Financial Risk Forecasting
Chapter 5
Implementing Risk Forecasts

Jon Danielsson ©2016
London School of Economics

To accompany
Financial Risk Forecasting
www.financialriskforecasting.com
Published by Wiley 2011

Version 2.1. November 2016

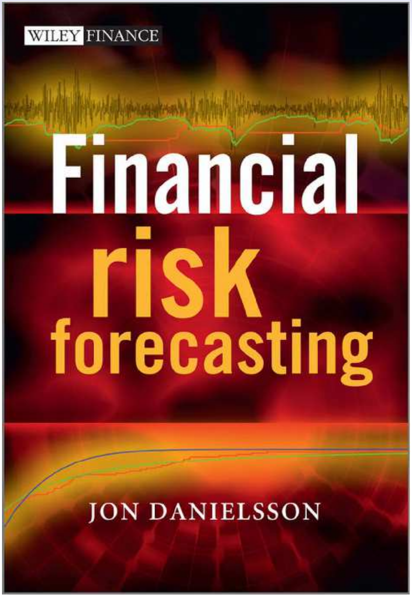
WILEY FINANCE
Financial
risk
forecasting
JON DANIELSSON

The focus of this chapter is on

- Techniques for implementing risk forecasting, especially
- Historical simulation
- Risk measures and parametric methods
- Expected returns
- VaR with time-dependent volatility

Notation

K	Number of assets
w	$K \times 1$ vector of portfolio weights
w_k	Portfolio weight on asset k
X and Y	Two different assets
$\varphi(.)$	Risk measure
ϑ	Portfolio value
$y_k = \{y_{t,k}\}_{t=1}^{T}$	$T \times 1$ vector of returns on asset k
$y = T \times K$	$T \times K$ matrix of historical returns
Σ	$K \times K$ covariance matrix

Historical simulation

Historical simulation

- For one asset,
- Assumes that one of the observations in the estimation window will be the next day return, therefore
- Assume history repeats itself
- VaR is one of the observations in the estimation window, multiplied by the monetary value of the asset holdings, the portfolio value ϑ_t. In the one asset case

$$\vartheta_t = \text{ number of stocks owned} \times P_t$$

- VaR_t is the negative of the $(W_E \times p)^{\text{th}}$ smallest return, times ϑ_r

500 days of the S&P 500

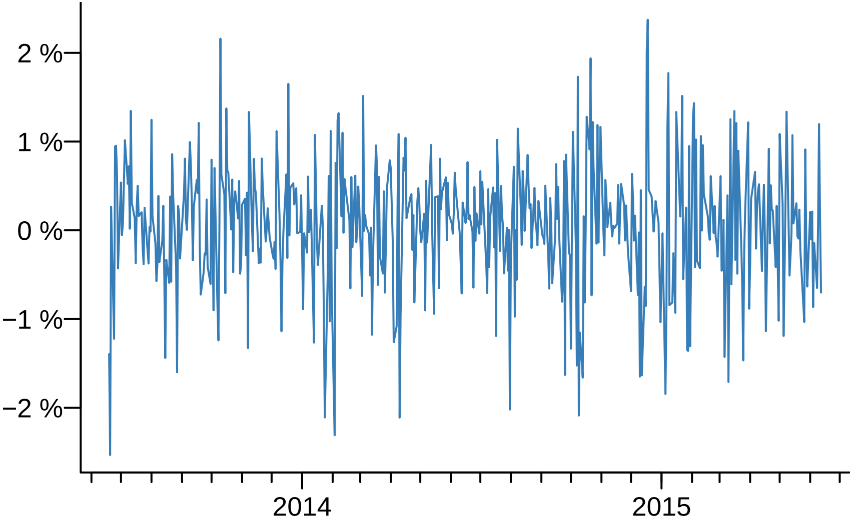

500 days of the S&P 500

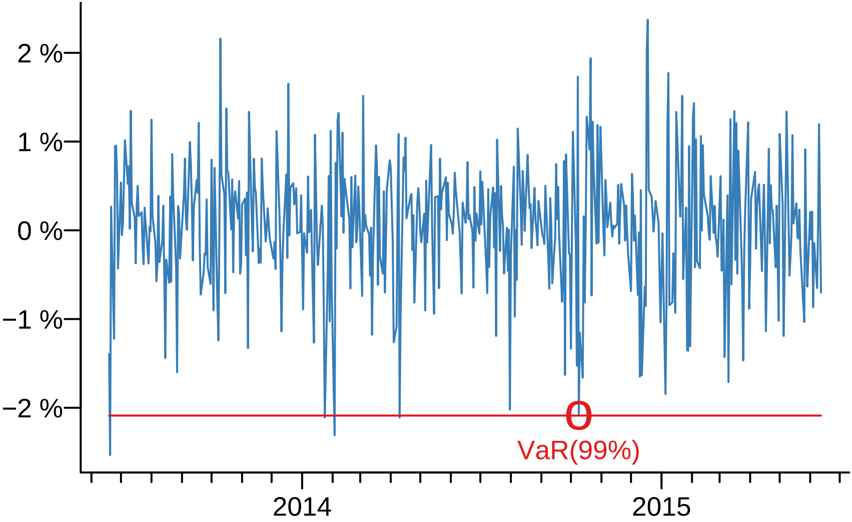

Sorted returns

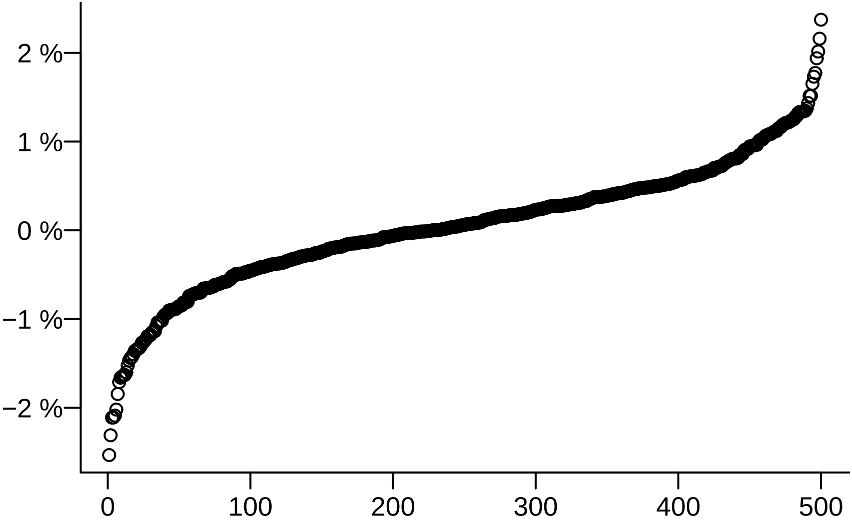

Zoom in

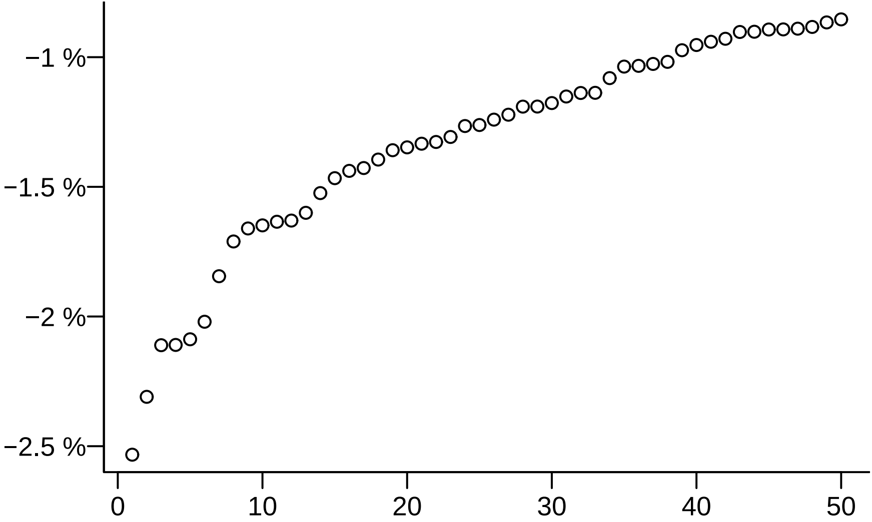

Zoom in

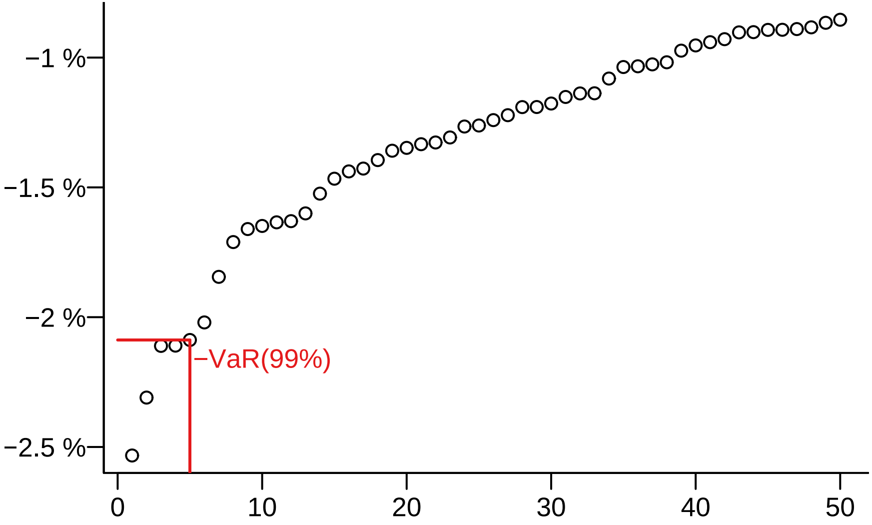

Zoom in

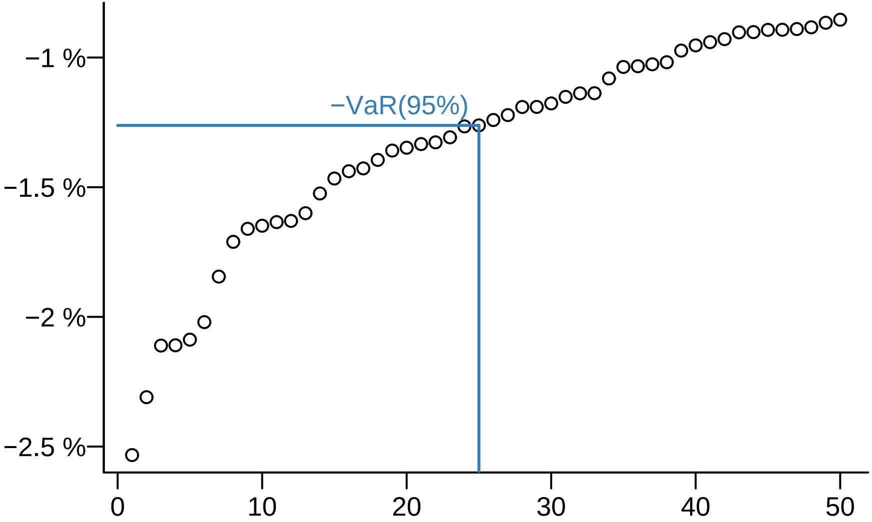

Procedure

- Identify probability, p, e.g. 1%
- Take W_E of the most recent returns (e.g. 1000)
- Sort them from the smallest to the largest, call that ys
- Take the $(W_E \times p)^{\text{th}} = (1000 \times 0.01)^{\text{th}}$ smallest value of ys, call that ys_{10}
- If you own one stock, and $P_{t-1} = 1$, then VaR is the 10^{th} smallest return, i.e. $\text{VaR} = -ys_{10}$
- Otherwise have to multiply that by the number of stocks you own and their $t-1$ price

Multiple assets

- Create a vector of historical portfolio returns
- y is a $T \times K$ matrix of returns
- w is a $K \times 1$ vector of portfolio weights
- We then get the timeseries vector of portfolio returns by

$$y_{\text{portfolio}} = yw$$

- And then you can simply treat the portfolio as if it were a single asset and apply HS

Expected shortfall estimation

- The expected losses conditional on VaR being violated
- Estimated by HS by taking the
- Mean of all observations less than or equal to -VaR
- Continuing with the single asset accent from above

$$\text{ES} = \frac{1}{10}\sum_{i=1}^{10} ys_i$$

Importance of sample size

- The most extreme observations fluctuate a lot more than observations that are less extreme
- Therefore, the bigger the sample the more precise the estimation of HS should be
- The downside is that old data may not be all that representative
- And if there is a structural break in the data (like in 2007) the VaR forecasts take longer to adjust to structural changes in risk
- As a general rule
- Minimum recommended sample size:

$$\frac{3}{p}$$

Issues

- No model assumptions needed
- In the absence of structural breaks HS tends to perform well
- It captures nonlinear dependence directly
- But performs badly when data has structural breaks
- This can be seen in the discussion about backtesting in Chapter 8

Parametric methods

Parametric and nonparametric

- HS does not use any distribution and therefore has no distributional parameters — it is *nonparametric*
- In all other cases, we have some distribution, and distributions have parameters — *parametric*
- For the remainder of this section assume we have one day, so we don't have to specify time, t
- The density of returns is then $f(x)$
- And the distribution is $F(x)$

Recall

- The definition of VaR is

$$p = \Pr[Q \leq -\text{VaR}(p)] = \int_{-\infty}^{-\text{VaR}(p)} f_q(x)dx$$

- Profit and loss when we own one stock

$$Q_t = P_t - P_{t-1}$$

VaR for simple returns

The definition of simple returns is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

VaR for simple returns

The definition of simple returns is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

We can then derive the VaR (we own one stock)

$$p = \Pr[Q_t \leq -\text{VaR}(p)]$$

VaR for simple returns

The definition of simple returns is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

We can then derive the VaR (we own one stock)

$$\begin{aligned} p =& \Pr[Q_t \leq -\text{VaR}(p)] \\ =& \Pr(P_t - P_{t-1} \leq -\text{VaR}(p)) \end{aligned}$$

VaR for simple returns

The definition of simple returns is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

We can then derive the VaR (we own one stock)

$$\begin{aligned} p &= \Pr[Q_t \leq -\text{VaR}(p)] \\ &= \Pr(P_t - P_{t-1} \leq -\text{VaR}(p)) \\ &= \Pr(P_{t-1} R_t \leq -\text{VaR}(p)) \end{aligned}$$

VaR for simple returns

The definition of simple returns is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

We can then derive the VaR (we own one stock)

$$\begin{aligned}
p &= \Pr[Q_t \leq -\mathrm{VaR}(p)] \\
&= \Pr(P_t - P_{t-1} \leq -\mathrm{VaR}(p)) \\
&= \Pr(P_{t-1} R_t \leq -\mathrm{VaR}(p)) \\
&= \Pr\left(\frac{R_t}{\sigma} \leq -\frac{\mathrm{VaR}(p)}{P_{t-1}\sigma}\right)
\end{aligned}$$

- Denote the distribution of standardized returns, R_t/σ, by

$$F_R(.)$$

- The inverse distribution is $F_R^{-1}(p)$

- Denote the distribution of standardized returns, R_t/σ, by

$$F_R(\cdot)$$

- The inverse distribution is $F_R^{-1}(p)$
- So

$$\text{VaR}_t(p) = -\sigma F_R^{-1}(p) P_{t-1}$$

VaR for continuously compounded returns

$$Y_t = \log P_t - \log P_{t-1}$$

VaR for continuously compounded returns

$$Y_t = \log P_t - \log P_{t-1}$$

$$p = \Pr\left(P_t - P_{t-1} \le -\mathrm{VaR}(p)\right)$$

VaR for continuously compounded returns

$$Y_t = \log P_t - \log P_{t-1}$$

$$\begin{aligned} p &= \Pr\left(P_t - P_{t-1} \leq -\mathrm{VaR}(p)\right) \\ &= \Pr\left(P_{t-1}\left(e^{Y_t} - 1\right) \leq -\mathrm{VaR}(p)\right) \end{aligned}$$

VaR for continuously compounded returns

$$Y_t = \log P_t - \log P_{t-1}$$

$$\begin{aligned} p &= \Pr\left(P_t - P_{t-1} \leq -\mathrm{VaR}(p)\right) \\ &= \Pr\left(P_{t-1}\left(e^{Y_t} - 1\right) \leq -\mathrm{VaR}(p)\right) \\ &= \Pr\left(\frac{Y_t}{\sigma} \leq \log\left(-\frac{\mathrm{VaR}(p)}{P_{t-1}} + 1\right)\frac{1}{\sigma}\right) \end{aligned}$$

when $-\mathrm{VaR}(p)/P_{t-1} \leq 1$

- Denote the distribution of standardized returns (Y_t/σ) by

$$F_y(.)$$

- The inverse distribution is $F_y^{-1}(p)$

- Denote the distribution of standardized returns (Y_t/σ) by

$$F_y(.)$$

- The inverse distribution is $F_y^{-1}(p)$
- So

$$\text{VaR}(p) = -(\exp(F_y^{-1}(p)\sigma) - 1)P_{t-1}$$

- Denote the distribution of standardized returns (Y_t/σ) by

$$F_y(\cdot)$$

- The inverse distribution is $F_y^{-1}(p)$
- So

$$\text{VaR}(p) = -(\exp(F_y^{-1}(p)\sigma) - 1)P_{t-1}$$

- For small $F_y^{-1}(p)\sigma$ the VaR for holding one unit of asset is:

$$\text{VaR}(p) \approx -\sigma F_y^{-1}(p)P_{t-1}$$

- VaR for continuously compounded returns is approximately the same as the VaR using simple returns

When there is more than one asset

- In the two assets case

$$\sigma^2_{\text{portfolio}} = \begin{pmatrix} w_1 & w_2 \end{pmatrix} \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix} \begin{pmatrix} w_1 \\ w_2 \end{pmatrix}$$

- And generally in the K asset case

$$\sigma^2_{\text{portfolio}} = w'\Sigma w$$

- Then as before

$$\text{VaR}(p) = -\sigma_{\text{portfolio}} F^{-1}(p) P_{t-1}$$

VaR when returns are normally distributed

- $\vartheta = 1$, $\sigma = 1$, $p = 0.05$

$$\text{VaR} = -\Phi^{-1}(0.05) = 1.64$$

VaR when returns are normally distributed

- $\vartheta = 1$, $\sigma = 1$, $p = 0.05$

$$\text{VaR} = -\Phi^{-1}(0.05) = 1.64$$

- If σ does not equal one, then the VaR is:

$$\text{VaR} = \sigma 1.64$$

VaR when returns are normally distributed

- $\vartheta = 1$, $\sigma = 1$, $p = 0.05$

$$\text{VaR} = -\Phi^{-1}(0.05) = 1.64$$

- If σ does not equal one, then the VaR is:

$$\text{VaR} = \sigma 1.64$$

- If portfolio value is not equal to 1:

$$\text{VaR} = \sigma 1.64 \vartheta$$

VaR under the Student-t distribution

- Advantage of Student-t VaR over normal is fat tails
- ν indicates how fat tails are
- When $\nu = \infty$ the Student-t becomes normal

Adjusting for variance

- Variance implied by ν of a *standard* Student-t distribution is:

$$\frac{\nu}{\nu - 2}$$

- If $\nu \leq 2$ variance of Student-t is not defined
- Volatility shows up in $F^{-1}(p)$ and $\hat{\sigma}$
- VaR would be overestimated if sample variance is used
- Volatility estimate *may* need to be scaled by ν

$$\sigma^2 = \frac{\nu}{\nu - 2}\tilde{\sigma}^2$$

- $\tilde{\sigma}^2$ is the variance in excess of that implied by the standard Student-t

Expected shortfall under normality

- One needs to obtain VaR and then calculate the conditional expectation:

$$
\begin{aligned}
\mathrm{ES} &= -\int_{-\infty}^{-\mathrm{VaR}(p)} x f_{\mathrm{VaR}}(x)dx \\
&= \frac{1}{p}\int_{-\infty}^{-\mathrm{VaR}(p)} x\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{1}{2}\frac{x^2}{\sigma^2}\right)dx \\
&= \frac{1}{p}\left[-\frac{\sigma^2}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{1}{2}\frac{x^2}{\sigma^2}\right)\right]_{-\infty}^{-\mathrm{VaR}(p)}
\end{aligned}
$$

- Given that bound is zero and normal density is

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right):$$

$$\text{ES} = -\frac{\sigma^2 \phi(-\text{VaR}(p))}{p}$$

- If portfolio value is ϑ:

$$\text{ES} = -\vartheta \frac{\sigma^2 \phi(-\text{VaR}(p))}{p}$$

Expected returns

Expected returns

- Is it reasonable to assume $\mu = 0$?
- Given that statistical uncertainty is more than 10% in most VaR calculations, VaR calculation is only significant to one digit
- Mean is smaller than that
- For S&P 500: $\mu = 0.019\%, \sigma = 1.15\%$

VaR with mean

- The definition of VaR is

$$p = \Pr[Q \leq -\text{VaR}(p)] = \int_{-\infty}^{-\text{VaR}(p)} f_q(x)dx$$

- Profit and loss when we own one stock, *and mean is zero,* $E(Q) = 0$

$$Q_t = P_t - P_{t-1}$$

- If mean is not zero, $\mathrm{E}(Q) \neq 0$, then the definition of VaR is

$$\Pr[Q + \mathrm{E}(Q) \leq -\mathrm{VaR}(p)] = p$$

so

$$\mathrm{VaR}(p) = -\sigma F^{-1}(p) - \mu$$

- Note how the mean has a minus in front of it
- A positive mean pulls the distribution to the right
- Making VaR smaller

Time aggregation of VaR with mean

- If the returns are IID, then both mean and variance aggregate at the same rate
- Mean and variance over T days is equal to T times mean and variance over one day
- Which implies that the volatility aggregates at the square root of time
- The T-period VaR is therefore:

$$\begin{aligned} \text{VaR}(T\text{day}) &= -\sigma(T\text{day})F^{-1}(p) - \mu(T\text{day}) \\ &= -\sqrt{T}\sigma(1\text{day})F^{-1}(p) - T\mu(1\text{day}) \end{aligned}$$

- And $\sqrt{T} < T$

so...

- The assumption $\mu = 0$ is relatively harmless as the error is small at the daily level.
- Daily S&P 500: $\mu = 0.019\%, \sigma = 1.15\%$
- Annual: $\mu = 4.87, \sigma = 18.2$

VaR when returns are not zero

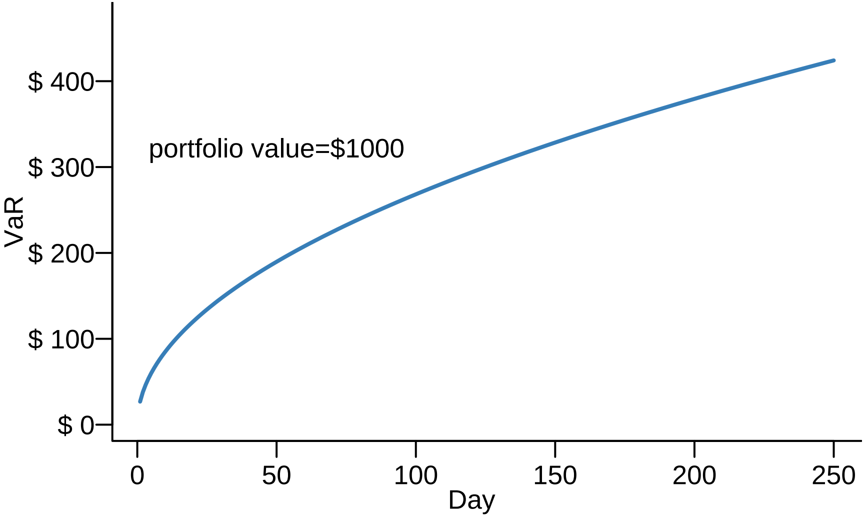

VaR when returns are not zero

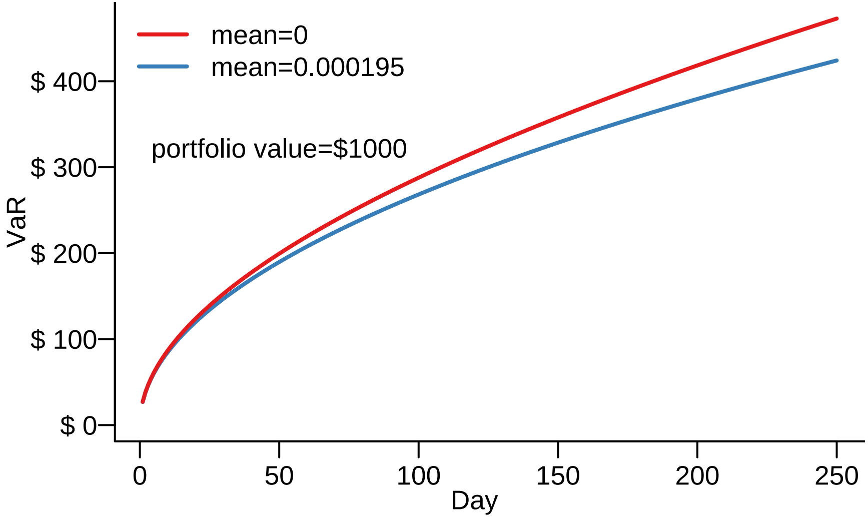

Issues in including the mean

- It is much more difficult to estimate the mean than the variance
- And unless necessary, should be avoided
- And for VaR over one day or 10 days is not necessary in most cases

VaR with time dependent volatility

VaR with time dependent volatility

- When using conditional volatility models, like EWMA and GARCH
- First update volatility
- And then calculate one day ahead VaR

GARCH

- GARCH(1,1) is:

$$\hat{\sigma}^2_{t+1} = \omega + \alpha Y_t^2 + \beta \sigma_t^2$$

- Software packages (R and Matlab) usually return the last volatility of the sample ($\hat{\sigma}_t$), and the parameters
- One then needs to update the volatility by using $\hat{\sigma}_t$ and the model parameters

$$\hat{\sigma}^2_{t+1} = \hat{\omega} + \hat{\alpha} y_t^2 + \hat{\beta} \hat{\sigma}_t^2$$

- So the VaR at $t+1$ is

$$\widehat{\text{VaR}}_{t+1} = -\hat{\sigma}_{t+1} F_y^{-1}(p) \vartheta t$$