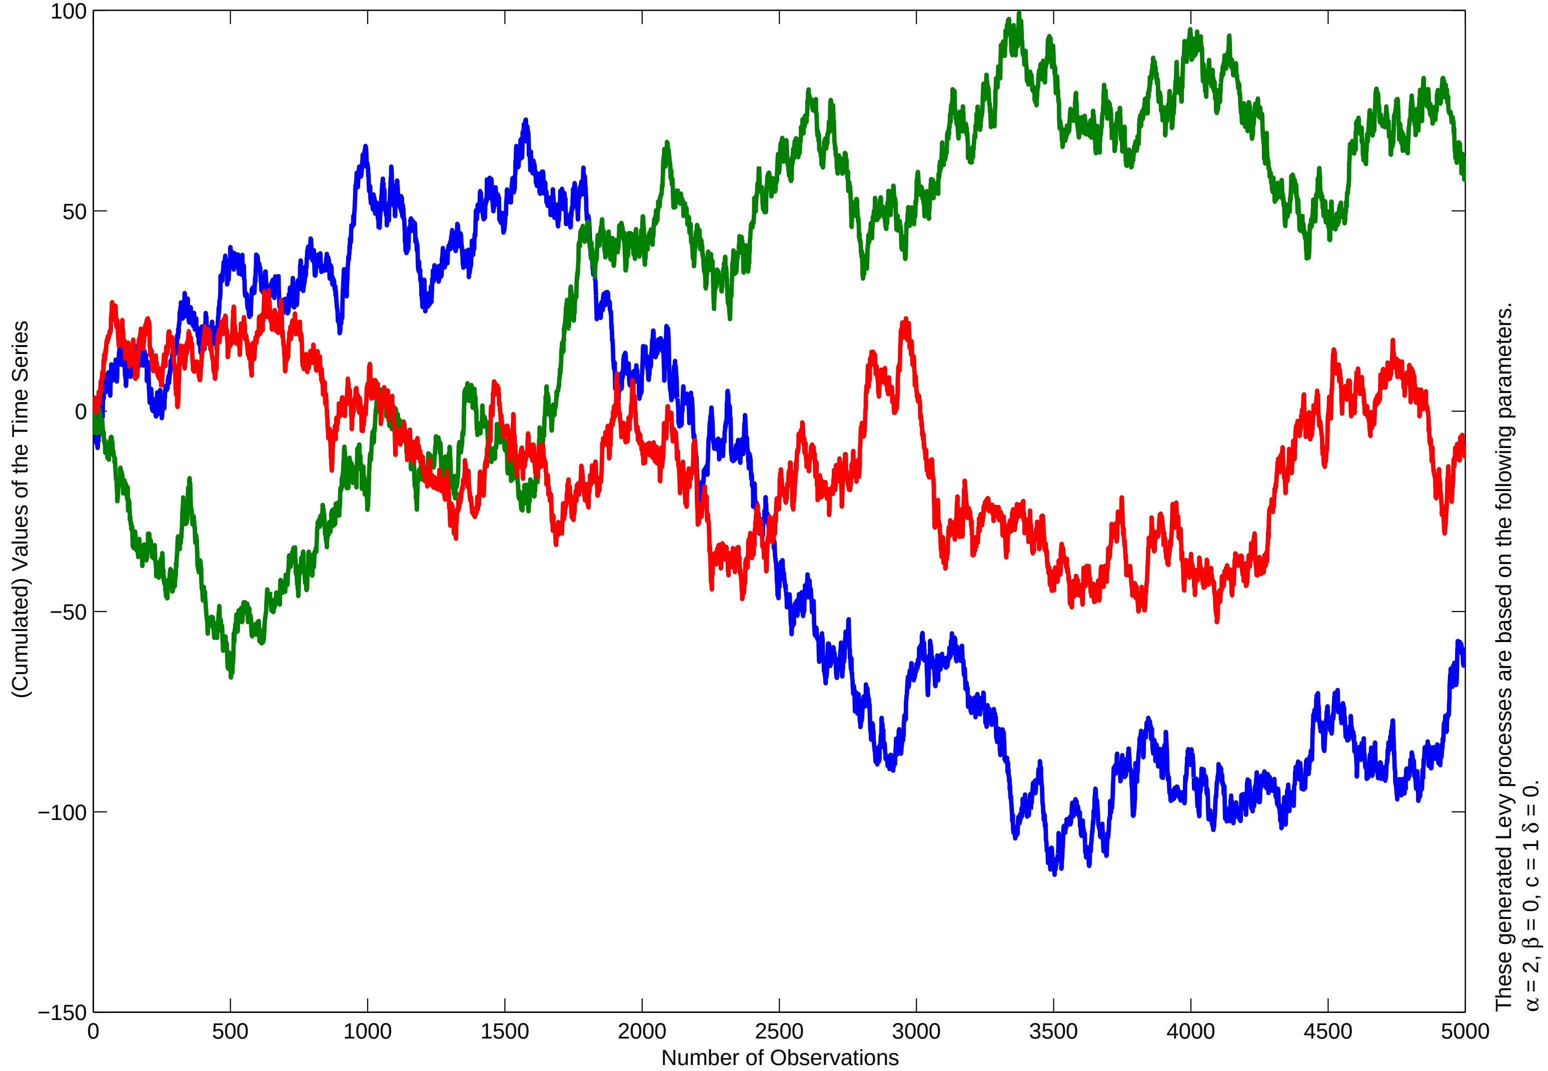

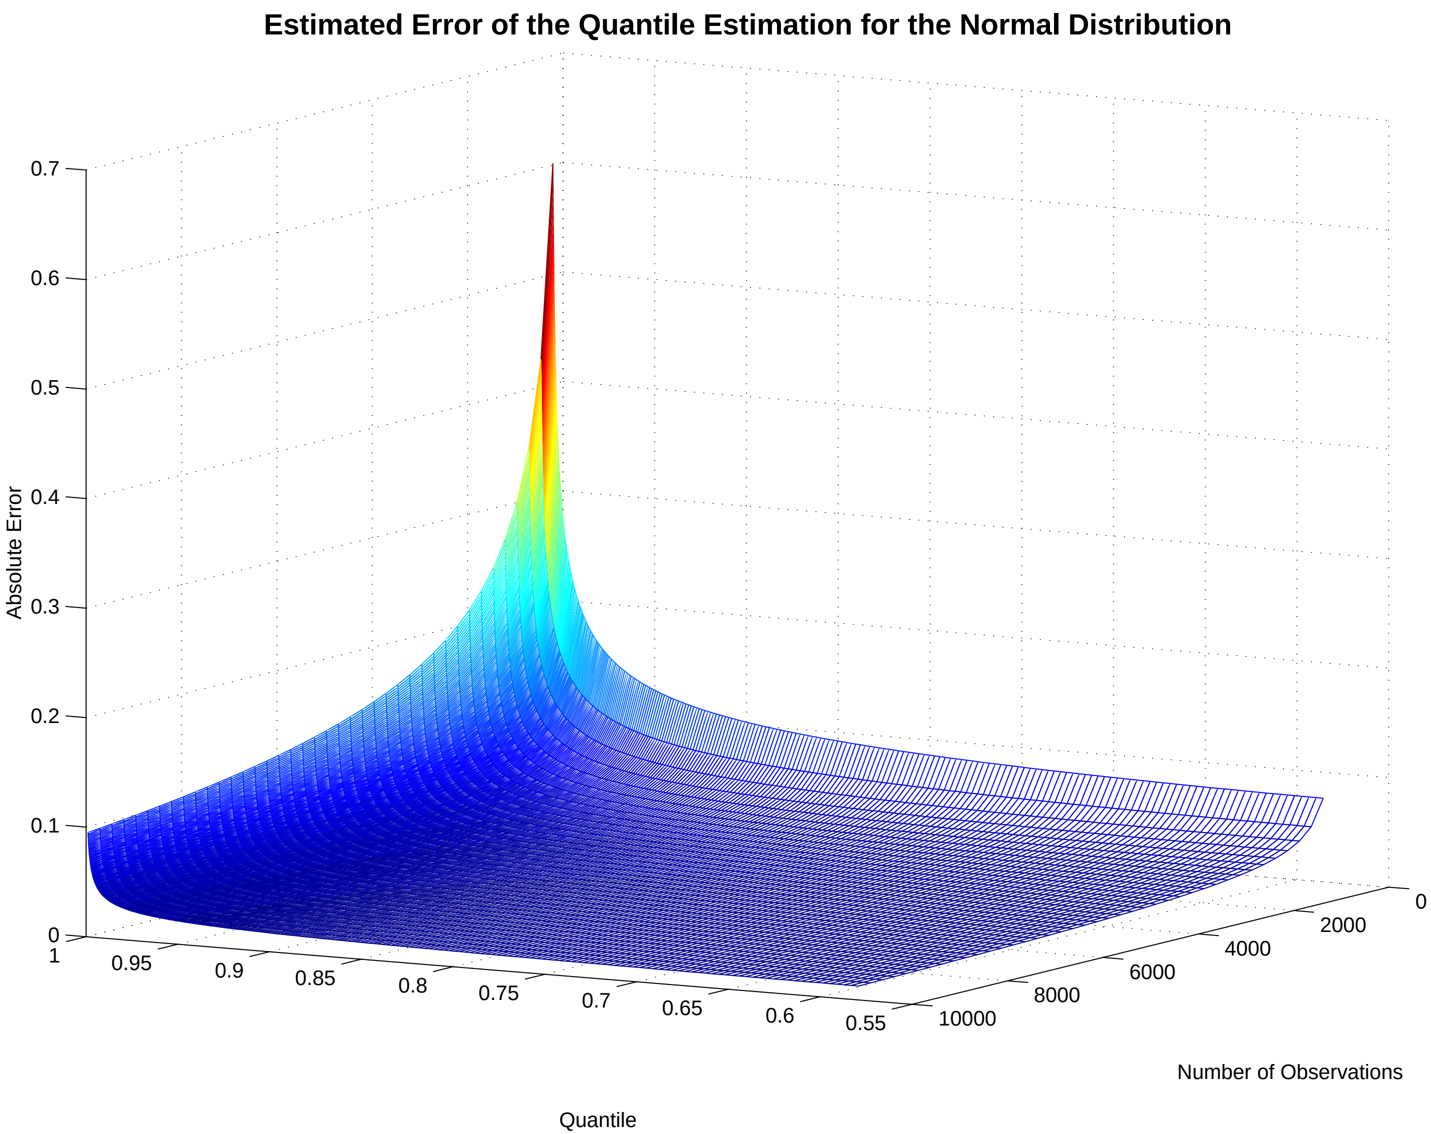
Estimated Error of the Quantile Estimation for the Normal Distribution
Absolute Error
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
0.55
Quantile
0
2000
4000
6000
8000
10000
Number of Observations

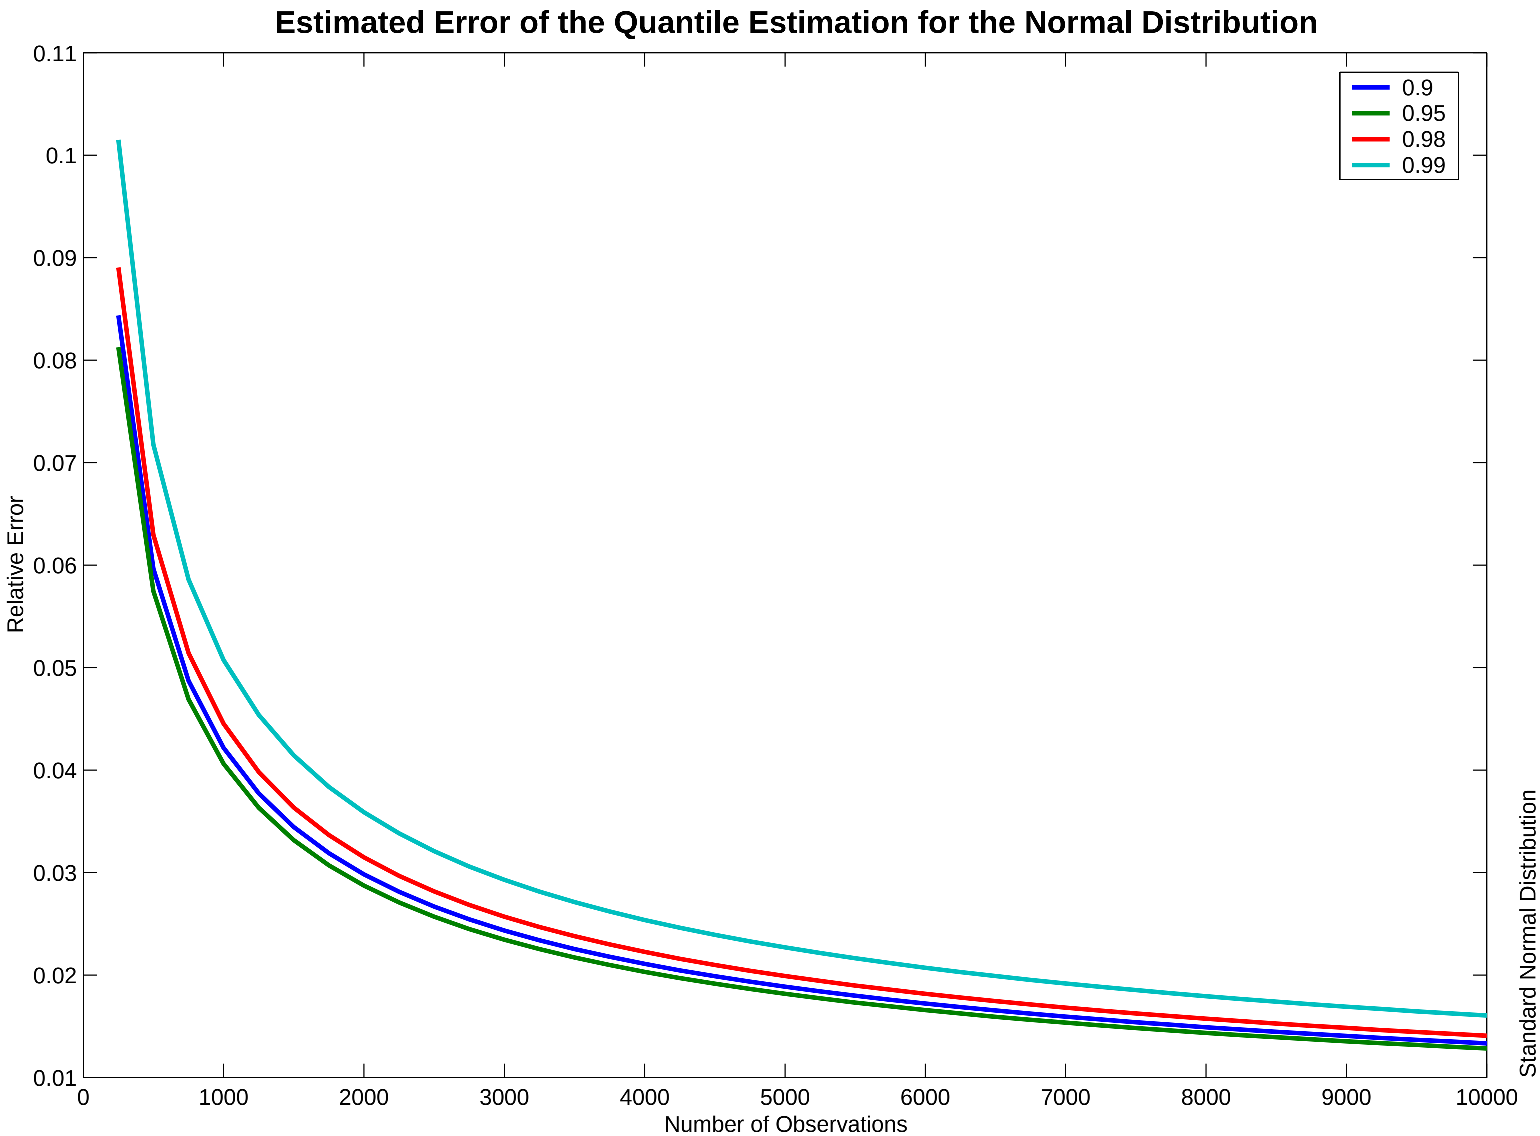
Estimated Error of the Quantile Estimation for the Normal Distribution
0.9
0.95
0.98
0.99
Relative Error
0.11
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
Number of Observations
Standard Normal Distribution

2. Example: **The Cauchy distribution**

Similarly, if X is Cauchy with mean zero, the propagation of the error can be shown to be

$$\begin{aligned}\sigma_x &= \frac{q\cdot(1-q)}{n}\cdot\frac{\pi^2\cdot(x^2+\gamma^2)^2}{\gamma^2}\\ &= \frac{q\cdot(1-q)}{n}\cdot\pi^2\cdot\gamma^2\cdot(y^2+1)\,,\end{aligned}$$

where the substitution $\gamma\cdot y = x$ has been used and γ is the scale parameter.

Some Typical Paths of the Cauchy Process

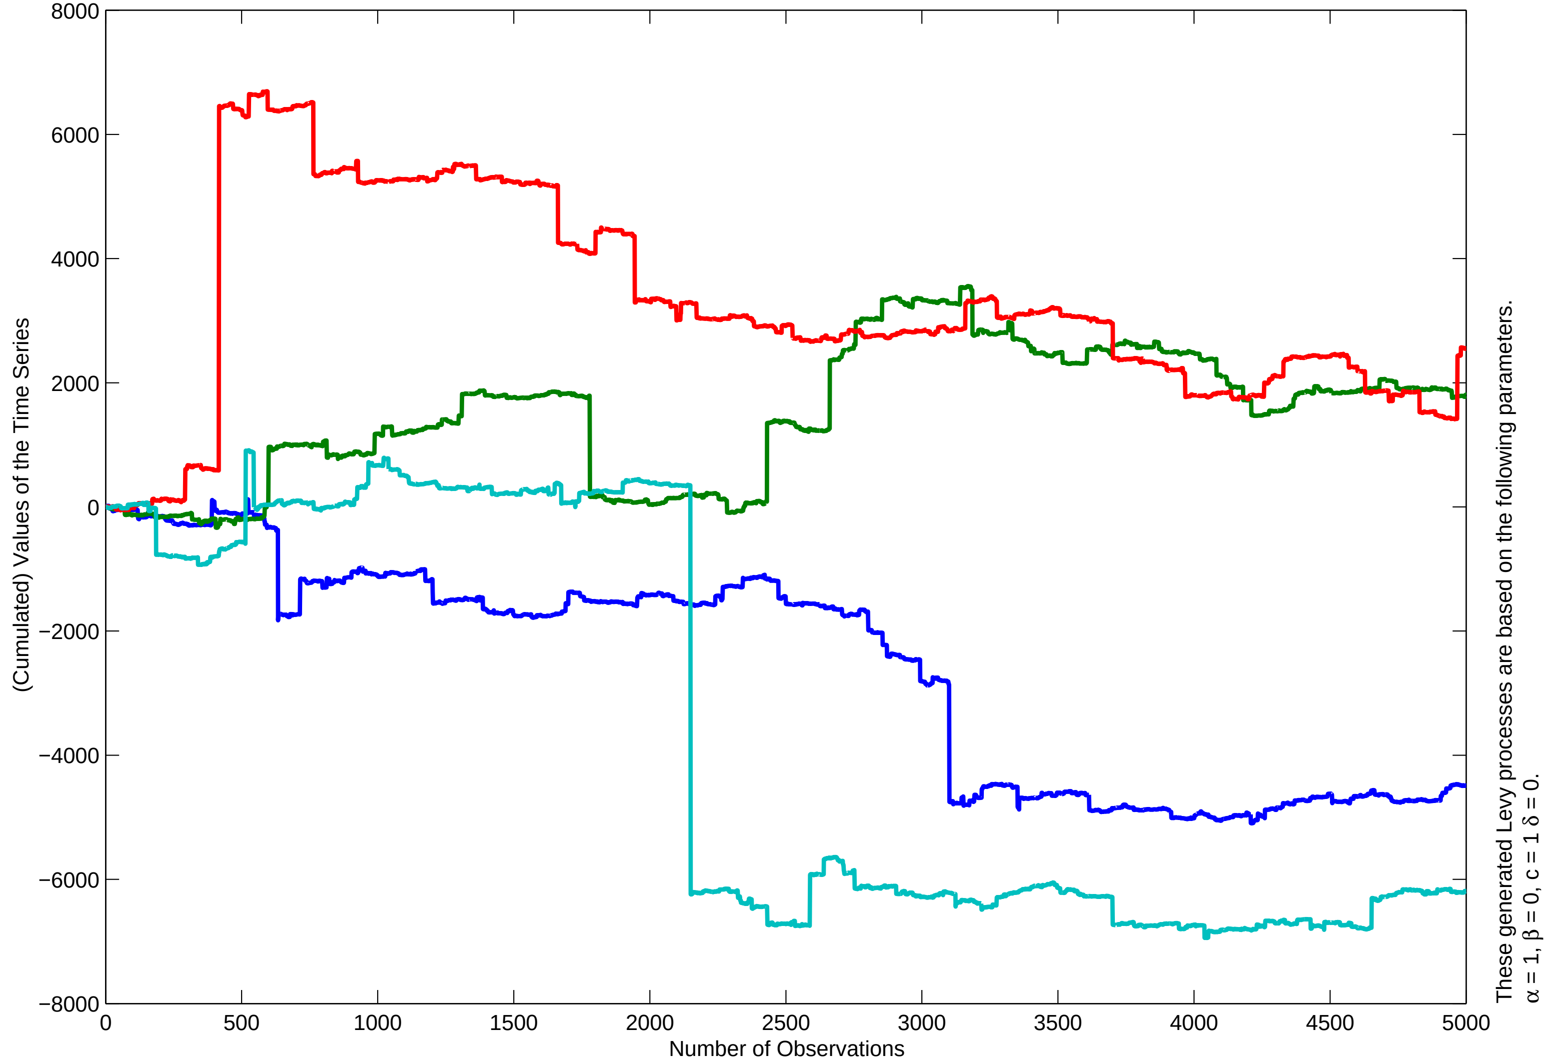

Some Other Typical Paths of the Cauchy Process

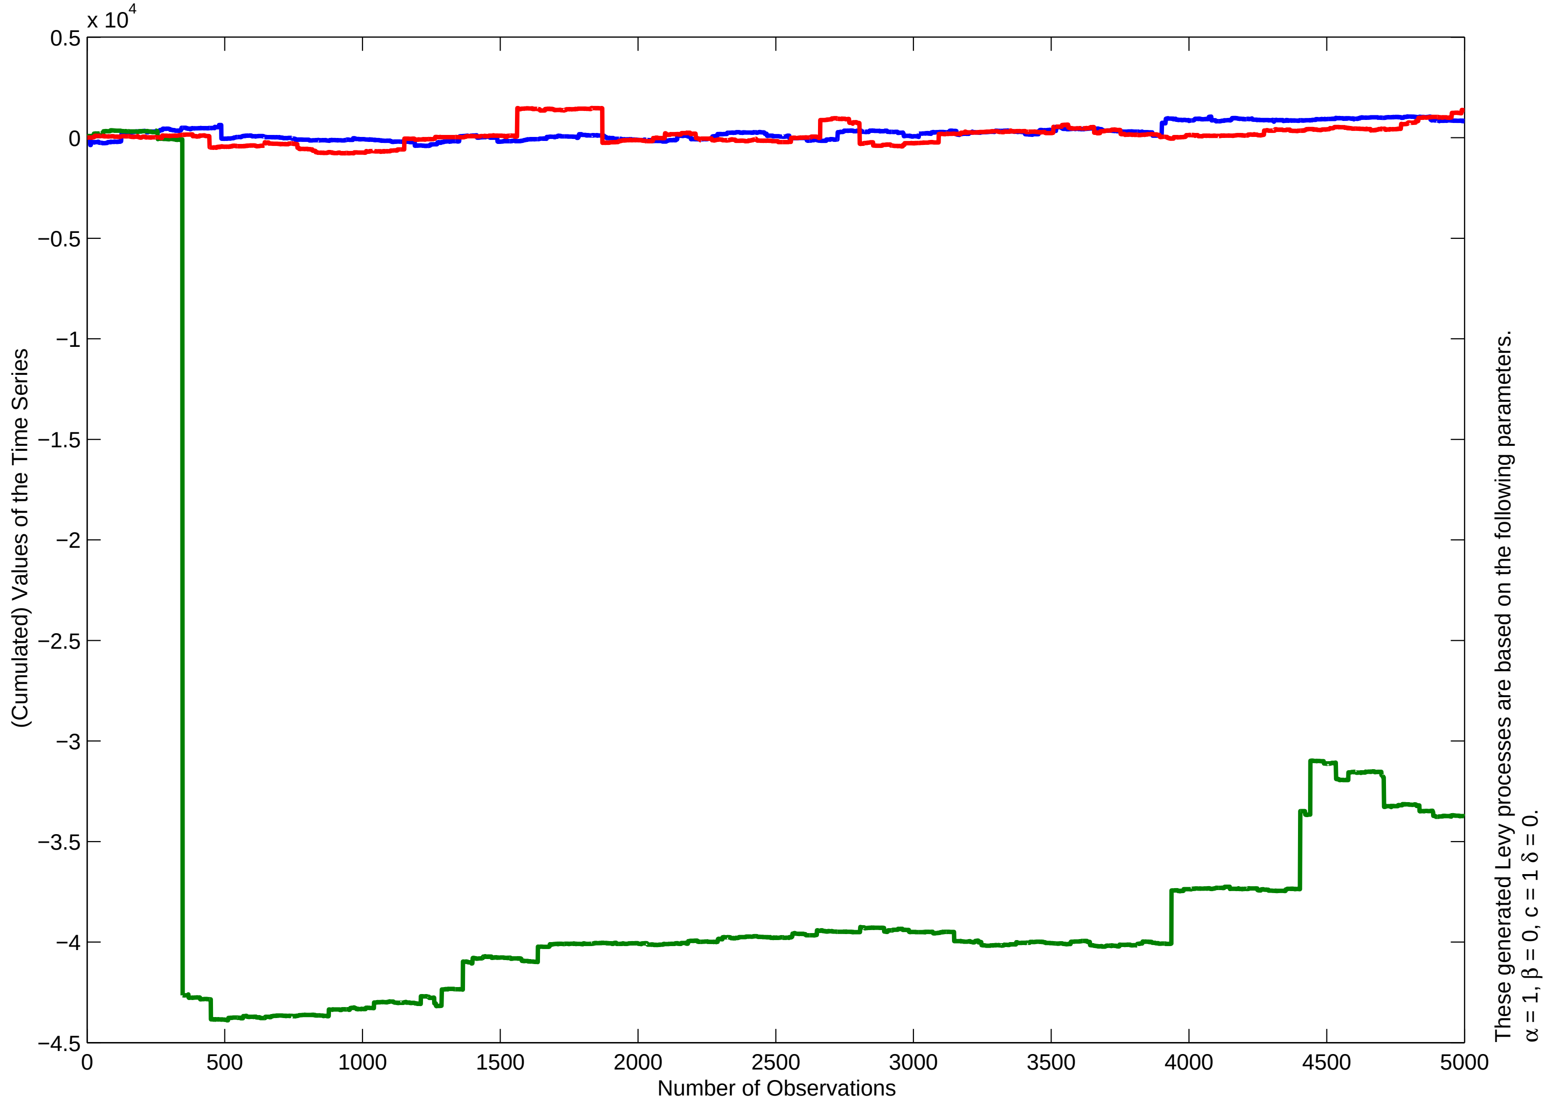

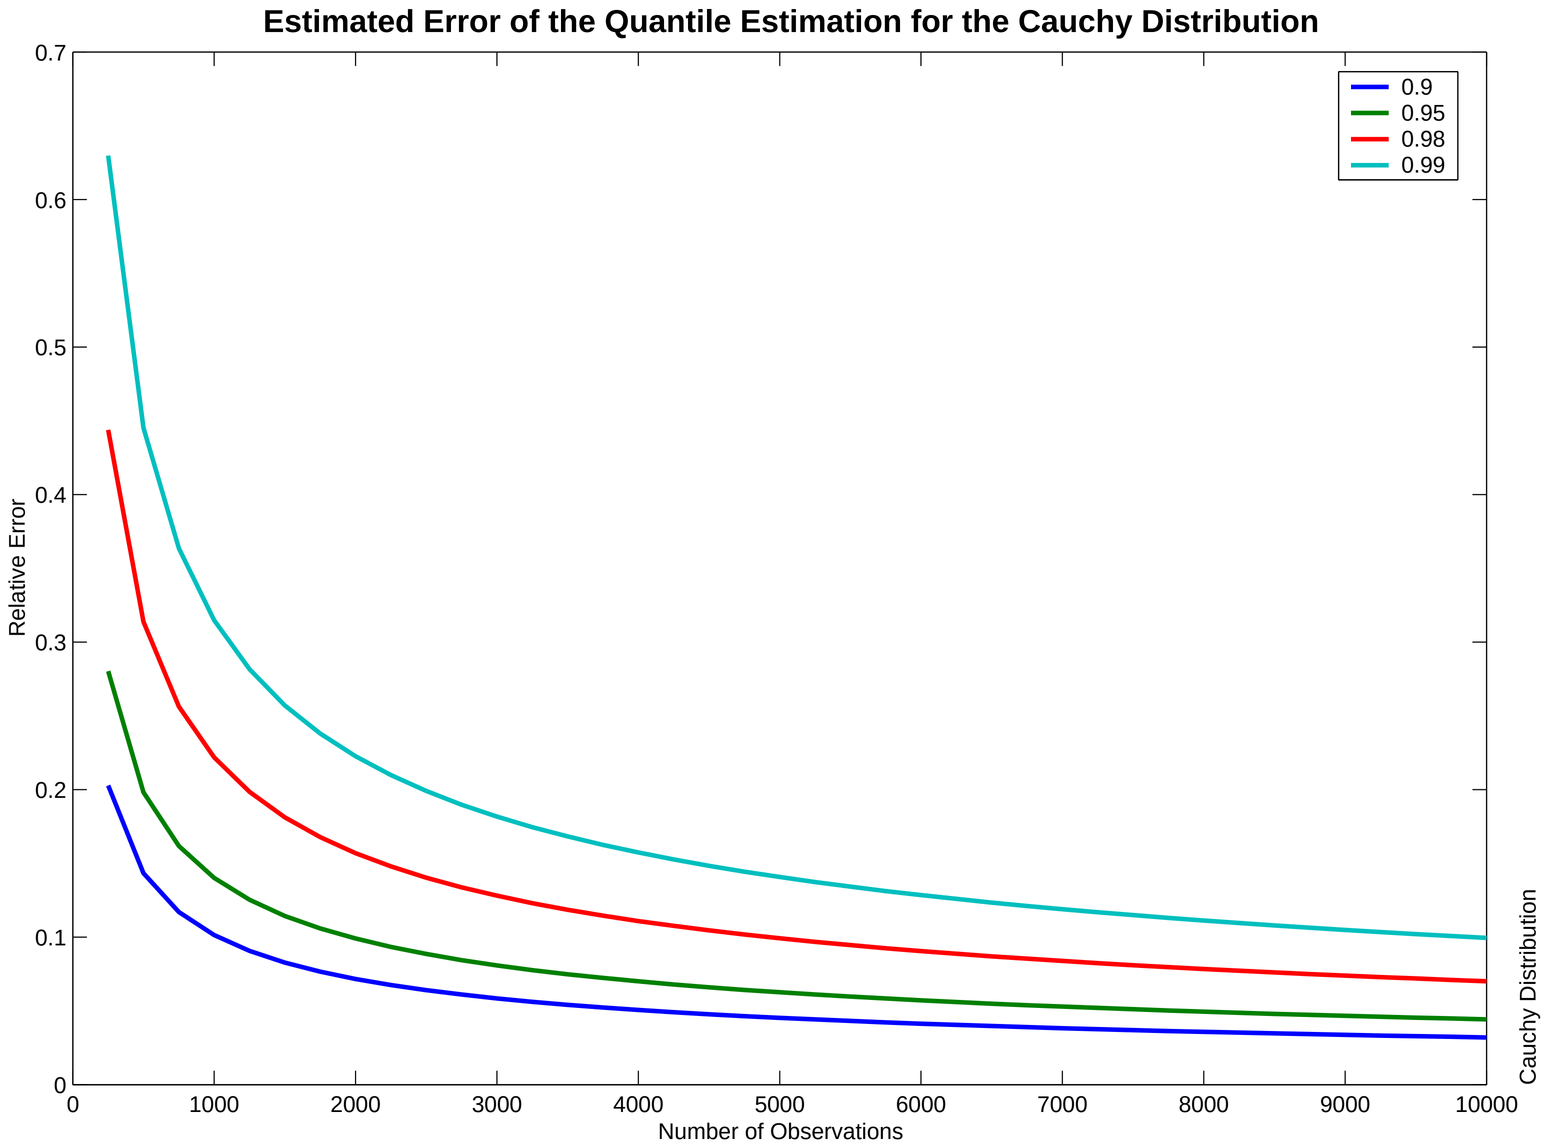
Estimated Error of the Quantile Estimation for the Cauchy Distribution
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Relative Error
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
Number of Observations
0.9
0.95
0.98
0.99
Cauchy Distribution

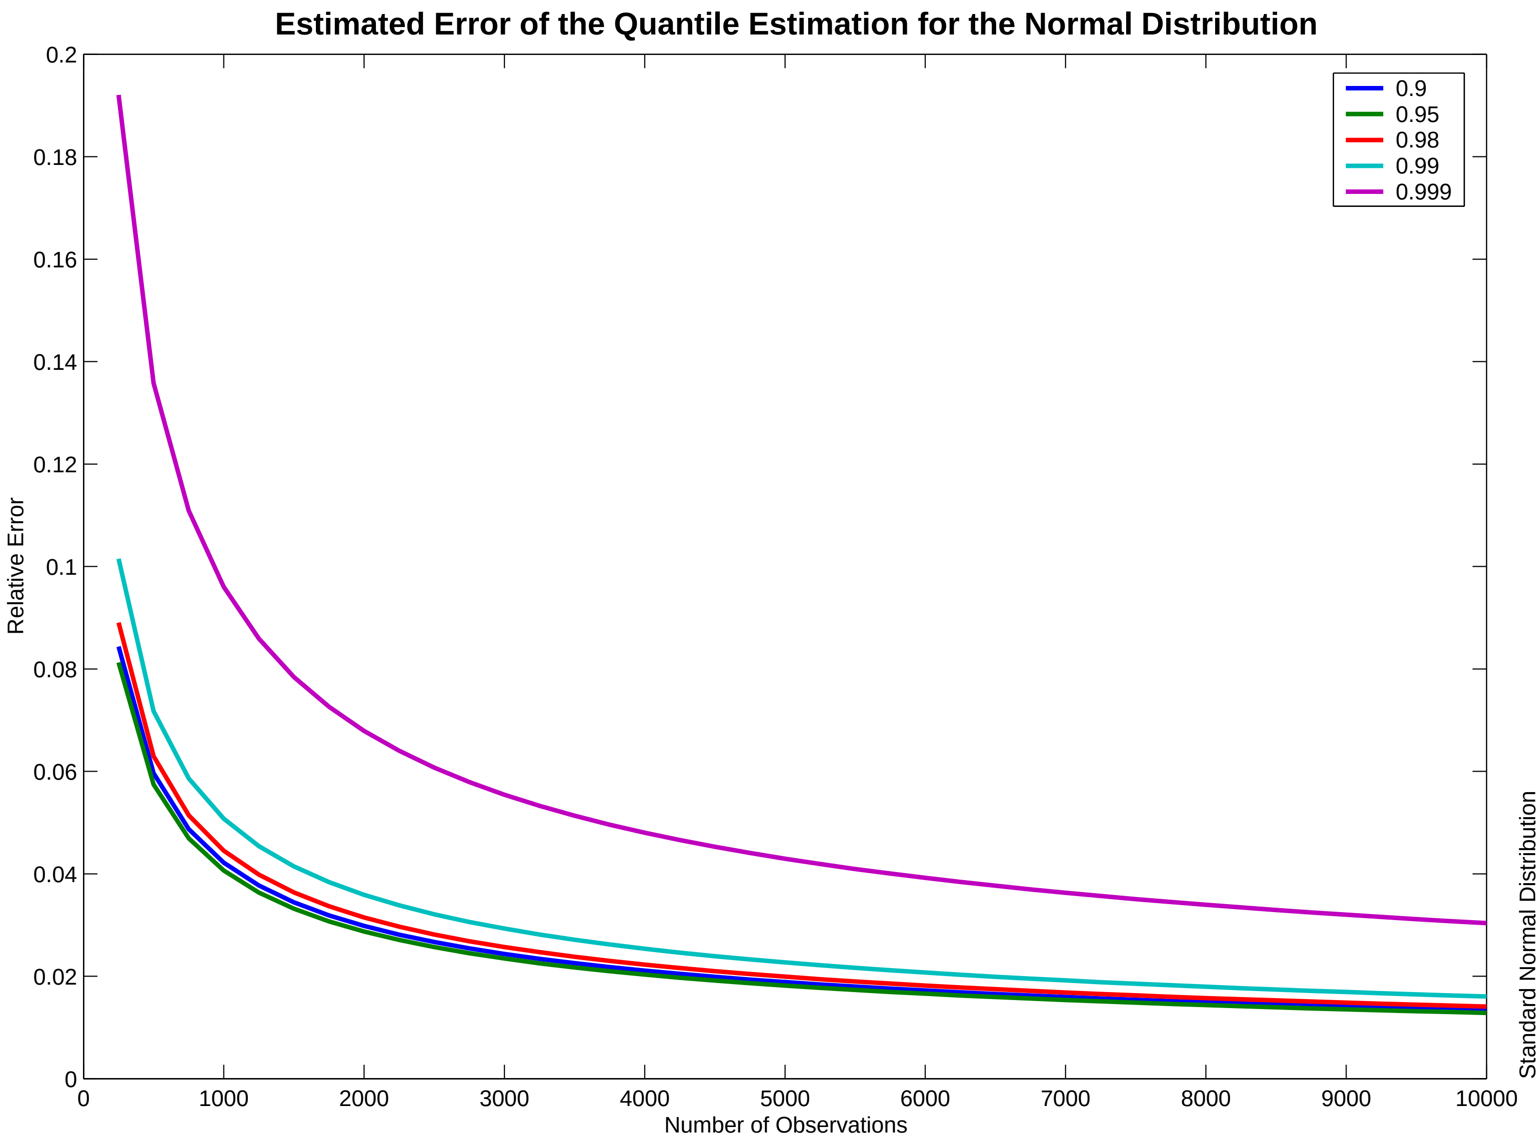
Estimated Error of the Quantile Estimation for the Normal Distribution
Relative Error
Number of Observations
Standard Normal Distribution
0.9
0.95
0.98
0.99
0.999
0
0.02
0.04
0.06
0.08
0.1
0.12
0.14
0.16
0.18
0.2
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

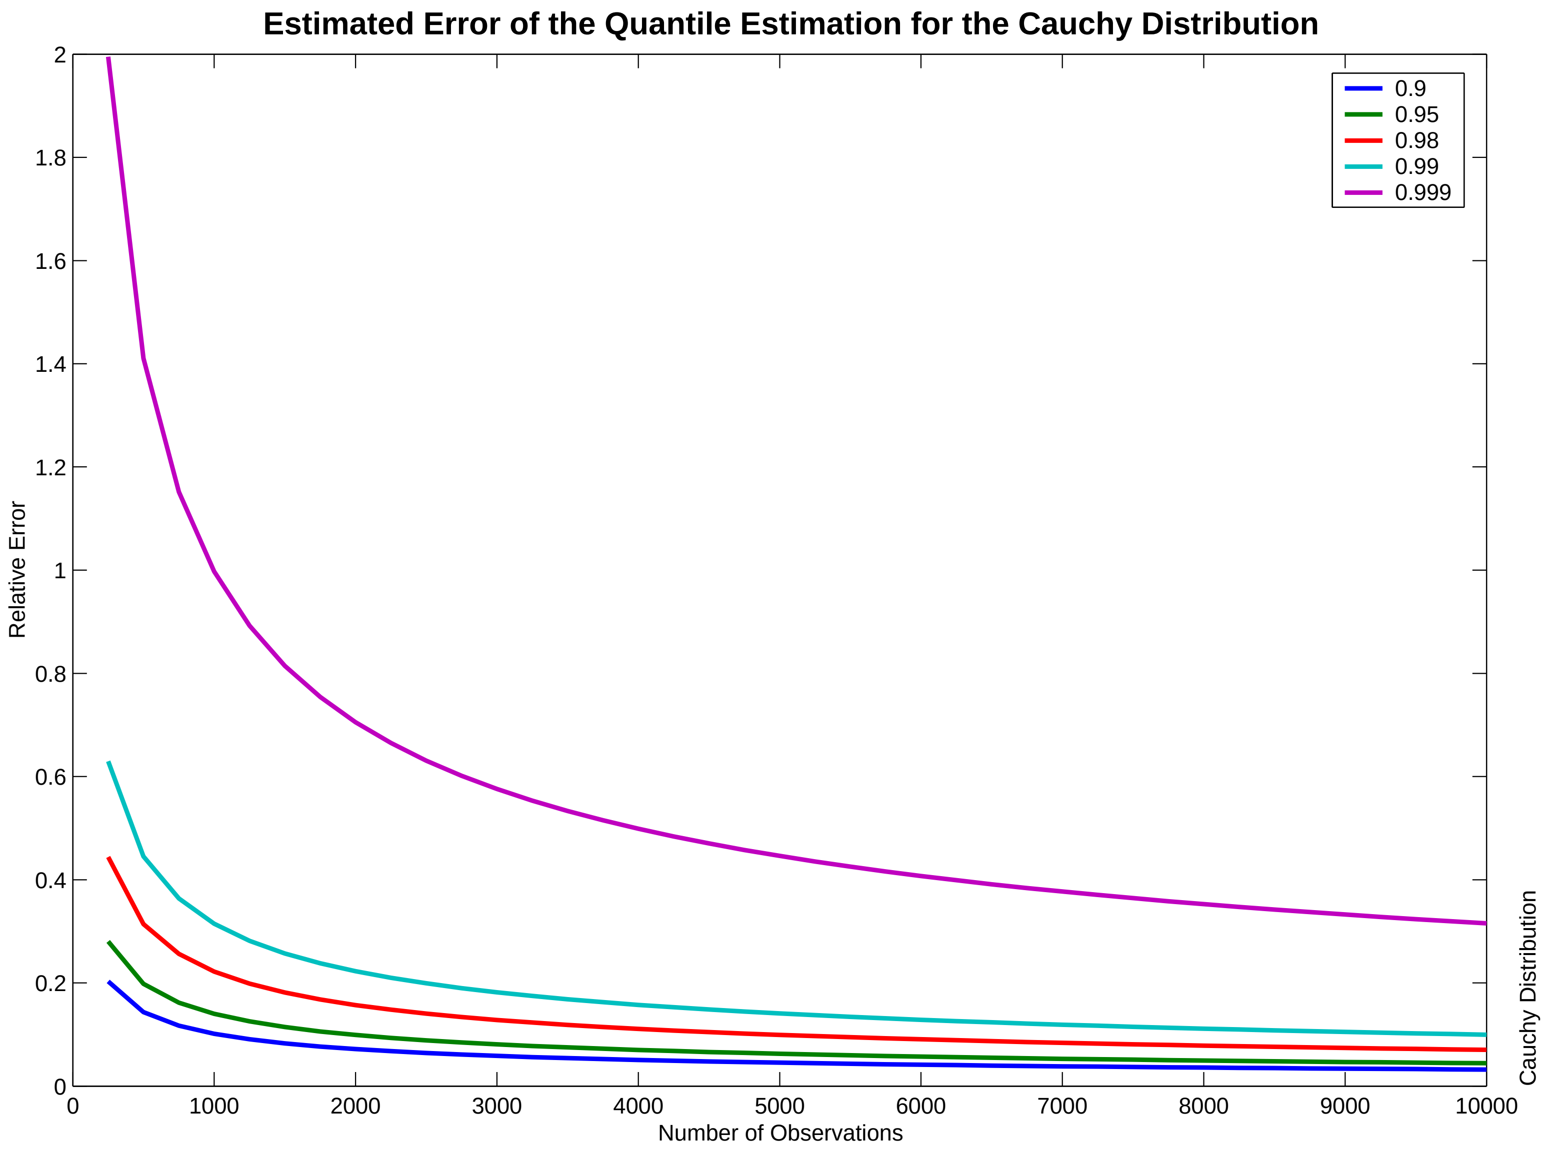
Estimated Error of the Quantile Estimation for the Cauchy Distribution
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
Relative Error
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
Number of Observations
0.9
0.95
0.98
0.99
0.999
Cauchy Distribution

Historical Simulation – Bootstrapping

Bootstrapping is a way to generate more observation than you have actually available. The procedure is as follows.

1. Number the n observations you have from 1 to n.
2. Draw a number from 1 to n, uniformly distributed.
3. Given that this number is i, take your i–th observation as your first observation for your new time series.
4. Repeat step 2. and 3. until you reached m, the desired length of your newly generated time series.
5. Repeat step 2. to 4. to create many more time series.

Monte Carlo Simulation

- is the generation of time series (such as distribution of returns or paths of asset prices) by the use of random numbers.
- draws numbers from a chosen distribution (e.g. normal, Student–t, or a diffusion) which is supposed to be the future distribution of the underlying to produce a time series – a future scenario.
- uses some price methodology to calculate the value of the portfolio and its VaR.

Principal Components Analysis

- can be used if the portfolio assets can be grouped into sup–portfolios which consists only of highly correlated assets.

- states that each portfolio asset $(S_i)_{i=1,\ldots,n}$ has a **factor representation**

$$\frac{dS_i}{S_i} = a_{i,0} + \sum_{j=1}^{m} a_{i,j} \frac{dF_j}{F_j} + \epsilon_i \,,$$

  with $m << n$ and $\left(F_j\right)_{j=1,\ldots,m}$ are pairwise independent.

- isolates the factors that are responsible for most of the variability.

- **Example:** three factors capture more than 95 percent of the variability for interst rates. The factors are the **shift**, the **twist**, and the **curvature**.

Stress Testing

Stress testing involves estimating how the portfolio would perform under some of the most extreme market moves seen in the past.

However, it should be estimated how the portfolio would perform under some **made–up** worst case scenario as well.

The **aim** of stress testing is to understand (or at least to get an idea of) the risk exposure of the portfolio.

Stress Testing – Examples

Examples are historical extreme movement such as

- October 19, 1987 when the S&P 500 moved by 22.3 standard deviations.
- January 8, 1988 when the S&P 500 moved by 6.8 standard deviations.
- April 10, 1992 when the 10–year UK–bond yields moved by 7.7 standard deviations.

or worst case scenario such as

- a sudden increase/decrease of volatility of ± 20 percent.
- a sudden increase (or devaluation) of a currency which is important for the portfolio.
- a default of a major customer.

Backtesting

- is a way to estimate the model risk.
- compares the d–day VaR estimation with the actually observed profit/loss over the next d days. If the actually observed profit/loss exceeds the VaR estimation too often, the model is not appropriate. For example, the 99 percent VaR–quantile estimation should be exceeded by the actually observed profit/loss "on average" 2.5 times given 250 observations.
- "clean" and "dirty" backtesting.
- see [10] and [11] for the requirements on backtesting by the regulators.

Model Risk

This is the risk that the model chosen is wrong.

In the Value at Risk framework the model risk is analysed by means of **backtesting**.

For example, one has the following rule of thumb:

Nonparametric Models

"Let the data speak for itself."

Parametric Models or Model Building Approach

lower

higher

Model Risk

Calculating VaR
– the Corresponding Model Risk

Variance–Covariance Approach

Historical Simulation

"History rhymes, it does not repeat." (Mark Twain)

Extreme Value Theory

Monte–Carlo Simulation

lower

higher

Model Risk

Backtesting – the Three Zone Approach

Zone	Number of Exceptions	Increase in Scaling Factor
Green	0 – 4	0
Yellow	5	0.4
	6	0.5
	7	0.65
	8	0.75
	9	0.85
Red	10 or more	1

The number of exceptions is to be read out of 250 observations. The initial scaling factor is 3.

Approach of the Banking Industry

The banking industry

- uses the VaR approach to measure the market risk in "normal" times.
- uses stress testing for estimating the impact of "crash" times to their portfolio.
- takes the liquidity risk into consideration by calculating the 10–day VaR.
- evaluates the model risk by doing backtesting.

Basic Requirements of the Regulators

- In order to calculate the **required capital** for the market risks, the banks calculate the 99 percent VaR–quantile for the 10–day returns. The bank must keep at least three times (plus a surplus depending on model [$\rightarrow$ backtesting] and data quality) of this VaR figure as capital for the market risk exposure.
- At least "clean" or "dirty" backtesting should be done, possibly both. The results of the backtesting will be evaluated according to the three zone approach.
- there are many more requirements such as on credit risk, operational risk, or data quality just to name some . . .
- Further information can be found on the website of the Basel Committee on Banking Supervision at the Bank for International Settlements (BIS): `http://www.bis.org/bcbs/index.htm` .

. . . and a Citation

> The problem with the math is that it adorned with certitude events that were inherently **uncertain**.
>
> ’You take Monica Lewinsky, who walks into Clinton’s office with a pizza. You have no idea where that’s going to go,’
>
> Conseco’s Max Bublitz, who had declined to invest in Long–Term, noted.
>
> ’Yet if you apply math to it, you come up with a thirty–eight percent chance she’s going to go down on him. It looks great, but it’s all a guess.’

(See Lowenstein [8], p. 75)

References

Essential:

1. Paul Wilmott, *Paul Wilmott on Quantitative Finance*, second edition, Volume 1–3, John Wiley & Sons, Ltd., Chichester, 2006. $\rightarrow$ **See Chapter 18, 19, 42, and 80** or for the first edition: $\rightarrow$ *See Chapter 51, 53, 57, and 66.*

2. Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk*, third edition, McGraw–Hill, 2007.

3. Kevin Dowd, *Measuring Market Risk*, second edition, John Wiley & Sons, Ltd., Chichester, 2005.

4. John C. Hull, *Risk Management and Financial Institutions*, Prentice Hall, Upper Saddle River, 2007.

Supplementary:

5. See http://www.riskmetrics.com for further information concerning RiskMetrics and CreditMetrics.

6. See http://www.gloriamundi.org for general information as well as actual news concerning VaR.

7. John C. Hull, *Options, Futures, and Other Derivatives*, sixth edition, Prentice Hall, Upper Saddle River, 2006. → **See Chapter 18. Value at Risk** or for the fifth edition:
→ *See Chapter 16.*

8. Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long–Term Capital Management*, paperback, Fouth Estate (a division of Harper–Collins), London, 2002.

9. Peter L. Bernstein, *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons, New York, 1996.

10. Basel Committee on Banking Supervision, *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*, Bank for International Settlement, June 2004. See http://www.bis.org/bcbs/index.htm and more specifically http://www.bis.org/publ/bcbs107.htm

11. Basel Committee on Banking Supervision, *Supervisory Framewok for the Use of "Backtesting" in Conjunction with the Internal Models Approach to Market Risk Capital Requirements*, Bank for International Settlement, January 1996. See http://www.bis.org/publ/bcbs22.htm

12. Joachim Hartung, Bärbel Elpelt, and Karl–Heinz Klösener, *Statistik*. Oldenbourg, München, tenth edition, 1995.

13. Alan Stuart and J. Keith Ord. Distribution Theory. In *Kendall's Advanced Theory of Statistics*, Volume 1. Arnold, London, sixth edition, 1994.