

Copulas in Risk Management

Introduction

The use of copulas in the study of Value-at-Risk in competitive markets is a research method which offers rich possibilities. In this post we compare VaR based on selected copulas with different marginals. We illustrate how to compute Value-at-Risk using Monte Carlo simulations in Wolfram Mathematica.

Background

A copula is a tool that describes dependence between variables.

Sklar's theorem proves the existence of a copula C_H that “couples” any joint distribution $H(x_1, \dots, x_n)$ with its univariate marginals $F_1(x_1), \dots, F_n(x_n)$ via the relation $H(x_1, \dots, x_n) = C_H(F_1(x_1), \dots, F_n(x_n))$ and thus demonstrates that copula distributions are ubiquitous in multivariate statistics.

Copula distributions date as far back as the 1940s, though much of the terminology and machinery used today were developed in the 1950s and 1960s. Since their inception, copulas have been used to model phenomena in areas including reliability theory, meteorology, and queueing theory, while specially purposed copulas and kernels have been developed to serve as tools in fields such as survival analysis (via survival copulas) and mathematical finance (via panic copulas). Copula distributions are also of independent theoretical interest in Monte Carlo theory and applied mathematics.

The Mathematica function `CopulaDistribution[ker, {dist1, ..., distn}]` represents a multivariate statistical distribution whose j^{th} marginal distribution is precisely dist_j , and for which the CDF of a dist_j -distributed random variate follows a uniform distribution. In this context, varying `ker` allows investigation of different types and degrees of dependence. For example, `{“FGM”, α }` best models weak variable dependence, whereas “Product” allows analysis of independent variables.

Value at Risk

Risk managers need to measure the exposure of the portfolios to different risk factors. In standard practice they use Value-at-Risk (VaR) and Expected Shortfall (ES). Both these measures are multivariate in the sense that they must account for correlation among the factors from which losses may arise. But this practice is obsolete with respect to structured finance products, according to Cherubini et al. (2012), because these products are non-linear, so that their value may change even though market prices do not move but their volatilities are changed.

Value-at-Risk describes the loss (amount of capital) that can occur over a given period, at a given confidence level, due to exposure to market risk. The Value-at-Risk measure is defined as quantile of a probability distribution of losses over a given period.

The Value-at-Risk measure is defined as quantile of a probability distribution of losses over a given period:

$$q_{\alpha}(X) = \inf \{x : P(X \leq x) \geq \alpha\}$$

where X is a random variable representing the value of the portfolio of assets or exposures to risk factors.

Then, the VaR of an exposure X at confidence level α will be defined as

$$\text{VaR}(X) = q_{\alpha}(-X) = F_{-X}^{-1}(\alpha).$$

where F_{-X}^{-1} denotes the generalized inverse function $F^{-1}: (0,1) \rightarrow \mathbb{R}$ as

$$F_{-X}^{-1}(x) = \inf \{l \in \mathbb{R} : F(l) \geq x\}.$$

Risk assessment, whether by VaR or other means, entails a series of assumptions, notably about the distributions of risk factors and their co-movement, generally assumed to be multivariate Normal. But these assumptions are simplifications of the complex way in which markets tend to operate. In general, if there are known two marginal continuous distributions we cannot derive their joint distribution (Sklar, 1959, 1996), (Embrechts et al., 2001), (Embrechts et al., 2002), but we can recover a joint distribution using copula function. One of the advantages of using copulas is that they isolate the dependence structure from the structure of the marginal distributions (Sklar, 1959), (Embrechts et al., 2001), (Embrechts et al., 2002), (Cherubini et al. 2004), (McNeil, 2005), (Alexander, 2008), etc. The marginal distribution may capture different types of symmetries, asymmetries, fat tails and structural breaks with strong influence in estimation results for modeling of the dependence structure.

Copulas and Correlation

Copula functions express joint distributions of random variable X . A copula enables us to separate the joint distribution into marginal distributions of each variable. Sklar's theorem (Sklar, 1959) states that any bivariate or multivariate distribution can be expressed as the copula function $C(u_1, \dots, u_n)$ evaluated at each of the marginal distributions. Using probability integral transform, each continuous marginal $u_i = F_i(x_i)$ has a uniform distribution on $I \in [0,1]$, where $F_i(x_i)$ is the cumulative integral of $f_i(x_i)$ for the random variable X_i , where X_i assume values on the extended real line $(-\infty, \infty)$ (Kumar, P., 2010). The n -dimensional probability distribution function H has a unique copula representation:

$$H(x_1, \dots, x_n) = C_H(F_1(x_1), \dots, F_n(x_n)) = C(u_1, \dots, u_n)$$

The joint probability density is

$$h(x_1, \dots, x_n) = \prod_{i=1}^n f_i(x_i) \times C(u_1, \dots, u_n)$$

$f_i(x_i)$ is each marginal density and coupling is provided by the copula probability density.

$$h(u_1, \dots, u_n) = \frac{\partial C(u_1, \dots, u_n)}{\partial u_1 \partial u_2 \dots \partial u_n}$$

When the random variables are independent, $C(u_1, \dots, u_n)$ is identically equal to one.

The key point is that the joint distribution can be segmented into an independent portion, expressed as the product of the marginals, and the copula function $C(u_1, \dots, u_n)$. This separation enables us to model the dependency between the variates directly.

Different copulas produce different joint distributions when applied to the same marginals. Consider two random variables and assume that we have calibrated their marginal distributions. Now suppose we apply two different copulas and so we obtain two different joint distributions. So if only one joint distribution exhibits strong lower tail dependence then this distribution should be regarded as more risky than the one with a weaker, symmetric dependence, at least according to a downside

risk metric.

Correlation is typically expressed as the Pearson correlation coefficient. But the Pearson correlation coefficient is a measure of linear dependence, and it is defined as covariance divided by the product of standard deviations, by assuming the data to be normalized to unit variance (Cherubini et al. 2012). Linear correlation does not work for different probability distributions. Therefore it is necessary to use non-parametric measures such as Kendall's τ (12) or Spearman's correlation coefficient (rank correlation) (Spearman's ρ) (13), see (Alexander, 2008), (Embrechts et al., 2002), etc.

Both these measures can be expressed in terms of copulas:

$$\tau = 4 \int \int C(u_1, u_2) dC(u_1, u_2) - 1$$

$$\rho = 12 \int \int u_1 u_2 dC(u_1, u_2) - 3$$

A great many forms have been applied in finance. Perhaps the most fundamental, if not the most useful, is the **bivariate normal copula** which takes the form:

$$C(u_1, u_2; \theta) = \Phi_N(\Phi^{-1}(u_1), \Phi^{-1}(u_2); \theta)$$

where Φ is the cdf of the standard normal distribution and Φ_N is the standard bivariate normal distribution with correlation parameter θ restricted to the interval $[-1, 1]$.

While the normal copula is flexible in that it allows for both positive and negative dependency, it has very weak tail dependency and is therefore not usually appropriate for modeling financial assets.

The correlation parameter θ is given as:

$$\theta = \sin\left(\frac{\pi}{2} \text{Kendall's } \tau\right)$$

The **Student t copula** is another implicit copula with ν degrees of freedom and correlation θ derived from a multivariate Student t distribution function

$$C_\nu(u_1, u_2; \theta) = t_\nu(t_{\nu_1}^{-1}(u_1), t_{\nu_2}^{-1}(u_2); \theta)$$

where t_ν and t_{ν_i} are multivariate and univariate Student t distribution functions with ν , or ν_i degrees of freedom that controls the heaviness of the tails and correlation parameter θ . This copula has symmetric tail dependency that is higher than those in normal copula, but the dependency is symmetric.

Copula parameters can be calibrated using Maximum Likelihood Estimate (MLE) (see (Alexander, 2008)), when both parameters are estimated together or it is calibrated only degree of freedom using MLE and as correlation parameter θ is used Spearman's correlation parameter:

$$\theta = 2 \sin\left(\frac{\pi}{6} \text{Spearman's } \rho\right)$$

The **Clayton copula** is widely used in finance, because it allows us to model stronger left tail dependence and weaker right tail dependence, characteristics that are typical of asset processes. It takes the form

$$C(u_1, u_2; \theta) = (u_1^{-\theta} + u_2^{-\theta} - 1)^{-\frac{1}{\theta}}$$

with the dependence parameter θ restricted on the region $(0, \infty)$. As θ approaches zero, the marginals become independent. It has been used to study correlated risks, it exhibits strong left tail dependence and relatively weak right tail dependence.

The relationship between Kendall's τ and θ is:

$$\text{Kendall's } \tau = \frac{\theta}{\theta + 2}$$

The **Frank copula** takes the form

$$C(u_1, u_2; \theta) = -\theta^{-1} \log \left\{ 1 + \frac{(e^{-\theta u_1} - 1)(e^{-\theta u_2} - 1)}{e^{-\theta} - 1} \right\}$$

where the dependence parameter θ may assume any real value from $(-\infty, \infty)$. It means that Frank copula can be used to model outcomes with strong positive or negative dependence. Frank copula is most appropriate for data that exhibits weak tail dependence compared to the normal copula and the strongest dependence in middle of the distribution.

Here, the relationship between Kendall's τ and θ is:

$$\text{Kendall's } \tau = 1 - \frac{1}{\theta}$$

Finally, the **Gumbel-Hougaard copula** takes the following form:

$$C(u_1, u_2; \theta) = \exp \left[- \left((-\ln u_1)^\theta + (-\ln u_2)^\theta \right)^{\frac{1}{\theta}} \right]$$

where $\theta \in [1, \infty]$. $\theta = 1$ corresponds to independence, $\theta \rightarrow \infty$ correspond to the perfect positive dependence. The Gumbel copula does not allow negative dependence, but in contrast to Clayton, Gumbel exhibits strong right tail dependence and relatively weak left tail dependence. The Gumbel copula is appropriate when risk factors are strongly correlated at high values but less correlated at low values.

Parameter θ is expressed using Kendall's τ as:

$$\text{Kendall's } \tau = 1 - \frac{1}{\theta}$$

Modeling VaR Using Copulas

To illustrate the application of copulas in a risk management context, we will construct a simple portfolio comprising the S&P 500 and NASDAQ Indices. We begin by computing daily log-returns for each series and examining their moments.

Data

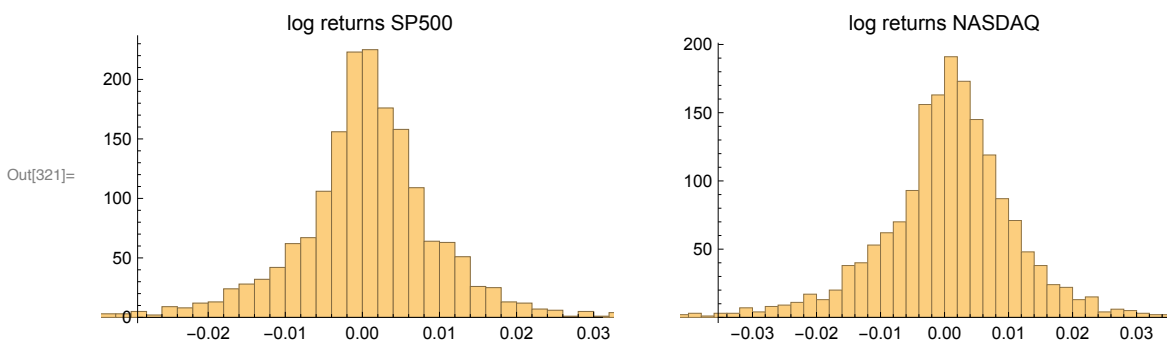
```
In[314]:= SP500prices = FinancialData["^GSPC", {{2010, 2, 1}, {2017, 1, 20}}];
NASDAQprices = FinancialData["^IXIC", {{2010, 2, 1}, {2017, 1, 20}}];
SP500returns =
  Log[Drop[SP500prices[[All, 2]], 1]] - Log[Drop[SP500prices[[All, 2]], -1]];
NASDAQreturns = Log[Drop[NASDAQprices[[All, 2]], 1]] -
  Log[Drop[NASDAQprices[[All, 2]], -1]];
```

```
In[318]:= TableForm[Through[{Mean, StandardDeviation, Skewness, Kurtosis}[
  Transpose[{SP500returns, NASDAQreturns}]]],
  TableHeadings → {"Mean", "St. Dev.", "Skewness", "Kurtosis"},
  {"SP500", "NASDAQ"}], TableAlignments → Right]
```

Out[318]//TableForm=

	SP500	NASDAQ
Mean	0.000418759	0.000535315
St. Dev.	0.00976428	0.0109706
Skewness	-0.44119	-0.422737
Kurtosis	7.30581	6.3333

```
In[319]:= p1 = Histogram[SP500returns, PlotLabel → "log returns SP500"];
p2 = Histogram[NASDAQreturns, PlotLabel → "log returns NASDAQ"];
GraphicsRow[{p1, p2}, ImageSize → Large]
```



Marginal Distribution Fitting

We begin by estimating the parameters for both normal and Student T distributions for each series by maximum likelihood and proceed to test each distribution for goodness of fit. In each case, we find that all of the fitness tests reject the null hypothesis of normally distributed returns, but that the Student T distribution appears to provide an adequate fit for both series.

```
In[322]:= params11 = FindDistributionParameters[SP500returns, NormalDistribution[ $\mu_{11}$ ,  $\sigma_{11}$ ]]
params12 =
  FindDistributionParameters[NASDAQreturns, NormalDistribution[ $\mu_{12}$ ,  $\sigma_{12}$ ]]
params21 = FindDistributionParameters[
  SP500returns, StudentTDistribution[ $\mu_{21}$ ,  $\sigma_{21}$ ,  $\nu_1$ ]]
params22 = FindDistributionParameters[NASDAQreturns,
  StudentTDistribution[ $\mu_{22}$ ,  $\sigma_{22}$ ,  $\nu_2$ ]]
```

Out[322]= { $\mu_{11} \rightarrow 0.000418759$, $\sigma_{11} \rightarrow 0.0097615$ }

Out[323]= { $\mu_{12} \rightarrow 0.000535315$, $\sigma_{12} \rightarrow 0.0109675$ }

Out[324]= { $\mu_{21} \rightarrow 0.00073051$, $\sigma_{21} \rightarrow 0.00627572$, $\nu_1 \rightarrow 2.98517$ }

Out[325]= { $\mu_{22} \rightarrow 0.000971617$, $\sigma_{22} \rightarrow 0.00760454$, $\nu_2 \rightarrow 3.49141$ }

```

In[326]:= H11 = DistributionFitTest[SP500returns,
      NormalDistribution[μ11, σ11], "HypothesisTestData"] /. params11;
H12 = DistributionFitTest[SP500returns, NormalDistribution[μ12, σ12],
      "HypothesisTestData"] /. params12;
Grid[{" ", "Hypothesis Tests - ", " "}, {" ", "Normal Distribution", " "},
      {"SP500", " ", "NASDAQ"}, {H11["TestDataTable", All],
      Spacer[100], H12["TestDataTable", All]}], Frame → True]

```

Out[328]=

Hypothesis Tests - Normal Distribution					
	SP500			NASDAQ	
	Statistic	P-Value		Statistic	P-Value
Anderson-Darling	22.3908	0.	Anderson-Darling	39.2412	0.
Baringhaus-Henze	34.2867	8.61533×10^{-14}	Baringhaus-Henze	54.5397	1.11022×10^{-16}
Cramér-von Mises	4.09209	2.9723×10^{-10}	Cramér-von Mises	7.09332	0.
Jarque-Bera ALM	1426.67	0.	Jarque-Bera ALM	1426.67	0.
Kolmogorov-Smirnov	0.0857067	1.12417×10^{-11}	Kolmogorov-Smirnov	0.100973	4.85456×10^{-16}
Kuiper	0.151578	8.98788×10^{-34}	Kuiper	0.193603	6.46092×10^{-56}
Mardia Combined	1426.67	0.	Mardia Combined	1426.67	0.
Mardia Kurtosis	36.8203	8.73457×10^{-297}	Mardia Kurtosis	36.8203	8.73457×10^{-297}
Mardia Skewness	56.9347	4.50524×10^{-14}	Mardia Skewness	56.9347	4.50524×10^{-14}
Pearson χ^2	234.63	9.21972×10^{-30}	Pearson χ^2	340.658	8.17712×10^{-50}
Shapiro-Wilk	0.946064	7.59518×10^{-25}	Shapiro-Wilk	0.946064	7.59518×10^{-25}
Watson U ²	4.02906	0.	Watson U ²	7.0883	0.

```

In[329]:= H21 = DistributionFitTest[SP500returns,
      StudentTDistribution[μ21, σ21, ν1], "HypothesisTestData"] /. params21;
H22 = DistributionFitTest[NASDAQreturns, StudentTDistribution[μ22, σ22, ν2],
      "HypothesisTestData"] /. params22;
Grid[{" ", "Hypothesis Tests - ", " "}, {" ", "Student T Distribution", " "},
      {"SP500", " ", "NASDAQ"}, {H21["TestDataTable", All],
      Spacer[200], H22["TestDataTable", All]}], Frame → True]

```

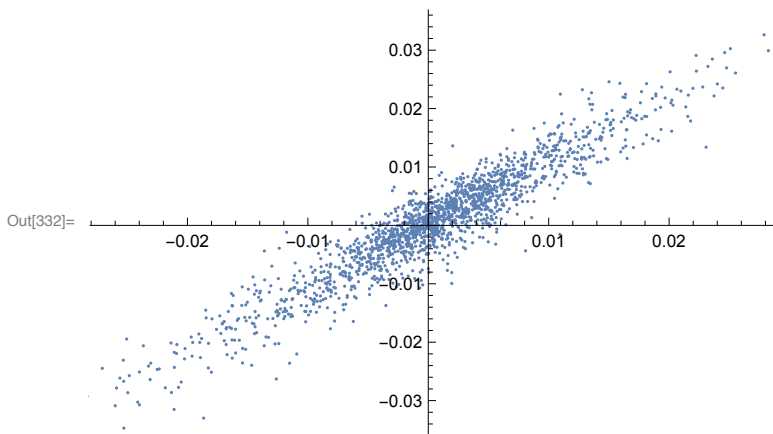
Out[331]=

Hypothesis Tests - Student T Distribution					
	SP500			NASDAQ	
	Statistic	P-Value		Statistic	P-Value
Anderson-Darling	1.29323	0.23454	Anderson-Darling	1.41866	0.197151
Cramér-von Mises	0.168883	0.336835	Cramér-von Mises	0.154309	0.376708
Kolmogorov-Smirnov	0.0192907	0.524611	Kolmogorov-Smirnov	0.0196333	0.501785
Kuiper	0.0366014	0.131103	Kuiper	0.034132	0.210365
Pearson χ^2	55.3932	0.0428145	Pearson χ^2	41.1709	0.375764
Watson U ²	0.151906	0.0996633	Watson U ²	0.129816	0.154314

Copula Calibration & Estimation

Below is a scatter plot of daily percentage log returns on the SP500 and NASDAQ indices. We use these data to calibrate normal, Clayton and Gumbel copulas. Based on the above analysis, in each case we assume the marginals are Student t distributed.

```
In[332]:= ListPlot[Transpose[{SP500returns, NASDAQreturns}]]
```



We next calibrate the parameters for each type of copula by maximum likelihood. There are several alternatives we can use, including the Method of Moments, for example, and these are listed in the documentation for the `EstimatedDistribution` function.

It is worth mentioning at this point that Alexander (2008) describes two possible procedures for estimating the copulas. The first is to use maximum likelihood to estimate all of the copula parameters, as we do here. That works well in this case because we are dealing with only two random variates. In a situation where we may be estimating a copula for several marginal distributions, MLE may fail to converge. So a second procedure can be adopted in which we first calibrate the copula correlation parameter, using the relationship with Spearman's Rho, and then use MLE to estimate the remaining copula parameter. See Alexander (2008) for details.

```
In[333]:= GaussianCopula = CopulaDistribution[{"Multinormal", {{1,  $\rho$ }, { $\rho$ , 1}}},
      {StudentTDistribution[ $\mu_{21}$ ,  $\sigma_{21}$ ,  $\nu_1$ ] /. params21,
      StudentTDistribution[ $\mu_{22}$ ,  $\sigma_{22}$ ,  $\nu_2$ ] /. params22}];
GaussianCopula = EstimatedDistribution[
  Transpose[{SP500returns, NASDAQreturns}], GaussianCopula]
```

```
Out[334]= CopulaDistribution[{Multinormal, {{1, 0.945884}, {0.945884, 1}}},
  {StudentTDistribution[0.00073051, 0.00627572, 2.98517],
  StudentTDistribution[0.000971617, 0.00760454, 3.49141]}]
```

```
In[336]:= ClaytonCopula = CopulaDistribution[
  {"Clayton", 0}, {StudentTDistribution[ $\mu_{21}$ ,  $\sigma_{21}$ ,  $\nu_1$ ] /. params21,
  StudentTDistribution[ $\mu_{22}$ ,  $\sigma_{22}$ ,  $\nu_2$ ] /. params22}];
ClaytonCopula = EstimatedDistribution[
  Transpose[{SP500returns, NASDAQreturns}], ClaytonCopula]
```

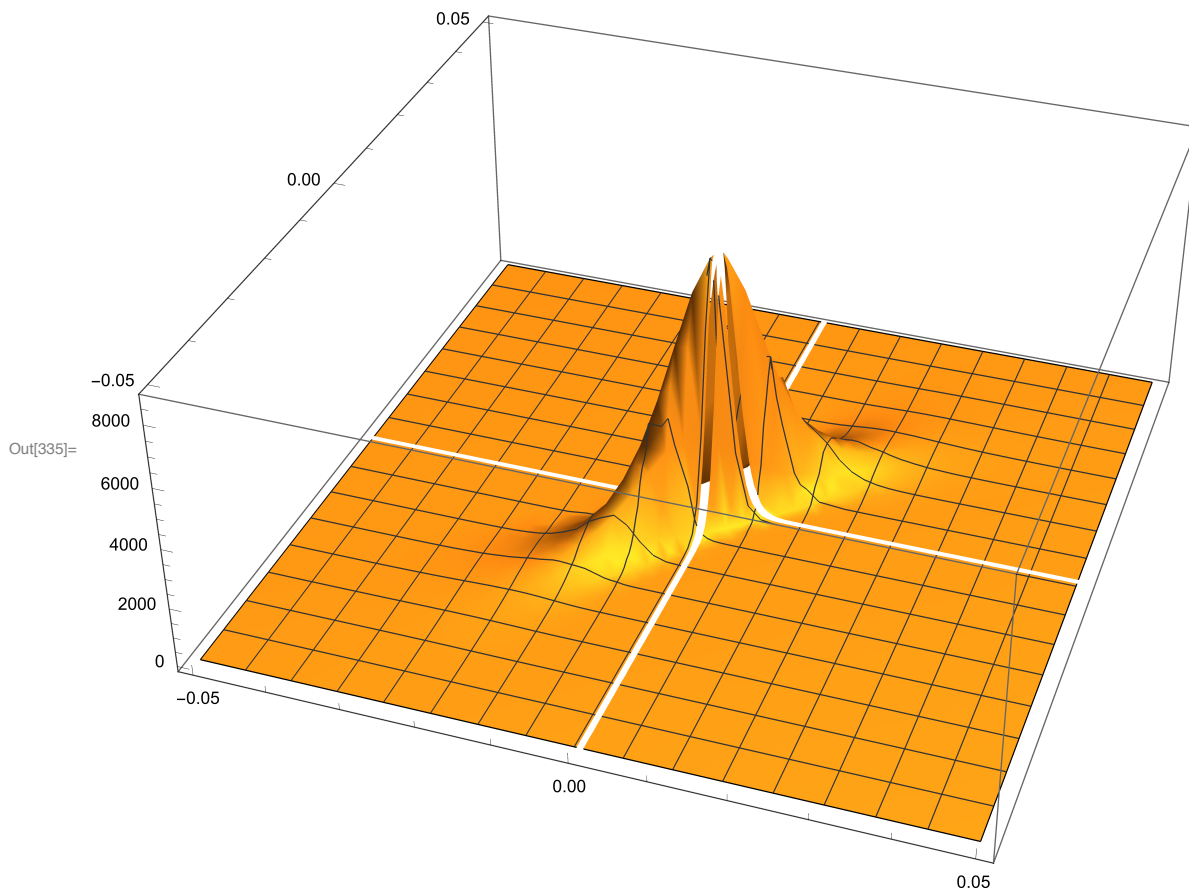
```
Out[336]= CopulaDistribution[{Clayton, 0.207943},
  {StudentTDistribution[0.00073051, 0.00627572, 2.98517],
  StudentTDistribution[0.000971617, 0.00760454, 3.49141]}]
```

```

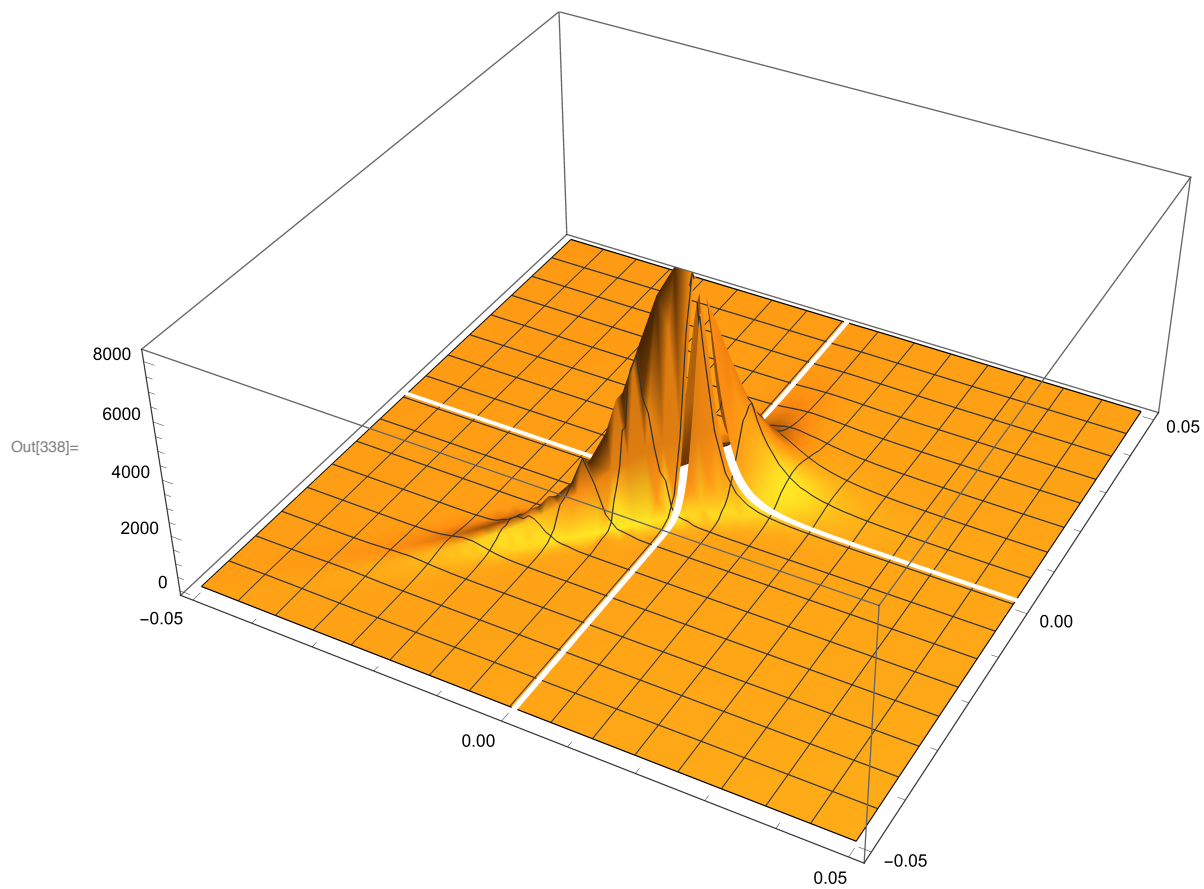
In[337]:= GumbelCopula = CopulaDistribution[{"GumbelHougaard",  $\theta$ },
      {StudentTDistribution[ $\mu_{21}$ ,  $\sigma_{21}$ ,  $\nu_1$ ] /. params21,
      StudentTDistribution[ $\mu_{22}$ ,  $\sigma_{22}$ ,  $\nu_2$ ] /. params22}];
GumbelCopula = EstimatedDistribution[
      Transpose[{SP500returns, NASDAQreturns}], GumbelCopula]
Out[337]= CopulaDistribution[{GumbelHougaard, 4.59593},
      {StudentTDistribution[0.00073051, 0.00627572, 2.98517],
      StudentTDistribution[0.000971617, 0.00760454, 3.49141]}]

In[335]:= Plot3D[PDF[GaussianCopula, {x, y}], {x, -0.05, 0.05},
      {y, -0.05, 0.05}, PlotRange -> All, ImageSize -> Large]

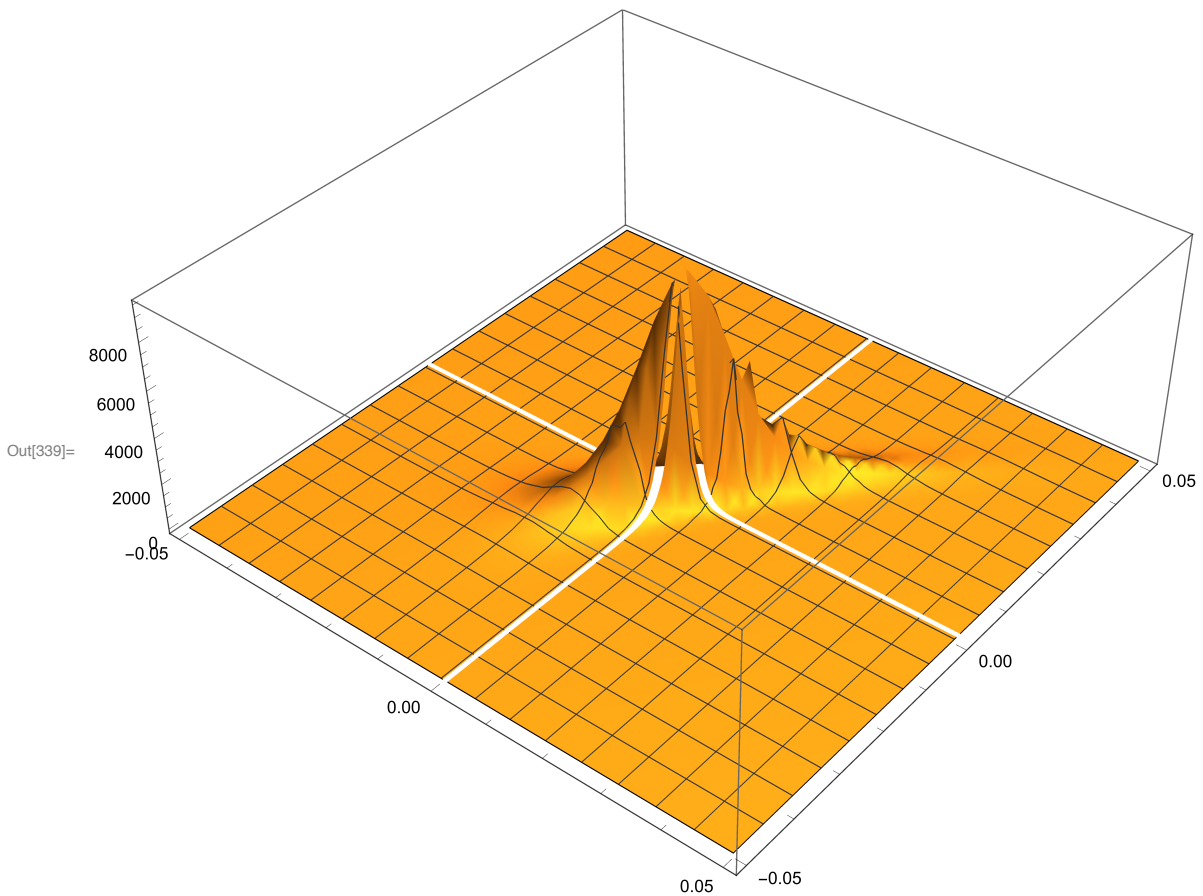
```



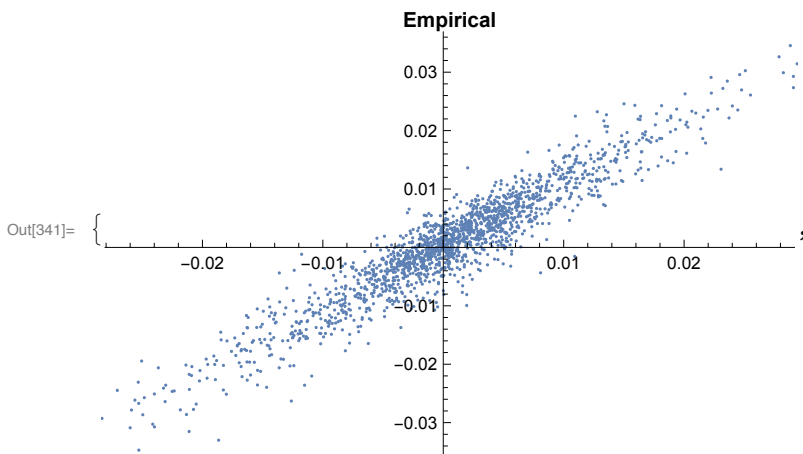

```
In[338]:= Plot3D[PDF[ClaytonCopula, {x, y}], {x, -0.05, 0.05},  
               {y, -0.05, 0.05}, PlotRange -> All, ImageSize -> Large]
```

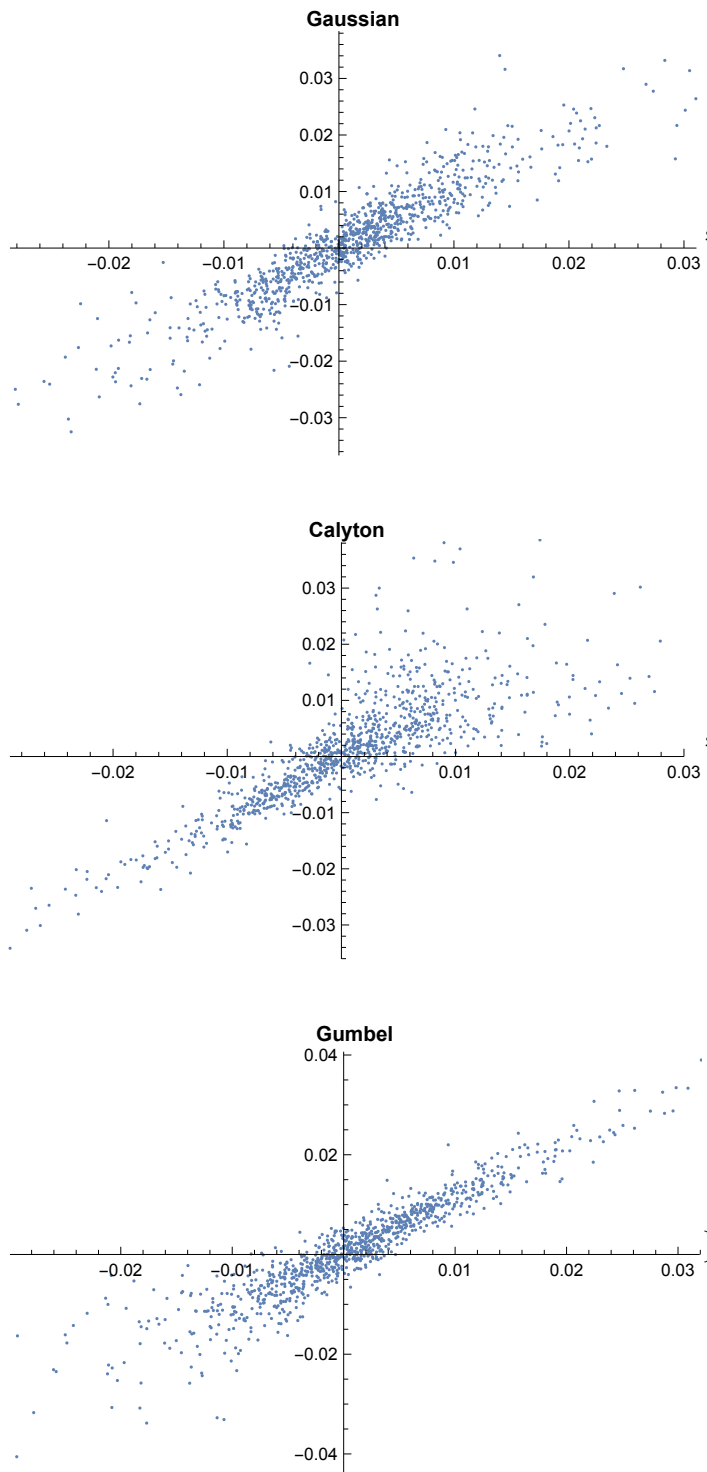


```
In[339]:= Plot3D[PDF[GumbelCopula, {x, y}], {x, -0.05, 0.05},
  {y, -0.05, 0.05}, PlotRange -> All, ImageSize -> Large]
```



```
In[341]:= {ListPlot[Transpose[{SP500returns, NASDAQreturns}],
  ImageSize -> Medium, PlotLabel -> Style["Empirical", Bold]],
  ListPlot[RandomVariate[GaussianCopula, 10^3],
  ImageSize -> Medium, PlotLabel -> Style["Gaussian", Bold]],
  ListPlot[RandomVariate[ClaytonCopula, 10^3],
  ImageSize -> Medium, PlotLabel -> Style["Calyton", Bold]],
  ListPlot[RandomVariate[GumbelCopula, 10^3],
  ImageSize -> Medium, PlotLabel -> Style["Gumbel", Bold]]}
```





Copula Selection

The distribution plots for the copulas illustrate the principal differences in their handling of tail dependency. The Gaussian copula shows symmetric dependency between the S&P500 and Nasdaq series, with higher dependency in the center of the distribution compared to the tails. The Clayton copula shows stronger left tail dependency, with greater dispersion in the right tail. The Gumbel copula shows the opposite - a higher degree of correlation in the right tail of the distribution, with lower dependency in the left tail.

One way to determine which of the models is most appropriate is to examine the likelihood. Of course, models with greater degrees of freedom (a larger number of parameters) will, other things being equal, tend to show a larger likelihood, so the standard procedure is to penalize the likelihood for the number of parameters employed in the model - the Akaike Information Criterion and Bayes Information Criterion are both examples of this approach.

```
In[342]:= TableForm[
  {LogLikelihood[GaussianCopula, Transpose[{SP500returns, NASDAQreturns}]],
    LogLikelihood[ClaytonCopula, Transpose[{SP500returns, NASDAQreturns}]],
    LogLikelihood[GumbelCopula, Transpose[{SP500returns, NASDAQreturns}]]},
  TableHeadings → {"Gaussian", "Claytom", "Gumbel"}]

Out[342]//TableForm=
Gaussian | 13 300.
Claytom  | 12 993.2
Gumbel   | 13 283.3
```

Based on the Log-Likelihood, the Gaussian copula appears to be the most suitable choice in this case. However, one should not neglect a straightforward comparison between the patterns of dependency seen in the empirical data, in comparison to the pattern of dependency depicted by the estimated copulas.

In the table below we show the estimated of the Pearson (linear) correlation, the Spearman ρ and Kendall τ for the data overall, and for the data comprising the lower and upper tails of the distribution of SP 500 Index returns.

It is noteworthy that the dependency appears to lessen in both left and right tails, with only slightly higher dependency in the right tail compared to the left. This would appear to suggest that a copula with lower, approximately symmetric dependency in the tails would be an appropriate choice, supporting the earlier finding in favor of the Gaussian copula (a Frank copula might be another appropriate choice).

```
In[343]:= LP05 = Flatten[Position[SP500returns, _? (# < Quantile[SP500returns, 0.05] &)]];
LP95 = Flatten[Position[SP500returns, _? (# > Quantile[SP500returns, 0.95] &)]];
TableForm[
  {Through[{Correlation, SpearmanRho, KendallTau}[SP500returns, NASDAQreturns]],
    Through[{Correlation, SpearmanRho, KendallTau}[
      SP500returns[[LP05]], NASDAQreturns[[LP05]]]},
    Through[{Correlation, SpearmanRho, KendallTau}[
      SP500returns[[LP95]], NASDAQreturns[[LP95]]]}],
  TableHeadings → {"Overall", "Lower 5%-tile", "Upper 95%-tile"},
  {"Pearson", "Spearman", "Kendall"}]

Out[345]//TableForm=
```

	Pearson	Spearman	Kendall
Overall	0.956041	0.927015	0.779359
Lower 5%-tile	0.924604	0.852245	0.674953
Upper 95%-tile	0.915731	0.774823	0.606522

We next compare the empirical and theoretical measures of dependency, Kendall's τ , for each of the fitted copulas:

```
In[346]:= TableForm[
  {KendallTau[SP500returns, NASDAQreturns], ArcSin[0.9458844026736891`]*2/π,
    0.19663268588974306`/(2+0.19663268588974306`), 1-(1/4.500333562276729`)},
  TableHeadings → {"Empirical", "Gaussian", "Clayton", "Gumbel"}, None}]
```

```
Out[346]//TableForm=
Empirical | 0.779359
Gaussian  | 0.789605
Clayton   | 0.0895155
Gumbel    | 0.777794
```

The Clayton copula appears mis-specified, which perhaps argues in favor of Alexander's two-step calibration procedure rather than relying a single-step MLE, but the values of Kendall's τ for the other copulas appear close to the empirical value.

Value at Risk

Let's suppose we want to estimate a 1% VaR for an investment portfolio comprising a 70% capital allocation to the S&P 500 Index and 30% allocation to the Nasdaq Index. Using Monte Carlo simulation we derive the 1% VaR levels for each of the different dependence assumptions, as follows:

```
In[347]:= TableForm[
  {Quantile[Dot[Transpose[{SP500returns, NASDAQreturns}], {0.7, 0.3}], 0.01],
    Quantile[Dot[RandomVariate[GaussianCopula, 10^3], {0.7, 0.3}], 0.01],
    Quantile[Dot[RandomVariate[ClaytonCopula, 10^3], {0.7, 0.3}], 0.01],
    Quantile[Dot[RandomVariate[GumbelCopula, 10^3], {0.7, 0.3}], 0.01]},
  TableHeadings → {"Empirical", "Gaussian", "Clayton", "Gumbel"}, None}]
```

```
Out[347]//TableForm=
Empirical | -0.0291767
Gaussian  | -0.0260779
Clayton   | -0.0250525
Gumbel    | -0.023972
```

The VaR estimates for the investment portfolio based on the copula models slightly under-estimate the 1% VaR estimated from the empirical data.

Conclusion

Depending on the characteristics of the portfolio the standard normal VaR model may misrepresent its riskiness. Typically this takes the form of underestimating the risk arising from extreme moves in the underlying assets, due to the dependency in the tails of the distributions, which may differ markedly from that assumed in a Gaussian framework.

Copulas are a flexible and effective tool with which to address these shortcomings, providing the risk manager with more reliable assessment of the risks across a wider span of possible outcomes, including those not yet manifest in the empirical data.

References

- Alexander, C., 2008. Market Risk Analysis. Chichester: John Wiley & Sons, Ltd.
- Duffie, D., Pan, J., 1997. An Overview of Value-at-Risk. In: Journal of Derivatives, Vol. 4, No. 3, 7-49.

- Embrechts, P, Dias, A., 2010. Modeling exchange rate dependence at different time horizons. In: Journal of International Money and Finance. Vol. 29, Issue 8, pp. 1687-1705
- Embrechts, P., Lindskog, F., McNeil, A. J., 2001. Modelling Dependence with Copulas and Applications to Risk Management. Zürich, <http://www.math.ethz.ch/finance>

```
In[349]:= DensityPlot[PDF[GaussianCopula, {x, y}], {x, -0.01, 0.01}, {y, -0.01, 0.01},  
PlotRange -> All, ImageSize -> Medium, ColorFunction -> ColorData["SolarColors"]]
```

