Multivariate Extremes and Dependence in Elliptical Distributions

Filip Lindskog, RiskLab, ETH Zürich
joint work with Henrik Hult, KTH Stockholm

E-mail: lindskog@math.ethz.ch
RiskLab: http://www.risklab.ch

Motivation

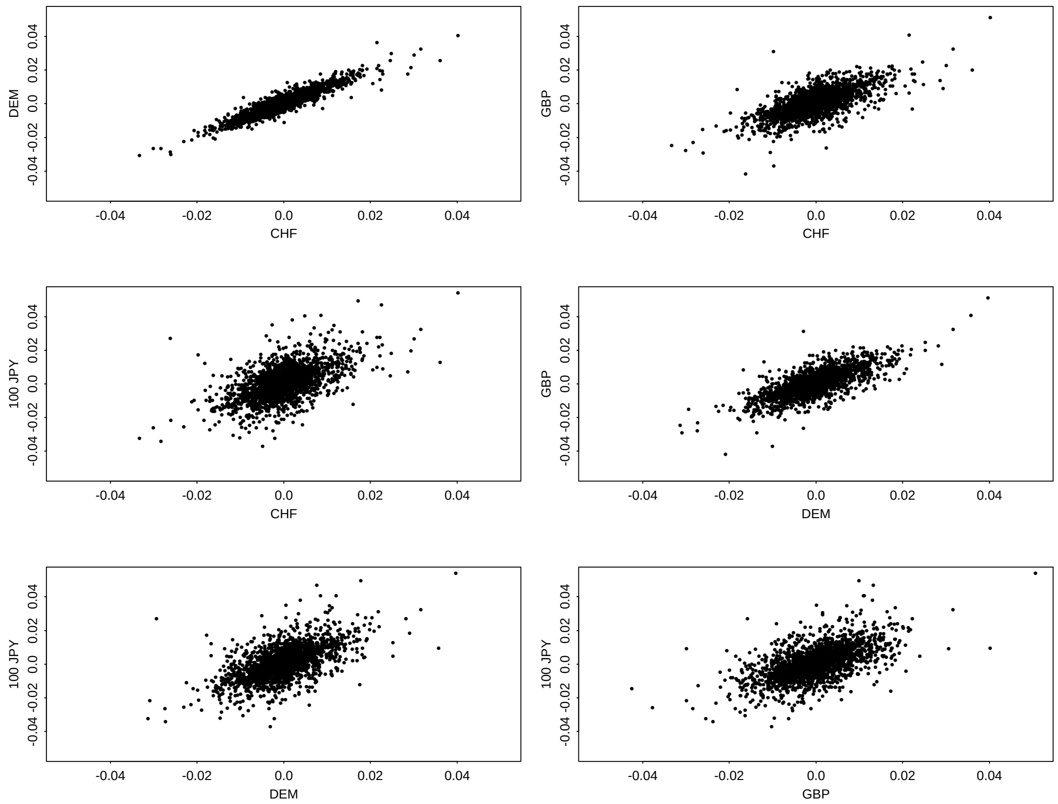

Daily FX log returns (quoted against the US Dollar)

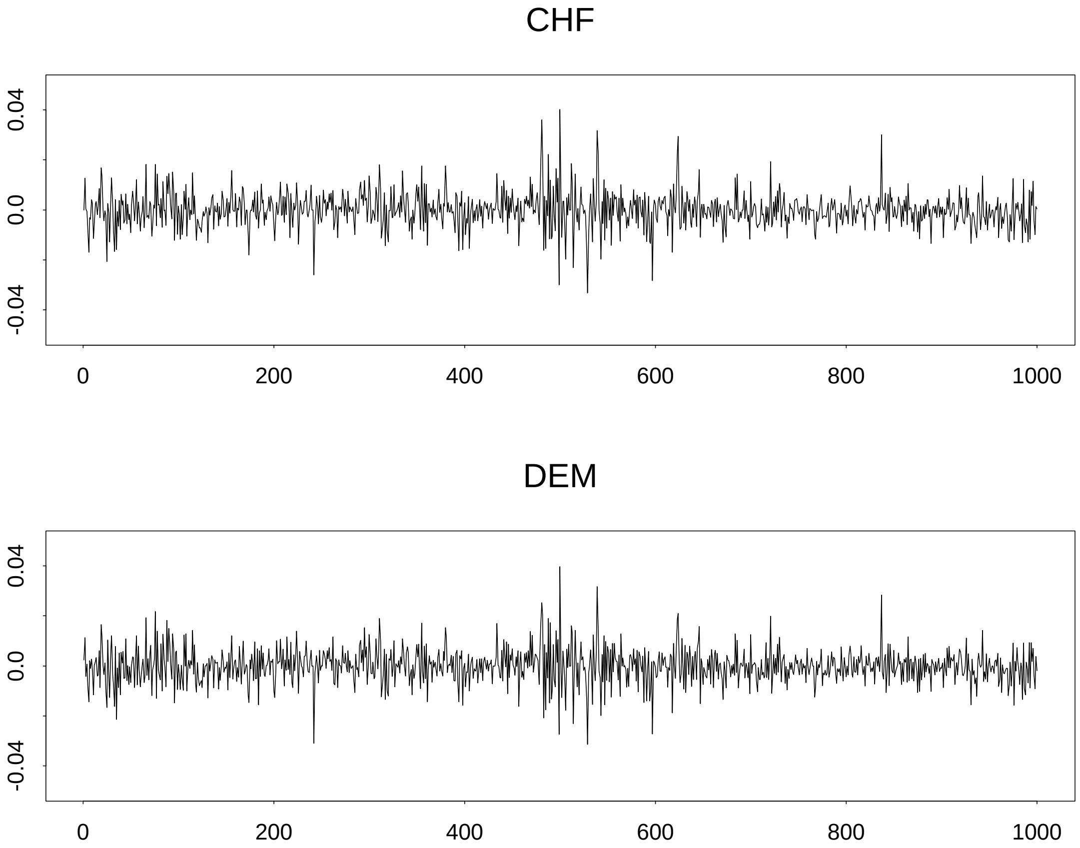

Daily FX log returns (quoted against the US Dollar)

Definition and Representation of Elliptical Distributions

Definition A d-dimensional random vector $\mathbf{X}$ has an *elliptical distribution* with parameters $\mu \in \mathbb{R}^d$, $\boldsymbol{\Sigma} \in \mathbb{R}^{d\times d}$ symmetric non-negative definite, and $\phi : [0, \infty) \to \mathbb{R}$ if the characteristic function $\varphi_{\mathbf{X}-\mu}$ of $\mathbf{X} - \mu$ is of the form

$$\varphi_{\mathbf{X}-\mu}(\mathbf{t}) = \phi(\mathbf{t}^\top \boldsymbol{\Sigma}\mathbf{t}).$$

We write $\mathbf{X} \sim E_d(\mu, \boldsymbol{\Sigma}, \phi)$.

Theorem $\mathbf{X} \sim E_d(\mu, \Sigma, \phi)$ with $\text{rank}(\Sigma) = k$ iff there exist a non-negative random variable R independent of $\mathbf{U}$, a random vector uniformly distributed on the unit hypersphere $\mathbb{S}_2^{k-1} = \{\mathbf{z} \in \mathbb{R}^k \,|\, \mathbf{z}^\top\mathbf{z} = 1\}$, and a $d \times k$ matrix A with $AA^\top = \Sigma$, such that

$$\mathbf{X} \overset{\text{d}}{=} \mu + RA\mathbf{U}.$$

Illustration of the Representation Theorem

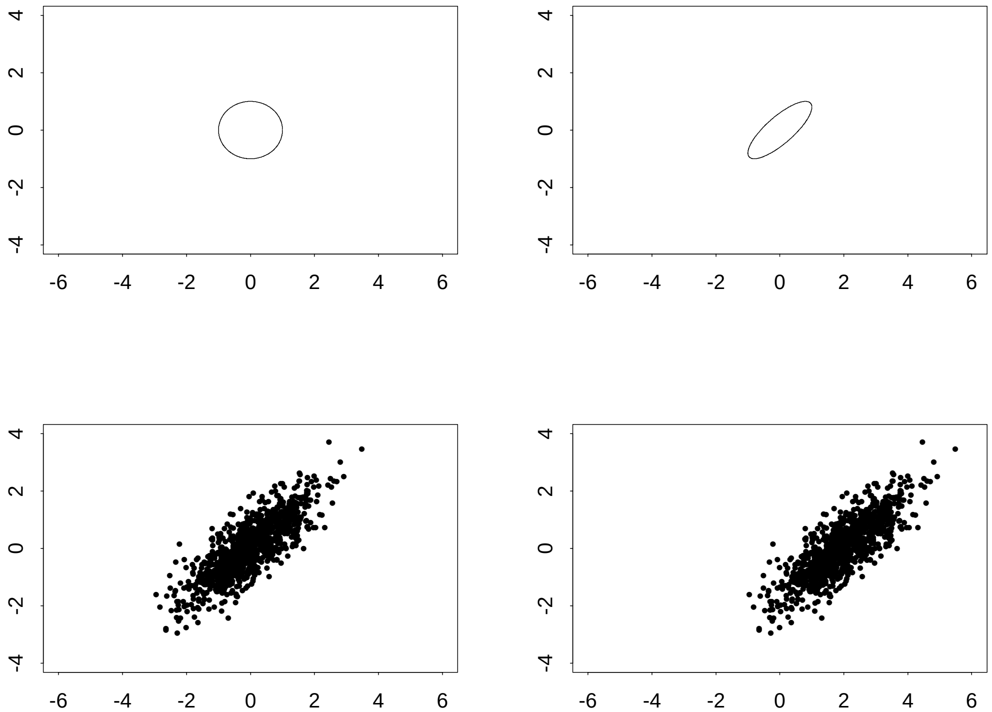

$$\mathbf{U} \longrightarrow A\mathbf{U} \longrightarrow RA\mathbf{U} \longrightarrow \mu + RA\mathbf{U}$$

Let $\mathbf{X} \stackrel{\text{d}}{=} \mu + RA\mathbf{U} \sim E_d(\mu, \Sigma, \phi)$.

Examples (with $\text{rank}(\Sigma) = d$)
$\mathbf{X}$ has a multivariate normal distribution if and only if R^2 has a χ^2_d-distribution.
$\mathbf{X}$ has a multivariate t-distribution with ν degrees of freedom if and only if R^2/d has an $F(d, \nu)$-distribution.

Theorem Let B be a $q \times d$ matrix and let $b \in \mathbb{R}^q$. Then

$$b + B\mathbf{X} \sim E_q(b + B\mu, B\Sigma B^\mathsf{T}, \phi).$$

Definition For $i, j \in \{1, \ldots, d\}$, if $\Sigma_{ii} > 0$ and $\Sigma_{jj} > 0$, then we call $\varrho_{ij} \triangleq \Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$ the *linear correlation coefficient* of $(X_i, X_j)^\mathsf{T}$.

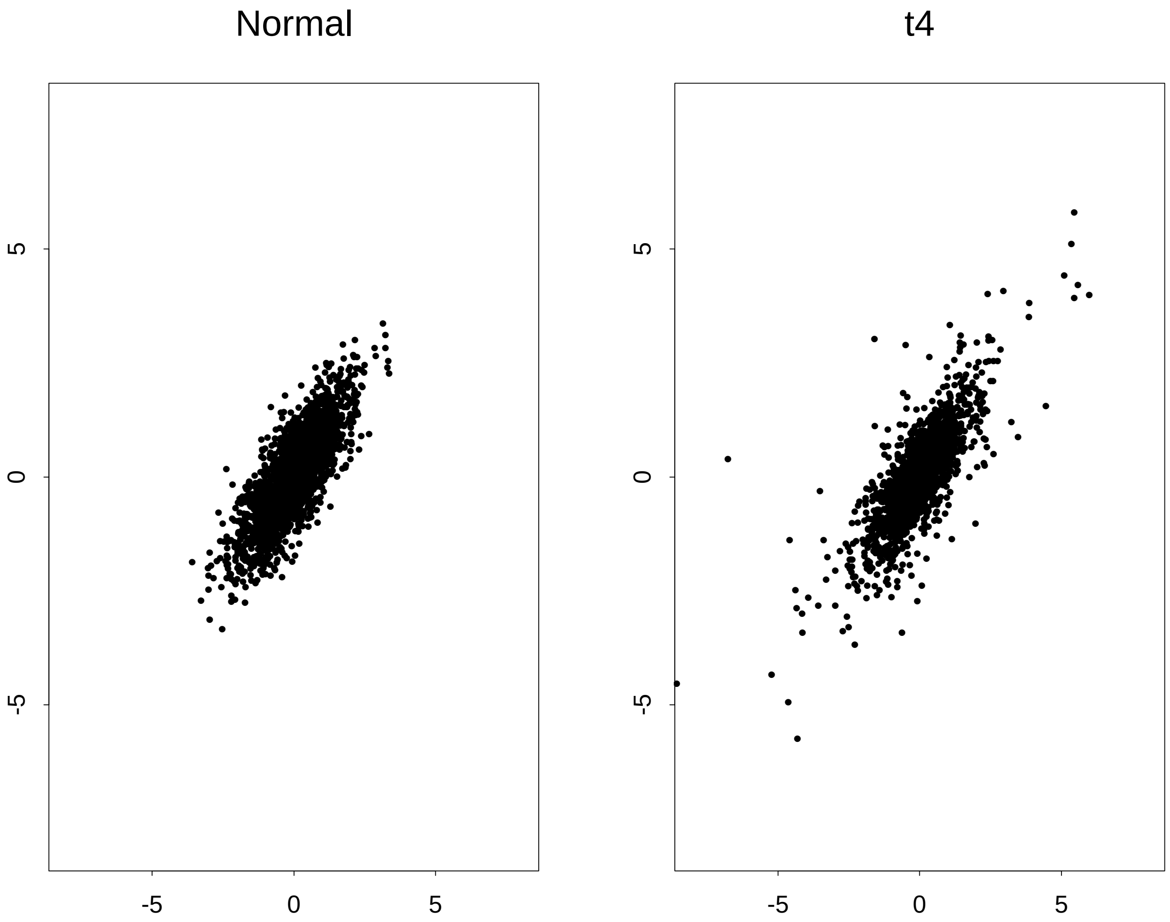

Two bivariate elliptical distributions with unit variances and $\varrho = 0.8$. Both plots are generated from the same sample.

A Class of Multivariate Time Series

For $t \in \mathbb{Z}$ let $\mathbf{X}_t = \sigma_t \mathbf{Z}_t$, where $\mathbf{Z}_t \triangleq R_t A \mathbf{U}_t \sim E_d(\mathbf{0}, \boldsymbol{\Sigma}, \phi_t)$ and σ_t is a non-negative random variable. Let the $\mathbf{Z}_t$s be independent, and for all t let $\sigma_t, \mathbf{Z}_t, \mathbf{Z}_{t+1}, \ldots$ be independent. *The σ_ts are allowed to be dependent.*

Then for every t, $\mathbf{X}_t$ is elliptically distributed with dispersion matrix Σ, i.e. with a constant linear correlation matrix (ϱ_{ij}).

Suppose further that $\{\sigma_t\}$ and $\{\mathbf{U}_t\}$ are independent.

Example: GARCH(1,1)-type

$$\sigma_t^2 = \alpha_0 + \alpha_1 |\mathbf{X}_{t-1}|_{\Sigma}^2 + \beta_1 \sigma_{t-1}^2 = \alpha_0 + (\alpha_1 R_{t-1}^2 + \beta_1)\sigma_{t-1}^2,$$

where $|\mathbf{x}|_{\Sigma} = (\mathbf{x}^{\mathsf{T}} \boldsymbol{\Sigma}^{-1} \mathbf{x})^{1/2}$.

Addition of Elliptically Distributed Random Vectors

What can we say about sums of the form

$$\mathbf{S}(t,k) \triangleq \mathbf{X}_{t+1} + \cdots + \mathbf{X}_{t+k}$$

and

$$\mathbf{S}(t,k,\mathbf{x}) \triangleq \mathbf{X}_{t+1} + \cdots + \mathbf{X}_{t+k} \mid \mathbf{X}_t = \mathbf{x} \quad ?$$

Theorem Let W and $\widetilde{W}$ be non-negative random variables and let

$$\mathbf{X} \triangleq \mu + W\mathbf{Z} \sim E_d(\mu, \Sigma, \phi)$$

and

$$\widetilde{\mathbf{X}} \triangleq \tilde{\mu} + \widetilde{W}\tilde{\mathbf{Z}} \sim E_d(\tilde{\mu}, \Sigma, \tilde{\phi}),$$

where $W, \mathbf{Z}, \tilde{\mathbf{Z}}$ are independent and $\widetilde{W}, \mathbf{Z}, \tilde{\mathbf{Z}}$ are independent. Then

$$\mathbf{X} + \widetilde{\mathbf{X}} \sim E_d(\mu + \tilde{\mu}, \Sigma, \phi^*).$$

Moreover, if W and $\widetilde{W}$ are independent, then $\phi^*(u) = \phi(u)\tilde{\phi}(u)$.

Hence, $\mathbf{S}(t,k)$ and $\mathbf{S}(t,k,\mathbf{x})$ are elliptical with dispersion matrix Σ.

Tail Dependence

Definition Let $(X_1, X_2)^\top$ be a random vector with marginal distribution functions F_1 and F_2. The *coefficient of upper tail dependence* of $(X_1, X_2)^\top$ is defined as

$$\lambda_U(X_1, X_2) \triangleq \lim_{u \nearrow 1} \mathbb{P}\{X_2 > F_2^{-1}(u) \mid X_1 > F_1^{-1}(u)\},$$

provided that the limit $\lambda_U \in [0, 1]$ exists. The *coefficient of lower tail dependence* is defined as

$$\lambda_L(X_1, X_2) \triangleq \lim_{u \searrow 0} \mathbb{P}\{X_2 \leq F_2^{-1}(u) \mid X_1 \leq F_1^{-1}(u)\},$$

provided that the limit $\lambda_L \in [0, 1]$ exists. If $\lambda_U > 0$ ($\lambda_L > 0$), then we say that $(X_1, X_2)^\top$ has upper (lower) tail dependence.

Illustration of Upper Tail Dependence

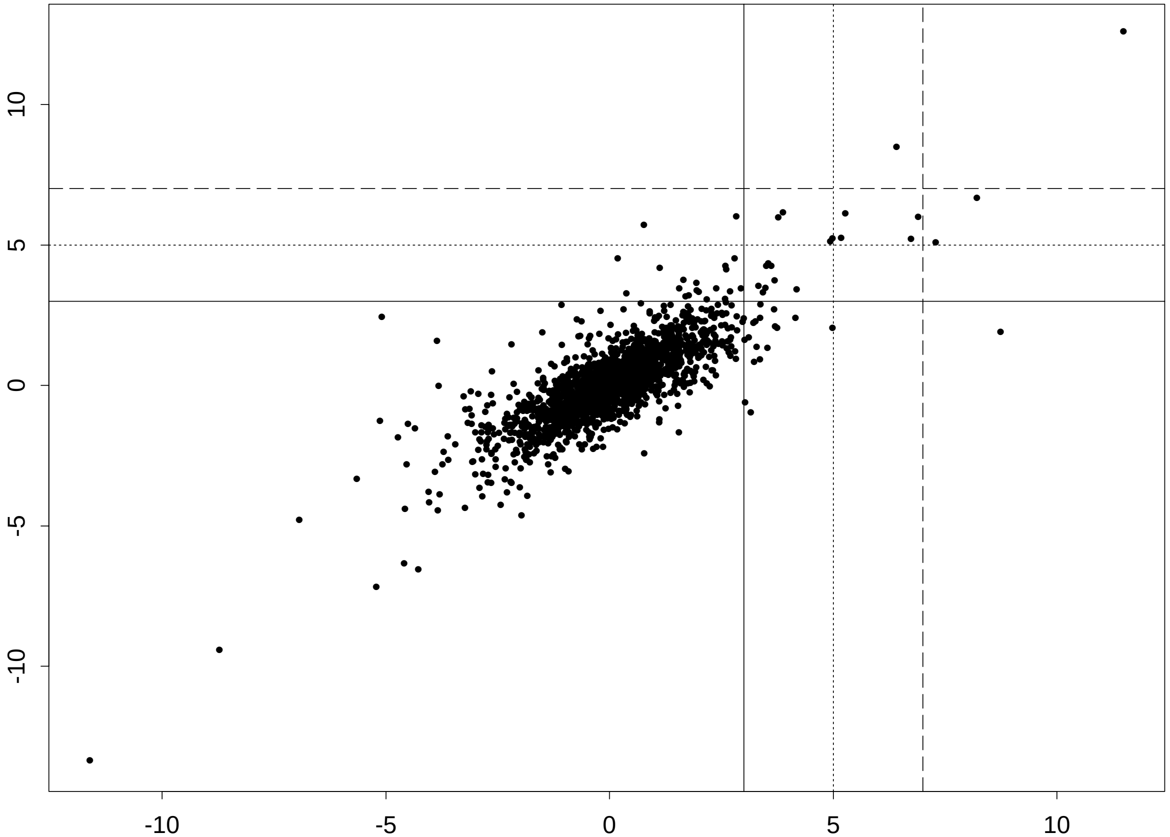

Regular Variation - Spectral Measure

Definition The random variable R is said to be *regularly varying at* ∞ *with index* $\alpha > 0$ if for all $x > 0$,

$$\lim_{t \to \infty} \frac{\mathbb{P}\{R > tx\}}{\mathbb{P}\{R > t\}} = x^{-\alpha}.$$

We denote by $\mathbb{S}^{d-1}$ the unit hypersphere in $\mathbb{R}^d$ with respect to a general norm $|\cdot|$, and by $\mathcal{B}(\mathbb{S}^{d-1})$ the Borel σ-algebra on $\mathbb{S}^{d-1}$.

Definition The d-dimensional random vector $\mathbf{X}$ is said to be *regularly varying with index* $\alpha > 0$ if there exists a random vector Θ with values in $\mathbb{S}^{d-1}$ a.s. such that for all $x > 0$ and $S \in \mathcal{B}(\mathbb{S}^{d-1})$ with $\mathbb{P}\{\Theta \in \partial S\} = 0$

$$\lim_{t \to \infty} \frac{\mathbb{P}\{|\mathbf{X}| > tx, \mathbf{X}/|\mathbf{X}| \in S\}}{\mathbb{P}\{|\mathbf{X}| > t\}} = x^{-\alpha} \mathbb{P}\{\Theta \in S\}.$$

The distribution of Θ is referred to as the *spectral measure* of $\mathbf{X}$ and α is referred to as the *tail index* of $\mathbf{X}$.

Does the choice of norm matter?

Theorem Let $|\cdot|_A$ and $|\cdot|_B$ be two norms on $\mathbb{R}^d$ and let $\mathbf{X}$ be a d-dimensional random vector. Then $\mathbf{X}$ is regularly varying with index $\alpha > 0$ with respect to the norm $|\cdot|_A$ if and only if $\mathbf{X}$ is regularly varying with index $\alpha > 0$ with respect to the norm $|\cdot|_B$.

Hence, whether a random vector $\mathbf{X}$ is regularly varying or not does not depend on the choice of norm in the definition of (multivariate) regular variation.

- Note that for different norms, the corresponding spectral measures do not in general coincide.
- Note also that, in general, the spectral measure depends on the tail index α.

For elliptical distributed random vectors, tail dependence and regular variation are closely related.

Theorem Let $\mathbf{X} \stackrel{\mathrm{d}}{=} \mu + RA\mathbf{U} \sim E_d(\mu, \Sigma, \phi)$, with $\Sigma_{ii} > 0$ for $i = 1, \ldots, d$ and $|\varrho_{ij}| < 1$ for all $i \neq j$. Then the following statements are equivalent.

(1) R is regularly varying with index $\alpha > 0$.

(2) $\mathbf{X}$ is regularly varying with index $\alpha > 0$.

(3) For all $i \neq j$, $(X_i, X_j)^{\mathsf{T}}$ has tail dependence.

Moreover, if R is regularly varying with index $\alpha > 0$, then for all $i \neq j$,

$$\lambda_U(X_i, X_j) = \lambda_L(X_i, X_j) = \frac{\int_{(\pi/2 - \arcsin \varrho_{ij})/2}^{\pi/2} \cos^\alpha t\,dt}{\int_0^{\pi/2} \cos^\alpha t\,dt}.$$

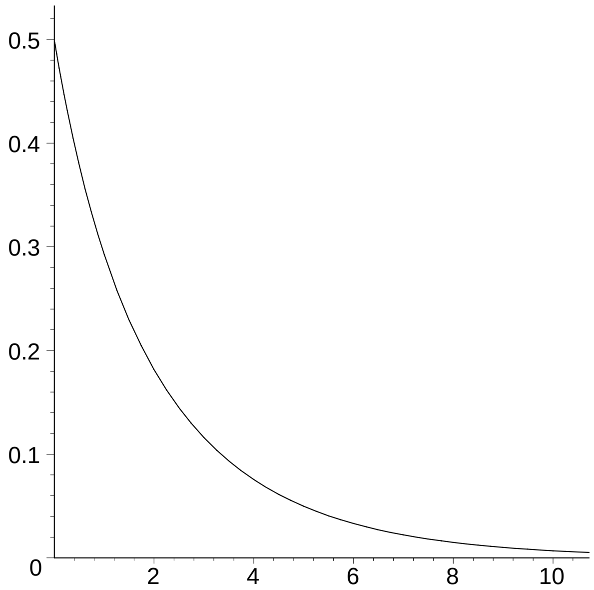

The coefficient of tail dependence for uncorrelated regularly varying bivariate elliptical distributions as a function of the tail index α.

How do we interpret the spectral measure?

For every choice of the norm the spectral measure is a measure of dependence between extreme values. However, the choice of norm becomes essential when interpreting the spectral measure.

The choice of norm must be related to the question we are trying to answer.

A natural question
What is the dependence between the components of a random vector given that at least one of its components is extreme?

Let $\mathbf{X} \sim E_d(\mathbf{0}, \mathbf{\Sigma}, \phi)$ be regularly varying with index $\alpha > 0$.

The Euclidean 2-norm $|\mathbf{X}|_2 \triangleq (X_1^2 + \cdots + X_d^2)^{1/2}$ gives the spectral measure

$$\begin{aligned}\mathbb{P}\{\Theta_2 \in \cdot\} &= \lim_{t\to\infty} \mathbb{P}\{\mathbf{X}/|\mathbf{X}|_2 \in \cdot \mid |\mathbf{X}|_2 > t\} \\ &= \lim_{t\to\infty} \mathbb{P}\{\mathbf{X}/|\mathbf{X}|_2 \in \cdot \mid (X_1^2 + \cdots + X_d^2)^{1/2} > t\}.\end{aligned}$$

The max-norm $|\mathbf{X}|_\infty \triangleq \max\{|X_1|, \ldots, |X_d|\}$ gives the spectral measure

$$\begin{aligned}\mathbb{P}\{\Theta_\infty \in \cdot\} &= \lim_{t\to\infty} \mathbb{P}\{\mathbf{X}/|\mathbf{X}|_\infty \in \cdot \mid |\mathbf{X}|_\infty > t\} \\ &= \lim_{t\to\infty} \mathbb{P}\{\mathbf{X}/|\mathbf{X}|_\infty \in \cdot \mid |X_1| > t \cup \cdots \cup |X_d| > t\},\end{aligned}$$

from which it is seen that the max-norm corresponds to the question posed.

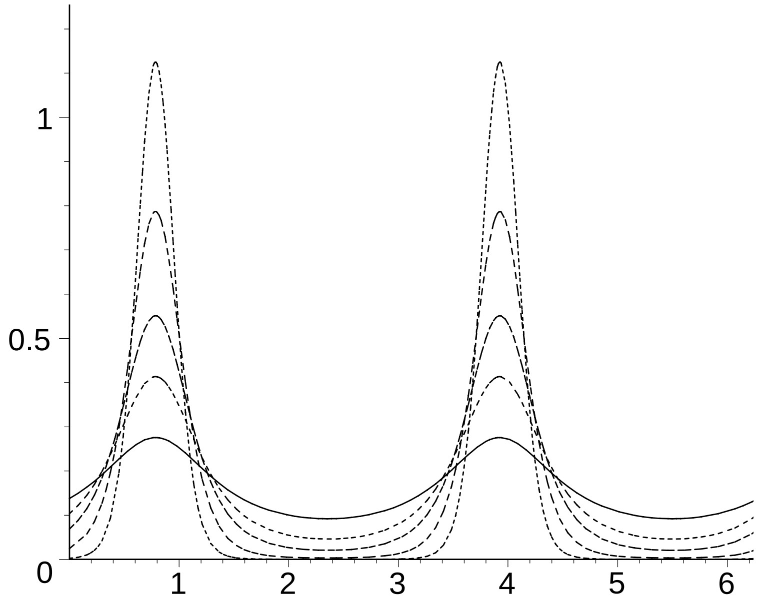

Densities of the spectral measure of $\mathbf{X} \sim E_2(\mu, \Sigma, \phi_\alpha)$ with respect to the Euclidean 2-norm.

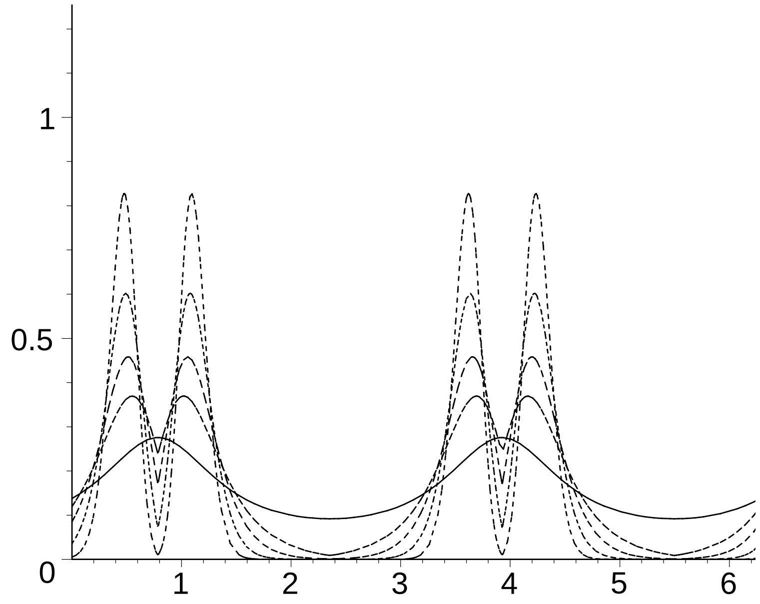

Densities of the spectral measure of $\mathbf{X} \sim E_2(\mu, \Sigma, \phi_\alpha)$ with respect to the max-norm.

Tail indices $\alpha = 0, 2, 4, 8, 16$. $\Sigma_{11} = \Sigma_{22} = 1$ and $\Sigma_{12} = \Sigma_{21} = 0.5$. Larger tail indices corresponds to higher peaks.

Estimation of Correlation Coefficients

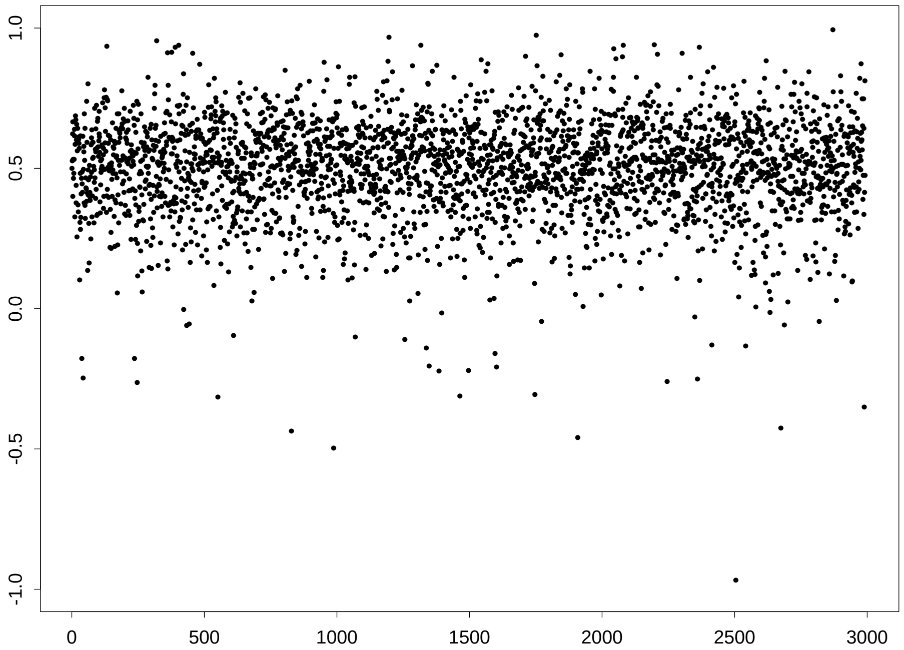

3000 independent linear correlation estimates with the standard estimator $\widehat{\varrho}_{\mathrm{SE}}$ from samples of size 90 from a bivariate t_3-distribution with $\varrho = 0.5$.

Definition Kendall's tau for the random vector $(X,Y)^T$ is defined as

$$\tau(X,Y) \triangleq \mathbb{P}\{(X-\widetilde{X})(Y-\widetilde{Y})>0\} - \mathbb{P}\{(X-\widetilde{X})(Y-\widetilde{Y})<0\},$$

where $(\widetilde{X},\widetilde{Y})^T$ is an independent copy of $(X,Y)^T$.

Theorem Let $\mathbf{X} \sim E_d(\mu,\Sigma,\phi)$, where for $i,j \in \{1,\ldots,d\}$, X_i and X_j are continuous. Then

$$\tau(X_i,X_j) = \frac{2}{\pi}\arcsin(\varrho_{ij}).$$

This relation provides the linear correlation estimator

$$\widehat{\varrho}_{\mathrm{KT}} = \sin\left(\frac{\pi}{2}\widehat{\tau}\right),$$

where $\widehat{\tau} = (c-d)/(c+d)$ with c and d denoting the number of concordant and discordant pairs respectively.

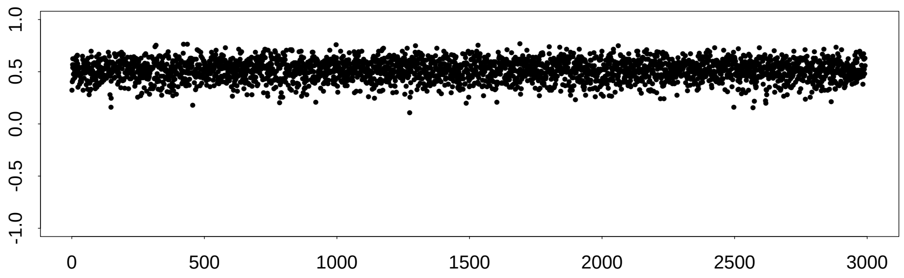

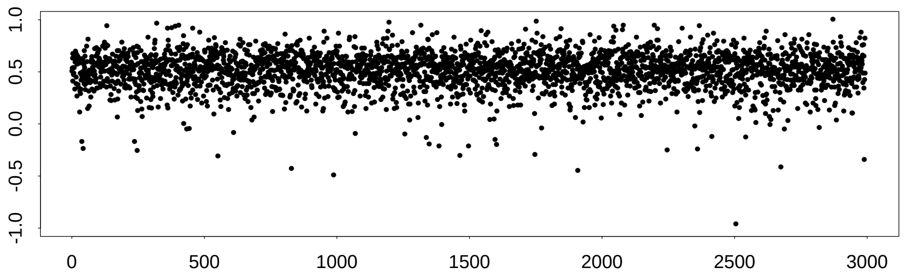

3000 independent linear correlation estimates with the estimator $\widehat{\varrho}_{\mathrm{KT}}$ (above) and with the estimator $\widehat{\varrho}_{\mathrm{SE}}$ (below) from (the same) samples of size 90 from a bivariate t_3-distribution with $\varrho = 0.5$.