

Goodfellow, Ch., & Salm, Ch. (2016). Risky Risk Measures: A note on underestimating financial risk under the Normal assumption. *Copernican Journal of Finance & Accounting*, 5(2), 85–108. <http://dx.doi.org/10.12775/CJFA.2016.017>

CHRISTIANE GOODFELLOW*

Jade University

CHRISTIAN SALM

**RISKY RISK MEASURES:
A NOTE ON UNDERESTIMATING FINANCIAL RISK
UNDER THE NORMAL ASSUMPTION**

Keywords: risk measurement, risk management, downside risk, value at risk, copula.

J E L Classification: C02, C18, G11, G28.

Abstract: This note compares three different risk measures based on the same stock return data: (1) the portfolio variance as in Markowitz (1952), (2) the value at risk based on the historical return distribution, and (3) the value at risk based on a t copula. Unless return series follow a Normal distribution, Normal-based risk measures underestimate risk, particularly so during periods of market stress, when accurate risk measurement is essential. Based on these insights, we recommend that supervisors discontinue to accept Normal-based value at risk estimations. We are happy to share our commented R-code with practitioners who wish to implement our methodology. Risk measurement is the foundation of risk management and hence of vital importance in any financial institution. Supervisory capital requirements according to Basel III or Solvency II are also derived from risk measures. Investors are interested in ratings which are based on risk assessments. This note is therefore relevant to practitioners and supervisors alike.

Date of submission: November 21, 2016; date of acceptance: January 30, 2017.

* Contact information: christiane.goodfellow@jade-hs.de, Jade University, Friedrich-Paffrath-Str. 101, D-26389 Wilhelmshaven, Germany, phone: +49 441 95726272.

■■■ INTRODUCTION

Risk measurement has come a long way since Markowitz (1952) where the risk associated with holding a stock or a portfolio is captured by its return variance, and the covariation between the instruments inside the portfolio is measured with the covariance. The focus on just mean and variance of a portfolio assumes that returns follow a Normal distribution, and the covariance only covers pairwise linear dependencies between the constituents of the portfolio. Furthermore, investors' prime interest is in the downside risk, rather than general variation. Hence, risk measurement should focus on the left tail of the return distribution, and it should not rely on the returns being Normally distributed. The value at risk is a risk measure that addresses these issues; it is a quantile of the multivariate statistical distribution of the portfolio value (or the portfolio return). More recently, copulas have been applied to risk measurement. Despite these improvements in risk measurement accuracy, both Basel III and Solvency II still rely on correlation coefficients to estimate interdependencies of different risks.

We calculate portfolio risk with (1) the portfolio variance as in Markowitz (1952), (2) the value at risk based on the historical return distribution, and (3) the value at risk based on copulas. We compare the estimation results of (2) and (3) to the value at risk based on the assumption that the returns follow a Normal distribution. Furthermore, we present a simple step-by-step procedure to derive an accurate risk measure with copulas.

Measuring risk reliably is relevant for practitioners in risk management and for supervisors alike. Both banks and insurance companies have an interest in not overestimating risks that could lead to capital requirements, and in not underestimating risks which might jeopardise the existence of the business. Likewise, investors and rating agencies are interested in risk assessments.

According to portfolio selection theory established by Markowitz (1952), the investor aims for a portfolio with maximum expected return and minimum risk, measured by the expected portfolio variance. This expected variance is minimised primarily by combining securities with low covariances among themselves. This leads to a well-diversified portfolio, where diversification refers to choosing securities whose gains will offset others' losses. Specifically, a well-diversified portfolio contains securities from companies of different, and dissimilar, industries, and investment decisions are not made by looking

at a particular security in isolation, but rather by evaluating expected return and variance of the overall portfolio when this particular security is added to it.

The investor seeks to minimise the expected portfolio variance σ_p . She can achieve this by choosing securities with low variances σ_i^2 , usually resulting in a low portfolio return which she is trying to maximise at the same time, and/or she can combine securities such that ρ_{ij} becomes minimal, which, by definition, is -1. The securities are then perfectly negatively correlated, i.e. when the return of security I increases, the return of security J decreases by the same amount. In such an extreme scenario, there is a set of weights (i.e. percentages of assets) that will completely eliminate expected portfolio risk, i.e. σ_p^2 . Sadly, in the real world, securities do not move in perfect lockstep, so that we do not actually observe perfectly negatively correlated return series. Still, a portfolio should be constructed from securities with dissimilar characteristics. The returns of these different securities will react differently to market movements, so that the gain in one security can (partially) compensate the loss in another, and thus the risk to the investor is substantially smaller than if she was holding just one security and be fully exposed to its fluctuations.

Portfolio selection is particularly relevant if you believe in actively managed funds. There is, however, empirical evidence suggesting that passive investment strategies outperform actively managed investments net of costs and fees and after controlling for risk (e.g. Malkiel, 2003). The investment advice would then be to buy an index and hold it. Still, institutional investors are widely trusted with active fund management, turning portfolio selection into a topic that is not just of academic interest, but also highly relevant to practitioners.

More recently, a number of deficiencies in traditional risk measures have been addressed:

- The portfolio value or the portfolio return are no longer assumed to follow a Normal distribution.
- Risk measures have been extended to also include non-linear dependencies.
- It has been recognised that return series exhibit different covariation patterns across the observations, e.g. extreme observations are more strongly dependent than less extreme observations. This risk of joint extreme values is called 'tail dependence'.
- Tail dependence can be symmetric (i.e. similar in both tails of the distribution) or asymmetric (e.g. far stronger in the lower tail of the distribution than in the upper).

- Downside risk has been separated from upside risk (Campbell et al., 2001; Benson et al., 2008).

These weaknesses in traditional risk measures will systematically bias the risk estimates and can lead to wrong decisions in risk management, in wrongly calculating supervisory capital requirements, in inaccurate risk assessments by rating agencies, etc. We present a simple and statistically correct methodology to derive comprehensive and accurate risk measures. Section 2 introduces the three risk measures we apply in this note, before section 3 presents the dataset and the empirical results, i.e. the extent to which the risk measures differ. Section 4 summarises our results and concludes.

METHODOLOGY

We present, and compare, three different risk metrics based on the same stock return data, thereby showing to what extent the empirical results differ across the different risk measures. From these differences, we derive recommendations for practitioners and supervisory authorities on how to correctly capture risk in order to accurately calculate, for example, capital requirements. Furthermore, we present a step-by-step procedure to calculate a comprehensive risk measure based on copulas.

Our analysis is based on continuously compounded returns in percent: $r_t = 100 \times \ln \frac{P_t}{P_{t-1}} = 100 \times (\ln P_t - \ln P_{t-1})$, where P_t is the price of a stock or a portfolio in period t . We apply three different risk measures. First, the portfolio variance based on Markowitz (1952). Second, the conventionally calculated value at risk (VaR). Third, the VaR based on copulas. The following subsections explain these three approaches.

According to the Basel III framework, and more specifically, according to the CRR, a Stressed Value at Risk (SVaR) has to be calculated in addition to the conventional VaR, because the latter is procyclical. The SVaR should, in principle, be calculated with a 99% confidence level. In practice, however, this can be achieved by calculating the risk measures for a period of financial stress, e.g. July 2007 to December 2009. We follow this approach and use these stressed results as a robustness check.

Portfolio variance

In the Markowitz (1952) world, the expected variance of the portfolio return (in short: portfolio variance) is the weighted sum of the expected variances of the securities in the portfolio plus the weighted sum of the expected covariances between all the possible pairs of the securities in the portfolio. The weights are the percentages of an investor's assets allocated to each security, and the covariances term has to be included twice because the covariance between stock I and stock J is the same as the covariance between stock J and stock I:

$$\sigma_p^2 = \sum_i a_i^2 \sigma_i^2 + 2 \sum_i \sum_{i>j} a_i a_j \sigma_{ij}$$

where σ_{ij} is the covariance between securities I and J. A standardised measure of the covariance σ_{ij} is the correlation coefficient ρ_{ij} , with $\rho_{ij} = \sigma_{ij} / \sigma_i \sigma_j$ and therefore $-1 \leq \rho_{ij} \leq +1$. The expected portfolio variance then becomes:

$$\sigma_p^2 = \sum_i a_i^2 \sigma_i^2 + 2 \sum_i \sum_{i>j} a_i a_j \rho_{ij} \sigma_i \sigma_j.$$

The covariances and hence correlation coefficients only cover linear dependencies between return series. As a result, the portfolio risk estimate is systematically biased. Moreover, returns often exhibit positive left-tail dependence, which means that one return will more likely perform poorly when another already is performing badly. This phenomenon is neither covered by the covariance nor by the implicit Normal assumption. Despite these insights, both Basel III and Solvency II rely on correlation coefficients to estimate interdependencies of different risks. Furthermore, investors (or, more generally, stakeholders) are concerned about downside risk, whereas they usually quite enjoy the results of upside risk. Portfolio risk treats both risk directions equally, whereas the value at risk is only concerned with the lower tail of the distribution.

Value at Risk based on historical return distribution

The value at risk (VaR) is defined as the minimum portfolio value (or maximum loss) over a given holding period and with a given level of statistical significance (i.e. probability). It can be derived based on the historical distribution of the portfolio value (or the portfolio return), in which case it captures the covariation of different portfolio components, including non-linear inter-

dependencies. Alternatively, it can be assumed that the portfolio value follows a multivariate Normal distribution with only linear dependencies among the underlying risk factors. This latter version is the original JP Morgan approach. However, given that return series generally violate the Normal assumption and exhibit non-linear dependencies, we favour the historical distribution approach, even though it requires a larger dataset and implicitly assumes that the past is an adequate predictor of the future.

The historical distribution-based VaR has three major advantages over the portfolio risk measure in the Markowitz (1952) world:

- 1) The VaR covers all dependencies among the portfolio components, including non-linear covariation.
- 2) The portfolio value (or return) does not have to follow a Normal distribution, which is implicitly assumed in the mean-variance-universe of Markowitz (1952). Empirically, return distributions have fatter tails than the Normal distribution (e.g. Mandelbrot, 1963; Fama, 1965; Peters, 1996), which is often addressed by using the Student's t distribution.
- 3) It evidently only captures downside risk, which is what investors wish to be compensated for, rather than including upside risk in the risk measure assuming equal investor aversion.

Given that the VaR is intuitive to understand, it has been popular especially among practitioners.

But there are also downsides to the VaR. First, it is just one figure, usually a loss that will only be exceeded with a very small probability over a certain period, which does not provide information on the shape of the distribution in the tail, i.e. if the low-probability scenario materialises and the loss happens to be larger, how quickly does it become how large? One solution to this problem is to rely on a supplement to the VaR which provides information on the tail. One of these supplements is the Expected Shortfall, which is the expected loss given that this loss is larger than a certain benchmark (Rachev et al., 2010). Another, rather similar, supplement to the VaR are Lower Partial Moments (LPM). These are characteristics of specifically the tail to the left of the VaR. LPM 1 is the average across all negative deviations from the pre-defined VaR, while LPM 2 is the variance of the negative deviations from the pre-defined VaR. LPM 2 thus places a heavier weight on larger deviations than on smaller ones, reflecting that investors perceive high losses on rare occasions as worse than smaller losses more often.

Related to this lack of information about the tail is the fact that one can construct two Normal distributions with different parameter values but the same α quantile and hence the same VaR, but differently fat tails. Furthermore, the VaR is based on a statistical distribution. This can either be assumed (e.g. Normal, which we know is wrong), or estimated from historical data (which requires a large dataset and assumes that the past can predict the future), or obtained through a Monte Carlo Simulation (requiring substantial calculations). If the return data are assumed to follow a Normal distribution while in fact they are Student's t distributed, for example, the Normal-based VaR will underestimate the potential loss in portfolio value substantially. There is thus a certain modelling risk associated with this risk metric.

Finally, for completeness, a few closely related risk metrics should be mentioned that also focus specifically on downside risk. None of these are calculated in this paper as we focus on the copula methodology. The Conditional Value at Risk (CVaR), also called mean shortfall or expected tail loss, is calculated as the average loss, conditional on the losses being at least as large as the losses under a Normal distribution VaR would be (Xiong et al., 2014; Rockafellar & Uryasev, 2000). This average is probability-weighted. A close relative to the CVaR is the Excess Conditional Value at Risk (ECVaR), which is the CVaR in excess of the implied Normal distribution CVaR. Another tail risk measure is co-skewness. Skewness in general captures how asymmetric the data are, while co-skewness addresses the question 'If a particular financial asset is added to the portfolio, how does this change the portfolio's skewness?'. The decision on whether to purchase a particular financial asset should hinge on how this marginal instrument affects the statistical distribution of the portfolio value (and its central moments), rather than on a stand-alone evaluation of that particular instrument's features (this is the Markowitz, 1952 idea). Xiong et al. (2014) apply both the ECVaR and the co-skewness with similar empirical results.

Value at Risk based on Copulas

The VaR is a quantile of the statistical distribution of the portfolio value or the portfolio return. Copulas (or 'copulae') are an alternative way to present a multivariate model, which in turn can be used to derive a value at risk.

Copulas are particularly useful in modelling tail dependencies. Tail dependence means that dependencies in the data vary across the observations, e.g. extremely low values tend to be realised jointly more often than other values.

If this is equally true for both tails of the distribution, the tail dependence is symmetric, and the Student's t distribution is, in general, a reasonable choice of marginal distribution. If, on the other hand, this dependence is stronger in the lower tail of the distribution (or vice versa), which is common in financial data, the tail dependence is asymmetric, and so an extreme value copula should be used (Gumbel is an example, although this only works for positive correlations). Our analysis is centred around the current supervisory requirement: Risk measures based on the Normal assumption. We therefore do not endlessly refine our copula calculation, but rather argue that a t copula risk calculation will be more accurate than any conventional, Normal distribution based measure that supervisors will accept.

Let X and Y be two random variables, e.g. the values of two stocks or their returns, with the joint continuous probability density function $f(X,Y)$ and the joint cumulative distribution function (CDF) $F(X,Y)$. The marginal probability density functions are $f(X)$ and $f(Y)$, respectively. Likewise, the individual cumulative distribution functions are $F(X)$ and $F(Y)$. According to Sklar (1959), there is a function:

$$C_{X,Y}(F(X), F(Y)) = F(X, Y)$$

with $C_{X,Y}: [0,1] \times [0,1] \rightarrow [0,1]$. In other words, the joint cumulative distribution function can be split into (i) the marginal cumulative distribution functions representing the univariate features, and (ii) the copula C capturing the joint characteristics of the joint cumulative distribution function. Thus, the copula only contains information on the dependencies between X and Y . The marginal distributions do not have to be identical. The copula is the joint cumulative distribution function with the individual quantile functions

$$C(u, v) = F_{X,Y}(F_X^{-1}(u), F_Y^{-1}(v)).$$

As marginal cumulative distribution functions, we choose Student's t distribution as well as the Normal distribution. We expect that the Normal distributions will underestimate risk, unless the return series are Normally distributed. We include the Normal based risk measures in our analysis, however, to show exactly by how much it underestimates risk compared to the t distributed marginal functions.

For the Normal distribution, the corresponding copula is the Gaussian copula:

$$C(u, v; \rho) = \Phi_\rho(\Phi^{-1}(u), \Phi^{-1}(v)) \\ = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\rho^2)} \exp\left(\frac{-(s^2 - 2\rho st + t^2)}{2(1-\rho^2)}\right) ds dt$$

with Φ being the cumulative distribution function of the univariate standard Normal distribution and Φ_ρ the cumulative distribution function of the joint, bivariate standard Normal distribution with parameter ρ ($-1 < \rho < 1$). The Gaussian copula is completely specified by the correlation matrix, which confirms the argument that the covariance or correlation matrix is a comprehensive enough risk metric only if the returns follow a Normal distribution.

For the Student's t distribution, the corresponding copula is the t copula:

$$C(u, v; \nu, \rho) = \int_{-\infty}^{F_\nu^{-1}(u)} \int_{-\infty}^{F_\nu^{-1}(v)} \frac{1}{2\pi\sqrt{1-\rho^2}} \left(1 + \frac{(s^2 - 2\rho st + t^2)}{\nu(1-\rho^2)}\right)^{\frac{-(\nu+2)}{2}} ds dt$$

with F_ν being the univariate cumulative distribution function of the Student's t distribution with $\nu > 0$ degrees of freedom and ($-1 < \rho < 1$). The number of degrees of freedom controls the probability of joint movements. For $\nu \rightarrow \infty$, the t copula approaches the Gaussian copula. Both copulas only model symmetric tail dependence, as they are both symmetric distributions.

Since the bivariate cumulative distribution function is estimated:

$$\hat{F}(x, y) = \frac{1}{T} \sum_{i=1}^T (x_i \leq x, y_i \leq y)$$

we also estimate the copula with $(\hat{u}_1, \hat{v}_1), \dots, (\hat{u}_T, \hat{v}_T)$:

$$\hat{C}(u, v) = \frac{1}{T} \sum_{i=1}^T (\hat{u}_i \leq u, \hat{v}_i \leq v)$$

with $(u, v) \in [0,1]$. Both this function and its contour lines can be visualised in a graph, and the shape of these isolines shows the relationship between X and Y . Perfectly circular lines are evidence in favour of independence, whereas an L shape or ellipses suggest dependencies. The stronger the dependence between two return series, the higher the portfolio risk, under otherwise equal circumstances.

Copulas are an abstract concept in that they do not directly convey the VaR. Instead, here is a step-by-step procedure to estimate the t copula and then the VaR based on this:

- 1) For the univariate margins, estimate the Kernel density. In addition, fit a Normal distribution and a t distribution.
- 2) Plot the marginal distributions from step (1). Compare, for each series, the Kernel density, the fitted Normal distribution, and the fitted t distribution. For the following steps, assume the t distribution achieves a good fit to the empirical Kernel function.
- 3) Based on t distributions, generate 1-dimensional marginal distributions over a uniform interval.
- 4) Estimate the t copula with Maximum Likelihood based on the transformed data from step (3).
- 5) Use the value of the maximised loglikelihood function and the Akaike information criterion (AIC) to evaluate the goodness of fit of this copula compared to, for example, the Gaussian copula. The higher the maximised loglikelihood function or the AIC, the better the fit.
- 6) Combine the best-fitting copula with the estimated marginal distributions from step (3) and generate artificial stock return data (simulation).
- 7) Determine the lowest percentile in this artificial distribution. This is the 99% copula-based VaR.

We use both the t copula and the Gaussian copula in order to measure the error of applying the Normal distribution to a non-Normally distributed return series.

Financial data often exhibit tail dependence: both variables tend to assume very low values or very high values at the same time. More formally, lower (upper) tail dependence refers to the probability that one portfolio component assumes very small (high) values given that the other component is already very low (high). Jointly normally distributed random variables will only have tail dependence when they are perfectly correlated (Rachev et al., 2010). Since tail dependence is a common phenomenon in stock returns, and since the correlation between stock returns is usually found to be less than 1, the Gaussian copula with Normal marginal distributions is a somewhat risky choice as it will not capture joint movements in extreme values. This further adds to the argument that risk measures based on the unrealistic assumption that returns follow a Normal distribution are likely to underestimate portfolio risk.

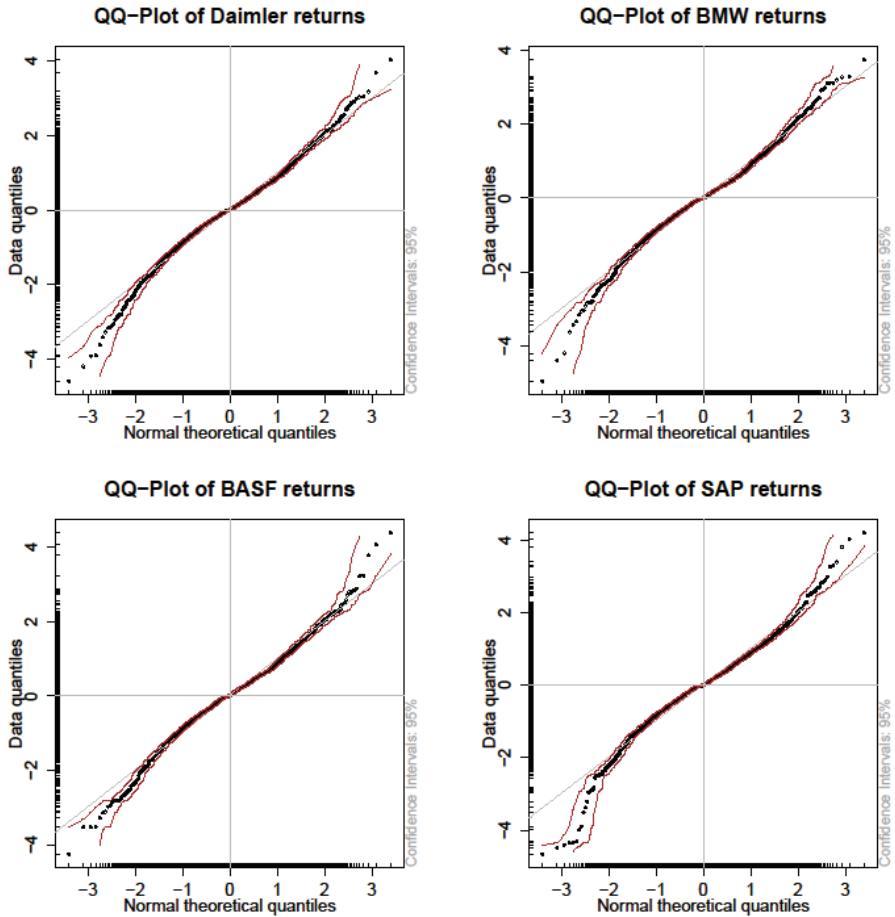
We are now acquainted with three risk measures: The Markowitz (1952) portfolio variance, the value at risk based on the historical return distribution, and the value at risk based on Gaussian and t copulas. We will now apply these risk measures to the same dataset and compare the results. We expect that the portfolio variance will underestimate risk as it ignores nonlinear dependencies. Likewise, the Normal assumption leads to underestimating risk with the value at risk and the Gaussian copula approaches. As a result, we expect the t copula based VaR to be the most accurate risk measure among the three investigated here.

DATA AND EMPIRICAL RESULTS

We use daily stock price data from yahoo finance from 30 December 2010 to 12 September 2016 in order to calculate continuously compounded returns, resulting in 1,484 return observations for the 4 German stocks Daimler, BMW, BASF, and SAP from 01 January 2011 to 12 September 2016. All 4 stocks are included in the DAX index, which is the bluechip index of the 30 largest German companies traded at Deutsche Börse in Frankfurt.

The Q-Q-plots of the return data for all 4 stocks are shown in figure 1. They indicate that (i) the return distributions are asymmetric in that the left tail is fatter than the right tail and that (ii) the Normal assumption along the diagonal line is violated at the 95% confidence level in the left tail. This confirms that risk measures based on the Normal assumption will underestimate the risk a financial institution faces. While the Normal assumption cannot be rejected in the upper tail of the Daimler return distribution, it is violated just like at the bottom end of the distribution for BMW, BASF, and SAP.

Figure 1. Q-Q-Plots for continuous daily returns of Daimler, BMW, BASF and SAP from 1 January 2011 to 12 September 2016



Source: developed by author.

The descriptive statistics for the 4 return series are given in Table 1.

Table 1. Descriptive Statistics for return series of Daimler, BMW, BASF, and SAP from 1 January 2011 to 12 September 2016

	Daimler	BMW	BASF	SAP
Number of Observations	1484	1484	1484	1484
Median	0,01728	0,00000	0,03981	0,06012
Mean	0,03105	0,02979	0,02648	0,05506
Variance	3,49495	3,50909	2,65169	1,84047
Standard Deviation	1,86948	1,87326	1,62840	1,35664
Skewness	-0,22993	-0,16826	-0,16072	-0,33708
Kurtosis	1,50169	1,55853	1,31826	2,41099
Correlation Coefficient	0,84		0,64	

Figures reported in this table are based on $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. In the Markowitz (1952) world, this table gives sufficient risk measures in that it lists standard deviations as well as correlation coefficients. The latter are needed to calculate the return covariances inside a portfolio.

Source: developed by author.

Table 1 shows that Daimler and BMW are riskier individual investments than BASF and SAP in the 2011 to 2016 period. Despite having the lowest standard deviation among the 4 stocks, SAP also scores the highest mean and median return. Combining BASF and SAP into a portfolio should be more risk-reducing than combining Daimler and BMW, as the correlation coefficient between the former is lower than for the latter. In other words, BASF and SAP absorb, respectively, more of each other's losses than is the case for Daimler and BMW. This is not surprising as 2 automobile stocks will move more closely together than a better diversified portfolio of BASF and SAP. It should be noted that the Markowitz (1952) risk measures can be calculated purely based on the data in Table 1. As these do not directly compare to the risk measures we derive from the estimated copulas, we do not report the Markowitz (1952) risk measures here.

Table 2. Conventional Conventional Value at Risk for return series of Daimler, BMW, BASF, and SAP as well as for a portfolio of Daimler and BMW and a portfolio of BASF and SAP

	VaR historical distribution	VaR Normal distribution	difference
Daimler	5,2391	4,3180	0,9211
BMW	4,9841	4,3281	0,6560
BASF	4,5047	3,7617	0,7430
SAP	3,7791	3,1010	0,6781
Daimler BMW	4,8897	4,1612	0,7285
BASF SAP	3,6310	3,1059	0,5251

Figures reported in this table are based on $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. Sampling period is 1 January 2011 to 12 September 2016, VaR holding period 1 day, 99% confidence level, portfolio weights 50%, difference in absolute values.

Source: developed by author.

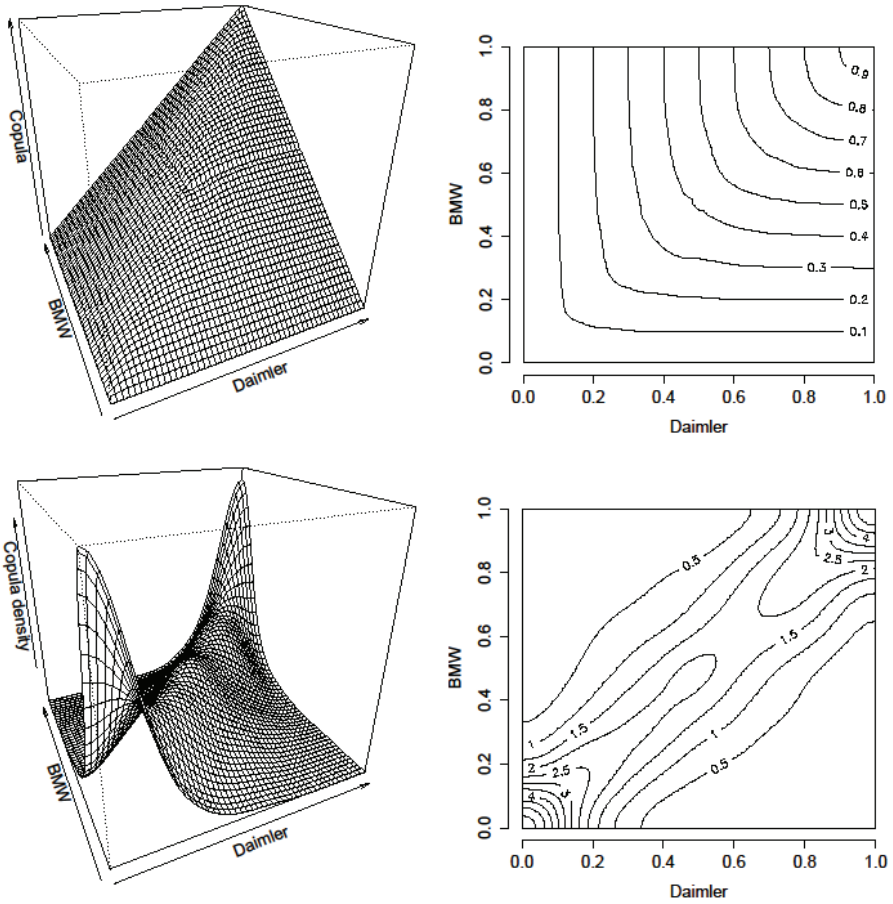
Table 2 reports the Value at Risk figures for all 4 stocks and for 2 portfolios which are constructed from 50% Daimler and 50% BMW in the first instance and 50% BASF and 50% SAP in the second instance, in order to validate the assumption that diversification is more beneficial in the second case. In fact, the value at risk for Daimler and BMW is higher than for BASF and SAP, suggesting that the expected loss is higher for the riskier portfolio. This is true for the VaR based on the empirical and the Normal distributions of the returns.

Indeed, the Normal assumption underestimates the risk (see column 'difference' in table 2). Not surprisingly, this holds true for individual stocks and for portfolios. Measured in Euros, for example, and based on a 1 million Euro portfolio in Daimler stocks, the maximum expected loss amounts to 52,391 Euros empirically or 43,180 Euros based on the Normal assumption, i.e. a difference in risk measure of 9,211 Euros or 18% of the empirical approach (or 21% under the Normal assumption, which is evidently wrong). The Normal assumption therefore produces an economically significant underestimate of risk. It should thus be avoided in practice, even though the original RiskMetrics methodology developed by JP Morgan relied on the Normal assumption. Since portfolio theory is also based on the Normal assumption, its estimation results based on the figures in Table 1 should be taken with a large grain of salt. These estimation results are not reported here because they cannot be directly compared to the VaR estimations based on copulas.

Interestingly, the return series whose bottom tail deviates the most from the Normal assumption, is SAP (see Figure 1). Despite this observation, the difference between the empirical VaR and the Normal VaR for SAP amounts to 0.6781 or 18% of the empirical approach, just like for Daimler. By contrast, BMW produces a difference between the empirical and the Normal VaRs of only 13% of the empirical value, while BASF produces 16.5% even though its bottom tail appears to be much closer to the Normal assumption than that of BMW or indeed SAP.

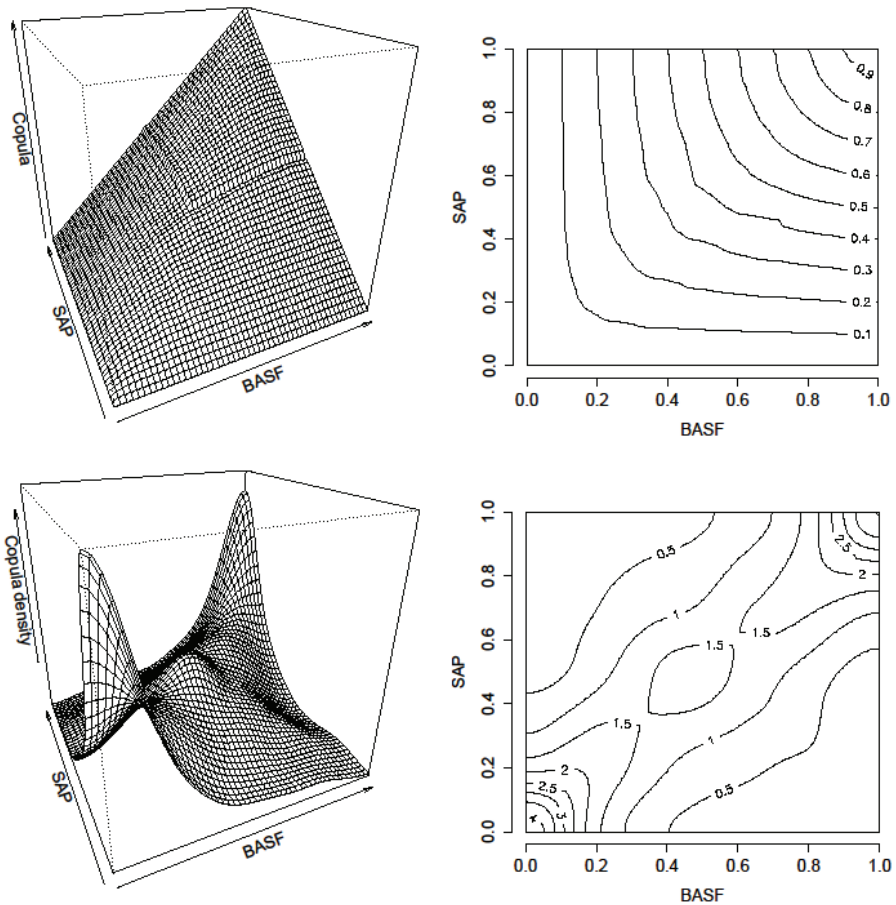
In order to determine whether the pairwise dependencies between the return series are strong enough to require a copula capturing this covariation at all, we plot the empirical copula cumulative density functions and the copula densities for the two portfolios of Daimler and BMW (figure 2) and BASF and SAP (figure 3). Both estimations are non-parametric, i.e. neither assumes that the returns follow a Normal or a t distribution. The contour lines indicate that the stocks are strongly dependent in that the isolines for the cumulative density functions are L shaped and the isolines for the copula densities are ellipses rather than circles. We conclude from this that it is worth proceeding with the copula estimation as there are substantial dependencies between the stock returns under investigation here.

Figure 2. Empirical (non-parametric) copula CDF and copula density for Daimler and BMW returns from 1 January 2011 to 12 September 2016



Source: developed by author.

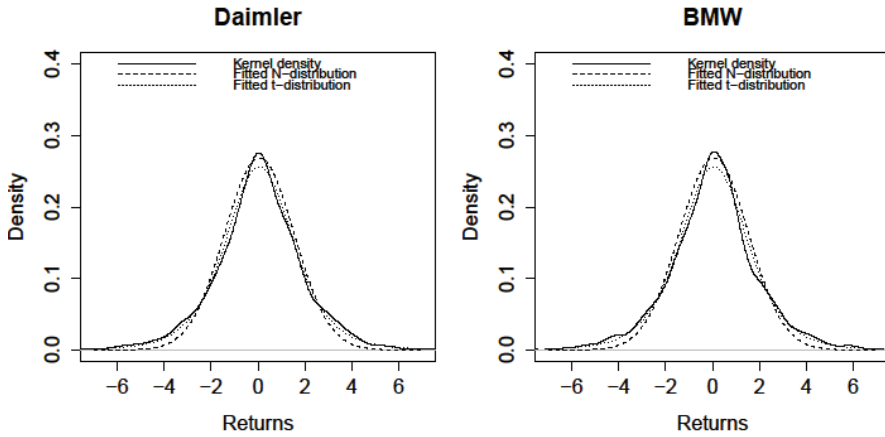
Figure 3. Empirical (non-parametric) copula CDF and copula density for BASF and SAP returns from 1 January 2011 to 12 September 2016



Source : developed by author.

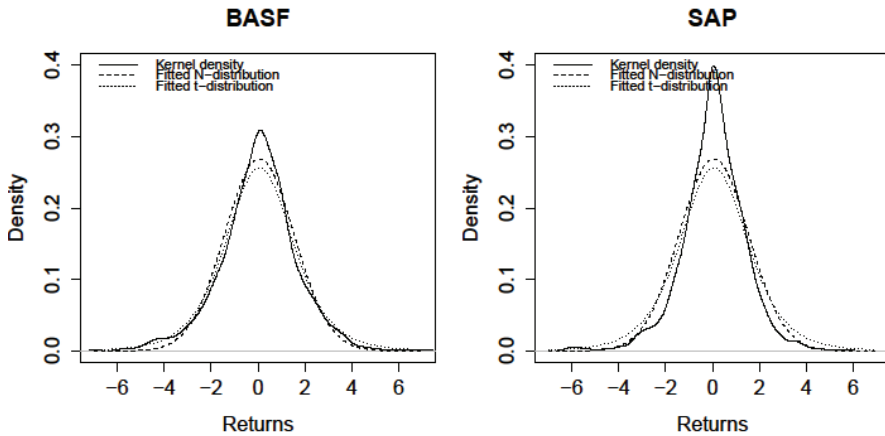
In light of these conclusions, we need to determine the suitable marginal distribution for each of the 4 stock returns. Figures 4 and 5 present the Kernel density, which is effectively the empirical frequency distribution, in comparison to the fitted Normal and *t* distributions. Not surprisingly, the *t* distribution is the best fit for the tails of the distributions. We therefore expect that the VaR based on the Gaussian copula will underestimate the risk, just like the Normal assumption for the conventional VaR calculation.

Figure 4. Fitted marginal distributions functions for Daimler and BMW returns from 1 January 2011 to 12 September 2016



Source: developed by author.

Figure 5. Fitted marginal distributions functions for BASF and SAP returns from 1 January 2011 to 12 September 2016



Source: developed by author.

Finally, Table 3 reports copula-based VaR figures.

Table 3. Copula-based Value at Risk for return series of Daimler, BMW, BASF, and SAP as well as for a portfolio of Daimler and BMW and a portfolio of BASF and SAP

	t-Copula VaR	Gaussian Copula VaR	difference
Daimler BMW	3,7277	3,2986	0,4291
BASF SAP	3,6121	2,4119	1,2002

Figures reported in this table are based on $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. Sampling period is 1 January 2011 to 12 September 2016, VaR holding period 1 day, 99% confidence level, portfolio weights 50%, difference in absolute values.

Source: developed by author.

Similarly to the conventional VaR estimates reported in table 2, we find that the Normal assumption leads to underestimating risk (see column ‘difference’ in table 3). We only report the results for the two portfolios as a copula does not make sense for individual return series. For the automobile portfolio, we underestimate risk with the Gaussian approach by 4,291 Euros or 11.5% of the *t* distribution VaR if 1 million Euros are invested into this portfolio. By contrast, for the better diversified portfolio of BASF and SAP, we underestimate risk by 12,002 Euros or 33% of the *t* distribution VaR for a 1-million Euro portfolio. With these results, we confirm our conjecture and our earlier conclusion that the Normal assumption substantially underestimates risk and should therefore be avoided in practice. While the empirical distribution is certainly the best fit, the *t* distribution works reasonably well for the 4 return series investigated here (see figures 4 and 5).

As a robustness check, we calculate the same indicators for the stress period of 02 July 2007 to 31 December 2009. During this time, the global financial crisis had its heaviest impact on financial markets. Therefore, the VaR calculated for this period is interpreted as the Stressed VaR in the Basel III framework. Tables 4 to 6 summarise the results.

Table 4. Descriptive statistics for stressed return series of Daimler, BMW, BASF, and SAP from 02 July 2007 to 31 December 2009

	Daimler	BMW	BASF	SAP
Number of Observations	635	634	637	635
Median	0,00000	-0,03806	0,00000	0,00000
Mean	-0,07830	-0,04896	-0,00651	-0,01690

Table 4. Descriptive statistics for stressed return series of Daimler...

	Daimler	BMW	BASF	SAP
Variance	10,52665	8,09299	7,27877	4,61677
Standard Deviation	3,24448	2,84482	2,69792	2,14867
Skewness	0,47567	0,31012	0,18209	-0,38797
Kurtosis	5,53263	3,47062	6,70795	13,16942
Correlation Coefficient	0,75		0,47	

Figures reported in this table are based $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. In the Markowitz (1952) world, this table gives sufficient risk measures in that it lists standard deviations as well as correlation coefficients. The latter are needed to calculate the return covariances inside a portfolio. This table compares directly to table 1, where the non-stressed results are summarised.

Source: developed by author.

Comparing tables 1 and 4, it becomes apparent that the mean and median returns were lower during the period of stress than across the later 5-year period from 2011 to 2016. If we were to measure risk with the standard deviation of the returns, the stress period would indeed be much riskier than the 2011 to 2016 period. Moreover, while the return data for the 5-year period are negatively skewed, the stressed return data are positively skewed (except for SAP), i.e. low observations occur with higher relative frequency than higher observations do during the stress period. This is plausible for a period of stress. The return series are leptokurtic both during the 2011 to 2016 period and under stress (the kurtosis is positive in both cases); however, during the stress period, the distributions are much thinner with heavier tails than under the Normal assumption, implying that extreme values are more likely to occur than during the 2011 to 2016 period. Finally, the return series become less correlated. This might be surprising at first sight; however, this only captures linear dependencies while a larger fraction of the overall co-movement may well be nonlinear during a period of stress.

Table 5. Conventional stressed Value at Risk for return series of Daimler, BMW, BASF, and SAP as well as for a portfolio of Daimler and BMW and a portfolio of BASF and SAP

	VaR historical distribution	VaR Normal distribution	difference
Daimler	9,3845	7,6323	1,7522
BMW	7,8666	6,6670	1,1996
BASF	8,9205	6,2979	2,6226
SAP	5,6164	5,0180	0,5984
Daimler BMW	8,4465	6,7732	1,6733
BASF SAP	7,2886	4,9716	2,3170

Figures reported in this table are based on $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. Sampling period is 02 July 2007 to 31 December 2009, VaR holding period 1 day, 99% confidence level, portfolio weights 50%, difference in absolute values. This table compares directly to table 2, where the non-stressed results are summarised.

Source: developed by author.

The conventional VaR figures are compared across tables 2 and 5. This comparison confirms our suspicion regarding linear and non-linear dependencies: The difference in VaR between the empirical and the Normal distributions becomes much larger in times of trouble (except for SAP which appears to function almost like an insurance policy). Furthermore, since the VaR measures an expected loss, we expect this to be higher in absolute terms during a period of stress compared to an ordinary state of the market. Indeed, for the stressed period of 2007 to 2009, we observe a VaR of 93,845 Euros for a 1 million Euro portfolio of Daimler stocks compared to 52,391 Euros in an ordinary state. Moreover, under ordinary circumstances, recall that the difference in VaR between the empirical and the Normal distributions for the same Daimler portfolio amounts to 9,211 Euros. By contrast, the difference between empirical and Normal-based VaRs amounts to 17,522 Euros in the stressed period, which is just under double the amount across the 5-year non-stressed period. We conclude that the error in using the Normal distribution to capture risk in financial markets becomes more severe as the markets experience a crisis. During such times, however, risk measures need to be extra-reliable and extra-prudent. This is a strong argument against the Gaussian copula to measure the dependence of return series.

Table 6. Stressed Copula-based Value at Risk for return series of Daimler, BMW, BASF, and SAP as well as for a portfolio of Daimler and BMW and a portfolio of BASF and SAP

	t-Copula VaR	Gaussian Copula VaR	difference
Daimler BMW	8,3938	4,5709	3,8229
BASF SAP	6,8408	2,6784	4,1624

Figures reported in this table are based on $r_t = 100 \times \ln\left(\frac{P_t}{P_{t-1}}\right)$. Sampling period is 02 July 2007 to 31 December 2009, VaR holding period 1 day, 99% confidence level, portfolio weights 50%, difference in absolute values. This table compares directly to table 3, where the non-stressed results are summarised.

Source: developed by author.

The copula-based VaR figures are reported in tables 3 and 6. During the stress period of 2007 to 2009, we observe much higher potential losses than during the 2011–2016 period. For example, a 1-million Euro Daimler-BMW portfolio loses with 99% no more than 37,277 Euros across 1 day, whereas the same portfolio loses 83,938 Euros under stress but with otherwise equal parameters. Furthermore, the difference between Normal and t distribution assumptions is much larger during the stress period than under ordinary circumstances. For a 1-million Euro Daimler-BMW portfolio, this difference in estimation approach amounts to 4,291 Euros under ordinary circumstances compared to 38,229 Euros during the stress period. This confirms our earlier finding that the mistake in underestimating risk using the Normal assumption is substantially bigger during a stress period compared to ordinary market circumstances. Since accurate risk measurement is particularly important during periods of market turbulence, both supervisors and practitioners should avoid using the Normal distribution as its tails are too thin.

In summary, we investigate continuous stock returns from 2011 to 2016, and during the stress period of 2007 to 2009. The return series deviate from the Normal distribution in both periods, but more severely so during the period of the financial crisis. This results in substantial errors in measuring risk based on the Normal assumption, and these errors are even larger during a period of market stress. In fact, this error is risky in that the Normal assumption underestimates risk. Based on this insight, we strongly recommend that financial institutions and supervisors adopt a non-Normal approach to capture the risk an institution faces. In order to facilitate the implementation of such an approach in practice, we are very happy to share our commented code that calculates

non-Normal risk measures with the open source software R. This code is available from the authors on request.

SUMMARY AND CONCLUSIONS

The goal of this paper is to compare risk measures that rely on the Normal assumption with those that do not, based on daily stock return data from 2011 to 2016. In fact, the Normal assumption substantially underestimates the risk an institution faces. During times of stressed markets from 2007 to 2009, this underestimation becomes even more severe. According to Basel III, however, a risk measure needs to adequately capture the risk facing an institution. It therefore appears that risk measures based on the Normal assumption should not be used. Despite this insight, regulators still accept Normal-based risk measures such as the value at risk based on the assumption that the return distribution is Normal. In order to help implement non-Normal risk measures in practice, we propose a simple methodology to comprehensively capture risk with a VaR based on the historical, empirical distribution or a VaR based on the t copula, and we are happy to share our code to produce accurate risk measures.

It is disadvantageous for a financial institution to overestimate risk because this may lead to equity requirements that are too high. It is, however, equally disadvantageous to underestimate risk since the institution has a vital interest in continuing to stay in business. This might be jeopardised if risks are systematically and substantially underestimated, resulting in too high an exposure. This concern is shared by supervisors. At the same time, it is still common practice to calculate the Value at Risk based on a Normal distribution assumption. In this paper, we have shown to what extent this method underestimates risk, especially so during times of stress, when risk exposure is particularly risky! In light of these findings, we strongly recommend that regulatory requirements are revised to exclude Normal-assumption based risk measures.

■■■ REFERENCES

- Campbell, R., Huismann, R., & Koedijk, K. (2001). Optimal portfolio selection in a Value-at-Risk framework. *Journal of Banking and Finance*, 25(9), 1789–1804. doi: [http://dx.doi.org/10.1016/S0378-4266\(00\)00160-6](http://dx.doi.org/10.1016/S0378-4266(00)00160-6).
- Fama, E. (1965). The Behavior of Stock Market Prices. *Journal of Business*, 38(1), 34–105. doi: <http://dx.doi.org/10.1086/294743>.

- Malkiel, B. (2003). Passive Investment Strategies and Efficient Markets. *European Financial Management*, 9(1), 1–10. doi: <http://dx.doi.org/10.1111/1468-036X.00205>.
- Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *Journal of Business*, 36(4), 394–419. doi: <http://dx.doi.org/10.1086/294632>.
- Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, 7(1), 77–91. doi: <http://dx.doi.org/10.2307/2975974>.
- Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments*. 16. New York: Cowles Foundation, Wiley. Second edition 1991, Blackwell, Oxford, UK.
- Peters, E.E. (1996). *Chaos and Order in the Capital Markets: A New View of Cycles, Prices, and Market Volatility*, 2nd edition. John Wiley and Sons.
- Rachev, S.T, Höchstötter, M., Fabozzi F.J., Focardi, S.M. (2010). *Probability and Statistics for Finance. The Frank Fabozzi Series*. New Jersey, Hoboken: John Wiley & Sons.
- Rockafellar, R., & Uryasev, S. (2000). Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2, 21–41. doi: <http://dx.doi.org/10.21314/JOR.2000.038>.
- Rubinstein, M. (2002). Markowitz' Portfolio Selection: A Fifty-Year Retrospective. *Journal of Finance*, 57(3), 1041–1045. doi: <http://dx.doi.org/10.1111/1540-6261.00453>
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges, *Publ. Inst. Statist. Univ. Paris*, 8, 229–231.
- Xiong, J.X, Idzorek, T.M., & Ibbotson, R.G. (2014). Volatility versus tail risk: Which one is compensated in equity funds? *Journal of Portfolio Management*, 40(2), 112–121.