Models for Dependent Risks Using Copulas

First Banco de Santander Financial Engineering School

Alexander J. McNeil, University of York

9th March 2021

Overview

1. Fundamentals of Copulas
   - Definition and properties
   - Examples and simulation
   - Estimation with maximum likelihood
   - Estimation with rank correlations

2. Attainability of Kendall Rank Correlations
   - Motivation of attainability problem
   - Extremal copulas and correlation matrices
   - Attainable Kendall correlation matrices
   - Testing for attainability

3. Copulas for Time Series
   - Introduction
   - ARMA copulas
   - D-vine copulas
   - V-transforms

Overview

Overview

Definition and Sklar's Theorem

A copula is a multivariate distribution function with standard uniform marginal distributions.

Theorem

- *Let F be a joint distribution function with margins $F_1, \ldots, F_d$. There exists a copula C such that for all $x_1, \ldots, x_d$ in $[-\infty, \infty]$*

$$F(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)).$$

  If the margins are continuous then C is unique.
- *And conversely, if C is a copula and $F_1, \ldots, F_d$ are univariate distribution functions, then F defined above is a multivariate df with margins $F_1, \ldots, F_d$.*

Probability and Quantile Transforms

Recall the following facts concerning stochastic simulation:

- **Probability Transform.**
  Let X be a random variable with continuous distribution function F. Then $F(X) \sim U(0,1)$ (standard uniform).

$$\mathbb{P}(F(X) \leq u) = u.$$

- **Quantile Transform.**
  Let U be uniform and F the distribution function of any rv X. Then

$$\mathbb{P}(F^{\leftarrow}(U) \leq x) = F(x).$$

These facts are essential in dealing with copulas.

Note, in general, we need generalized inverse $F^{\leftarrow}(u) = \inf\{x : F(x) \geq u\}$.

Basic Properties

- Extracting the copula. Given F the copula C is given by

$$C(u_1, \dots, u_d) = F\left(F_1^{\leftarrow}(u_1), \dots, F_d^{\leftarrow}(u_d)\right),$$

  where $F_1, \dots, F_d$ are the marginal distributions of F.
- Invariance. The copula C of a random vector is invariant under strictly increasing transformations of the variables.
- Independence copula. The copula representing independence is

$$C(u_1, \dots, u_d) = \prod_{i=1}^{d} u_i.$$

- Comonotonicity copula. The copula representing perfect positive dependence (all variables increasing functions of a single variable) is $M(\boldsymbol{u}) = \min\{u_1, \dots, u_d\}$.
- Countermonotonicity copula. For two variables, the copula representing perect negative dependence is $W(u_1, u_2) = \max\{u_1 + u_2 - 1, 0\}$.

Overview

Parametric Copulas

There are basically two possibilities:

1. Copulas implicit in well-known parametric distributions. Sklar's Theorem states that we can always find a copula in a parametric distribution function. Denoting the df by F and assuming the margins $F_1, \ldots, F_d$ are continuous, the implied copula is

$$C(u_1, \ldots, u_d) = F\left(F_1^{\leftarrow}(u_1), \ldots, F_d^{\leftarrow}(u_d)\right) .$$

   Such a copula may not have a simple closed form.

2. Closed form parametric copula families generated by some explicit construction that is known to yield copulas. The best example is the well-known Archimedean copula family. These generally have limited numbers of parameters and limited flexibility; the standard Archimedean copulas are dfs of exchangeable random vectors (distribution invariant under permutations).

Examples of Implicit Copulas

- Gaussian Copula

$$C_P^{\text{Ga}}(\mathbf{u}) = \boldsymbol{\Phi}_P\left(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_d)\right)$$

  - Φ denotes the standard univariate normal df.
  - $\boldsymbol{\Phi}_P$ denotes the joint df of $\mathbf{X} \sim N_d(\mathbf{0}, P)$and P is a correlation matrix.
  - Write C_ρ^{Ga} when $d = 2$.
  - $P = I_d$ gives independence.
  - $P = J_d$ (matrix of 1's) gives comonotonicity.
- t Copula

$$C_{\nu,P}^{\text{t}}(\mathbf{u}) = \mathbf{t}_{\nu,P}\left(t_\nu^{-1}(u_1), \ldots, t_\nu^{-1}(u_d)\right)$$

  - t_ν is the df of a standard univariate t distribution.
  - $\mathbf{t}_{\nu,P}$ is the joint df of a multivariate Student t distribution with correlation matrix P and degree of freedom ν.

Archimedean Copulas

- Gumbel Copula

$$C_\theta^{\mathrm{Gu}}(u_1, \dots, u_d) = \exp\left(-\left((-\log u_1)^\theta + \dots + (-\log u_d)^\theta\right)^{1/\theta}\right).$$

  $\theta \geq 1$: $\theta = 1$ gives independence; $\theta \to \infty$ gives comonotonicity.

- Clayton Copula

$$C_\theta^{\mathrm{Cl}}(u_1, \dots, u_d) = \left(u_1^{-\theta} + \dots + u_d^{-\theta} - d + 1\right)^{-1/\theta}.$$

  $\theta > 0$: $\theta \to 0$ gives independence ; $\theta \to \infty$ gives comonotonicity.

- Other common examples include the Frank and Joe copulas.

Simulating Copulas

Simulating Gaussian copula C_P^{Ga}

1. Simulate $\mathbf{X} \sim N_d(\mathbf{0}, P)$
2. Set $\mathbf{U} = (\Phi(X_1), \ldots, \Phi(X_d))'$ (probability transformation)

The other three copulas are also straightforward to simulate but we do not go into details.

How do you simulate a normal vector $\mathbf{X} \sim N_d(\mathbf{0}, P)$?

1. Simulate d independent standard normals $Z_1, \ldots, Z_d$.
2. Compute Cholesky decomposition $P = AA'$ for lower-triangular matrix A.
3. Return $\boldsymbol{X} = A\boldsymbol{Z}$ where $\boldsymbol{Z} = (Z_1, \ldots, Z_d)'$.

Simulating Copulas II

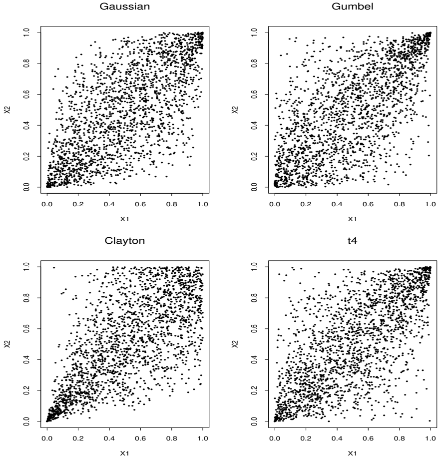

Gauss: $\rho = 0.7$, Gumbel: $\theta = 2$, Clayton: $\theta = 2.2$, t: $\rho = 0.71, \nu = 4$

Meta-Distributions and Their Simulation

- By the converse of Sklar's Theorem we know that if C is a copula and $F_1, \ldots, F_d$ are univariate dfs, then $F(\mathbf{x}) = C(F_1(x_1), \ldots, F_d(x_d))$ is a multivariate df with margins $F_1, \ldots, F_d$.
- We refer to F as a meta-distribution with the dependence structure represented by C. For example, if C is a Gaussian copula we get a meta-Gaussian distribution and if C is a t copula we get a meta-t distribution.
- If we can sample from the copula C, then it is easy to sample from F.

1. Generate a vector $(U_1, \ldots, U_d)$ with df C.
2. Return $(F_1^{\leftarrow}(U_1), \ldots, F_d^{\leftarrow}(U_d))$.

Simulating Meta Distributions

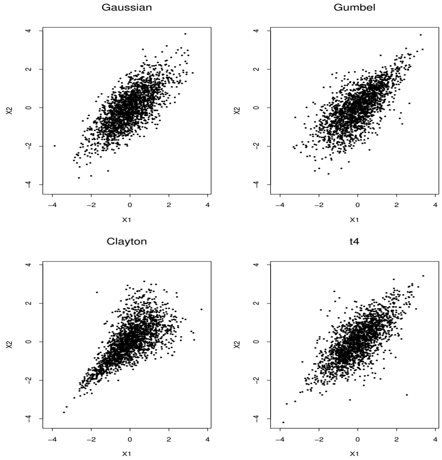

Linear correlation $\rho(X_1, X_2) \approx 0.7$ in all cases.

Overview

The Set-Up

- We have data vectors $\mathbf{X}_1, \dots, \mathbf{X}_n$ with identical distribution function F.
- We assume that this distribution function F has continuous marginal distributions $F_1, \dots, F_d$ and thus by Sklar's theorem a unique representation $F(x) = C(F_1(x_1), \dots, F_d(x_d))$.
- We assume that C belongs to a parametric copula family $C(\cdot; \boldsymbol{\theta})$.
- In order to estimate $C(\cdot; \boldsymbol{\theta})$ we first need to estimate $F_1, \dots, F_d$. There are two main approaches:

  Inference functions for margins (IFM) -
  : margins estimated with parametric distributions. (Joe, 1997)

  Pseudo-maximum likelihood -
  : margins estimated using version of empirical distribution function, e.g. $\widehat{F}_j(x) = \frac{1}{n+1} \sum_{i=1}^{n} 1_{\{X_{i,j} \leq x\}}$. (Genest and Rivest, 1993)

Stage 2: estimating the copula

- Having estimated the margins, we form a pseudo-sample from copula

$$\widehat{\mathbf{U}}_i = \left(\widehat{U}_{i,1}, \ldots, \widehat{U}_{i,d}\right)' = \left(\widehat{F}_1(X_{i,1}), \ldots, \widehat{F}_d(X_{i,d})\right)', \quad i = 1, \ldots, n.$$

  and fit parametric copula C by maximum likelihood.
- Copula density is $c(u_1, \ldots, u_d; \boldsymbol{\theta}) = \frac{\partial}{\partial u_1} \cdots \frac{\partial}{\partial u_d} C(u_1, \ldots, u_d; \boldsymbol{\theta})$.
- The log-likelihood is

$$l(\boldsymbol{\theta}; \widehat{\mathbf{U}}_1, \ldots, \widehat{\mathbf{U}}_n) = \sum_{i=1}^{n} \log c(\widehat{U}_{i,1}, \ldots, \widehat{U}_{i,d}; \boldsymbol{\theta}).$$

- Independence of vector observations is assumed here for simplicity.

Overview

Why rank correlation?

Why not ordinary (Pearson/linear) correlation?

- The correlation of a bivariate distribution F depends on both copula C and margins F_1 and F_2.
- Thus the correlation is not necessarily invariant under one-to-one transformations of the margins.
- For given margins, the range of attainable correlation need not be equal to $[-1, 1]$ and can be quite limited.

In contrast:

- Rank correlations depend only on copula C.
- They are thus invariant under one-to-one transformations of the margins
- For given margins, the range of attainable rank correlations is $[-1, 1]$.

This is important for Monte Carlo scenario generation under specification of margins and correlations (e.g. @RISK software).

Rank Correlation

Spearman's rho

$$
\begin{aligned}
\rho_S(X_1, X_2) &= \rho(F_1(X_1), F_2(X_2)) \\
\rho_S(X_1, X_2) &= 12 \int_0^1 \int_0^1 \{C(u_1, u_2) - u_1 u_2\} du_1 du_2 .
\end{aligned}
$$

Kendall's tau

Take an independent copy of (X_1, X_2) denoted (X_1^*, X_2^*).

$$
\begin{aligned}
\tau(X_1, X_2) &= \mathbb{P}\left((X_1 - X_1^*)(X_2 - X_2^*) > 0\right) - \mathbb{P}\left((X_1 - X_1^*)(X_2 - X_2^*) < 0\right) \\
&= \mathbb{P}\{\text{points concordant}\} - \mathbb{P}\{\text{points discordant}\} \\
&= 2\mathbb{P}\{\text{points concordant}\} - 1 \\
\tau(X_1, X_2) &= 4 \int_0^1 \int_0^1 C(u_1, u_2) dC(u_1, u_2) - 1 .
\end{aligned}
$$

Concordance

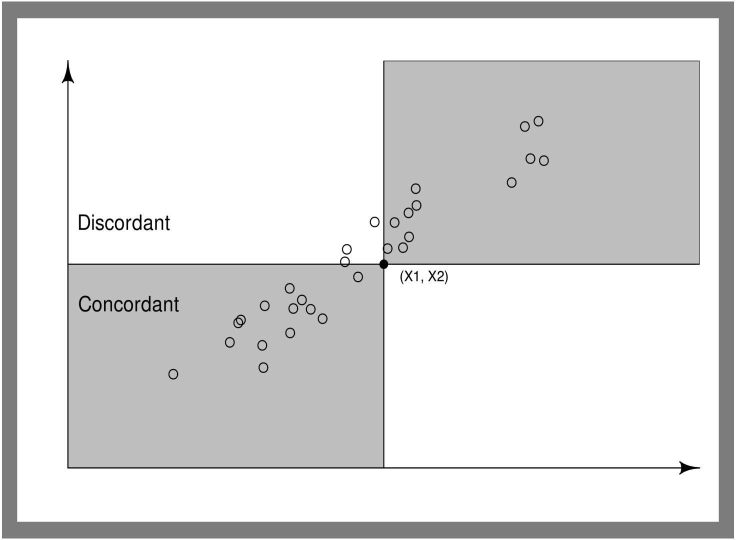

Rank correlations for certain copulas

- Let (X_1, X_2) be a bivariate random vector with copula C_ρ^{Ga} and continuous margins. Then the rank correlations are

$$\begin{aligned} \tau(X_1, X_2) &= \frac{2}{\pi} \arcsin \rho, \\ \rho_S(X_1, X_2) &= \frac{6}{\pi} \arcsin \frac{\rho}{2}. \end{aligned}$$

- The first formula also holds when (X_1, X_2) has the t copula $C_{\nu,\rho}^{\text{t}}$ or indeed the copula of a bivariate elliptical distribution with correlation ρ.
- Explicit formulas in terms of parameter θ available for many Archimedean copulas.
- Kendall's tau tends to be more widely available than Spearman's rho,
- This permits method-of-moments estimation by equating theoretical values of rank correlation with sample values.

Sample Rank Correlations

Consider iid bivariate data $\{(X_{1,1}, X_{1,2}), \ldots, (X_{n,1}, X_{n,2})\}$. The standard estimator of $\tau(X_1, X_2)$ is

$$\frac{1}{\binom{n}{2}} \sum_{1 \leq i < j \leq n} \operatorname{sgn}\left[(X_{i,1} - X_{j,1})(X_{i,2} - X_{j,2})\right],$$

and the standard estimator of $\rho_S(X_1, X_2)$ is

$$\frac{12}{n(n^2-1)} \sum_{i=1}^{n} \left(\operatorname{rank}(X_{i,1}) - \frac{n+1}{2}\right)\left(\operatorname{rank}(X_{i,2}) - \frac{n+1}{2}\right).$$

Overview

Overview

When only partial information is available

- In a risk analysis or stress test, we often need to generate scenarios from a model with prescribed rank correlations between variables.
- To obtain these rank correlations, we may have limited or non-existent data and may need to incorporate expert knowledge.
- To fix ideas, suppose we want to find the Kendall correlation matrix for returns on the cryptocurrencies Bitcoin, Ethereum, Litecoin and Ripple.
- Suppose we only have data on the first three and an expert opinion on the correlation between Bitcoin and Ripple.
- In a recent paper McNeil et al. (2020) have developed methods for determining the attainability of supposed Kendall correlation matrices and for completing missing values.

Illustration with crypto currencies

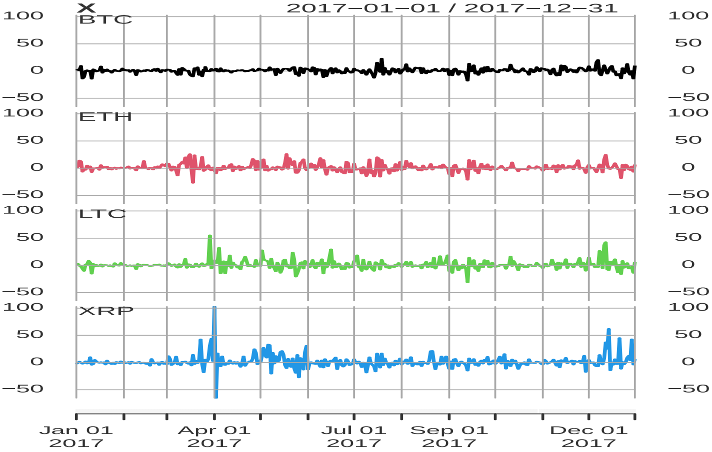

Log-returns on Bitcoin, Ethereum, Litecoin and Ripple prices in USD.

The attainability problem

Consider a continuous random vector $\boldsymbol{X} = (X_1, \ldots, X_d)$. The $d \times d$ Kendall rank correlation matrix P_τ is given by

$$(P_\tau)_{ij} = \tau(X_i, X_j).$$

What properties characterize P_τ?

P_τ is (1) symmetric and (2) its diagonal entries equal 1 and off-diagonal entries are elements of $[-1, 1]$.

It is also (3) positive semi-definite (no negative eigenvalues). These three properties are necessary and sufficient for Pearson correlation matrices.

However these properties are necessary but NOT sufficient for Kendall correlation matrices.

Overview

Extremal copulas

Consider an index set $J \subseteq \mathcal{D} = \{1, \dots, d\}$ and a random vector $\boldsymbol{U}$ with

$$U_j = \begin{cases} U & \text{if } j \in J, \\ 1 - U & \text{if } j \notin J, \end{cases}$$

where U is uniform on $[0, 1]$.

The vector $\boldsymbol{U}$ has uniform margins and spreads its mass uniformly along a main diagonal of the unit hypercube $[0, 1]^d$. Its cdf is the extremal copula

$$C(u_1, \dots, u_d) = \left(\min_{j \in J} u_j + \min_{j \in J^\complement} u_j - 1 \right)^+ .$$

When $J = \mathcal{D}$ we get the comonotonicity copula, as encountered earlier.

Why extremal?

The Kendall rank correlation matrix P_τ of the copula C with index set $J \subseteq \{1, \ldots, d\}$ satisfies

$$(P_\tau)_{ij} = \begin{cases} 1 & \text{if } i,j \in J \text{ or } i,j \in J^\complement, \\ -1 & \text{otherwise.} \end{cases}$$

- In other words, P_τ is an extremal correlation matrix consisting exclusively of entries 1 and -1.
- Pairs of variables are either perfectly positive dependent or perfectly negative dependent.
- In fact, P_τ is also the Spearman and Pearson correlation matrix of C.
- In dimension d there are 2^{d-1} extremal copulas, each associated with a diagonal of the d-dimensional unit cube.
- There are various ways of developing a unique enumeration $C^{(1)}, \ldots, C^{(2^{d-1})}$ of these copulas.

Illustration in 3D

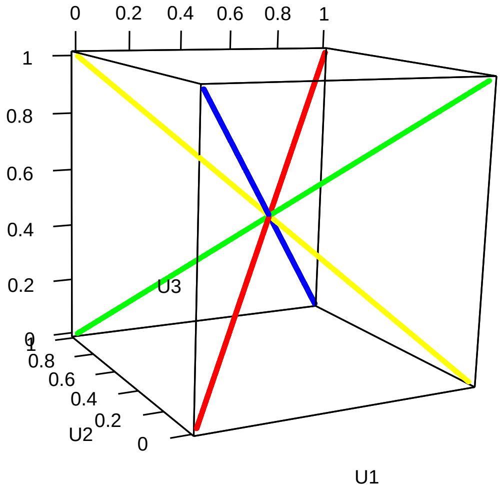

Samples of size 2000 from extremal copulas $C^{(1)}$, $C^{(2)}$, $C^{(3)}$ and $C^{(4)}$ in $d = 3$.

Overview

Attainability of Kendall rank correlation matrices

Let $P^{(k)}$ be the extremal correlation matrix of the kth extremal copula $C^{(k)}$ for $k \in \{1, \dots, 2^{d-1}\}$.

Characterization of Kendall's tau matrices

The $d \times d$ correlation matrix P is a Kendall's tau rank correlation matrix if and only if P can be represented as a convex combination of the extremal correlation matrices in dimension d, i.e.,

$$P = \sum_{k=1}^{2^{d-1}} w_k P^{(k)} \quad w_k \geq 0, \forall k, \quad \sum_{k=1}^{2^{d-1}} w_k = 1.$$

The convex hull of the matrices $P^{(k)}$ is called the cut-polytope.

Illustration in 3D

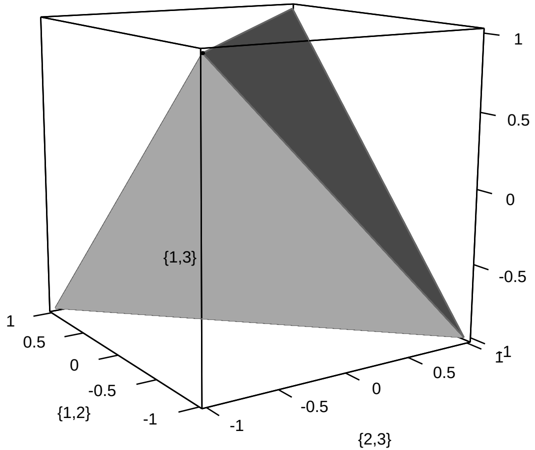

The cut-polytope of $(\tau_{\{1,2\}}, \tau_{\{1,3\}}, \tau_{\{2,3\}})$.

Not all correlation matrices are attainable

Consider the matrix

$$\frac{1}{12}\begin{pmatrix} 12 & -5 & -5 \\ -5 & 12 & -5 \\ -5 & -5 & 12 \end{pmatrix}.$$

This matrix is symmetric and positive definite and therefore a Pearson correlation matrix.

However, it is not a Kendall rank correlation matrix. The only possible weights would be

$$(w_1, w_2, w_3, w_4) = \frac{1}{24} \times (-27, 17, 17, 17)$$

Good news: when Kendall rank correlations are estimated from data without ties using the standard estimator, the resulting matrix is always attainable.

Overview

Linear programming problem

- The theory of testing for attainability of a $d \times d$ Kendall correlation matrix involves looking for a vector of non-negative weights $\boldsymbol{w} = (w_1, \dots, w_{2^{d-1}})$ to solve a linear system of the form

$$A\boldsymbol{w} = \begin{pmatrix} 1 \\ (1 + \tau_{\{1,2\}})/2 \\ \vdots \\ (1 + \tau_{\{d-1,d\}})/2 \end{pmatrix}$$

- A is a known matrix with $1 + d(d-1)/2$ rows and 2^{d-1} columns. The first row consist of 1's because of the sum constraint on the weights.
- For $d \leq 3$ there is either a unique solution or no solution.
- For $d \geq 4$ there may be multiple solutions or no solution.
- When there are multiple solutions, they form a convex set.
- When Kendall's tau values are missing, the corresponding rows of A are deleted.

Range of unobserved correlations

Imagine we have not observed Ripple prices and we only know Kendall's tau between Bitcoin, Ethereum and Litecoin: $(0.278, 0.333, 0.361)$.

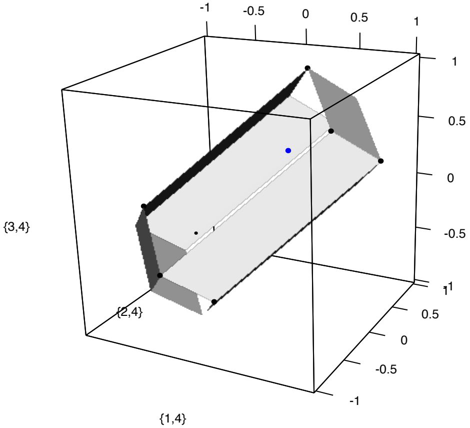

Range of attainable correlations with Ripple.

Adding more information

Now we are told that the correlation between Ripple and Bitcoin is 0.196.

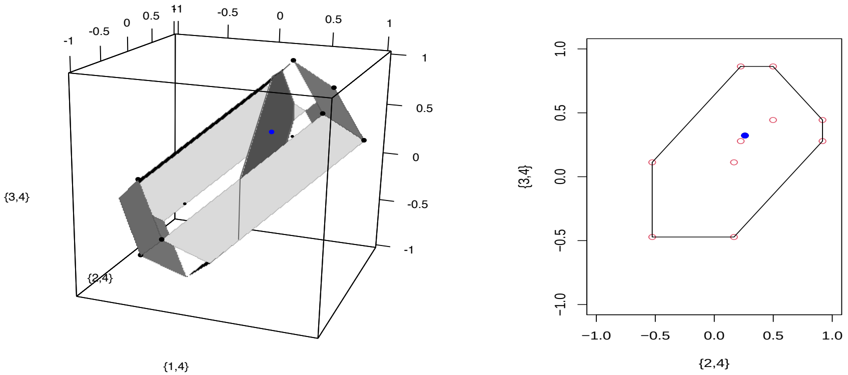

Range of attainable correlations with Ripple.

Finding a copula to match an attainable matrix

Suppose P_τ has been found to be an attainable Kendall rank correlation matrix. How do we find a copula C that has Kendall matrix P_τ?

1. For some matrices P_τ we can choose an elliptical copula like the Gauss copula or the t copula. Let P be the matrix given by the componentwise transformation

$$P = \sin(\tfrac{\pi}{2} P_\tau).$$

   If P is positive semi-definite, then the Gauss copula C_P^{Ga} or the t copula $C_{\nu,P}^{\mathrm{t}}$ have Kendall correlation matrix P_τ.

2. For all matrices P_τ we can use a mixture of extremal copulas. If $P_\tau = \sum_{i=1}^{2^{d-1}} w_k P^{(k)}$ then the copula

$$C = \sum_{i=1}^{2^{d-1}} w_k C^{(k)}$$

   has Kendall correlation matrix P_τ.

Overview

Overview

Aim of section

- In this section we consider strictly stationary time series models for modelling financial time series.
- There is a growing literature on copulas and time series; see survey article by Fan and Patton (2014).
- The main emphasis has been on cross-sectional dependencies between series; however, we are interested in serial dependence.
- The idea is to build strictly stationary copula processes (U_t) and to combine these with arbitrary, continuous marginal distributions F_x to create processes (X_t) such that $U_t = F_X(X_t)$ for all t.
- We will concentrate on two core classes for (U_t):
  1. ARMA copula processes as proposed in McNeil (2021).
  2. D-vine copula processes as developed in Chen and Fan (2006) (first-order) and Smith et al. (2010); these are based on the vine copula concept (Joe, 1997; Bedford and Cooke, 2002; Aas et al., 2009).
- To capture stochastic volatility we combine these models with v-transforms as proposed in McNeil (2021) and Bladt and McNeil (2020).

Overview

ARMA copula process

Definition

Let $(Z_t)_{t\in\mathbb{Z}}$ be a causal and invertible Gaussian ARMA(p,q) process such that $\mathbb{E}(Z_t) = 0$ and $\mathrm{var}(Z_t) = 1$ for all t. Then the process $(U_t)_{t\in\mathbb{Z}}$ such that $U_t = \Phi(Z_t)$ for all t is an ARMA(p,q) copula process.

- The distribution of any finite-dimensional random vector $(U_{t_1}, \ldots, U_{t_d})$ is a Gaussian copula with correlation matrix P depending on acf $\rho(k)$ of (Z_t).
- For example, $(U_1, \ldots, U_d) \sim C_P^{\mathrm{Ga}}$ where $P_{ij} = \rho(|i-j|)$.
- The process has $p+q$ parameters.
- It is possible to estimate a model combining an ARMA copula process with a continuous parametric marginal distribution using the method of maximum likelihood.
- The finite-dimensional marginal distributions of the resulting model are meta-Gaussian.

Overview

Generalizing the AR copula process

- The d-vine copula process can be thought of as a generalization of the ARMA(p,0) or AR(p) copula process.
- The AR(p) copula process can be fully described by p bivariate Gaussian copulas $C_1, \ldots, C_p$.
- For all t,
  - C_1 is the copula of (U_t, U_{t+1});
  - C_2 is the copula of (U_t, U_{t+2}) given U_{t+1};
  - C_3 is the copula of (U_t, U_{t+3}) given U_{t+1} and U_{t+2};
  - and so on.
- The copula C_i could be called the ith partial copula.
- Its parameter ρ_i satisfies $\rho_i = \alpha(i)$ where $\alpha(\cdot)$ is the partial autocorrelation function of the underlying AR process (Z_t).
- What would happen if we swapped these bivariate Gaussian copulas for non-Gaussian copulas?

D-vine copula process

Definition

A strictly stationary copula process $(U_t)_{t\in\mathbb{Z}}$ is a d-vine(p) process if, for any $t \in \mathbb{Z}$ and $d \geq 2$, the joint density of $(U_{t+1}, \ldots, U_{t+d})$ is given by

$$\prod_{i=1}^{\max(p,d-1)} \prod_{j=1+i}^{d} c_i\big(R_i^*(u_{j-i}, (u_{j-1}, \ldots, u_{j-i+1})), R_i(u_j, (u_{j-1}, \ldots, u_{j-i+1}))\big)$$

where $c_1, \ldots, c_p$ are bivariate copula densities and

$$\begin{aligned} R_i(x, (u_1, \ldots, u_{i-1})) &= \mathbb{P}(U_t \leq x \mid U_{t-1} = u_1, \ldots, U_{t-i+1} = u_{i-1}) \\ R_i^*(x, (u_1, \ldots, u_{i-1})) &= \mathbb{P}(U_{t-i} \leq x \mid U_{t-1} = u_1, \ldots, U_{t-i+1} = u_{i-1}). \end{aligned}$$

and where by convention $R_1(x, \cdot) = R_1^*(x, \cdot) = x$.

Remarks on d-vine copula process

- Such processes do exist and can be constructed and simulated.
- If $c_1, \ldots, c_p$ are bivariate Gaussian copula densities, the joint density of the d-vine copula process is identical to that of an AR(p) copula process.
- The joint density is easier to evaluate than it first appears; the functions R_i and R_i^* can be evaluated recursively.
- This means that maximum likelihood inference is possible.
- We generally use a mixture of Archimedean, Gauss and t copulas for the pair copula densities.
- Combined with non-Gaussian marginal distributions, these models can improve the fit in many situations where AR and ARMA processes are standard models.
- However, they cannot capture stochastic volatility (at least using common copula choices).

Overview

Bitcoin log return data

U and ***V*** scales:

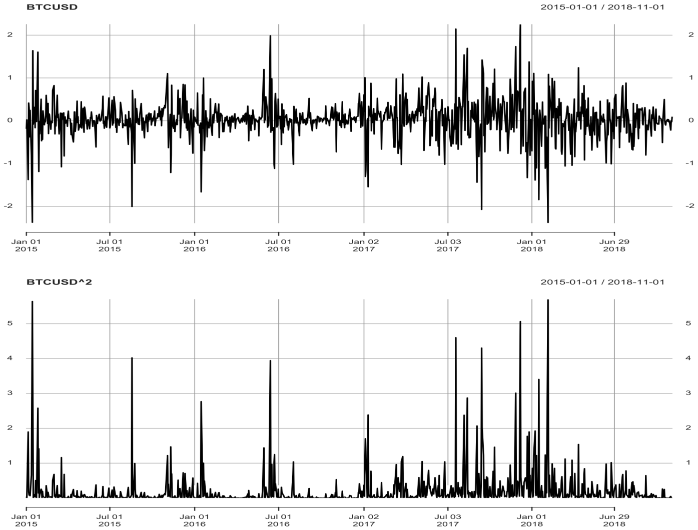

Empirical v-transform

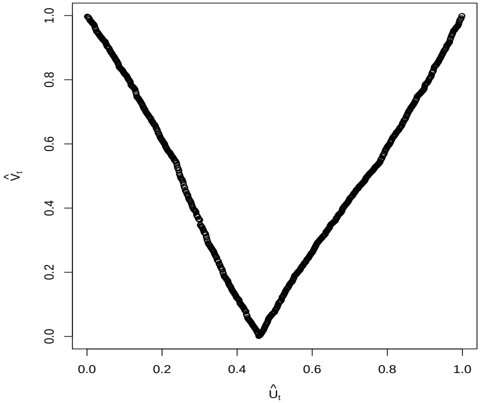

$$\widehat{U}_t = F_n^{(X)}(X_t) \quad \widehat{V}_t = F_n^{(X^2)}(X_t^2)$$

Volatility-revealing transformations

- Volatile financial return series (X_t) seldom show serial correlation, but the series of absolute vales $(|X_t|)$ or squared values (X_t^2) do.
- If we attempt to transform the raw data and the squared data to uniform (e.g. with the empirical distribution function) then the latter show a v-shaped relationship to the former.
- This is not suprising: if X has cdf F_X symmetric around 0 and we define $U = F_X(X)$ and $V = F_{X^2}(X^2)$ then $V = \mathcal{V}(U)$ where $\mathcal{V}(u) = |2u - 1|$.
- $\mathcal{V}(u) = |2u - 1|$ belongs to a class of v-shaped, uniformity-preserving transformations that we call v-transforms.
- To model the serial dependence in real data we construct copula processes (U_t) such that, under some componentwise v-transform $V_t = \mathcal{V}(U_t)$, the copula process (V_t) follows one of our models for serial dependence (ARMA copula or d-vine copula).
- The model for (U_t) is called a vt-ARMA or vt-d-vine copula process.
- They are alternatives to GARCH models.

Summary

- Copulas are particularly useful for scenario analyses or stress tests where we consider risk factors with different marginal distributions.
- When full data are available, they can be estimated.
- When data are sparse or missing, they can be elicited from expert opinion. Rank correlations are particularly helpful in this case.
- There have been recent advances in understanding the conditions for matrices of rank correlations to be attainable (McNeil et al., 2020; Wang et al., 2019).
- Copula models for time series is also an area of contemporary research interest (Bladt and McNeil, 2020; Nagler et al., 2020).

For Further Reading

Aas, K., Czado, C., Frigessi, A., and Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2):182–198.

Bedford, T. and Cooke, R. M. (2002). Vines–a new graphical model for dependent random variables. *Annals of Statistics*, 30(4):1031–1068.

Bladt, M. and McNeil, A. (2020). Time series copula models using d-vines and v-transforms: an alternative to GARCH modelling. arXiv:2006.11088.

Chen, X. and Fan, Y. (2006). Estimation of copula-based semiparametric time series models. *Journal of Econometrics*, 130(2):307–335.

Fan, Y. and Patton, A. (2014). Copulas in econometrics. *Annual Review of Economics*, 6:179–200.

Genest, C. and Rivest, L. (1993). Statistical inference procedures for bivariate Archimedean copulas. *Journal of the American Statistical Association*, 88:1034–1043.

Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.

For Further Reading (cont.)

McNeil, A. (2021). Modelling volatility with v-transforms and copulas. *Risks*, 9(1):14.

McNeil, A., Nešlehová, J., and Smith, A. (2020). On attainability of kendall's tau matrices and concordance signatures. arXiv preprint 2009.08130.

Nagler, T., Krüger, D., and Min, A. (2020). Stationary vine copula models for multivariate series. Preprint.

Smith, M., Min, A., Almeida, C., and Czado, C. (2010). Modeling Longitudinal Data Using a Pair-Copula Decomposition of Serial Dependence. *Journal of the American Statistical Association*, 105(492):1467–1479.

Wang, B., Wang, R., and Wang, Y. (2019). Compatible matrices of Spearman's rank correlation. *Statistics & Probability Letters*, 151:67–72.