Unbiased estimation of risk measures

Albert-Ludwigs-Universität Freiburg

Thorsten Schmidt
joint work with Marcin Pitera
Abteilung für Mathematische Stochastik

www.stochastik.uni-freiburg.de
thorsten.schmidt@stochastik.uni-freiburg.de
Vienna, Mar 2017

Problem

- A standard risk measure is the **value-at-risk**. This risk measure is given by a quantile of the profit & loss distribution.

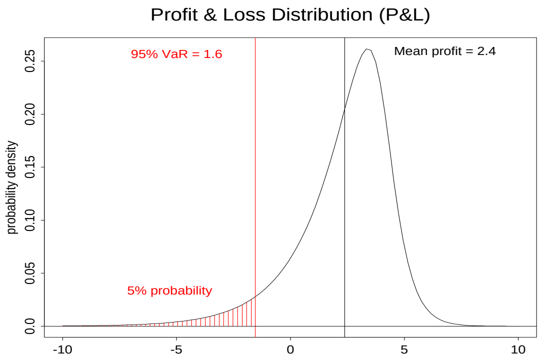

Density of X together with VaR at the level $\alpha = 0.05$. Source: McNeil, Frey, Embrechts **Quantitative Risk Management**.

- Assume that the **profit** of a portfolio, say X, is normally distributed with mean μ and variance σ^2. Then the value-at-risk at level α is given by

$$\mathrm{VaR}_\alpha(X) = -\Big(\mu + \sigma\Phi^{-1}(\alpha)\Big).$$

- But in practice, μ and σ are unknown and have to be estimated. In this regard, let us consider the simplest case: we have an i.i.d. sample $X_1, \dots, X_n =: \boldsymbol{X}$ at hand.
- **Efficient** estimators of μ and σ are at hand:

$$\hat{\mu}_n = \bar{\boldsymbol{X}}, \qquad \hat{\sigma}_n = \bar{\sigma}(\boldsymbol{X}) = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{\boldsymbol{X}})^2}. \tag{1}$$

- **Common practice** is to use the plug-in estimator

$$\mathrm{VaR}_\alpha^{plugin} := -\Big(\hat{\mu}_n + \hat{\sigma}_n\Phi^{-1}(\alpha)\Big).$$

- Can this be efficient?

Motivation from Statistics

- In the normal case, for **known** σ, the likelihood-ratio test turns out to be the Gauss-test, or, equivalently, the confidence-interval is a normal distribution.
- If σ is **unknown**, one utilizes the t-distribution to obtain an efficient test: consider w.l.o.g. the test for $\mu = 0$ versus $\mu \neq 0$. The standardized test statistic is

$$T(X_1, \ldots, X_n) =: T(\boldsymbol{X}) = \frac{\sqrt{n}\bar{\boldsymbol{X}}}{\bar{\sigma}(\boldsymbol{X})}$$

  and the test rejects the null hypothesis if

$$T(X) > t_n(1 - \alpha).$$

- Shouldn't there be a similar adjustment towards the t-distribution in the estimator for VaR?

Motivation from Backtesting

- Let us perform a standard backtesting-procedure, i.e. we run several simulations, estimate the value-at-risk and check if the percentage of unsufficient capital does not exceed 5%.

Tabelle: Estimates of $\mathrm{VaR}_{0.05}$ for NASDAQ100 (first column) and for a sample from normally distributed random variable with mean and variance fitted to the NASDAQ data (second column), both for 4.000 data points. **Exceeds** reports the number of exceptions in the sample, where the actual loss exceeded the risk estimate.

Estimator		NASDAQ		Simulated	
		exceeds	percentage	exceeds	percentage
Plug-in	$\hat{\mathrm{VaR}}_\alpha^{\text{plugin}}$	241	0.061	221	0.056

- Our findings suggest that the estimator is biased. In a statistical sense !
- Our goal is to analyse this problem and give a new notion of unbiasedness in an economic sense.

Measuring risk

We begin with well-known results on the measurement of risk, see McNeil et al. (2005).

- Let $(\Omega, \mathscr{A})$ be a measurable space and $(P_\theta : \theta \in \Theta)$ be a family of probability measures.
- For simplicity, we assume that the measures P_θ are equivalent, such that their null-sets coincide.
- For the estimation, we assume that we have a sample $X_1, X_2, \ldots, X_n$ of observations at hand.
- A risk measure ρ is a mapping from L^0 to $\mathbb{R} \cup \{+\infty\}$.
- The value $\rho(X)$ is a quantification of risk for a future position: it is the amount of money one has to add to the position X such that the position becomes acceptable.

A priori, the definition of a risk measure is formulated without any relation to the underlying probability. However, in most practical applications one typically considers law-invariant risk-measures. Denote by $\mathscr{D}$ the convex space of cumulative distribution functions of real-valued random variables.

Definition

The family of risk-measures $(\rho_\theta)_{\theta\in\Theta}$ is called **law-invariant**, if there exists a function $R:\mathscr{D}\to\mathbb{R}\cup\{+\infty\}$ such that for all $\theta\in\Theta$ and $X\in L^0$

$$\rho_\theta(X) = R(F_X(\theta)), \tag{2}$$

$F_X(\theta) = P_\theta(X\le\cdot)$ denoting the cumulative distribution function of X under the parameter θ.

Estimation

We aim at estimating the risk of the future position when $\theta \in \Theta$ is unknown and needs to be estimated from a data sample $x_1, \dots, x_n$.

Definition

An **estimator** of a risk measure is a Borel function $\hat{\rho}_n : \mathbb{R}^n \to \boldsymbol{R} \cup \{+\infty\}$.

Sometimes we will call $\hat{\rho}_n$ also risk estimator.

The following definition introduces an economically motivated formulation of unbiasedness.

Definition

The estimator $\hat{\rho}_n$ is called **unbiased** for $\rho(X)$, if for all $\theta \in \Theta$,

$$\rho_\theta(X + \hat{\rho}_n) = 0. \tag{3}$$

- If the estimator is unbiased, adding the estimated amount of risk capital $\hat{\rho}_n$ to the position X makes the position $X + \hat{\rho}_n$ acceptable under all possible scenarios $\theta \in \Theta$.
- Requiring equality in Equation (3) ensures that the estimated capital is not too high.
- Except for the i.i.d. case, the distribution of $X + \hat{\rho}_n$ does also depends on the dependence structure of $X, X_1, \ldots, X_n$ and not only on the (marginal) laws.

Relation to the statistical definition of unbiasedness

Our Definition differs from unbiasedness in the statistical sense!

- The estimator $\hat{\rho}_n$ is called **statistically unbiased**, if

$$E_\theta[\hat{\rho}_n] = \rho_\theta(X), \qquad \text{for all } \theta \in \Theta, \tag{4}$$

- One point why the statistical unbiasedness is not reasonable here is that it does not behave well in various backtesting or stress-testing procedures.

Relation to probability unbiasedness

- In Francioni and Herzog (2012), the authors introduced the concept of a **probability unbiased** estimation: denote by $F_X(\theta,t) = P_\theta(X \leq t), t \in \mathbb{R}$ umulative distribution function of X under P_θ. Then the estimator $\hat{\rho}_n$ is called **probability unbiased**, if

$$E_\theta[F_X(\theta,-\hat{\rho}_n)] = F_X(\theta,-\rho_\theta(X)), \qquad \text{for all } \theta \in \Theta. \tag{5}$$

- This approach coincides for value-at-risk in the strongly restricted setting of the i.i.d. Example with our definition of unbiasedness: indeed, assume that $F_X(\theta)$ is continuous and that $X_1,\dots,X_n,X$ are i.i.d. Then $\hat{\rho}_n$ and X are independent and hence

$$E_\theta[F_X(\theta,-\hat{\rho}_n)] = P_\theta[X+\hat{\rho}_n < 0].$$

On the other hand we know that, for ρ_θ being value-at-risk at level α, we obtain $F_X(\theta,-\rho_\theta(X)) = \alpha$, so (5) is equivalent to

$$P_\theta[X+\hat{\rho}_n < 0] = \alpha. \tag{6}$$

Now it follows this is equivalent to

$$\rho_\theta(X+\hat{\rho}_n) = \inf\{x \in \boldsymbol{R} \colon P_\theta[X+\hat{\rho}_n + x < 0] \leq \alpha\} = 0.$$

Relation to level adjustment

- A further alternative is to adjust the level α, see Frank (2016) and Francioni and Herzog (2012).
- An existing estimator depending continuously on α can always be trimmed to match exactly the unbiased estimator. However, the adjusted α will typically depend on n and the sample !

Unbiased estimation of value-at-risk under normality

- Let $X \sim \mathcal{N}(\theta_1, \theta_2^2)$ and denote $\theta = (\theta_1, \theta_1) \in \Theta = \mathbb{R} \times \mathbb{R}_{>0}$.
- The value-at-risk is

$$\rho_\theta(X) = \inf\{x \in \boldsymbol{R} \colon P_\theta[X + x < 0] \leq \alpha\}, \qquad \theta \in \Theta, \tag{7}$$

- Unbiasedness as defined in Equation (3) is equivalent to

$$P_\theta[X + \hat{\rho} < 0] = \alpha, \qquad \text{for all } \theta \in \Theta. \tag{8}$$

- We define estimator $\hat{\rho}$, as

$$\hat{\rho}(x_1, \dots, x_n) = -\bar{x} - \bar{\sigma}(\boldsymbol{x}) \sqrt{\frac{n+1}{n}} t_{n-1}^{-1}(\alpha), \tag{9}$$

- The this estimator is unbiased: using the fact that X, $\bar{X}$ and $s(\boldsymbol{X})$ are independent for any $\theta \in \Theta$, we obtain

$$T := \sqrt{\frac{n}{n+1}} \cdot \frac{X - \bar{X}}{\bar{\sigma}(\boldsymbol{X})} = \frac{X - \bar{X}}{\sqrt{\frac{n+1}{n}}\theta_2} \cdot \sqrt{\frac{n-1}{\sum_{i=1}^n (\frac{X_i - \bar{X}}{\theta_2})^2}} \sim t_{n-1}.$$

Thus, the random variable T is a pivotal quantity and

$$P_\theta[X + \hat{\rho} < 0] = P_\theta[T < q_{t_{n-1}}(\alpha)] = \alpha.$$

Unbiased estimation of expected shortfall under normality

- We continue in the previous setting,
- The expected shortfall at level α under a continuous distribution is

$$\rho_\theta(X) = E_\theta[-X | X \le q_X(\theta, \alpha)],$$

  where $q_X(\theta, \alpha)$ is α-quantile of X under P_θ.
- We consider estimators of the form

$$\hat{\rho}(x_1, \dots, x_n) = -\bar{x} - \bar{\sigma}(x) a_n, \tag{10}$$

  for some $(a_n)_{n \in \boldsymbol{N}}$, where $a_n \in \boldsymbol{R}$.
- We can show that there exists a_n which makes $\hat{\rho}$ unbiased. This a_n can be computed numerically.

Asymptotically unbiased estimators

Definition

A sequence of risk estimators $\hat{\rho} = (\hat{\rho}_n)_{n \in \mathbb{N}}$ will be called **unbiased** at $n \in \mathbb{N}$, if $\hat{\rho}_n$ is unbiased. If unbiasedness holds for all $n \in \boldsymbol{N}$, we call the sequence $\hat{\rho}$ unbiased. The sequence $\hat{\rho}$ is called **asymptotically unbiased**, if

$$\rho_\theta(X + \hat{\rho}_n) \xrightarrow{n \to \infty} 0, \qquad \text{for all } \theta \in \Theta.$$

- The proposed definition of asymptotical unbiasedness has similarities to the notion of consistency suggested in Davis (2016). This notion of consistency requires that averages of the calibration errors converge suitable fast to 0 when the time period tends to infinity. Hence, asymptotically unbiased risk estimators will be consistent when the calibration error is measured with the risk measure itself. On the other side, it should be noted that our main goal is to obtain the optimal risk estimator without averaging out under- or overestimates as they have an asymmetric effect on the portfolio performance.

For empircal estimators we obtain asymptotic unbiasedness.
We obtain the following result. Recall that we study an i.i.d. sequence $X, X_1, X_2, \dots$ Let $\alpha \in (0,1)$ and consider the negative of emprical α-quantile

$$\hat{\rho}_n(x_1, \dots, x_n) = -x_{(\lfloor n\alpha \rfloor + 1)}, \quad n \in \mathbb{N}, \tag{11}$$

which we call empirical estimator of value-at-risk at level α.

Proposition

Assume that X is absolutely continuous under P_θ for any $\theta \in \Theta$. The sequence of empirical estimators of value-at-risk given in (11) is asymptotically unbiased.

Empirical study

- It is the aim of this section to analyse the performance of selected estimators on various sets of real market data (Market) as well as on simulated data (Simulated). Our focus is on the practically most relevant risk measures, VaR and ES.
- The market data we use are returns from the data library Fama and French (2015), containing returns of 25 portfolios formed on book-to-market and operating profitability in the period from 27.01.2005 to 01.01.2015. We obtain exactly 2500 observations for each portfolio.
- The sample is split into 50 separate subsets, each consisting of 50 consecutive trading days. For $i = 1, 2, \ldots, 49$, we estimate the risk measure using the i-th subset and test it's adequacy on $(i + 1)$-th subset.
- The simulation study uses i.i.d. normally distributed random variables whose mean and variance was fitted to each of the 25 portfolios. The sample size was set to 2500 for each set of parameters. In this way we are able to exclude difficulties due to dependencies in the data or bad model fit.

Backtesting VaR

- We considered the unbiased estimator $\hat{\mathrm{VaR}}_\alpha^{\mathrm{u}}$, the empirical sample quantile $\hat{\mathrm{VaR}}_\alpha^{\mathrm{emp}}$, the modified Cornish-Fisher estimator $\hat{\mathrm{VaR}}_\alpha^{\mathrm{CF}}$ and the plug-in estimator $\hat{\mathrm{VaR}}_\alpha^{\mathrm{norm}}$ and the GPD plug-in estimator[1] $\hat{\mathrm{VaR}}_\alpha^{\mathrm{GPD}}$.

$$\hat{\mathrm{VaR}}_\alpha^{\mathrm{emp}}(x) := -\Big(x_{(\lfloor h\rfloor)} + (h-\lfloor h\rfloor)(x_{(\lfloor h+1\rfloor)} - x_{(\lfloor h\rfloor)})\Big), \tag{12}$$

$$\hat{\mathrm{VaR}}_\alpha^{\mathrm{CF}}(x) := -\Big(\bar{x} + \bar{\sigma}(x)\bar{Z}_{CF}^{\alpha}(x)\Big), \tag{13}$$

$$\hat{\mathrm{VaR}}_\alpha^{\mathrm{norm}}(x) := -\Big(\bar{x} + \bar{\sigma}(x)\Phi^{-1}(\alpha)\Big), \tag{14}$$

$$\hat{\mathrm{VaR}}_\alpha^{\mathrm{GPD}} := -u + \frac{\hat{\beta}}{\hat{\xi}}\left(\left(\frac{\alpha n}{k}\right)^{-\hat{\xi}} - 1\right), \tag{15}$$

$$\hat{\mathrm{VaR}}_\alpha^{\mathrm{u}}(x_1,\dots,x_n) := -\left(\bar{x} + \bar{\sigma}(x)\sqrt{\frac{n+1}{n}}\,t_{n-1}^{-1}(\alpha)\right), \tag{16}$$

where $x_{(k)}$ is the k-th order statistic of $x = (x_1,\dots,x_n)$, the value $\lfloor z\rfloor$ denotes the integer part of $z \in \boldsymbol{R}$, $h = \alpha(n-1)+1$, Φ denotes the cumulative distribution function of the standard normal distribution and $\bar{Z}_{CF}^{\alpha}$ is a standard Cornish-Fisher α-quantile estimator.

[1] For each portfolio, we set the threshold value u to match the 0.7-empirical quantile of the corresponding sample.

Tabelle: First we show the results for portfolios in the period from 27.01.2005 to 01.01.2015 from the Fama & French dataset. Second, we show the results on simulated Gaussian data. We perform the standard backtest, splitting into intervals of length 50 and computing average rate of exceptions.

Type of data:	MARKET				
Portfolio	Estimator type				
	$\hat{VaR}_\alpha^{emp}$	$\hat{VaR}_\alpha^{norm}$	$\hat{VaR}_\alpha^{CF}$	$\hat{VaR}_\alpha^{GPD}$	$\hat{VaR}_\alpha^{u}$
LoBM.LoOP	0.071	0.073	0.067	0.067	0.069
BM1.OP2	0.076	0.070	0.069	0.069	0.065
BM1.OP3	0.071	0.064	0.063	0.064	0.061
BM1.OP4	0.069	0.071	0.067	0.067	0.068
LoBM.HiOP	0.071	0.071	0.070	0.067	0.068
...	...	...	...	...	
mean	0.073	0.071	0.068	0.067	0.067

Type of data:	SIMULATED				
	$\hat{VaR}_\alpha^{emp}$	$\hat{VaR}_\alpha^{norm}$	$\hat{VaR}_\alpha^{CF}$	$\hat{VaR}_\alpha^{GPD}$	$\hat{VaR}_\alpha^{u}$
LoBM.LoOP	0.065	0.057	0.055	0.056	**0.051**
BM1.OP2	0.064	0.053	0.053	0.053	**0.050**
BM1.OP3	0.069	0.058	0.058	0.060	**0.052**
BM1.OP4	0.069	0.057	0.058	0.062	**0.053**
LoBM.HiOP	0.060	0.054	0.053	0.056	**0.047**
...	...	...	...	...	
mean	0.066	0.057	0.057	0.058	**0.051**

A further test

To gain further insight on the performance of the Gaussian unbiased estimator, we have replicated the results from the simulations in Table 2 for $N = 10.000$ times and the first portfolio LoBM.LoOP. We consider three statistics: the **exceedance rate** $\mathrm{ER}_i(\hat{\rho})$, the **relative deviation**

$$\mathrm{RD}_i(\hat{\rho}) := \frac{\mathrm{ER}_i(\hat{\rho}) - \mathrm{ER}_i(\hat{\mathrm{VaR}}^{\mathrm{u}}_{\alpha})}{\mathrm{ER}_i(\hat{\mathrm{VaR}}^{\mathrm{u}}_{\alpha})}$$

and, the **outperformance rate** of the unbiased estimator in the sense that the exceedance rate is closer to $\alpha = 0.05$,

$$\mathrm{OR}_i(\hat{\rho}) := \begin{cases} 1 & \text{if } |\mathrm{ER}_i(\hat{\rho}) - \alpha| > |\mathrm{ER}_i(\hat{\mathrm{VaR}}^{\mathrm{u}}_{\alpha}) - \alpha|, \\ 0 & \text{otherwise.} \end{cases}$$

In Table 3 we state mean and standard deviations (sd) of these statistics.

Tabelle: We fit a normal distribution to the first portfolio from the Fama & French dataset, i.e. LoBM.LoOP portfolio, compare Table 2. We show average exception (ER) rate, relative deviation (RD) and outperformance rate (OR).

Estimator		ER mean	ER sd	RD mean	RD sd	OR mean
Percentile	$\hat{\text{VaR}}^{\text{emp}}_{\alpha}(x)$	0.067	0.004	29.2%	8.9%	100%
Modified C-F	$\hat{\text{VaR}}^{\text{CF}}_{\alpha}(x)$	0.057	0.003	11.2%	5.0%	91.7%
Gaussian	$\hat{\text{VaR}}^{\text{norm}}_{\alpha}(x)$	0.057	0.004	9.8%	3.0%	88.2%
GPD	$\hat{\text{VaR}}^{\text{GPD}}_{\alpha}(x)$	0.058	0.003	12.5%	6.4%	93.3%
Gaussian unbiased	$\hat{\text{VaR}}^{\text{u}}_{\alpha}(x)$	0.052	0.003	-	-	-

Backtesting Expected Shortfall

In this example we will use the same dataset, but instead of VaR at level 5% we consider ES at level 10%. Following the notation in Equations (16)–(14) we obtain the estimators

$$\hat{\mathrm{ES}}_\alpha^{\mathrm{emp}}(x) := -\left(\frac{\sum_{i=1}^n x_i \mathbb{1}_{\{x_i + \hat{\mathrm{VaR}}_\alpha^{\mathrm{emp}}(x) < 0\}}}{\sum_{i=1}^n \mathbb{1}_{\{x_i + \hat{\mathrm{VaR}}_\alpha^{\mathrm{emp}}(x) < 0\}}} \right), \tag{17}$$

$$\hat{\mathrm{ES}}_\alpha^{\mathrm{CF}}(x) := -\left(\bar{x} + \bar{\sigma}(x) C(\bar{Z}_{CF}^\alpha(x)) \right), \tag{18}$$

$$\hat{\mathrm{ES}}_\alpha^{\mathrm{norm}}(x) := -\left(\bar{x} + \bar{\sigma}(x) \frac{\phi(\Phi^{-1}(\alpha))}{1-\alpha} \right), \tag{19}$$

$$\hat{\mathrm{ES}}_\alpha^{\mathrm{GPD}}(x) := \frac{\hat{\mathrm{VaR}}_\alpha^{\mathrm{emp}}(x)}{1-\hat{\xi}} + \frac{\hat{\beta} - \hat{\xi} u}{1-\hat{\xi}}, \tag{20}$$

Let the Gaussian unbiased Expected Shortfall estimator be

$$\hat{\mathrm{ES}}_\alpha^{\mathrm{u}}(x) := -(\bar{x} - \bar{\sigma}(x) a_n), \tag{21}$$

where a_n was computed numerically. Note that the **non-elicitability of ES** is directly reflected in (17) and (20). The joint elicitability of ES together with VaR is also visible: the estimator for the ES also makes use of an estimator for VaR.

- For the backtest we follow *Test 2* suggested in Acerbi and Székely (2014) utilizing the 50 separate subsets of our data denoted by $(x_1^i, \dots, x_{50}^i)$.
- The test statistic for the backtest is given by

$$Z := \frac{1}{49} \sum_{i=1}^{49} \left(\frac{1}{50} \sum_{j=1}^{50} \frac{x_j^{i+1} \, \mathbb{1}_{\{x_j^{i+1} + \hat{\mathrm{VaR}}_\alpha^i < 0\}}}{\alpha \, \hat{\mathrm{ES}}_\alpha^i} \right) + 1. \tag{22}$$

- The results of our backtest are presented in Table 4.
- Note that while 0 would be optimal, negative values of the test statistic Z correspond to underestimtion of risk of the considered estimator.

The unbiased estimator clearly outperforms the biased estimators, both on the market data and the simulated data.

Tabelle: The results of our backtest. Note that while 0 would be optimal, negative values of the test statistic *Z* correspond to **underestimtion** of risk.

Type of data:	MARKET				
Portfolio	Estimator type				
	ES_α^{emp}	ES_α^{norm}	ES_α^{CF}	ES_α^{GPD}	ES_α^{u}
LoBM.LoOP	-0.357	-0.393	-0.325	-0.302	-0.331
BM1.OP2	-0.428	-0.303	-0.338	-0.335	-0.235
BM1.OP3	-0.327	-0.322	-0.336	-0.295	-0.254
BM1.OP4	-0.326	-0.354	-0.348	-0.282	-0.272
LoBM.HiOP	-0.424	-0.421	-0.371	-0.335	-0.331
mean	-0.374	-0.363	-0.339	-0.308	-0.290

Type of data:	SIMULATED				
	ES_α^{emp}	ES_α^{norm}	ES_α^{CF}	ES_α^{GPD}	ES_α^{u}
LoBM.LoOP	-0.177	-0.073	-0.077	-0.104	**-0.005**
BM1.OP2	-0.143	-0.083	-0.069	-0.074	**-0.014**
BM1.OP3	-0.220	-0.084	-0.100	-0.157	**-0.019**
BM1.OP4	-0.224	-0.086	-0.101	-0.150	**-0.012**
LoBM.HiOP	-0.183	-0.082	-0.072	-0.098	**-0.016**
mean	-0.174	-0.101	-0.103	-0.109	**-0.030**

Conclusion

- We studied the estimation of risk, with a particular view on **unbiased** estimators and backtesting.
- The new notion of unbiasedness introduced is motivated from economic principles rather than from statistical reasoning, which links this concept to a better performance in backtesting.
- Some unbiased estimators, for example the unbiased estimator for value-at-risk in the Gaussian case, can be computed in closed form while for many other cases numerical methods are available.
- A small empirical analysis underlines the outperformance of the unbiased estimators with respect to standard backtesting measures.

The paper is available on SSRN: `https://ssrn.com/abstract=2890034`

Many thanks for your attention !

Acerbi, C. and Székely, B. (2014), 'Back-testing expected shortfall', **Risk magazine** (November).

Davis, M. (2016), 'Verification of internal risk measure estimates', **Statistics & Risk Modeling 33**, 67 – 93.

Fama, E. F. and French, K. R. (2015), `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`. accessed 20.10.2015.

Francioni, I. and Herzog, F. (2012), 'Probability-unbiased value-at-risk estimators', **Quantitative Finance 12**(5), 755–768.

Frank, D. (2016), 'Adjusting var to correct sample volatility bias', **Risk magazine** (October).

McNeil, A., Frey, R. and Embrechts, P. (2005), **Quantitative Risk Management: Concepts, Techniques and Tools**, Princeton University Press.