Aalborg University

10TH SEMESTER

Mathematics-Economics

Master thesis

Dependence Modelling in Portfolio Optimization

A GARCH-Copula Approach under Mean-CVaR Optimization

27/05-2026

AALBORG
UNIVERSITY

AALBORG UNIVERSITY

The Department of Mathematical Sciences
Mathematics-Economics
Thomas Manns Vej 23
9220 Aalborg Øst
https://math.aau.dk

Title:

Dependence Modelling in Portfolio Optimization

Subtitle:

A GARCH-Copula Approach under Mean-CVaR Optimization

Project period:

1. February 2026 - 27. May 2026

Authors:

Cecilie Stoklund Nørnberg

Supervisor:

Esben Høg

Pagecount: 82 + front page
Done: 27/5-2026

Abstract:

This thesis investigates portfolio performance under different dependence structures using a mean–CVaR optimisation framework. In particular, three copula models are considered, namely the Gaussian, Student's t, and R-vine copula, in order to assess how the modelling of dependence affects the trade-off between return and downside risk. The analysis is based on ten stocks from the OMX C25 index, representing different sectors, where optimal portfolio allocations are constructed under varying levels of risk aversion and evaluated against a naive equally weighted benchmark. The results show that all copula-based portfolios consistently outperform the benchmark portfolio in terms of both return and risk-adjusted performance. Among the copula specifications, Student's t delivers the strongest overall performance, followed by R-vine, while Gaussian generally underperforms. However, despite differences in realised returns, the model-implied downside risk dynamics are highly similar across all models, with nearly identical responses to periods of market stress. Overall, the findings suggest that the choice of dependence structure primarily affects portfolio performance through realised returns rather than through differences in downside risk.

Preface

The project period has elapsed from 01/02-2027 - 27/05-2026. Data processing and calculations are done using the programming language R.

Reader's guide

Literature and source code references are specified at the start of every section, where it is used. Literature references refer to the bibliography at the end of the project, and source code references refer to the appendix. Literature references follow the Vancouver method and are therefore given by [Source number].

Acknowledgements

The work on the project has been supervised by Esben Høg during the entire work process, I therefore thank him for input and feedback.

Contents

1 Introduction

In highly volatile and unpredictable financial markets, uncertainty is inevitable, meaning that investors are continuously faced with difficult decisions when allocating capital across different assets. Having an efficient portfolio optimization strategy that balances the trade-off between expected return and downside risk is therefore essential, as the rapidly changing financial markets offer both the potential for substantial gains as well as significant losses.

A commonly used approach within Modern Portfolio Theory is *Mean-Variance Portfolio Optimization*, introduced by the economist Harry Markowitz. This approach formalizes the portfolio optimization problem as a trade-off between maximizing the expected portfolio return and minimizing the variance. Risk is therefore solely characterised by the covariance matrix, which assumes linear dependence between the assets and treats positive and negative deviations from the mean equally. In practice, these assumptions are often too restrictive for financial markets, where stock returns typically exhibit skewness, heavy tails, and non-linear dependence structures, which cannot be fully described by variance alone.

To address these limitations an alternative approach, known as *Mean-Conditional Value-at-Risk (Mean-CVaR)*, is implemented in this thesis. Here, risk is measured by CVaR, which captures the expected loss in the lower tail of the return distribution, thereby focusing explicitly on downside risk rather than the overall variability of the returns. This optimization problem is therefore particularly sensitive to extreme negative outcomes, which makes it well suited for financial markets, where large losses tend to co-occur more frequently than large gains. However, implementing this approach requires an appropriate specification of the joint distribution of the stock returns

There a several ways of modelling the joint distribution of stock returns. A relatively straightforward method is to use a multivariate distribution, describing both the stocks' individual behaviour and their dependence structure. However, this approach quite restrictive in practice, as it typically requires all marginal distributions to belong to the same distributional family, while also imposing a fixed dependence structure. A more flexible approach is to separate the modelling of the marginal distributions from the dependence structure, which is the core of the copula framework. In particular, copulas allow the dependence structure between stocks to be modelled independently of the marginal distributions, making it possible to capture non-linear and tail dependence commonly observed in financial markets. In this thesis, a GARCH–copula framework is therefore used, where the marginal distributions are modelled using GARCH processes to capture time-varying volatility, while the dependence structure is specified through different copula families.

In view of this, the problem statement for this thesis is as follows:

Problem statement *How does the specification of dependence structure in a GARCH-copula framework influence portfolio performance compared to a naive equal-weight strategy.*

Furthermore, how does portfolio performance vary across different levels of downside risk aversion?

2 Copula Theory

A typical mathematical framework for portfolio optimization is *Modern Portfolio Theory (MPT)*, which assumes that portfolio returns can be fully characterised by a mean vector and a covariance matrix, typically estimated from historical data. Consequently, the dependence structure between the assets in the portfolio is entirely determined by the covariance matrix, meaning that only linear dependence is captured. This limitation is particularly important in portfolio optimization, as asset returns typically exhibit complex dependence structures, including non-linear relationships and extreme co-movements in the tails of the distribution. (See Allen and Luciano [2019])

Another approach is to fit a multivariate distribution directly to the data, and use it to model the portfolio returns. However, in this case both the marginal behaviour of the assets and the dependence structure are fully determined by the chosen multivariate distribution. For instance, if the portfolio returns are modelled by a multivariate normal distribution, the dependence structure is restricted to linear dependence, and the marginal distributions are all assumed to be normal, where neither the marginals nor the dependence structure can be adjusted or changed without altering the entire multivariate distribution.

To address these limitations, *copulas* are introduced as a flexible framework for modelling the multivariate distribution of a portfolio. By separating the modelling of the marginal distributions from the dependence structure, copulas allow both linear and non-linear co-movements to be captured, including extreme co-movements in the tails of the distribution, which is crucial for accurate risk assessment and effective portfolio management. In the following sections, the fundamental concepts of copulas, including *Sklar's theorem*, will be presented, along with different *dependence measures* and some of the most commonly used *copula families*.

2.1 Foundations of Copula Theory

To establish the proper mathematical framework, it is first necessary to introduce a few fundamental concepts. Central to these is the extended real line $[-\infty, +\infty]$, denoted by $\bar{\mathbb{R}}$, along with its d-dimensional Cartesian product, $\bar{\mathbb{R}}^d$, which will be used throughout this section. (See Durante and Sempi [2016])

The first concept, which will be introduced, is the *d-box*.

Definition 2.1.1. d-Box
Let $d \geq 2$ and let $\mathbf{a}, \mathbf{b} \in \bar{\mathbb{R}}^d$, where $\mathbf{a} \leq \mathbf{b}$. A *d-box*, denoted as $[\mathbf{a}, \mathbf{b}]$, is then defined as

$$[\mathbf{a}, \mathbf{b}] = [a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_d, b_d].$$

So intuitively, when $d = 1$ the d-box represents a closed interval, when $d = 2$ it represents a rectangle, and when $d = 3$ it forms a rectangular box. The d-box is therefore a generalization of a rectangle to d-dimensions, also known as a hyperrectangle. The vertices of the d-box correspond to the points $\mathbf{c} = c_1, \ldots, c_d$, where each c_i is equal to either a_i or b_i. These vertices play a central role in defining the *H-volume of an d-box*.

Definition 2.1.2. H-Volume of an d-Box
Let $S_1, \ldots, S_d$ be non-empty subsets of $\bar{\mathbb{R}}$, let $H : S_1 \times \cdots \times S_d \to \mathbb{R}$ and let $B = [\mathbf{a}, \mathbf{b}]$ be a d-box, where $\mathbf{a}, \mathbf{b} \in S_1 \times \cdots \times S_d$. The *H-volume of* B is then given by

$$V_H(B) = \sum_{\mathbf{c} \in \{a_1, b_1\} \times \cdots \times \{a_d, b_d\}} \operatorname{sign}(\mathbf{c}) H(\mathbf{c}),$$

where the sign of the corners $\mathbf{c}$ of B is defined as

$$\operatorname{sign}(\mathbf{c}) = \begin{cases} 1, & \text{if } c_i = a_i \text{ for an even number of indices } i \\ -1 & \text{if } c_i = a_i \text{ for an odd number of indices } i. \end{cases}$$

The H-volume measures the total change of the function H over the d-box. For instance, if H is a bivariate distribution function, then the H-volume of a 2-box B represents the probability that two random variables assume values within B, which is calculated as

$$V_H(B) = H(b_1, b_2) - H(a_1, b_2) - H(b_1, a_2) + H(a_1, a_2). \tag{2.1}$$

The altering signs in (2.1) account for the contributions of the corners of the box. Each corner contributes positively or negatively depending on number of coordinates equal to the lower bound a_i, for example corners with an odd number of lower bounds contribute negatively, while corners with an even number contribute positively. For instance, in (2.1), the corner $\mathbf{c} = [a_1, b_2]$ contains one lower bound a_1, meaning that it contributes negatively, while the corner $[a_1, a_2]$ contains two lower bounds a_1 and a_2, meaning that it contributes positively. (See Jaworski et al. [2010])

Now that the H-volume of a d-box has been introduced, it is possible to define a copula.

Definition 2.1.3. Copula
Let $d \geq 2$. A function $C : [0,1]^d \to [0,1]$ is called a *copula* it it satisfies the following properties:

1. C is *grounded*, i.e., for each $i \in \{0, \ldots, d\}$

$$C(u_1, \ldots, u_{i-1}, 0, u_{i+1}, \ldots, u_d) = 0, \quad \text{for all } u_j \in [0,1], \text{ where } j \neq i,$$

2. C has *uniform marginals* $C_i : [0,1] \to [0,1]$ for every $i \in \{1, \ldots, d\}$, i.e.,

$$C_i(u_i) = C(1, \ldots, 1, u_i, 1, \ldots, 1) = u_i, \quad \text{for all } u_i \in [0,1],$$

3. C is *d-increasing*, i.e., for every $\mathbf{a}, \mathbf{b} \in [0,1]^d$, where $\mathbf{a} \leq \mathbf{b}$:

$$V_C([\mathbf{a}, \mathbf{b}]) \geq 0.$$

Remark. *Notice, that a copula is sometimes referred to as a d-copula to emphasise its domain* $[0,1]^d$ *for* $d \geq 2$.

As stated in Definition 2.1.3, a copula is both grounded, has uniform marginals and is d-increasing. These properties ensure that it also satisfies the defining properties of a d-dimensional distribution function, which are provided in Theorem A.0.2 (Properties of d-Dimensional Distribution Functions). In other words, a copula is a multivariate distribution function that models the dependence structure between variables by joining their marginal distributions, which is exactly what Sklar's theorem states. (See Nelsen [2006])

Theorem 2.1.4. Sklar's Theorem
Let $d \geq 2$ and let $F : \bar{\mathbb{R}}^d \to [0,1]$ be a d-dimensional distribution function with univariate marginal distribution functions $F_i : \bar{\mathbb{R}} \to [0,1]$ for $i \in \{1, \ldots, d\}$, then there exists a d-copula $C : [0,1]^d \to [0,1]$, such that for all $(x_1, \ldots, x_d) \in \bar{\mathbb{R}}^d$

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

If all F_i are continuous, then C is unique, otherwise, C is uniquely determined on $\text{ran}(F_1) \times \cdots \times \text{ran}(F_d)$. Conversely, if C is a d-dimensional copula and $F_1, \ldots, F_d$ are univariate marginal distribution functions, then F is a d-dimensional distribution function with univariate marginal distribution functions $F_1, \ldots, F_d$.

Proof. This proof is omitted, but it can be found in [Durante and Sempi, 2016, p. 54–56]. □

The decomposition of the d-dimensional distribution function in Theorem 2.1.4 separates the individual behaviour of the random variables from the dependence structure captured

by the copula. This makes it possible to analyse the relationships between the random variables independently of their individual distributions, which is particular useful in portfolio optimization.

Under the assumptions of Theorem 2.1.4 (Sklar's Theorem), if the d-dimensional distribution F is also absolutely continuous, the joint density of $(X_1, \ldots, X_d)$ with marginal densities $f_1, \ldots, f_d$ can be expressed as

$$f(x_1, \ldots, x_d) = c_{1,\ldots,d}\left(F_1(x_1), \ldots, F_d(x_d)\right) \cdot f_1(x_1) \cdots f_d(x_d), \tag{2.2}$$

for almost all $(x_1, \ldots, x_d) \in \bar{\mathbb{R}}^d$, where $c_{1,\ldots,d}$ denotes the *Copula Density* given by

$$c_{1,\ldots,d}\left(F_1(x_1), \ldots, F_d(x_d)\right) = \frac{\partial^d C\left(F_1(x_1), \ldots, F_d(x_d)\right)}{\partial F_1(x_1) \cdots \partial F_d(x_d)}.$$

It is thus possible to extend Theorem 2.1.4 (Sklar's Theorem) to densities, when the d-dimensional distribution function F is absolute continuous. The joint density can then be decomposed into a product of copula densities and marginal densities, which allows the dependence structure and the marginal behaviour to be modelled separately, while maintaining the d-dimensional distribution function. (See Durante and Sempi [2016])

2.2 Dependence Measures

As states in Theorem 2.1.4 (Sklar's Theorem), the joint distributions of a set of random variables can be decomposed into its marginal distributions and a copula capturing their dependence structure. While marginal distributions can often be determined directly from the data, choosing a suitable copula is less straightforward. A naive approach is to just select a copula family, from the families introduced in Section 2.3, fit it to the data and then evaluate its goodness-of-fit. However, this approach may be misleading, as different copula families impose different characteristics, such as symmetry or particular tail behaviours, which may not reflect the actual dependence structure observed in the data. (See Jaworski et al. [2010])

A more robust approach is to first analyse the pairwise dependence between the variables. In a d-dimensional setting, this involves examining the dependence for each unique pair of variables individually, where each generic pair will be denoted as (X_1, X_2) throughout this section for simplicity. Three commonly used measures of dependence are *Tail Dependence*, which quantifies the probability that extreme values occur simultaneously in a pair of variables, and *Kendall's Tau*, which measures the tendency that the pair move in the same direction, and *Spearman's Rho*, which captures the strength and direction of association between the pair. Collectively, these dependence measures provide a nuanced insight into the dependence structure and will be discussed in more detail in the following sections.

2.2.1 Spearman's Rho

Spearman's rho is a rank-based dependence measure, meaning that it relies on the relative ordering of the observations of the random variables rather than their numerical values. It is therefore necessary to formally introduce the concept of the *rank of an observation*, before defining Spearman's rho.

Definition 2.2.1. Rank of an Observation
Let X_1 be a random variable with $n \in \mathbb{N}$ observations. The *rank of an observation* $X_{1,k}$, denoted $\text{Rank}\left(X_{1,k}\right)$, is defined as

$$\text{Rank}\left(X_{1,k}\right) = \sum_{j=1}^{n} \mathbf{1}_{\left\{X_{1,j} \leq X_{1,k}\right\}},$$

where $k \in \{1, \ldots, n\}$.

In other words, the rank of an observation $X_{i,k}$ in the sample $X_{1,1}, X_{1,2}, \ldots, X_{1,n}$ is equal to 1 if the observation is the smallest in the sample, 2 if it is the second smallest and so on, up to n if it is the largest observation. By converting the observations into ranks, Spearman's rho can be interpreted as the usual Pearson correlation coefficient just based on the relative ordering of the observations instead of their actual numerical values, allowing it to measure the strength of the monotonic relationship between two random variables. (See Ruppert and Matteson [2015])

To provide a more intuitive understanding of how Spearman's rho relates to the Pearson correlation, consider a pair of variables (X_1, X_2), each with n observations, then the empirical Spearman's rho, $\hat{\rho}_s \in [-1, 1]$ for the pair is given by

$$\begin{aligned}
\hat{\rho}_s &= \hat{\rho}\left(\text{Rank}(X_1), \text{Rank}(X_2)\right) \\
&= \frac{\sum_{k=1}^{n}\left(\text{Rank}(X_{1,k}) - \overline{\text{Rank}(X_1)}\right)\left(\text{Rank}(X_{2,k}) - \overline{\text{Rank}(X_2)}\right)}{\sqrt{\sum_{k=1}^{n}\left(\text{Rank}(X_{1,k}) - \overline{\text{Rank}(X_1)}\right)^2 \sum_{k=1}^{n}\left(\text{Rank}(X_{2,k}) - \overline{\text{Rank}(X_2)}\right)^2}},
\end{aligned}$$

where $\hat{\rho}$ denotes Pearson's correlation, and where $\overline{\text{Rank}(X_i)} = \frac{1}{n}\sum_{k=1}^{n} \text{Rank}\left(X_{i,k}\right)$ for $i = 1, 2$, meaning that, rather than centering the observations around their mean as in the Pearson correlation, Spearman's rho centers them around the mean of the ranks. If the empirical Spearman's rho, $\hat{\rho}_s$, is equal to -1 the pair (X_1, X_2) have a perfect decreasing monotonic relationship, if it is equal to 1, they have a perfect increasing monotonic relationship, and if it is 0, where is no monotonic association. (See Czado [2019])

So while the empirical Spearman's rho is based on observed data, the *Population Spearman's rho*, which will be referred to as Spearman's rho throughout the rest of the thesis,

is defined in terms of the underlying distributions of the variables.

Definition 2.2.2. Spearman's Rho
Let (X_1, X_2) be a pair of continuous random variables. *Spearman's Rho*, $\rho_s \in [-1, 1]$, for (X_1, X_2) is defined as

$$\rho_s(X_1, X_2) = \rho\left(F_2(X_1), F_2(X_2)\right),$$

where $\rho \in [-1, 1]$ denotes the Pearson correlation coefficient, and $F_i(X_i)$ is the marginal distribution function of X_i for $i = 1, 2$.

By this definition, Spearman's rho of a pair of variables (X_1, X_2) can be interpreted as as the Pearson correlation applied to the marginal distribution functions of X_1 and X_2.

Now, since the two random variables X_1 and X_2 are continuous, their marginal distribution functions are also continuous, meaning that by Theorem A.0.3 (Probability Integral Transform) it follows that $F_i(X_i) \sim \text{Uniform}(0, 1)$ for $i = 1, 2$. It is therefore possible to write

$$\rho_s(X_1, X_2) = \rho\left(F_2(X_1), F_2(X_2)\right) = \rho\left(U_1, U_2\right),$$

where $F_i(X_i) = U_i$ and $U_i \sim \text{Uniform}(0, 1)$ for $i = 1, 2$. So basically, by transforming both marginal distribution functions into standard uniform variables, Spearman's rho reflects only the dependence structure between X_1 and X_2, and is independent of their specific marginal distributions. Furthermore, by Theorem 2.1.4 (Sklar's Theorem) the joint distribution of U_1 and U_2 is completely determined by the copula C of (X_1, X_2). It is therefore possible to express Spearman's rho entirely in terms of C as defined in Theorem 2.2.3. (See Czado [2019])

Theorem 2.2.3. Copula Representation of Spearman's Rho
Let (X_1, X_2) be a pair of continuous random variables with copula C. *Spearman's Rho*, $\rho_s \in [-1, 1]$, for (X_1, X_2) can be expressed as

$$\rho_s(X_1, X_2) = 12 \int_{[0,1]^2} u_1\, u_2 \,\mathrm{d}C(u_1, u_2) - 3$$

where $u_1 = F_{X_1}(x_1)$ and $u_2 = F_{X_2}(x_2)$ denote the marginal distribution functions of X_1 and X_2, respectively, so that $u_i \sim \text{Uniform}(0, 1)$ for $i = 1, 2$.

Proof. Firstly, Spearman's rho can be expressed as

$$\rho_s(X_1, X_2) = \rho\left(F_2(X_1), F_2(X_2)\right) = \rho\left(U_1, U_2\right),$$

by Definition 2.2.2 (Spearman's Rho), where $U_i = F_i(X_i)$ for $i = 1, 2$.

Then recall that the formula for Pearson's correlation coefficient is given by

$$\rho\left(U_1, U_2\right) = \frac{\mathrm{Cov}(U_1, U_2)}{\sqrt{\mathrm{Var}(X_1)\mathrm{Var}(X_2)}} = \frac{\mathbb{E}[U_1U_2] - \mathbb{E}[U_1]\mathbb{E}[U_2]}{\sqrt{\mathrm{Var}(X_1)\mathrm{Var}(X_2)}}.$$

Since U_1 and U_2 are both uniformly distributed on the interval $[-1, 1]$, their variance and expected value can be expressed as

$$\mathrm{Var}(U_i) = \frac{1}{12}, \quad \text{and} \quad \mathbb{E}[U_i] = \frac{1}{2}, \quad \text{for } i = 1, 2,$$

which means that

$$\rho\left(U_1, U_2\right) = \frac{\mathbb{E}[U_1U_2] - \frac{1}{4}}{\frac{1}{12}} = 12 \cdot \mathbb{E}[U_1U_2] - 3$$

Finally, by Theorem 2.1.4 (Sklar's Theorem), the joint distribution of (U_1, U_2) can be expressed in terms of their copula C, allowing the expected value to be written as

$$\mathbb{E}[U_1U_2] = \int_{[0,1]^2} u_1\, u_2\, \mathrm{d}C(u_1, u_2).$$

The expression for Spearman's rho therefore becomes

$$\rho\left(U_1, U_2\right) = 12 \int_{[0,1]^2} u_1u_2\mathrm{d}C(u_1, u_2) - 3,$$

concluding the proof. □

2.2.2 Kendall's Tau

Kendall's tau is another rank-based correlation coefficient, measuring the strength of association between two variables based on the ordering of their observations rather that their actual values. This rank-based association is formalized through the notions of *Concordant* and *Discordant Pairs*. (See Nelsen [2006])

Definition 2.2.4. Concordant and Discordant Pairs
Let (X_1, X_2) be a pair of random variables, each with $n \in \mathbb{N}$ observations. Consider two distinct observation from this pair, $(x_{1,k}, x_{2,k})$ and $(x_{1,l}, x_{2,l})$, where $k, l \in \{1, \ldots, n\}$ and $k \neq l$. This observation pair is classified as

- *Concordant* if both variables move in the same direction, i.e., either $x_{1,k} < x_{1,l}$ and $x_{2,k} < x_{2,l}$ or $x_{1,k} > x_{1,l}$ and $x_{2,k} > x_{2,l}$.
- *Discordant* if the variables move in opposite directions, i.e., either $x_{1,k} < x_{1,l}$ and $x_{2,k} > x_{2,l}$ or $x_{1,k} > x_{1,l}$ and $x_{2,k} < x_{2,l}$.

Concordant pairs are therefore determined as those observation pairs that preserve the same ordering across both variables, whereas discordant pairs are those for which the ordering is reversed. Kendall's tau quantifies the overall tendency of the pair by comparing the total number of concordant and discordant observation pairs. (See Czado [2019])

Specifically, for a pair of variables (X_1, X_2) based on a sample of n observations, the empirical Kendall's tau, $\hat{\tau} \in [-1, 1]$, is given as

$$\hat{\tau}\left(X_1, X_2\right) = \frac{N_c - N_d}{\binom{n}{2}} = \binom{n}{2}^{-1} \sum_{1 \le k < l \le n} \operatorname{sign}\left((x_{1,k} - x_{1,l})(x_{2,k} - x_{2,l})\right), \tag{2.3}$$

where $\binom{n}{2}$ denotes the total number of unique observation pairs, and $N_c, N_d \in \mathbb{N}$ denote the number of concordant and discordant observation pairs, respectively.

The first expression in (2.3) is the normalized difference between the number of concordant and discordant observation pairs, which can be interpreted as the overall tendency for the pair (X_1, X_2) to move together. The second expression is an equivalent measure, just with the numerator rewritten as a sum over all observation pairs using the sign function, which is defined as

$$\operatorname{sign}\left((x_{1,k} - x_{1,l})(x_{2,k} - x_{2,l})\right) = \begin{cases} +1, & \text{if the pair is concordant,} \\ -1, & \text{if the pair is discordant,} \\ 0, & \text{otherwise.} \end{cases}$$

This empirical version of Kendall's tau thereby provides a concrete way of estimating the dependence between a pair of variables directly from observed data. (see Ruppert and Matteson [2015])

This concept can also be extended to a theoretical setting, where Kendall's tau for a pair of random variables (X_1, X_2) is defined as the probability that two independent observations drawn from the same joint distribution are concordant or discordant.

Definition 2.2.5. Kendall's Tau
Let (X_1, X_2) be a pair of random variables. *Kendall's Tau*, $\tau \in [-1, 1]$, for (X_1, X_2) is defined as

$$\tau(X_1, X_2) = \mathbb{P}\left((X_{1,1} - X_{2,1})(X_{1,2} - X_{2,2}) > 0\right) - \mathbb{P}\left((X_{1,1} - X_{2,1})(X_{1,2} - X_{2,2}) < 0\right),$$

where $(X_{1,1}, X_{1,2})$ and $(X_{2,1}, X_{2,2})$ are i.i.d. copies of (X_1, X_2).

In other words, Kendall's Tau measures the probability that a pair of i.i.d. observations move in the same or opposite directions. Specifically, when $\tau(X_1, X_2) > 0$ concordant

pairs are more likely, when $\tau(X_1, X_2) < 0$ discordant pairs are more likely, and when $\tau(X_1, X_2) = 0$ there is no tendency for the variables to move together.

From Definition 2.2.5 (Kendall's Tau) it also becomes clear, that Kendall's tau is based solely on the dependence structure of the variables and not their marginal distributions, which allows it to be expressed directly in terms of the copula associated with (X_1, X_2), as stated in the following theorem. (See Czado [2019])

Theorem 2.2.6. Copula Representation of Kendall's Tau
Let (X_1, X_2) be a pair of continuous random variables with copula C. *Kendall's Tau*, $\tau \in [-1, 1]$, for (X_1, X_2) can be expressed as

$$\tau(X_1, X_2) = 4 \int_{[0,1]^2} C(u_1, u_2) \, \mathrm{d}C(u_1, u_2) - 1,$$

where $u_1 = F_{X_1}(x_1)$ and $u_2 = F_{X_2}(x_2)$ denote the marginal distribution functions of X_1 and X_2, respectively, so that $u_i \sim \text{Uniform}(0, 1)$ for $i = 1, 2$.

Proof. This proof is omitted, but it can be found in [Czado, 2019, p. 32–33]. □

It has now been established that Kendall's tau can be expressed either in terms of probabilities of concordant and discordant pairs or directly through the copula of the variables, while remaining invariant to their marginal distributions. This makes it a robust measure of the overall dependence structure between variables, reflecting how they move together. However, Kendall's tau primarily captures the general tendency of association and does not provide any information about tail behaviour. To address this limitation, it is therefore useful to also introduce tail dependence.

2.2.3 Tail Dependence

Tail Dependence is a difference type of dependence measure compared to Spearman's rho and Kendall's tau, as it assesses the strength of association between a pair of variables in the tail of their distributions. One way of defining tail dependence is in terms of the pairs' marginal distribution functions.

Definition 2.2.7. Tail Dependence
Let X_1 and X_2 be two continuous random variables with marginal distribution functions F_1 and F_2. The *Lower-Tail Dependence Coefficient* of the pair (X_1, X_2), $\lambda_L \in [0,1]$, is defined as

$$\lambda_L(X_1, X_2) = \lim_{q \downarrow 0} \mathbb{P}\left(X_2 \leq F_2^{-1}(q) \mid X_1 \leq F_1^{-1}(q)\right),$$

while the *Upper-Tail Dependence Coefficient*, $\lambda_U \in [0,1]$, is defined as

$$\lambda_U(X_1, X_2) = \lim_{q \uparrow 1} \mathbb{P}\left(X_2 > F_2^{-1}(q) \mid X_1 > F_1^{-1}(q)\right),$$

with $q \in [0,1]$, if the limits exist.

Hence, the lower tail dependence coefficient, $\lambda_L(X_1, X_2)$, quantifies the limiting probability that X_2 takes on an extremely low value given that X_1 is already very low, while the upper tail dependence coefficient, $\lambda_U(X_1, X_2)$, captures the limiting probability that X_2 takes on an extremely high value given that X_1 is already high. In terms of interpretation, values close to 1 indicate a strong tail dependence, meaning that extreme events tend to occur together, while values close to 0 indicate a weak tail dependence, where extreme events are unlikely to coincide.

While Definition 2.2.7 (Tail Dependence) provides a definition for tail dependence in terms of marginal distribution functions of the variables, it is also possible to define it in terms of the copula of the pair (X_1, X_2).

Theorem 2.2.8. Copula Representation of Tail Dependence
Let X_1 and X_2 be two continuous random variables with marginal distribution functions F_1 and F_2 and let C be the copula of (X_1, X_2). The *Lower-Tail Dependence Coefficient* of the pair (X_1, X_2), $\lambda_L \in [0,1]$, is defined as

$$\lambda_L(X_1, X_2) = \lim_{q \downarrow 0} \frac{C(q,q)}{q},$$

while the *Upper-Tail Dependence Coefficient*, $\lambda_U \in [0,1]$, is defined as

$$\lambda_U(X_1, X_2) = \lim_{q \uparrow 1} \frac{1 - 2q + C(q,q)}{1-q},$$

where $q \in [0,1]$, if the limits exist.

Proof. In order to prove Theorem 2.2.8 (Copula Representation of Tail Dependence), begin with the definition of the lower-tail dependence coefficient from Definition 2.2.7 (Tail

Dependence), which for $q \in [0,1]$ is given as

$$\lambda_L(X_1, X_2) = \lim_{q \downarrow 0} \mathbb{P}\left(X_2 \leq F_2^{-1}(q) \mid X_1 \leq F_1^{-1}(q)\right), \tag{2.4}$$

Then by the definition of conditional probability,

$$\mathbb{P}\left(X_2 \leq F_2^{-1}(q) \mid X_1 \leq F_1^{-1}(q)\right) = \frac{\mathbb{P}\left(X_2 \leq F_2^{-1}(q), X_1 \leq F_1^{-1}(q)\right)}{\mathbb{P}\left(X_1 \leq F_1^{-1}(q)\right)} \tag{2.5}$$

Now since X_1 and X_2 are both continuous, their marginal distribution functions, F_1 and F_2, are also continuous. It then follows from Theorem 2.1.4 (Sklar's Theorem) that the joint probability in the numerator of (2.5) can be expressed in terms of the copula C. In particular,

$$\mathbb{P}\left(X_2 \leq F_2^{-1}(q), X_1 \leq F_1^{-1}(q)\right) = C\left(F\left(F_1^{-1}(q)\right), F_2(F_2^{-1}(q))\right) = C(q,q),$$

where the last equality follows Definition A.0.4 (Quasi-Inverse).

Similarly, the denominator can be written as

$$\mathbb{P}\left(X_1 \leq F_1^{-1}(q)\right) = F_1\left(F_1^{-1}(q)\right) = q.$$

Combining these expressions, the lower-tail dependence coefficient in (2.4) can be written as

$$\lambda_L(X_1, X_2) = \lim_{q \downarrow 0} \frac{C(q,q)}{q}. \tag{2.6}$$

The proof for the upper-tail dependence coefficient follows a similar approach, where it can be expressed as

$$\lambda_U(X_1, X_2) = \lim_{q \uparrow 1} \mathbb{P}\left(X_2 > F_2^{-1}(q) \mid X_1 > F_1^{-1}(q)\right) = \lim_{q \uparrow 1} \frac{\mathbb{P}\left(X_2 > F_2^{-1}(q), X_1 > F_1^{-1}(q)\right)}{\mathbb{P}\left(X_1 > F_1^{-1}(q)\right)},$$

for $q \in [0,1]$ according to Definition 2.2.7 (Tail Dependence) and the definition of conditional probability.

Then using the laws of probability, the numerator can be written as

$$\begin{aligned}\mathbb{P}\left(X_2 > F_2^{-1}(q), X_1 > F_1^{-1}(q)\right) &= 1 - \mathbb{P}\left(X_1 \leq F_1^{-1}(q)\right) - \mathbb{P}\left(X_2 \leq F_2^{-1}(q)\right) \\ &\quad + \mathbb{P}\left(X_1 \leq F_1^{-1}(q), X_2 \leq F_2^{-1}(q)\right) \\ &= 1 - q - q + C(q,q),\end{aligned}$$

and using the complement rule, the denominator becomes

$$\mathbb{P}\left(X_1 > F_1^{-1}(q)\right) = 1 - \mathbb{P}\left(X_1 \leq F_1^{-1}(q)\right) = 1 - q.$$

Now combining these expressions, the upper-tail dependence coefficient can be written as

$$\lambda_U(X_1, X_2) = \lim_{q \uparrow 1} \frac{1 - 2q + C(q, q)}{1 - q}, \tag{2.7}$$

which completes the proof. □

In summary, all dependence measures introduced in this section can be expressed either in terms of the joint probability of the pair (X_1, X_2) or equivalently in terms of their copula. Spearman's rho and Kendall's tau captures the overall dependence between the variables, while tail dependence focuses specifically on the co-movements in the tails of the distributions, so together, these measures provide a complementary insight into the dependence structure.

2.3 Families of Copulas

As discussed in Section 2.2 (Dependence Measures), choosing an appropriate copula family when modelling the dependence structure between random variables is crucial, as different copula families impose specific assumptions on the dependence structure. The two main classes of copula families, included in this project, are the *Elliptical* and the *Archimedean* copulas, which differ in their flexibility and their ability to capture asymmetric dependence and extreme co-movements.

2.3.1 Elliptical Copulas

Elliptical copulas are derived from multivariate elliptical distributions, such as the multivariate Gaussian or Student's t distributions, which are both included in this section.

The key idea behind elliptical copulas is to isolate the dependence structure of an elliptical distribution by transforming the marginals into uniform variables. Specifically, given a d-dimensional elliptical distribution function F with marginal distributions $F_1, \ldots, F_d$, each marginal is transformed using Theorem A.0.3 (Probability Integral Transform), such that $F_i(x_i) = u_i$, where $U_i \sim \text{Uniform}(0, 1)$, for all $i \in \{1, \ldots, d\}$. Once the margins are uniform, Theorem 2.1.4 (Sklar's Theorem) can be applied to construct the elliptical copula by evaluating the joint distribution at the inverse marginal distribution functions of these uniform variables, i.e.,

$$C(u_1, \ldots, u_d) = F\left(F_1^{-1}(u_1), \ldots, F_d^{-1}(u_d)\right).$$

This formula illustrates how the elliptical copula can be used to separate the dependence structure form the marginal distributions, providing a flexible framework for modelling the joint behaviour, while preserving the behaviour of each individual variable.

The Gaussian Copula

The most well-known elliptical copula is the *Gaussian copula*, which is constructed from the multivariate normal distribution and is defined the following definition.

Definition 2.3.1. Gaussian Copula
Let $d \geq 2$ and let $\Phi_P : \mathbb{R}^d \to [0,1]$ denote the standardized d-dimensional normal distribution function with correlation matrix $P \in \mathbb{R}^{d \times d}$, where P is symmetric, positive definite and satisfies $\text{diag}(P) = (1, \ldots, 1)^T$. The *Gaussian Copula* is then defined as

$$C_P^{\mathbf{Ga}}(u_1, \ldots, u_d; P) = \Phi_P\left(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_d)\right),$$

for all $(u_1, \ldots, u_d) \in [0,1]^d$, with the corresponding copula density

$$c_P^{\mathbf{Ga}}(u_1, \ldots, u_d; P) = |P|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\boldsymbol{\varsigma}^T\left(P^{-1} - I_d\right)\boldsymbol{\varsigma}\right),$$

where $\boldsymbol{\varsigma} = \left(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_d)\right)^T$.

So basically, the Gaussian copula captures the dependence structure of a multivariate normal distribution while allowing the marginals to follow any continuous distribution. (See Cherubini et al. [2004])

The dependence structure is entirely characterized by the correlation matrix $P = [\rho_{i,j}]$, where $\rho_{i,j} \in [-1,1]$ for $i, j \in \{1, \ldots, d\}$ with $i \neq j$ describes the pairwise linear dependence between the variables. In particular, values of $\rho_{i,j}$ close to 1 or -1 indicate a strong positive or negative dependence, respectively, while values of $\rho_{i,j} = 0$ indicate independence between the two variables. (See Czado [2019])

In the special case of two variables, i.e., the bivariate Gaussian copula, the dependence is fully described by $\rho_{1,2} = \rho_{2,1}$ as the correlation matrix is symmetric. This dependence structure can be further analysed by implementing a dependence measure, such as Kendall's tau, $\tau\left(C_P^{\mathbf{Ga}}\right)$, or Spearman's rho, $\rho_s\left(C_P^{\mathbf{Ga}}\right)$, which are given as

$$\tau\left(C_P^{\mathbf{Ga}}\right) = \frac{2}{\pi}\arcsin(\rho_{1,2}), \quad \text{and} \quad \rho_s\left(C_P^{\mathbf{Ga}}\right) = \frac{6}{\pi}\arcsin\left(\frac{\rho_{1,2}}{2}\right)$$

or by the tail dependence coefficients, which is expressed as

$$\lambda_L\left(C_P^{\mathbf{Ga}}\right) = \lambda_U\left(C_P^{\mathbf{Ga}}\right) = 0,$$

Notice, that the Gaussian copula exhibits no tail dependence, meaning that extreme values in one variable do not increase the probability of extreme values in the other variable. This might be a limitation in many financial applications, where extreme values often occur

simultaneously and need to be accounted for properly. An alternative elliptical copula is therefore the *Student's t copula*, which allows for tail dependence and is presented in the next section. (See Cherubini et al. [2004])

The Student's t Copula

The Student's t Copula is a natural extension of the Gaussian copula within the class of elliptical copulas. It is derived from the multivariate Student's t distribution, whose heavier tails allow for greater flexibility, as it is able to capture extreme co-movements in the tails. A formal definition of the Student's t Copula is provided below.

Definition 2.3.2. Student's t Copula
Let $d \geq 2$ and let $\mathbf{t}_{P,\nu} : \mathbb{R}^d \to [0,1]$ denote the standardized d-dimensional Student's t distribution function with $\nu > 0$ degrees of freedom and correlation matrix $P \in \mathbb{R}^{d\times d}$, where P is symmetric, positive definite and satisfies $\text{diag}(P) = (1,\dots,1)^T$. The *d-Dimensional Student's t Copula* is then defined as

$$C^{\mathbf{t}}_{P,\nu}(u_1,\dots,u_d;P,\nu) = \mathbf{t}_{P,\nu}\left(t_\nu^{-1}(u_1),\dots,t_\nu^{-1}(u_d)\right),$$

for all $(u_1,\dots,u_d) \in [0,1]^d$, with the corresponding copula density is

$$c^{\mathbf{t}}_{P,\nu}(u_1,\dots,u_d;P) = |P|^{-\frac{1}{2}} \frac{\Gamma\left(\frac{\nu+d}{2}\right)\left(\Gamma\left(\frac{\nu}{2}\right)\right)^d\left(1+\frac{1}{\nu}\boldsymbol{\varsigma}^T P^{-1}\boldsymbol{\varsigma}\right)^{-\frac{\nu+d}{2}}}{\left(\Gamma\left(\frac{\nu+1}{2}\right)\right)^d\Gamma\left(\frac{\nu}{2}\right)\prod_{i=1}^d\left(1+\frac{\varsigma_i^2}{\nu}\right)^{-\frac{\nu+1}{2}}}$$

where $\boldsymbol{\varsigma} = \left(t_\nu^{-1}(u_1),\dots,t_\nu^{-1}(u_d)\right)^T$ and $t_\nu^{-1} : \mathbb{R} \to [0,1]$ denotes the inverse of the univariate standard student's t distribution with ν degrees of freedom.

In contrast to the Gaussian copula, the dependence structure of the Student's t Copula is characterised by both the correlation matrix P, and the degrees of freedom parameter ν. Here, the correlation matrix P plays the same role as in the Gaussian copula Subsection 2.3.1, capturing the linear dependence between the variables, while the degrees of freedom parameter ν controls the heaviness of the tails, where smaller values of ν correspond to heavier tails and increases the probability of extreme values occurring simultaneously in the variables. Consequently, the Student's t copula provides a more flexible framework for modelling dependence in cases where extreme co-movements are present. (See Bouye et al. [2000])

Similarly to the Gaussian copula, the dependence measures for the Student's t copula depend only on the correlation parameter $\rho_{1,2} = \rho_{2,1}$, and can be expressed as

$$\tau\left(C^{\mathbf{t}}_{P,\nu}\right) = \tau\left(C^{\mathbf{Ga}}_P\right) \quad \text{and} \quad \rho_s\left(C^{\mathbf{t}}_{P,\nu}\right) = \rho_s\left(C^{\mathbf{Ga}}_P\right),$$

meaning that they are identical to corresponding measures for the Gaussian copula.

However, while the Gaussian copula exhibits no tail dependence, the Student's t-copula allows for symmetric tail dependence, as captured by the tail dependence coefficients, which are given by

$$\lambda_L\left(C^{\mathbf{t}}_{P,\nu}\right) = \lambda_U\left(C^{\mathbf{t}}_{P,\nu}\right) = 2\,t_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\rho)}{1+\rho}}\right),$$

where $t_{\nu+1}$ denotes the univariate standard student's t distribution function with $\nu+1$ degrees of freedom. (See Czado [2019])

2.3.2 Archimedean Copulas

Archimedean copulas are another widely used class of copulas, particularly known for their flexibility and their ability to model asymmetric tail dependencies. Unlike elliptical copulas, Archimedean copulas are able to capture asymmetric dependence structures, allowing for different behaviour in the lower and the upper tails. This makes them especially useful in situations where joint extreme events are more likely to occur in one tail than the other.

A key feature of Archimedean copulas is that they can be written in an explicit form. Rather than being constructed through Sklar's Theorem, as in the case of elliptical copulas, they are defined via a function, known as the *additive generator*, which is introduced in the following definition. (See Jaworski et al. [2010])

Definition 2.3.3. Additive Generator
A function $\varphi : [0,\infty[\to [0,1]$ is said to be an *Additive Generator* if

1. it is continuous,
2. it is decreasing,
3. it satisfies
$$\varphi(0) = 1, \quad \text{and} \quad \lim_{t\to\infty}\varphi(t) = 0,$$
4. it is strictly decreasing on $[0, t_0]$, where $t_0 = \inf\{t > 0 : \varphi(t) = 0\}$.

Remark. *The Pseudo-inverse of the additive generator is defined as*

$$\varphi^{(-1)}(t) = \begin{cases} \varphi^{-1}(t), & t \in]0,1], \\ t_0, & t = 0. \end{cases}$$

The additive generator and its pseudo-inverse fully determines the structure of an Archimedean copula, which is defined in the following definition.

Definition 2.3.4. Archimedean Copula
Let $d \geq 2$ and let C be a d-copula. If there exists an additive generator φ, such that

$$C(\mathbf{u}) = \varphi\left(\varphi^{(-1)}(u_1) + \cdots + \varphi^{(-1)}(u_d)\right)$$

for every $\mathbf{u} \in [0,1]^d$, then C is called a *Archimedean copula*.

Archimedean copulas include some of the most well-known copula families, such as the *Gumbel copula*, the *Clayton Copula* and the *Frank Copula*, which will be the focus of the remainder of this section. (See Durante and Sempi [2016])

The Gumbel Copula

The Gumbel copula is particularly well-suited for situations where extreme joint events are more likely to occur in the upper tail, as it effectively captures upper tail dependence. Being a Archimedean copula, it can be expressed in a simple, explicit form, which is introduced in the following definition.

Definition 2.3.5. Gumbel Copula
Let $d \geq 2$ and let $\alpha \geq 1$. A d-copula defined as

$$C_\alpha^{\mathbf{Gum}}(u_1, \ldots, u_d) = \exp\left(-\left(\sum_{i=1}^{d}(-\log(u_i))^\alpha\right)^{1/\alpha}\right) \tag{2.8}$$

for $(u_1, \ldots, u_d) \in [0,1]^d$ is referred to as a *Gumbel copula* with the Archimedean generator $\varphi_\alpha(t) = \exp\left(-t^{1/\alpha}\right)$, where $t \in [0, \infty)$.

The Gumbel copula can therefore be defined either through the explicit formula in (2.8) or by it Archimedean generator.

The parameter α governs the strength of dependence between the variables, where a higher value of α corresponds to stronger dependence and more pronounced upper tail behaviour. In the bivariate case, this dependence can be quantified using the measures introduced in Section 2.2 (Dependence Measures). For the bivariate Gumbel copula, Kendall's tau, $\tau(C_\alpha^{\mathbf{Gum}}) \in [0,1]$, is given by

$$\tau(C_\alpha^{\mathbf{Gum}}) = \frac{\alpha - 1}{\alpha},$$

whereas Spearman's rho does not admit a closed-form expression, but can be evaluated numerically using Theorem 2.2.3 (Copula Representation of Spearman's Rho). The lower-

and upper-tail dependence coefficients are expressed as

$$\lambda_L\left(C_\alpha^{\mathbf{Gum}}\right) = 0, \quad \lambda_U\left(C_\alpha^{\mathbf{Gum}}\right) = 2 - 2^{1/\alpha}.$$

Notice, that the Gumbel copula is asymmetric, exhibiting a positive dependence in the upper-tail, while the lower tail is asymptotically independent, with $\lambda_L = 0$. This means, that it is best suited for variables where extreme high values tend to occur simultaneously, and where extreme low values are unlikely to occur together. (See Czado [2019])

The Clayton Copula

As opposed to the Gumbel copula, the Clayton copula is more well-suited for modelling situations in which extreme joint events occur more frequently in the lower tail. Like other Archimedean copulas, it can be expressed either through an explicit formula or via its Archimedean generator, which are both provided in the following definition.

Definition 2.3.6. Clayton Copula
Let $d \geq 2$ and let $\alpha \geq \frac{-1}{d-1}$, where $\alpha \neq 0$. A d-copula defined as

$$C_\alpha^{\mathbf{Clay}}(u_1, \ldots, u_d) = \max\left\{\left(\sum_{i=1}^{d}(-u_i^{-\alpha} - (d-1)\right)^{-1/\alpha}, 0\right\}$$

for $(u_1, \ldots, u_d) \in [0,1]^d$, is referred to as a *Clayton copula* with the Archimedean generator $\varphi_\alpha(t) = (\max\{1 + \alpha t, 0\})^{-1/\alpha}$, where $t \in [0, \infty)$.

Again, the parameter α determines the level of dependence between the variables, with higher values of α corresponding to stronger lower-tail dependence. In the bivariate case, Kendall's tau $\tau(C_\alpha^{\mathbf{Clay}}) \in [0,1]$, is given by

$$\tau(C_\alpha^{\mathbf{Clay}}) = \frac{\alpha}{\alpha + 1},$$

while Spearman's rho does not admit a closed-form expression. The lower- and upper tail dependence coefficients are given as

$$\lambda_L\left(C_\alpha^{\mathbf{Clay}}\right) = \begin{cases} 2^{-1/\alpha}, & a > 0, \\ 0, & \alpha \in [-1, 0], \end{cases} \quad \lambda_U\left(C_\alpha^{\mathbf{Clay}}\right) = 0.$$

These two tail dependence coefficients emphasize that the Clayton copula is asymmetric, showing pronounced lower-tail dependence. It is therefore most appropriate for modelling variables where extreme low values tend to occur simultaneously. (See Durante and Sempi [2016])

The Frank Copula

Unlike the Gumbel and Clayton copulas, the Frank copula is symmetric and best suited for variables in which extreme co-movements are unlikely in either of the tails. The Frank copula is introduced in the following definition.

Definition 2.3.7. Frank Copula
Let $d \geq 2$ and let $\alpha > 0$. A d-copula defined as

$$C_\alpha^{\mathbf{Frank}}(u_1, \ldots, u_d) = -\frac{1}{\alpha} \log \left(1 + \frac{\prod_{t=1}^{d} \left(\mathrm{e}^{-\alpha u_i} - 1 \right)}{\left(\mathrm{e}^{-\alpha} - 1 \right)^{d-1}} \right)$$

for $(u_1, \ldots, u_d) \in [0,1]^d$, is referred to as a *Frank copula* with the Archimedean generator $\varphi_\alpha(t) = \frac{1}{\alpha} \log \left(1 - \left(1 - \mathrm{e}^{-\alpha} \right) \mathrm{e}^{-t} \right)$, where $t \in [0, \infty)$.

The parameter α governs the overall dependence between the variables, with larger values indicating a stronger dependence.

In the bivariate case, Kendall's tau, $\tau(C_\alpha^{\mathbf{Frank}}) \in [-1, 1]$, is given by

$$\tau(C_\alpha^{\mathbf{Frank}}) = 1 - \frac{4}{\alpha} \left(1 - D_1(\alpha) \right)$$

and Spearman's rho is given by

$$\rho_s(C_\alpha^{\mathbf{Frank}}) = 1 - \frac{12}{\alpha} \left(D_1(\alpha) - D_2(\alpha) \right),$$

where D_n is the *Debye function*, which for any $n \in \mathbb{N}$ is defined as

$$D_n(x) = \frac{n}{x^n} \int_0^x \frac{t^n}{\mathrm{e}^t - 1} \mathrm{d}t.$$

As stated in the beginning, the Frank copula exhibits no tail dependence, meaning that both tail dependence coefficients are

$$\lambda_L \left(C_\alpha^{\mathbf{Frank}} \right) = \lambda_U \left(C_\alpha^{\mathbf{Frank}} \right) = 0.$$

Consequently, the Frank copula is best suited for modelling moderate, symmetric dependence, rather then extreme co-movements. (See Durante and Sempi [2016])

Rotated Copulas

Rotated copulas can be seen as an extension of bivariate copula families, particularly within the class of Archimedean copulas. In many cases, such as for the Gumbel and Clayton copulas, Archimedean copulas exhibit asymmetric tail dependence, meaning that dependence is present in either the lower- or the upper-tail, but not in both.

This asymmetry naturally motivates the use of transformations, which provide more flexibility and allow the dependence structure to be adapted to different tail behaviours. One such transformation is rotations, which provide a simple method for preserving the copula structure, as defined in Definition 2.1.3 (Copula), while also modifying the way dependence is allocated between the tails. The notation of rotated copulas is formally introduced in the following definition.

Definition 2.3.8. Rotated Copulas
Let C be a 2-dimensional Archimedean copula. The 90, 180 and 270 degree rotations of C are then given as

$$\begin{aligned} C^{90}(u_1, u_2) &= u_2 - C(1 - u_2, u_1), \\ C^{180}(u_1, u_2) &= u_1 + u_2 - 1 + C(1 - u_1, 1 - u_2), \\ C^{270}(u_1, u_2) &= u_1 - C(u_2, 1 - u_1) \end{aligned}$$

respectively, where $(u_1, u_2) \in [0, 1]^2$.

Basically, a rotated copula can be seen as a transformation of the original copula, where the tail dependence structure is altered. In particular, a 180 degree rotated copula, C^{180}, corresponds to an interchange of the lower- and upper-tail dependence. For instance, the 180 degree rotated Gumbel copula, also referred to as the *Survival Gumbel copula*, is a reflection of the Gumbel copula, exhibiting lower tail dependence instead of upper tail dependence. This means that the Survival Gumbel copula can be expressed as

$$C_\alpha^{\mathbf{Gum},180} = u_1 + u_2 - 1 + C_\alpha^{\mathbf{Gum}}(1 - u_1, 1 - u_2),$$

with the tail dependence coefficients

$$\lambda_L \left(C_\alpha^{\mathbf{Gum},180} \right) = 2 - 2^{1/\alpha}, \quad \lambda_U \left(C_\alpha^{\mathbf{Gum},180} \right) = 0,$$

highlighting that the rotation do not change the copula structure, but only the direction of the tail dependence. Rotated copulas are therefore a useful tool for modelling asymmetric dependence structures within the class of Archimedean copulas. (See Czado [2019])

3 Vine Copulas

While copulas provide a flexible method for modelling dependence between random variables by separating their individual behaviour from the dependence structure, limitations tend to arise as the dimensionality increases. In particular, many copula families impose restrictive assumptions on the dependence structure. For instance, the Archimedean copulas, presented in Subsection 2.3.2 (Archimedean Copulas), rely on a single parameter, α, to describe the dependence between all variables, implying that the strength of dependence must be the same across all variable pairs. On the other hand, elliptical copulas, such as those presented in Subsection 2.3.1 (Elliptical Copulas), allow for different levels of dependence between the variables by incorporating a correlation matrix P. However, these copulas still exhibit some limitations, as the Gaussian copula cannot capture tail dependence, while the Student's t copula relies on the degrees of freedom parameter ν to model the tail dependence, enforcing the same type of tail behaviour across all variables.

To overcome these limitations, *vine copulas* are introduced as a far more flexible approach in which multivariate copulas are constructed from a sequence of bivariate copulas by implementing a technique called the *pair-copula constructions (PCC)*. This technique will be presented in the following sections together with the construction of vine and R-vine copulas, which provide a framework for organizing these bivariate building blocks.

3.1 Pair-Copula Construction

Pair-copula constructions (PCCs) provide a flexible framework for modelling complex dependence structures in high dimensions by decomposing a joint density into a sequence of bivariate copulas and marginal densities. Using such a decomposition provides greater flexibility, as each bivariate copula can be chosen independently, allowing the dependence between pairs of variables to reflect their individual relationships rather than enforcing the same dependence structure across all pairs. In order to implement PCC it is assumed that the joint distribution, F, of a random vector $\mathbf{X}$ is absolutely continuous. This guarantees the existence of a joint density, f, which forms the basis for the PCC decomposition.

For an intuitive understanding of PCC, consider a random vector $\mathbf{X} = (X_1, \dots, X_d)$ with marginal densities $f_1(x_1), \dots, f_d(x_d)$ and joint density $f(x_1, \dots, x_d)$. The joint density can then be decomposed into a product of conditional densities using the chain rule, such that

$$f(x_1, \dots, x_d) = f(x_d|x_1, \dots, x_{d-1}) \cdot f(x_{d-1}|x_1, \dots, x_{d-2}) \cdots f(x_2|x_1) \cdot f_1(x_1).$$

This decomposition is the starting point of PCC, as it allows the joint density to be expressed as a sequence of conditional densities. The idea is then to express each conditional

density in terms of bivariate copulas, thereby reducing the modelling of high-dimensional dependence structures into a collection of bivariate copulas, which are considerably easier to estimate and interpret.

To exemplify this approach, consider a bivariate case with two random variables X_1 and X_2 with marginal density functions f_1 and f_2. Their joint density can then be expressed as

$$f(x_1, x_2) = f(x_2|x_1) \cdot f_1(x_1). \tag{3.1}$$

Using (2.2), the same joint density can be expressed using a bivariate copula $c_{1,2}$, such that

$$f(x_1, x_2) = c_{1,2}\left(F_1(x_1), F_2(x_2)\right) \cdot f_1(x_1) \cdot f_2(x_2), \tag{3.2}$$

where F_1 and F_2 denote the marginal distributions of X_1 and X_2, respectively. Using (3.2), it is possible to rewrite the conditional density, $f(x_1|x_2)$, in (3.1) as

$$\begin{aligned} f(x_2|x_1) &= \frac{f(x_1, x_2)}{f_1(x_1)} \\ &= \frac{c_{1,2}\left(F_1(x_1), F_2(x_2)\right) \cdot f_1(x_1) \cdot f_2(x_2)}{f_1(x_1)} \\ &= c_{1,2}\left(F_1(x_1), F_2(x_2)\right) \cdot f_2(x_2). \end{aligned} \tag{3.3}$$

This shows how the conditional density can be expressed by the bivariate copula density $c_{1,2}$ and the marginal density f_2, and thereby how the dependence structure can be separated from the marginal behaviour, which is the essence of PCC. (See Jaworski et al. [2010])

To understand how this extends to higher dimensions, consider the three-dimensional case with the random variables X_1, X_2 and X_3 and their marginal densities f_1, f_2 and f_3. The joint density can be expressed as

$$f(x_1, x_2, x_3) = f(x_3|x_1, x_2) \cdot f(x_2|x_1) \cdot f_{x_1}(x_1).$$

In contrast to the two-dimensional case, this expression contains the conditional density, $f(x_3|x_1, x_2)$, which involves multiple conditioning variables. This is where the recursive structure becomes apparent, because instead of modelling $f(x_3|x_1, x_2)$ directly, the idea of PCC is to decompose it into simpler components, such that

$$f(x_3|x_1, x_2) = c_{1,3|2}\left(F(x_1|x_2), F(x_3|x_2)\right) \cdot f(x_3|x_2),$$

where

$$f(x_3|x_2) = c_{2,3}\left(F_2(x_2), F_3(x_3)\right) \cdot f_3(x_3).$$

Using the expression for $f(x_2|x_1)$, derived in (3.3), the joint density can be written as

$$\begin{aligned} f(x_1, x_2, x_3) = c_{1,3|2}\big(F(x_1|x_2), F(x_3|x_2)\big) \cdot c_{2,3}\big(F_2(x_2), F_3(x_3)\big) \\ \cdot c_{1,2}\big(F_1(x_1), F_2(x_2)\big) \cdot f_3(x_3) \cdot f_2(x_2) \cdot f_1(x_1). \end{aligned}$$

This decomposition illustrates the idea of PCC in higher dimensions, as the joint density is expressed as a product of marginal densities and bivariate copulas. The dependence structure is therefore fully captured through the copulas $c_{1,2}$ $c_{2,3}$ and $c_{1,3|2}$, rather than through a single copula.

An important note is that this decomposition is not unique. In particular, the representation of the conditional density $f(x_3|x_1, x_2)$ can be decomposed in different ways depending on the choice of the conditioning variable in the copula density. For example, an alternative decomposition of $f(x_3|x_1, x_2)$ is

$$f(x_3|x_1, x_2) = c_{2,3|1}\left(F(x_2|x_1), F(x_3|x_1)\right) \cdot f(x_3|x_1),$$

where the role of the conditioning variable is shifted from X_2 to X_1, resulting in a different but equivalent decomposition.

Formalizing the structure of PCC in a d-dimensional setting, the conditional density of X_i given a set of variables $X_v = \left(X_{v_1}, \ldots, X_{v_{d-1}}\right)$ can be written as

$$f(x_i|x_v) = c_{i,v_j|v_{-j}}\left(F(x_i|x_{v_{-j}}), F(x_{v_j}|x_{v_{-j}})\right) \cdot f(x_i|x_{v_{-j}}), \tag{3.4}$$

for any $i \in \{1, \ldots, d\}$, where $v_j \in v$ and $v_{-j} = v \setminus \{v_j\}$. The choice of v_j is arbitrary, which reflects the non-uniqueness of the decomposition.

The copula densities appearing in (3.4) can be expressed as

$$c_{i,v_j|v_{-j}}\left(F\left(x_i|x_{v_{-j}}\right), F\left(x_{v_j}|x_{v_{-j}}\right)\right) = \frac{\partial^2 C_{i,v_j|v_{-j}}\left(F\left(x_i|x_{v_{-j}}\right), F\left(x_{v_j}|x_{v_{-j}}\right)\right)}{\partial F\left(x_i|x_{v_{-j}}\right) \partial F\left(x_{v_j}|x_{v_{-j}}\right)},$$

whereas the corresponding conditional distribution functions can be obtained by

$$F(x_i|x_v) = \frac{\partial C_{i,v_j|v_{-j}} \left(F\left(x_i|x_{v_{-j}}\right), F\left(x_{v_j}|x_{v_{-j}}\right)\right)}{\partial F\left(x_{v_j}|x_{v_{-j}}\right)},$$

where $C_{i,v_j|v_{-j}}$ denotes the underlying bivariate copula. (See Czado [2019])

PCC therefore provides a flexible approach for modelling high-dimensional distributions through bivariate copulas. However, it does not specify how these bivariate copulas should be organised. This motivates the introduction of vine copulas, which provide a hierarchical structure for arranging these bivariate copulas.

3.2 Vine Structures

In order to construct a vine copula, it is essential to understand the underlying structure of *vines*. Vines are graphical models that provide a systematic framework for organising the bivariate building blocks obtained from the PCC introduced in Section 3.1. However, before presenting a formal definition of vines, some fundamental concepts within graph theory must be introduced. First, is the definition of a *graph*, which formalizes the relationships between *nodes* and *edges*.

Definition 3.2.1. Graph
A *graph* is defined as a pair (N, E), where

- N is a non-empty set of *nodes*,
- $E \subseteq \{\{a, b\} : a, b \in N, a \neq b\}$ is a set of *edges*,
- $d(v)$ is the number of nodes connected to the node $v \in N$ by an edge and is referred to as the *degree* of v.

The definition of a graph provided in Definition 3.2.1 is typically referred to as an *undirected* graph, meaning that the connection between two nodes does not have any specific direction. In contrast, if the connection between the nodes have a clear direction, the graph is referred to as a *directed* graph. (See Czado [2019])

Building on to this framework, it is natural to consider how multiple nodes can be connected through a sequence of edges. This concept is referred to as a *path* and is formally defined in the following definition together with a *cycle*.

Definition 3.2.2. Path and Cycle
Let $k \in \mathbb{N}_0$, and consider the node set $N = \{v_0, v_1, \ldots, v_k\}$ and the edge set $E = \{\{v_0, v_1\}, \{v_1, v_2\}, \ldots, \{v_{k-1}, v_k\}\}$. A *path* is defined as a graph $P = (N, E)$ and a *cycle* is defined as a path where $v_0 = v_k$.

A path can thereby be interpreted as a route through the graph, where the nodes are connected sequentially by the edges, while a cycle represents a path that starts and ends at the same node. The presence or absence of such a cycle is important when characterizing different types of graphs. In particular, a *tree* is a special type of graph with no cycles. (See Czado [2019])

Definition 3.2.3. Tree
A *tree* is a graph $T = (N, E)$ in which any two nodes $a, b \in N$ are connected by a unique path.

So basically, a tree is *connected*, meaning that every node can be reached from any other node through a sequence of edges, and also that it contains no cycles. This acyclic property of trees provides the foundation for the definition of a *vine*, presented below.

Definition 3.2.4. Vine
Let $d, m \in \mathbb{N}$, with $d \geq 2$. A set of trees $\mathcal{V} = (T_1, \ldots, T_m)$ is called a *vine* on d elements if

- T_1 is a tree with node set $N_1 = \{v_1, \ldots, v_d\}$ and edge set E_1.
- for $i = 2, \ldots, m$, the tree T_i has a node set N_i and an edge set E_i, where

$$N_i \subset N_1 \cup E_1 \cup \ldots \cup E_{i-1}.$$

Essentially, a vine is a hierarchical sequence of trees, where the edges of T_i serve as the nodes of T_{i+1} for all $i = 2, \ldots, m$. This recursive structure allows complex multivariate dependencies to be decomposed into bivariate building blocks, which may be conditioned on other variables. It also provides a graphical representation of the dependence structure in the joint distribution, making it easier to specify and interpret the relationships between variables. Two commonly used vine structures are the *Canonical vine (C-vine)*, where one variable is chosen as the central node connected to all other nodes, and the *Drawable vine (D-vine)*, in which the nodes are arranged in a single line structure. Both of these vine structures are presented in greater detail in Appendix A.1. The focus of this thesis is the *Regular vine (R-vine)*, which generalizes both C-vines and D-vines by allowing a more flexible tree structure.

Definition 3.2.5. Regular Vine (R-Vine)
Let $d \geq 2$. A vine $\mathcal{V} = (T_1, \ldots, T_{d-1})$ is called a *regular vine (R-vine)* on d elements if

- the tree T_1 has node set $N_1 = \{v_1, \ldots, v_d\}$ and edge set E_1.
- for $j = 2, \ldots, d-1$
  - the node set of the tree T_j is $N_j = E_{j-1}$.
  - the *proximity condition* holds, i.e., any two nodes $\{a, b\} \in E_j$ satisfy $|a \cap b| = 1$.

Remark. *The notation $|a \cap b|$ denotes the cardinality of the set $a \cap b$, i.e., the number of shared elements between the two nodes.*

In practice, an R-vine is a sequence of trees, where the first tree, T_1, represents the original set of nodes, N_1, with edges representing the connecting between them. Each subsequent tree, T_j then uses the edges from the previous tree, T_{j-1}, as its nodes. An important rule for these higher-level trees is the proximity condition, which requires that any two nodes in T_j can only be connected by an edge if they share exactly one node from the previous tree. To formally describe how the nodes and edges are organized across different tree-levels, the concepts of *complete union*, *conditioning set* and *conditioned set* must be defined. (See Bedford and Cooke [2002])

Definition 3.2.6. Complete Union, Conditioning and Conditioned Sets
Let $d \geq 2$ and let $\mathcal{V} = (T_1, \ldots, T_{d-1})$ be a R-vine, then for any edge $e \in E_i$, where $i = 1, \ldots, d-1$, the *Complete Union* of e is the subset defined by

$$A_e = \begin{cases} e, & i = 1 \\ \left\{ n \in N_1 \;\middle|\; \exists\, e_1 \in E_1, \ldots, e_{i-1} \in E_{i-1}, \text{ such that } n \in e_1 \in \cdots \in e_{i-1} \in e \right\}, & i \geq 2. \end{cases}$$

For $e = \{a, b\}$ the *Conditioning Set* of e is defined as

$$D_e = A_a \cap A_b,$$

and the *Conditioned Sets* of e are defined by

$$\mathcal{C}_{e,a} = A_a \setminus D_e, \quad \mathcal{C}_{e,b} = A_b \setminus D_e,$$

where $\mathcal{C}_e = \mathcal{C}_{e,a} \cup \mathcal{C}_{e,b}$

The sets defined in Definition 3.2.6 describe the structure of an edge in an R-vine, meaning that the complete union identifies all nodes that are included in the edge, the conditioning set specifies the nodes shared between the connected nodes, and the conditioned sets contain the nodes that are unique to each. Using these sets, any edge $e = \{a, b\}$ can be compactly written as $e = (\mathcal{C}_{e,a}, \mathcal{C}_{e,b} \mid D_e)$, clearly distinguishing the nodes in the condi-

tioned sets from those in the conditioning set. To give a further understanding of these concepts an example see Example 3.2.7.

Example 3.2.7. Consider an R-vine on $d = 6$ elements, whose structure is illustrated in Figure 3.1.

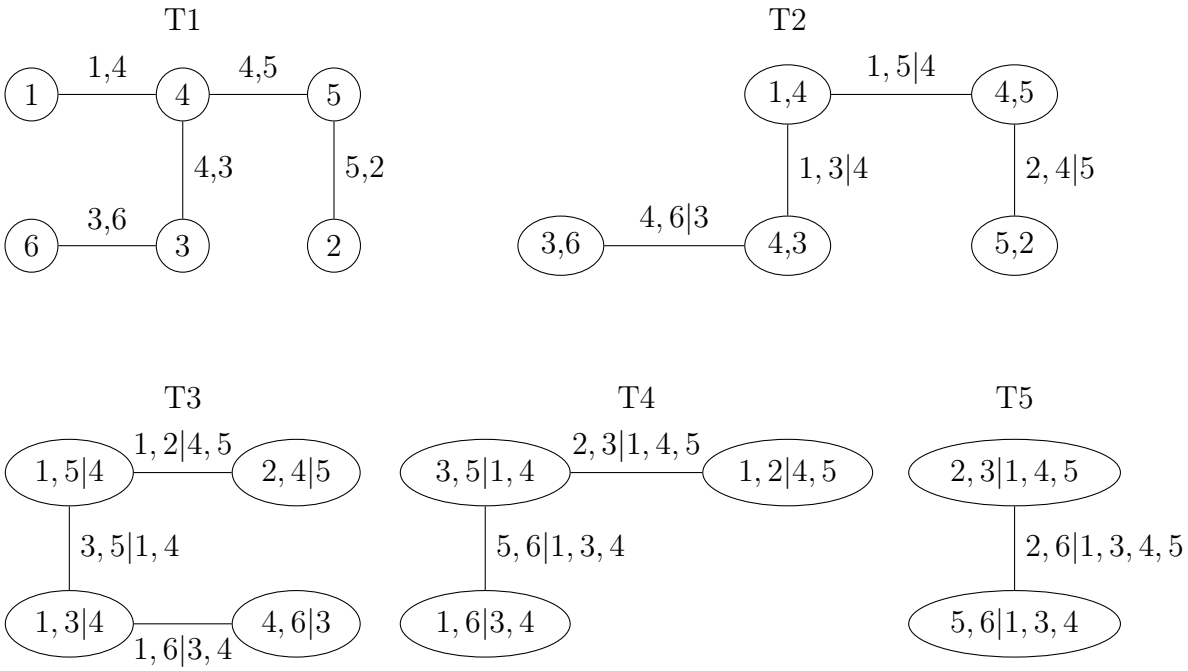

Figure 3.1: Example of an R-vine.

For the first tree, T_1, the node set is $N_1 = \{1, 2, 3, 4, 5, 6\}$ and the edge set is

$$E_1 = \{\{1, 4\}, \{4, 5\}, \{3, 6\}, \{4, 3\}, \{5, 2\}\}.$$

These edges in T_1 form the foundation of the R-vine structure, because once they are fixed, all edges in the subsequent trees are systematically determined according to the proximity condition in Definition 3.2.5 (Regular Vine (R-Vine)).

Continuing to the second tree, T_2, its nodes is defined as the edges from T_1, such that $N_2 = E_1$ and it edge set is

$$E_2 = \{\{\{1, 4\}, \{4, 5\}\}, \{\{4, 5\}, \{5, 2\}\}, \{\{4, 5\}, \{4, 3\}\}, \{\{4, 3\}, \{3, 6\}\}\}.$$

Now consider as an example the edge $e = \{\{1, 4\}, \{4, 5\}\} \in E_2$, the complete union is then defined as the set $A_e = \{1, 4, 5\}$, whereas the conditioning set is $D_e = \{4\}$ and the conditioned set is $\mathcal{C}_e = \{1, 5\}$.

In the third tree, T_3, the node set is $N_3 = E_2$ and the edge set consists of pairs of nodes that satisfy the proximity condition. For example, the edge $\{\{1,5|4\},\{2,4|5\}\}$ is included in E_3 because these two nodes share exactly one node from T_2, namely $\{4,5\}$. In contrast, there is no edge between the nodes $\{2,4|5\}$ and $\{4,6|3\}$ as they do not share a node from T_2.

3.3 R-Vine Copulas

Having established the graphical structure of R-vines, it is possible decompose a multivariate distribution into a sequence of bivariate building blocks using the PCC framework introduced in Section 3.1. Each of these bivariate building blocks corresponds to an edge in the R-vine and can be modelled by a bivariate copula, capturing the dependence between the variables associated with that edge. Thus, by arranging these bivariate copulas according to the hierarchical R-vine structure, the resulting multivariate distribution can be formally described as an *R-vine specification*. (See Bedford and Cooke [2002])

Definition 3.3.1. R-Vine Specification
Let $d \geq 2$. The d-dimensional distribution function F of a random vector $\mathbf{X} = (X_1, \ldots, X_d)$ is said to admit an *R-vine specification* if there exists a triplet $(\mathcal{F}, \mathcal{V}, \mathcal{B})$ such that

1. $\mathcal{F} = (F_1, \ldots, F_d)$ is a vector of continuous invertible marginal distribution functions of $X_1, \ldots, X_d$.
2. $\mathcal{V} = (T_1, \ldots, T_{d-1})$ is an R-vine on d elements.
3. $\mathcal{B} = \{C_e \mid e \in E_i\,;\, i = 1, \ldots, d-1\}$ is a set of bivariate copulas, such that for each edge $e = \{a, b\} \in E_i$:
   - C_e is the copula associated with the conditional distribution of $X_{\mathcal{C}_{e,a}}$ and $X_{\mathcal{C}_{e,b}}$ given $\mathbf{X}_{D_e} = \mathbf{x}_{D_e}$,
   - C_e does not depend on the specific value of $\mathbf{x}_{D_e}$.

As stated in Definition 3.3.1, a multivariate distribution function F of a random vector $\mathbf{X} = (X_1, \ldots, X_d)$ can be represented by an R-vine structure through the triplet $(\mathcal{R}, \mathcal{V}, \mathcal{B})$, if X_i has a continuous, invertible marginal distribution for all $i = 1, \ldots, d$, the vine $\mathcal{V}$ is a valid R-vine, and the set $\mathcal{B}$ consists of bivariate copulas associated with the edges of the vine.

In this framework, each edge $e = \{a, b\}$ is associated with a bivariate copula C_e, often referred to as a *pair copula*, capturing the dependence between the two sets of variables $\mathbf{X}_{\mathcal{C}_{e,a}}$ and $\mathbf{X}_{\mathcal{C}_{e,b}}$ conditioned on the set $\mathbf{X}_{D_e}$, determined by the lower-level trees. For example, taking the edge $e = \{\{1,5|4\},\{2,4|5\}\}$ from the third tree, T_3, in Example 3.2.7,

then $\mathbf{X}_{\mathcal{C}_{e,a}} = X_1$, $\mathbf{X}_{\mathcal{C}_{e,b}} = X_2$ and $\mathbf{X}_{D_e} = \{X_4, X_5\}$, so that the corresponding pair copula can be written as $C_{1,2|4,5}$. An commonly imposed assumption of each pair copula, C_e, is that it does not depend on the specific values of the conditioning variables $\mathbf{x}_{D_e}$, which is referred to as the *simplifying assumption*. With these notations, each pair copula can be written as $C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$, with corresponding density $c_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$.

Based on Definition 3.3.1 (R-Vine Specification), it is now possible to define its corresponding density, which is expressed in the following theorem.

Theorem 3.3.2. R-vine Density
Let $d \geq 2$ and assume that the triplet $(\mathcal{F}, \mathcal{V}, \mathcal{B})$ satisfy the properties $(1) - (3)$ of Definition 3.3.1 (R-Vine Specification). Then there exists a unique d-dimensional distribution function F that realizes this R-vine copula specification with density

$$f(x_1, \ldots, x_d) = \prod_{j=1}^{d} f_j(x_j) \prod_{i=1}^{d-1} \prod_{e \in E_i} c_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e} \left(F_{\mathcal{C}_{e,a}|D_e}(\mathbf{x}_{\mathcal{C}_{e,a}} \mid \mathbf{x}_{D_e}), F_{\mathcal{C}_{e,b}|D_e}(\mathbf{x}_{\mathcal{C}_{e,b}} \mid \mathbf{x}_{D_e}) \right),$$

where $e = \{a, b\}$, f_i denotes the density for F_i for $i = 1, \ldots, d$, and $\mathbf{x}_{D_e} = \{x_i | i \in D_e\}$.

Proof. This proof is omitted, but it can be found in [Bedford and Cooke, 2001, p. 259 – 260]. □

The density provided in Theorem 3.3.2 is fully consistent with the formula for the joint density introduced in Section 3.1 (Pair-Copula Construction). In particular, the structure of the R-vine density is based on the same idea as in PCC, where a d-dimensional density is decomposed into a product of marginal densities and bivariate copulas.

The main difference is that in this R-vine framework, the bivariate copulas are organised in a systematic and hierarchical way. More precisely, each tree T_i in the R-vine $\mathcal{V} = (T_1, \ldots, T_{d-1})$ consists of an edge set E_i, where every edge $e = \{a, b\} \in E_i$ corresponds to a bivariate copula density $c_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$, Here, the conditioning set, D_e, specifies the variables that have already been accounted for in previous trees, while the conditioned sets, $\mathcal{C}_{e,a}$ and $\mathcal{C}_{e,b}$, specify the variables whose conditional dependence is modelled at the current stage. As an example, consider tree T_3 in Example 3.2.7. Here the edge connecting the nodes $\{1, 5|4\}$ and $\{1, 3|4\}$ is associated with the conditional copula density $c_{3,5|1,4}$, modelling the dependence between X_3 and X_5 given X_1 and X_4. This specific copula density can be denoted as

$$c_{3,5|1,4} \left(F_{5|1,4}(x_5|x_1, x_4), F_{3|1,4}(x_3|x_1, x_4) \right),$$

illustrating how the copula density should be interpreted within the R-vine structure. (See Dißmann et al. [2013])

One of the main difficulties in evaluating the R-vine density in Theorem 3.3.2 lies in computing the conditional distributions $F_{\mathcal{C}_{e,a}|D_e}(\mathbf{x}_{\mathcal{C}_{e,a}} \mid \mathbf{x}_{D_e})$ and $F_{\mathcal{C}_{e,b}|D_e}(\mathbf{x}_{\mathcal{C}_{e,b}} \mid \mathbf{x}_{D_e})$, which are not specified directly in the model, but are derived recursively, as described in Theorem 3.3.1 (Sequential Estimation).

It is now possible to define an *R-vine copula*, as a special case of the R-vine specification, where the marginals are uniformly distributed.

Definition 3.3.3. R-Vine Copula
A *R-vine copula* is an R-vine specification $(\mathcal{F}, \mathcal{V}, \mathcal{B})$, whose marginal distributions are all uniform on $[0, 1]$.

Thus, this definition implies that an R-vine copula can be constructed by applying the transformation defined in Theorem A.0.3 (Probability Integral Transform), such that the transformed variables are defined by $F_i(x_i) := u_i$ for $i = 1, \ldots, d$, where $u_i \sim \text{Uniform}(0, 1)$. The corresponding copula density is then given by

$$c(u_1, \ldots, u_d) = \prod_{i=1}^{d-1} \prod_{e \in E_i} c_{\mathcal{C}_{e,a}, \mathcal{C}_{e,b}|D_e} \left(u_{\mathcal{C}_{e,a}|D_e}, u_{\mathcal{C}_{e,b}|D_e} \right).$$

OVERVEJ OM RESTEN HER SKAL MED! something of a overgang. *R-vine matrix*, which provides a compact representation of an R-vine structure.

Definition 3.3.4. R-Vine Matrix
Let $d \geq 2$ and let $M \in \mathbb{R}^{d \times d}$ be an upper triangular matrix with entries $m_{i,j} \in \{1, \ldots, d\}$ for $i \leq j$. The matrix M is called a *R-vine matrix* if it satisfies

1. $\{m_{1,i}, \ldots, m_{i,i}\} \subset \{m_{1,j}, \ldots, m_{j,j}\}$ for all $i < j$.
2. $m_{i,i} \notin \{m_{1,i-1}, \ldots, m_{i-1,i-1}\}$.
3. For all $i = 3, \ldots, d$ and $k = 1, \ldots, i-1$, there exists indices $j < i$ and $l < j$, such that

$$\{m_{k,i}\} \cup \{m_{1,i}, \ldots, m_{k-1,i}\} = \{m_{j,j}\} \cup \{m_{1,j}, \ldots, m_{l,j}\} \text{ or}$$
$$\{m_{k,i}\} \cup \{m_{1,i}, \ldots, m_{k-1,i}\} = \{m_{l,j}\} \cup \{m_{1,j}, \ldots, m_{l-1,j}, m_{j,j}\}.$$

Condition 3. in Definition 3.3.4 is equivalent to the proximity condition from Definition 3.2.5 (Regular Vine (R-Vine)) and ensures that edges in higher tree levels are only formed between nodes sharing exactly one node from the previous tree level. This guarantees that only nodes with exactly one element in common can be connected by an edge, thereby preserving the structure of an R-vine.

The R-vine matrix is an convenient tool for representing an R-vine structure in a compact

way, as it provides a systematic way of identifying both the conditioned and conditioning sets associated with each edge. This can be shown by considering a given column k, the diagonal element $m_{k,k}$ together with an entry $m_{i,k}$, where $i < k$, then defines the conditioned set, while the entries above $m_{i,k}$ in the same column determine the corresponding conditioning set. More formally, for an edge $e \in E_j$, where $j = 1, \ldots, d-1$, the conditioned set is given by

$$C_e = \{m_{k,k}, m_{i,k}\},$$

and the conditioned set is given by

$$D_e = \{m_{i-1,k}, \ldots, m_{1,k}\},$$

where $i < k$. $D_e = \{m_{i-1,k}, \ldots, m_{1,k}\}$. Thus, each column of the R-vine matrix represents a collection of conditional pair copulas, where the dependence structure is constructed recursively from previously tree levels.

Finally, it should be noted that the R-vine matrix is not unique, meaning that many different matrices may represent the same R-vine structure. For a more practical understanding, see Example 3.3.5.

Example 3.3.5. This example is a continuation of Example 3.2.7. Given the tree structures of the R-vine in Figure 3.1, it is possible to construct multiple R-vine matrices representing this structure. Two examples of possible R-vine matrices are seen in (3.5).

$$M_1 = \begin{bmatrix} 4 & 4 & 5 & 4 & 1 & 1 \\ & 5 & 4 & 5 & 4 & 3 \\ & & 2 & 2 & 5 & 4 \\ & & & 1 & 2 & 5 \\ & & & & 3 & 2 \\ & & & & & 6 \end{bmatrix}, \quad M_2 = \begin{bmatrix} 6 & 6 & 4 & 3 & 1 & 1 \\ & 4 & 6 & 4 & 3 & 3 \\ & & 3 & 6 & 4 & 4 \\ & & & 1 & 6 & 5 \\ & & & & 5 & 6 \\ & & & & & 2 \end{bmatrix} \tag{3.5}$$

Based on these two matrices, it is possible to reconstruct the underlying R-vine structure together with its associated edges, copulas and copula densities.

As an example, consider M_1 column 4, here the diagonal entry is $m_{4,4} = 1$, and the entry above is $m_{3,4} = 2$, meaning that the conditioned set is

$$C_e = \{1, 2\}.$$

The conditioning set is constructed from the entries above $M_{3,4}$ in the same column, meaning that

$$D_e = \{4, 5\}$$

This corresponds to the conditional pair copula $C_{1,2|4,5}$ with the corresponding density

$$c_{1,2|4,5}\left(F_{1|4,5}(x_1|x_4,x_5), F_{2|4,5}(x_2|x_4,x_5)\right).$$

3.3.1 Estimation

The practical implementation of R-vine copulas for modelling dependence structures can be described by the following three steps:

1. Determine the R-vine structure.
2. Select a bivariate copula family for each pair copula.
3. Estimate the copula parameters.

The first step is to choose the R-vine structure, defining how the variables are connected and specifying which pairs are conditioned on others. In practice, this involves determining the edges in each tree of the R-vine, which for a d-dimensional R-vine corresponds to $\frac{d(d-1)}{2}$ edges in total. Once, the R-vine structure has been determined, a suitable bivariate copula family is selected for each pair, followed by an estimation of the associated parameters. Each of these steps is essential when constructing a model the adequately describes the dependence structure among the variables.

In principle, one way to identify the best model could be by performing steps (2) and (3) for all possible R-vine constructions. However, this quickly becomes computationally infeasible, as the number of possible constructions grows extremely rapidly with the number of variables, reaching approximately $\left(\frac{d!}{2}\right) \cdot 2^{\left(\frac{d-2}{2}\right)}$ possible combinations (See Morales-Nápoles et al. [2010]). Furthermore, some methods used in step (2) rely on visual interpretation of diagnostic plots , such as K- or Chi-plots, which further limits their practical application. (See Dißmann et al. [2013])

To address these challenges, a sequential heuristic method is proposed as an alternative. This method is described in more detail in the next section.

Sequential Estimation

A more feasible estimation approach is *sequential estimation*, where the R-vine copula is constructed stepwise, tree by tree, starting from the first tree and proceeding to higher level trees. To formalize this approach, a full R-vine copula specification can be identified by the triplet

$$\left(\mathcal{V}, \mathcal{B}(\mathcal{V}), \Theta(\mathcal{B}(\mathcal{V}))\right),$$

where $\mathcal{V} = (T_1, \ldots, T_{d-1})$ denotes the R-vine structure, $\mathcal{B}(\mathcal{V})$ is the set of all bivariate copula

families associated with the edges of the R-vine and $\Theta(\mathcal{B}(\mathcal{V}))$ denotes the corresponding set of copula parameters.

This triplet is constructed sequentially, starting with the first tree, T_1. As explained in Subsection 3.3.1 (Estimation), the first step is to determine the structure of T_1. This is done by computing the pairwise dependence between all variable pairs (X_i, X_j), for $1 \leq i < j \leq d$, using empirical Kendall's tau, defined in Subsection 2.2.2, and then selecting the maximum spanning tree based on the absolute values of these empirical Kendall's taus.[1] Formally, T_1 is given by

$$T_1 = \arg\max_{T} \sum_{(i,j)\in E_{P,1}} |\hat{\tau}\left(X_i, X_j\right)|, \tag{3.6}$$

where $E_{P,1}$ denotes all of the admissible edges in T_1 satisfying the proximity condition from Definition 3.2.5 (Regular Vine (R-Vine)), and $\hat{\tau}\left(X_i, X_j\right)$ denotes the empirical Kendall's tau between the pair $\left(X_i, X_j\right)$. Once (3.6) have been solved, each selected pair (X_i, X_j) is represented as an edge $e = (a, b)$ and stored in E_1, ensuring that the final edges in the first tree, E_1, correspond to the strongest pairwise dependencies.

The second step combines steps (2) and (3) from Subsection 3.3.1 (Estimation), meaning that both a bivariate copula family and the corresponding parameters are determined for each edge in E_1. This procedure is best explained by considering an arbitrary edge $e = (a, b) = (\mathcal{C}_{e,a}, \mathcal{C}_{e,b}) \in E_1$, where $\mathcal{B}_e$ denotes the set of candidate bivariate copula families for this edge.

For each edge e, initial *pseudo observations* are constructed form the marginal distributions as

$$u_{k,\mathcal{C}_{e,a}} = F_{\mathcal{C}_{e,a}}\left(x_{k,\mathcal{C}_{e,a}}\right) \quad \text{and} \quad u_{k,\mathcal{C}_{e,b}} = F_{\mathcal{C}_{e,b}}\left(x_{k,\mathcal{C}_{e,b}}\right),$$

for $k = 1, \ldots, n$, where $n \in \mathbb{N}$ denotes the number of observations in the dataset. Given these pseudo-observations, the parameter vector θ_B for each candidate copula family $C^B \in \mathcal{B}_e$ is then estimated by *maximum likelihood (ML)*. The ML estimator $\hat{\theta}_B$ is defined as

$$\hat{\theta}_B = \arg\max_{\theta_B} \mathcal{L}\left(\theta_B; u_e\right),$$

where the log-likelihood is given as

$$\mathcal{L}\left(\theta_B; u_e\right) = \sum_{k=1}^{n} \log\left(c^B\left(u_{k,\mathcal{C}_{e,a}}, u_{k,\mathcal{C}_{e,b}}; \theta_B\right)\right),$$

where c^B denotes the copula density associated with C^B.

[1] Note that Kendall's tau is used in this section as it is the most commonly used dependence measure in the literature, however, alternative measures such as Spearman's rho or tail dependence coefficients can also be used.

These ML estimates yield the best-fitting parameters within each family, allowing for model comparison across all candidate families C^B using a information criteria such as the *Akaike Information Criterion (AIC)* or the *Bayesian Information Criterion (BIC)*. The AIC is defined as

$$\mathrm{AIC}(C^B, \hat{\theta}_B; u_e) = -2 \sum_{k=1}^{n} \log \left(c^B \left(u_{k,\mathcal{C}_{e,a}}, u_{k,\mathcal{C}_{e,b}}; \hat{\theta}_B \right) \right) + 2k^B, \tag{3.7}$$

where $k^B = \mathrm{Dim}\left(\theta_B\right)$, whereas the BIC is defined as

$$\mathrm{BIC}(C^B, \hat{\theta}_B; u_e) = -2 \sum_{k=1}^{n} \log \left(c^B \left(u_{k,\mathcal{C}_{e,a}}, u_{k,\mathcal{C}_{e,b}}; \hat{\theta}_B \right) \right) + \log(n) k^B, \tag{3.8}$$

where $n \in \mathbb{N}$ denotes the number of observations. (See Nagler et al. [2026])

The copula family C^e and the corresponding parameter estimate $\hat{\theta}_B$ are then selected as the combination that minimizes the chosen information criteria the most across all candidate families.

For the remaining trees T_i, where $i \geq 2$, the estimation procedure changes slightly, as the pair copulas are now defined conditionally on previously estimated variables. Consequently, the initial observations are replaced by recursively updated pseudo-observations, which are obtained from previous trees.

More precisely, assuming that all pair copula parameters in tree $T_1, \ldots, T_{i-1}$ have already been estimated and collected in $\hat{\theta}(T_1, \ldots, T_{i-1})$, the pseudo-observations for an edge $e = (a, b) = (\mathcal{C}_{e,a}, \mathcal{C}_{e,b} | D_e) \in T_i$ are given by

$$\begin{aligned} u_{k,\mathcal{C}_{e,a}|D_e} &:= F_{\mathcal{C}_{e,a}|D_e} \left(u_{k,\mathcal{C}_{e,a}} | u_{k,D_e}; \hat{\theta}(T_1, \ldots, T_{i-1}) \right), \\ u_{k,\mathcal{C}_{e,b}|D_e} &:= F_{\mathcal{C}_{e,b}|D_e} \left(u_{k,\mathcal{C}_{e,b}} | u_{k,D_e}; \hat{\theta}(T_1, \ldots, T_{i-1}) \right), \end{aligned}$$

for $k = 1, \ldots, n$. These values are then used as inputs for tree T_i in the same way as in T_1. The remaining of the estimation and selection procedure is otherwise identical to that used for T_1, including the ML estimates of the copula parameters for each candidate family and the subsequent model comparison via chosen information criteria. This procedure is applied sequentially to each pair copula in the R-vine until every edge is specified.

Using this sequential estimation approach, the edges in the first tree E_1 typically capture the strongest dependence structure in the data, implying that T_1 has the largest impact on the overall model fit. Consequently, it is crucial that these dependencies are modelled accurately, particularly as many copula families exhibit very similar behaviour under weak dependence, for instance when $\tau \approx 0$, families such as Frank, Gaussian, Clayton and Gumbel act quite similar, whereas their behaviour becomes more distinct under strong dependence.

Another point is that although this approach is computationally efficient and widely used in literature (See Czado [2019] and Dißmann et al. [2013]), it does not guarantee a global optimal R-vine specification in terms of the likelihood. This is a direct consequence of its sequential construction, where each tree is constructed based on local optimums rather than the overall global optimum of the R-vine. For this reason, this procedure is commonly classified as a greedy algorithm.

An well-known realisation of this idea is the Dißmann algorithm, where the R-vine copula is constructed sequentially tree by tree. An example of this algorithm is presented in Appendix A.2 using Kendall's tau as the dependence measure.

3.3.2 Sampling

Once an R-vine copula specification $\left(\mathcal{V}, \mathcal{B}(\mathcal{V}), \Theta(\mathcal{B}(\mathcal{V}))\right)$ has been estimated, it is possible to generate random samples from the corresponding multivariate distribution.

The key idea is to generate a random vector $(u_1, \ldots, u_d)$ from the R-vine copula specification, using a vector of i.i.d. random variables $(w_1, \ldots, w_d)$, where $w_i \sim \text{Uniform}(0,1)$ for $i = 1, \ldots, d$. This is conducted by sequentially obtaining each u_j by transforming w_j through the inverse of its conditional distribution function, conditioned on all previously generated u_{j-1}. Here the random variables $(w_1, \ldots, w_d)$ are transformed as

$$
\begin{aligned}
u_1 &= w_1 \\
u_2 &= F^{-1}_{2|1}(w_2|u_1) \\
u_3 &= F^{-1}_{3|1,2}(w_3|u_1, u_2) \\
&\vdots \\
u_d &= F^{-1}_{d|1,\ldots,d-1}(w_d|u_1, \ldots, u_{d-1}),
\end{aligned}
$$

where $F^{-1}_{j|1,\ldots,j-1}$ denotes the inverse conditional distribution function of U_j given $U_1, \ldots U_{j-1}$. These conditional distributions are not specified directly, but are obtained recursively from the R-vine specification via the bivariate copulas associated with the edges of the R-vine. In particular, for each edge $e = (\mathcal{C}_{e,a}, \mathcal{C}_{e,b}|D_e)$, the corresponding conditional distribution functions are obtained from the associated bivariate copula distribution function $C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$, such that

$$
F_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b};D_e}\left(u_{\mathcal{C}_{e,a}}|u_{\mathcal{C}_{e,b}}, u_{D_e}\right) = \frac{\partial C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}\left(u_{\mathcal{C}_{e,a}}, u_{\mathcal{C}_{e,b}}\right)}{\partial u_{\mathcal{C}_{e,b}}} := h_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b};D_e}\left(u_{\mathcal{C}_{e,a}}, u_{\mathcal{C}_{e,b}}\right),
$$

or equivalently

$$
F_{\mathcal{C}_{e,b}|\mathcal{C}_{e,a};D_e}\left(u_{\mathcal{C}_{e,b}}|u_{\mathcal{C}_{e,a}}, u_{D_e}\right) = \frac{\partial C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}\left(u_{\mathcal{C}_{e,a}}, u_{\mathcal{C}_{e,b}}\right)}{\partial u_{\mathcal{C}_{e,a}}} := h_{\mathcal{C}_{e,b}|\mathcal{C}_{e,a};D_e}\left(u_{\mathcal{C}_{e,a}}, u_{\mathcal{C}_{e,b}}\right).
$$

Here, the choice of which h-function to use for each edge depends on how the conditioning is specified within the R-vine structure. Hence, sequentially applying the h-functions to all edges allows for the construction of the conditional distributions used to generate the random vector $(u_1, \ldots, u_d)$. (See Czado [2019])

4 Portfolio Optimization Framework

In this chapter, the theoretical framework for portfolio optimization is introduced, establishing the fundamental concepts of portfolio construction and the portfolio optimization problem.

A portfolio is defined as a collection of investments in different financial assets such as stocks and bonds. Bonds are often refereed to as riskless assets, as they represent a debt security issued by the government or a corporation, where the investor lends money to the issuer over a agreed upon period and interest rate. Stocks, on the other hand, are typically considered more risky assets, as they represents shares of ownership in a company. Consequently, their returns depend on both the performance of the company, as well as market conditions and geopolitical events. (See Morah [2026])

When constructing a portfolio, an investor is not only concerned with which assets to include, but also how to allocate the portfolio weights. This leads to the portfolio optimization problem, which determines how capital is allocated across the assets in the portfolio in order achieve the desired trade-off between risk and return. For this purpose, the *Mean-CVaR* portfolio optimization framework is introduced, focusing on maximizing the expected return, while also reducing the worst-case losses.

To formalize this portfolio optimization problem, consider a portfolio of $d \in \mathbb{N}$ assets, the portfolio return is then defined through the function $R : (0,1)^d \times \mathbb{R}^d \to \mathbb{R}$, given by

$$R(\omega, r) = \sum_{i=1}^{d} \omega_i r_i,$$

where $r = (r_1, \ldots, r_d) \in \mathbb{R}^d$ contains the asset returns and $\omega = (\omega_1, \ldots, \omega_d) \in (0,1)^d$ contains the portfolio weights. These weights must satisfy

$$\sum_{i=1}^{d} \omega_i = 1,$$

in order to ensure, that the entire capital is allocated across the assets in the portfolio.

To assess the risk associated with the portfolio, it is convenient to consider the portfolio losses rather then the returns. The loss of the portfolio is defined through the *Loss Function* $L : (0,1)^d \times \mathbb{R}^d \to \mathbb{R}$, given by

$$L(\omega, r) = -R(\omega, r)$$

Based on this loss function, the probability that the portfolio loss does not exceed a given threshold $\alpha \in \mathbb{R}$, is given by

$$\Psi(\omega, \alpha) = \int_{L(\omega,r)\leq\alpha} p(r)\,\mathrm{d}r,$$

where $p(r)$ denotes the probability density of the asset returns and $\Psi(\omega, \alpha) : (0,1)^d \times \mathbb{R} \to (0,1)$ denotes the distribution function of the portfolio loss associated with the weights ω.

Now that the loss function has been introduced, the *Value at Risk (VaR)*, which forms the basis for the *Conditional Value-at-Risk (CVaR)*, can be introduced. For a given confidence level $\beta \in (0,1)$, the risk measure $\text{VaR}_\beta : (0,1)^d \to \mathbb{R}$ is defined as

$$\text{VaR}_\beta(\omega) = \min\{\alpha \in \mathbb{R} : \Psi(\omega, \alpha) \geq \beta\}.$$

The value of VaR_β therefore represents the smallest threshold α, that the loss of the portfolio will not exceed with probability β. However, although VaR is able to provide a threshold, that the losses are expected to remain below with a certain confidence level, it provides no information about the magnitude of the losses exceeding this threshold.

To address this limitation, the CVaR is introduced. The risk measure $\text{CVaR} : (0,1)^d \to \mathbb{R}$ measures the expected portfolio loss in the tail of the loss distribution beyond the VaR threshold. Formally, CVaR is defined as

$$\text{CVaR}_\beta = \frac{1}{1-\beta} \int_{L(\omega,r) \geq \text{VaR}_\beta(\omega)} L(\omega, r)\, p(r)\, \mathrm{d}r.$$

This means, that CVaR is able to capture the expected severity of the losses beyond the threshold determined by VaR, making it a more comprehensive measure of tail risk in portfolio optimization. (see Tyrrell Rockafellar and Uryasev [1999])

With the CVaR defined, it is now possible to introduce the mean-CVaR optimization problem, which is formulated as

$$\max_{\omega} (1-\lambda)\, \mathbb{E}\left[R(\omega, r)\right] - \lambda \text{CVaR}_\beta(\omega), \tag{4.1}$$

where the term $\mathbb{E}\left[R(\omega, r)\right]$ denotes the expected portfolio return, and $\lambda \in (0,1)$ denotes the risk aversion parameter controlling the trade-off between the expected return and downside risk. Here, a higher value of λ implies a greater degree of risk aversion, where the emphasis is on minimizing extreme losses, whereas a lower value of λ is associated with a lower degree of risk aversion, placing a greater emphasis on maximizing the expected returns. This parameter can therefore be interpreted as reflecting different investors' preferences towards risk. (See Luan et al. [2022])

Other commonly used frameworks in portfolio optimization include mean-variance and *mean-VaR* optimization. Mean-variance optimization is a framework where the portfolio weights are based on the trade-off between the expected return and risk, where the risk is measured as the variance of the portfolio returns. However, since the variance treats positive and negative deviations from the mean equally, it fails to account for skewness and heavy tails, which are often observed in financial data. Mean-VaR optimization, on the other hand, defines the risk through the risk measure VaR, meaning that the portfolio

weights are determined based on the trade-off between the expected return and the VaR of the portfolio. When minimizing the risk, it therefore only focuses on reducing the threshold α for a given confidence level β, but does not account for the severity of the losses beyond this threshold. These limitations therefore motivates the use of mean-CVaR, which focuses directly on the expected tail losses and thereby captures their severity.

5 Empirical Framework

The aim of this chapter is to present the empirical framework used to construct different portfolios by implementing the theory from Chapter 2 and Chapter 3. In this thesis, three copula-based portfolio strategies are constructed based on the Gaussian copula, the Student's t copula, and the R-vine copula. These copula specifications are chosen, based on their ability to capture different dependence characteristics, ranging from linear dependence to tail dependence and more flexible pairwise dependence structures. However, before these portfolios can be constructed, several components of the empirical framework must be specified. The following sections will therefore provide a description of the dataset used in this thesis, including the selected stocks and the sample period, the specification of the marginal models used to describe the individual stock behaviour, as well as the out-of-sample portfolio construction framework.

5.1 Data Description

The portfolios are all constructed using an identical set of ten constituent stocks selected from the *OMX Copenhagen 25 (OMX C25)*. The OMX C25 is a stock index, tracking the 25 largest and most traded stocks on Nasdaq Copenhagen, covering a broad range of sectors including financials, healthcare, consumer goods and industrials. (See Nasdaq) It is therefore widely regarded as a reliable benchmark for the Danish stock market.

The ten stocks included in this thesis are picked from different sectors within the OMX C25, deliberately selected to ensure sectoral diversification. This reduces the portfolio's exposure to adverse shocks within a single sector, while also introducing a greater variation in the dependence structure between stocks, which is important when assessing the ability of different copulas to capture dependence. The selected stocks, along with their sector, ticker symbol and industry classification, are listed in Table 5.1.

Sector	Stock	Ticker Symbol	Industry
Healthcare	Nova Nordisk	NOVO-B.CO	Drug Manufacturers, General
	Demant A/S	DEMANT.CO	Medical Devices
	Coloplast A/S	COLO-B.CO	Medical Instruments & Supplies
Industrials	DSV A/S	DSV.CO	Integrated Freight & Logistics
	Vestas Wind Systems A/S	VWS.CO	Specialty Industrial Machinery
Financial	Danske Bank A/S	DANSKE.CO	Banks - Regional
	Jyske Bank A/S	JYSK.CO	Banks - Regional
Basic Materials	Novozymes A/S	NSIS-B.CO	Specialty Chemicals
Consumer Defensive	Carlsberg A/S	CARL-B.CO	Beverages, Brewers
Consumer Cyclical	Pandora A/S	PNDORA.CO	Luxury Goods

Table 5.1: Providing the sector, ticker symbol and industry of the ten selected stocks from OMX C25. (See Nasdaq)

Throughout the remainder of this thesis, each stock is referred to by its ticker symbol.

For each of the selected stocks, listed in Table 5.1, the daily adjusted closing prices are obtained from *Yahoo Finance* (see Yahoo Finance [b]) for all trading days during the period spanning from January 1st 2017 to December 31st 2025. This choice to use adjusted closing prices, rather then raw closing prices, is based on their ability to accurately reflect the value of the underlying stock, as adjusted closing prices account for corporate actions such as dividends and stock splits. (See Yahoo Finance [a]) The sample period from January 1st 2017 to December 31st 2025 is chosen to provide a sufficiently long sample, including periods of varying market volatility. This sample is then split into an in-sample period from January 1st 2017 to December 31st 2018, used for the selection of the marginal models, and an out-of-sample period from January 1st 2019 to December 31st 2025, used for the construction of the portfolios. The adjusted closing prices for all ten stocks over the entire sample are plotted in Figure 5.1, where the dashed vertical line marks the split between the two periods.

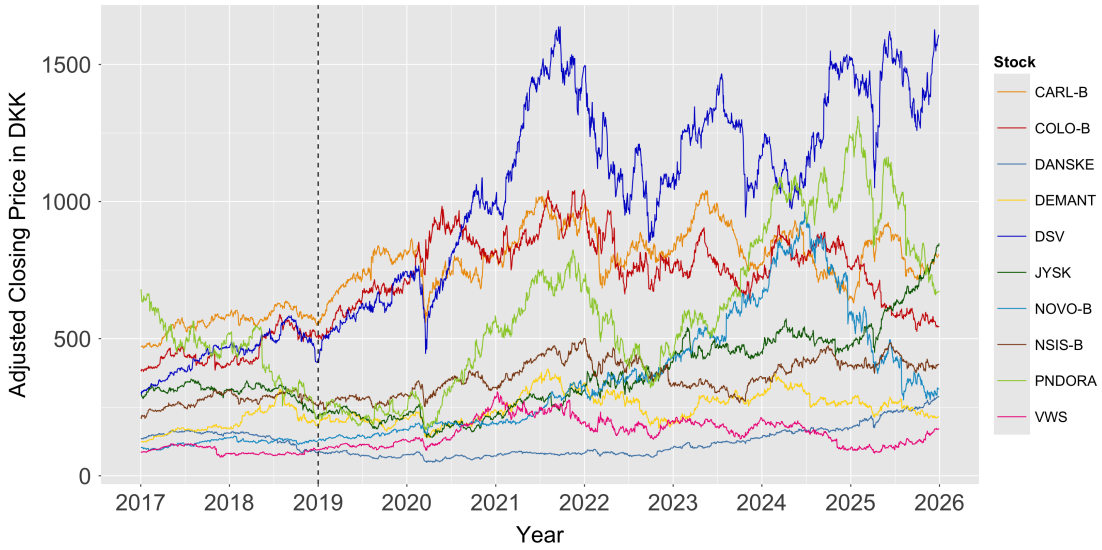

Figure 5.1: Plot of the adjusted closing prices for the selected stocks from the OMX C25 index over the period from January 1st 2017 to December 31st 2025, where the black, dashed, vertical line separates the in-sample and out-of-sample periods.

From Figure 5.1 substantial differences in both price dynamics and volatility across the selected stocks can be observed. Here some stocks exhibit relatively stable prices, such as VWS, DANSKE and DEMANT, while others display stronger upward trends and larger fluctuations, particularly DSV and JYSK. These differences indicate varying return and volatility characteristics across stocks, motivating the use of copulas to capture the dependence structure between the stocks.

Before modelling the dependence structure, an appropriate marginal distribution must be specified for each stock. Since stock prices are typically non-stationary, as also illustrated

in Figure 5.1 where both trends and fluctuations are present, they are not consistent with the assumption of stationarity typically required in time series analysis. The adjusted closing prices are therefore transformed into daily log returns, which for each stock are computed as

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right), \quad t = 1, \ldots, T,$$

where $T \in \mathbb{N}$ denotes the number of observations in the sample period, and P_t denotes the price of the stock at time t. (See Franke et al. [2019]) Having transformed the adjusted closing prices into log returns, the next step is to specify suitable marginal distributions, which is done in the following section.

5.2 Model selection for the Marginal Distributions

The specification of the marginal distribution for each selected stock is based on the in-sample period from January 1st 2017 to December 31st 2018. Selecting an appropriate marginal distribution for each stock is essential, as they form the basis for the dependence structure, which will be modelled using copulas as described in Theorem 3.3.1.

A commonly used approach is to model the marginal distributions of the log returns using an *Autoregressive Moving Average (ARMA)* combined with a *Generalized Autoregressive Conditional Heteroscedasticity (GARCH)* model, commonly referred to as an *ARMA-GARCH* model. This framework is particularly well suited for financial time series, as is captures both potential short-term dynamics in the mean, as well as time-varying volatility. However, before specifying the marginal models, the statistical properties of the log returns are examined in order to assess whether they are suitable for the ARMA-GARCH framework. This is done by computing descriptive statistics and conducting diagnostic tests over the in-sample period. The results are presented in Table 5.2.

Looking at Table 5.2, the log returns of all stocks have mean values close to zero, providing no evidence of a clear trend over time. The standard deviations are also relative small, however, they differ notable across stocks. In particular, the standard variation of PNDORA is more than three times higher that the standard deviation of CARL-B, highlighting the differences in volatility across the stocks. So, even though the level of volatility is relative low, it varies considerably across the selected stocks.

Stock	Mean	SD	Skew	Kurt	JB[2]	ADF	LB(15)
NOVO-B	0.00039	0.015	-0.263	5.793	0	0.01	0.167
DEMANT	0.00079	0.0178	-1.318	9.674	0	0.01	0.028
COLO-B	0.00055	0.014	-1.684	16.078	0	0.01	0.921
DSV	0.00061	0.013	0.073	3.027	0	0.01	0.050
VWS	0.00013	0.021	-1.904	22.238	0	0.01	0.126
DANSKE	-0.00089	0.015	-0.602	5.709	0	0.01	0.033
JYSK	-0.00064	0.014	-1.263	10.262	0	0.01	0.012
NSIS-B	0.00040	0.0132	-0.649	2.889	0	0.01	0.132
CARL-B	0.00032	0.009	-0.035	1.421	0	0.01	0.161
PNDORA	-0.00233	0.028	-3.005	24.031	0	0.01	0.350

Table 5.2: The results from the descriptive statistics and diagnostic tests for the log returns over the in-sample period from January 1st 2017 to December 31st 2018.

In terms of skewness, most stocks deviate from zero, indicating asymmetry in their log return distributions. Only DSV exhibits positive skewness, while the remaining stocks are negatively skewed, suggesting a higher likelihood of extreme negative returns. In addition to this asymmetry, most stocks also exhibit pronounced tail behaviour, as all assets, except for DSV, NSIS-B and CARL-B, have kurtosis values exceeding that of a Gaussian distribution. This indicates that they have heavy tailed distributions, where extreme returns occur more frequently than what would be expected under a Gaussian distribution. Overall, while most stocks exhibit noticeable deviations from the Gaussian distribution in terms of asymmetry and heavy tails, the distribution of DVS appears to be relatively close to a Gaussian distribution. However, this is rejected by the *Jarque-Bera (JB)* test, which strongly rejects the null hypothesis that the log returns of any of the stocks are normally distributed.

In addition to the distributional properties of the log returns, it is also necessary to examine whether these are actually stationary. This is assessed using the *Argumented Dickey-Fuller (ADF)* test, where the null hypothesis assumes the presence of a unit root, implying non-stationarity. This null hypothesis is strongly rejected for all stocks with p-values of 0.01, indicating that the log returns can be treated as stationary.

The next step is then to examine whether the log returns exhibit autocorrelation, as this

[2]Note that the Jarque Bera Test yields p-value $< 2.2 \cdot 10^{-16}$ for all stocks except CARL-B where p-value $< 4.458 \cdot 10^{-10}$, meaning that their p-values are effectively close to zero.

indicates whether an ARMA specification is necessary. Autocorrelation is analysed by first performing a *Ljung–Box (LB)* test, where the null hypothesis assumes that no autocorrelation is present up to a given lag. This test is conducted using 15 lags, which corresponds to 15 trading days. As reported in Table 5.2, the null hypothesis cannot be rejected for the majority of the stocks, suggesting limited evidence of autocorrelation. This is then further analysed through a visual assessment of the plots of the *autocorrelation function (ACF)*, which measures the linear dependence between log returns across different lags. Note that only the ACF plots for COLO-B, DSV and JYSK are included in Figure 5.2, while the remaining ACF plots can be found in Appendix A.3. These three stocks are selected based on their results from the LB test in Table 5.2, as they represent different degrees of evidence against the null hypothesis.

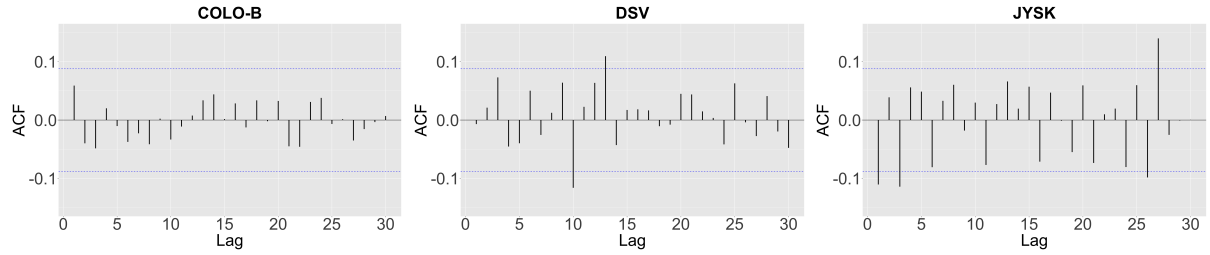

Figure 5.2: ACF plots over 30 lags of the log returns for COLO-B (left), DSV (middle) and JYSK (right) based on the in-sample period from January 1st 2017 to December 31st 2018.

COLO-B exhibits the highest p-value of 0.921, implying that the null hypothesis of no autocorrelation cannot be rejected. This is also consistent with its ACF plot in Figure 5.2, where the autocorrelation across all lags remain small, fluctuate around zero and show no clear systematic pattern. DSV, with a p-value of 0.05, barely rejects the null hypothesis and shows slightly more variation in the autocorrelation across lags, including two spikes outside the confidence bounds. Finally, JYSK, with the lowest p-value of 0.012, strongly rejecting the null hypothesis. However, despite this, its ACF plot does not exhibit any systematic pattern, but rather a few isolated spikes outside the confidence bounds.

Even though the evidence of autocorrelation is generally weak, an ARMA specification is still considered in order to capture any remaining linear dependence in the log returns.

In practice, the log returns of each stock can be decomposed as

$$r_t = \mu_t + \varepsilon_t, \tag{5.1}$$

where $\mu_t = \mathbb{E}[r_t|\mathcal{F}_{t-1}]$ denotes the conditional mean at time t based on the available information at time $t-1$, which is contained in the set $\mathcal{F}_{t-1}$ and ε_t represents the innovations.

The conditional mean in (5.1) is modelled using an ARMA(m, n) specification defined as

$$\mu_t = \mu + \sum_{j=1}^{m} \phi_j r_{t-j} + \sum_{k=1}^{n} \theta_k \varepsilon_{t-k}, \tag{5.2}$$

where μ denotes the unconditional mean, ϕ_j is the autoregressive coefficient, capturing the linear dependence of past log returns and θ_k is the moving average coefficient, capturing the dependence on past innovations.

For each stock, ARMA(m, n) are estimated for $m, n \in \{0, 1, 2, 3\}$, where the final ARMA specification is selected based on the lowest value of the *Bayesian Information Criterion (BIC)*. The BIC is calculated as

$$\text{BIC} = -2 \ln \left(\hat{L} \right) + p \ln(N),$$

where $\hat{L}$ denotes the maximum likelihood of the ARMA model, p is the number of parameters, and N is the number of observations.

Unlike the *Akaike Information Criterion (AIC)*, which primarily focuses on finding the model with the best in-sample fit, BIC penalizes complex models, meaning that it favours simpler models. This helps reducing the risk of overfitting and improves model stability over time, which is particular important in the out-of-sample portfolio framework presented in Section 5.3. Based on the result from the BIC ranking, an ARMA(0,0) specification is chosen for all stocks, implying that the conditional mean is sufficiently captured by the unconditional mean, μ, in (5.2). The the log returns in (5.1) is therefore simplified to

$$r_t = \mu + \varepsilon_t. \tag{5.3}$$

The attention is now shifted to the modelling of the innovations, which are decomposed as

$$\varepsilon_t = \sigma_t z_t, \quad \text{where } z_t \sim \text{i.i.d.}(0, 1),$$

where z_t is an independent and identically distributed sequence with zero mean and unit variance. The innovations also satisfy

$$\mathbb{E}[\varepsilon_t | \mathcal{F}_{t-1}] = 0, \quad \text{Var}(\varepsilon_t | \mathcal{F}_{t-1}) = \sigma_t^2,$$

where the conditional variance σ_t^2 is modelled using a GARCH(p, q) specification, defined as

$$\sigma_t^2 = \omega + \sum_{j=1}^{p} \alpha_j \, \varepsilon_{t-j}^2 + \sum_{k=1}^{q} \beta_k \, \sigma_{t-k}^2,$$

where $\omega > 0$, $\alpha_j \geq 0$ and $\beta_k \geq 0$. The parameter α_j captures the impact of past squared innovations, while β_k captures the persistence in the conditional variance.

Before specifying the GARCH models for the conditional variances of the innovations, it is first examined whether volatility clustering is present in the data. This is done by examining the ACF plots of the squared log returns over 30 lags. Notice that only the plots for COLO-B, DSV and DANSKE are included in Figure 5.3, while the remaining plots can be found in Appendix A.4.

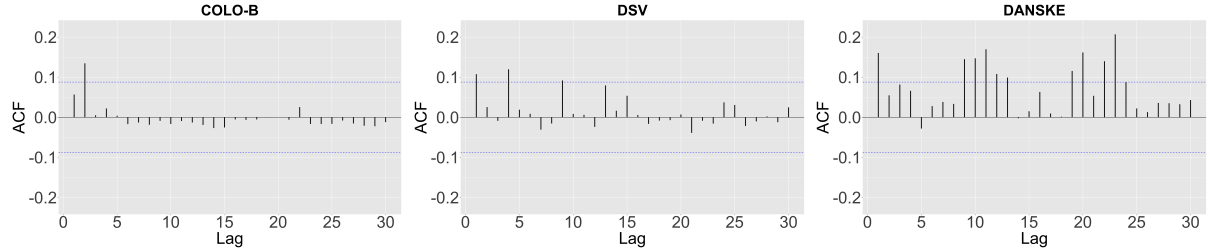

Figure 5.3: ACF plots over 30 lags of the squared log returns for COLO-B (left), DSV (middle) and DANSKE (right) based on the in-sample period from January 1st 2017 to December 31st 2018.

Looking at the ACF plots in Figure 5.3, COLO-B exhibits only weak autocorrelation in the squared log returns, which is mainly concentrated in the first few lags, providing limited evidence of volatility clustering. DSV displays a slightly higher degree of volatility clustering with several positive spikes exceeding the confidence bounds across multiple lags. DANSKE, on the other hand, shows a more pronounced and persistent pattern of positive autocorrelation with several spikes beyond the confidence bounds over a larger number of lags, suggesting a stronger degree of volatility clustering compared to the other stocks.

To formally identify the most appropriate GARCH specification for each stock, GARCH(p, q) are estimated for $p, q \in \{1, 2, 3\}$ and selected based on the BIC. Specifications with $p = 0$, $q = 0$ or $p = q = 0$ are excluded, as the focus of this thesis is on modelling the joint distribution of stock returns for a portfolio using a GARCH-copula framework. Based on the BIC ranking, all stocks are modelled using a GARCH(1,1) specification, except for DSV and VWS, which are better described by a GARCH(1,3). This suggests that, despite DANSKE showing stronger evidence of volatility clustering in the ACF plots in Figure 5.3, a GARCH(1,1) specification is still sufficient to capture its volatility dynamics. Conversely, DSV, which shows a more moderate degree of volatility clustering, is better described by a GARCH(1,3) specification, indicating more persistent volatility over time.

Given the selected GARCH specifications, the estimated conditional variance $\hat{\sigma}^2_{i,t}$ for each stock $i = 1, \ldots, 10$ is used to obtain the standardized residuals, defined as

$$z_{i,t} = \frac{\hat{\varepsilon}_{i,t}}{\hat{\sigma}_{i,t}},$$

where $\hat{\varepsilon}_{i,t}$ denotes the estimated innovations. These standardized residuals are going to be used as inputs for the copulas in Section 5.3. However, since the copula framework is formulated in terms of uniformly distributed variables on the interval $[0, 1]$, as defined in

Definition 2.1.3 (Copula), the standardized residuals are transformed using Theorem A.0.3 (Probability Integral Transform), such that

$$u_{i,t} = F_i(z_{i,t}),$$

where $u_{i,t} \sim \text{Uniform}(0, 1)$ and F_i denotes the assumed distribution function of the standardized residuals for stock i.

Since the true distribution of the standardized residuals is unknown, F_i is selected based on the empirical characteristics of the standardized residuals for each stock. These are analysed by first performing a visual inspection comparing the empirical distribution of the standardized residuals to the standard normal distribution using QQ-plots. In Figure 5.4, the QQ-plots for the standardized residuals of DSV, DANSKE and PNDORA are shown, while the remaining plots can be found in Appendix A.5.

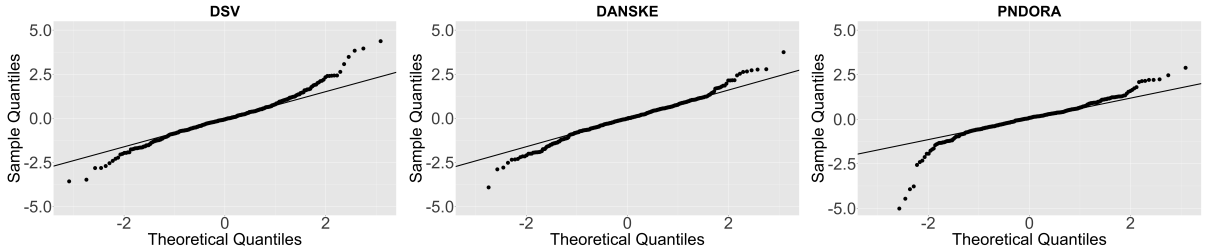

Figure 5.4: QQ-plots of the sample quantiles of the standardized residuals against the theoretical quantiles of the standard normal distribution for DSV (left), DANSKE (middle) and PNDORA (right) based on the in-sample period from January 1st 2017 to December 31st 2018.

According to the QQ-plots of the sample quantiles of the standardized residuals against the theoretical quantiles in Figure 5.4, all three stocks exhibit clear deviations form the standard normal distribution. This is particular evident in the lower tails, which lie below the reference line, and the upper tails, which lie above it, indicating the presence of heavier tails compared to the normal distribution.

This visual analysis is supported by the JB test, which is performed on the standardized residuals, and strongly rejects the null hypothesis of normality for all stocks, with p-values equal to zero. These results are consistent with the findings from the QQ-plot analysis, providing statistical evidence of non-normal behaviour in the standardized residuals. It is therefore necessary to consider alternative distributions for standardized residuals. This is done by fitting them to a range of candidate distributions, including the standardized normal distribution, the standardized Student's t distribution and the standardized *Normal Inverse Gaussian (NIG)* distribution, where the final distribution for each stock is selected based on the lowest value of the BIC.

The selected marginal models for each stock, are presented in Table 5.3, reporting the estimated ARMA-GARCH specifications, the distributions and the corresponding parameter estimates.

Stock	Model	μ	ω	α	β	β_2	β_3	Distribution	Shape	Skewness
NOVO-B	ARMA$(0,0)$-GARCH$(1,1)$	0.0003	0	0	0.999	0	0	std	3.067	0
DEMANT	ARMA$(0,0)$-GARCH$(1,1)$	0.0011	0	0.025	0.974	0	0	std	2.689	0
COLO-B	ARMA$(0,0)$-GARCH$(1,1)$	0.0009	0	0	0.999	0	0	std	3.093	0
DSV	ARMA$(0,0)$-GARCH$(1,3)$	0.0008	0	0.199	0.214	0	0.434	std	4.184	0
VWS	ARMA$(0,0)$-GARCH$(1,3)$	0.0003	0	0.100	0.059	0	0.768	std	4.454	0
DANSKE	ARMA$(0,0)$-GARCH$(1,1)$	-0.0001	0	0.037	0.962	0	0	std	4.078	0
JYSK	ARMA$(0,0)$-GARCH$(1,1)$	-0.0002	0	0.021	0.966	0	0	std	3.413	0
NSIS-B	ARMA$(0,0)$-GARCH$(1,1)$	0.0001	0	0.141	0.745	0	0	nig	1.379	-0.344
CARL-B	ARMA$(0,0)$-GARCH$(1,1)$	0.0003	0	0	0.999	0	0	std	5.031	0
PNDORA	ARMA$(0,0)$-GARCH$(1,1)$	-0.0005	0	0.004	0.993	0	0	std	2.778	0

Table 5.3: The selected marginal models and their parameter estimates for each stock based on the in-sample period from January 1st 2017 to December 31st 2018.

5.3 Out-Of-Sample Portfolio Construction

This section presents the rolling window approach used to construct the copula-based portfolios over the out-of-sample period. This framework is designed to mimic realistic investment strategies, where the portfolios are updated and reallocated daily using only the information available that day. To ensure that the subsequent analysis focuses solely on the how the different copulas affect portfolio performance, the portfolios are constructed under a frictionless market assumption, where costs such as transaction costs, fees, taxes and other trading frictions are ignored.

The rolling window approach is set up such that for each trading day t, a fixed-length estimation window of the most recent historical log returns up to $t-1$ is used to update and reallocate the portfolio weights before the market opens. The resulting portfolio weights are then implemented at the beginning of day t, whereafter their performance is evaluated by the end of the same day using the realized log returns. Within this framework, the parameters of the selected marginal models are re-estimated each trading day t, as stock prices can change rapidity in response to market events, trends or geopolitical events. In contrast, the copula parameters are only updated every 21 trading days, corresponding to approximately one month, as the dependence structure between the stocks is assumed to evolve at a slower pace.

In order to analyse how these copula-based portfolios perform under different levels of risk-aversion in the mean-CVaR framework, presented in Chapter 4, they are all constructed for $\lambda \in \{0.2, 0.5, 0.8\}$. These values of the risk-aversion parameter are chosen to represent

different investor preferences with respect to the trade-off between the expected return and downside risk.

Formally, the out-of-sample portfolio construction is implemented according to the following procedure, where for each trading day t, all portfolio strategies are evaluated simultaneously across the selected copula specifications and risk aversion values.

Step 1. Construct a rolling estimation window of length $W = 498$ using realized log returns from $\{t - W, \ldots, t - 1\}$.

Step 2. Fit the selected marginal models from Section 5.2 on the rolling estimation window and obtain one-step-ahead forecasts of the conditional mean $\hat{\mu}_{i,t}$ and volatility $\hat{\sigma}_{i,t}$ and compute the standardized residuals $\hat{z}_{i,t}$.

Step 3. Transform the standardized residuals into uniform variables using Theorem A.0.3 (Probability Integral Transformation), yielding the corresponding pseudo-observations

$$\hat{u}_{i,t} = \hat{F}_{i,t}(\hat{z}_{i,t}), \quad \hat{u}_{i,t} \sim \text{Uniform}(0, 1).$$

The following steps are implemented separately for each of the considered copula specifications.

Step 4. Model the dependence structure using the copula specification, where Student's t and Gaussian are estimated using maximum likelihood, while R-vine is estimated using the approach introduced in Theorem 3.3.1, with BIC as the information criteria.

Step 5. Simulate $N = 5000$ uniform random samples from the fitted copula, yielding

$$(\hat{v}_{1,k}, \ldots, \hat{v}_{10,k}), \quad k = 1, \ldots, N.$$

Step 6. Transform the simulated uniform samples into standardized residuals using the inverse marginal distributions, such that

$$\hat{z}_{i,k} = \hat{F}_{i,t}^{-1}\left(\hat{v}_{i,k}\right), \quad i = 1, \ldots, 10, \ k = 1, \ldots, N.$$

Step 7. Obtain the simulated returns as

$$\hat{r}_{i,t}^{(k)} = \hat{\mu}_{i,t} + \hat{\sigma}_{i,t}\hat{z}_{i,k}, \quad i = 1, \ldots, 10, \ k = 1, \ldots, N.$$

The following steps are implemented separately for each of the considered levels of risk aversion.

Step 8. Use the simulated returns to obtain the expected return as

$$\mathbb{E}_t[R] = \frac{1}{N}\sum_{k=1}^{N}\hat{R}_{t,k}, \quad \text{where } \hat{R}_{t,k} = \sum_{i=1}^{10}\omega_{i,t}\hat{r}_{i,t}^{(k)},$$

and the CVaR as

$$\text{CVaR}_{\beta,t} = \hat{\alpha}_t + \frac{1}{(1-\beta)N}\sum_{k=1}^{N}\hat{u}_{t,k},$$

subject to

$$\hat{u}_{t,k} \geq -\hat{R}_{t,k} - \hat{\alpha}_t,$$
$$\hat{u}_{t,k} \geq 0$$

where $\hat{\alpha}_t$ denotes the $\text{VaR}_{\beta,t}$ threshold with $\beta = 0.95$, and $\hat{u}_{t,k}$ captures the simulated portfolio losses exceeding this threshold.

Step 9. Use the expected return and the CVaR to solve the mean-CVaR portfolio optimization problem, presented in Chapter 4, and obtain the portfolio weights $(\omega_{1,t}, \ldots, \omega_{10,t})$.

Step 10. Compute the realized out-of-sample portfolio return for t using the realized log returns, as

$$R_t = \sum_{i=1}^{10}\omega_{i,t}r_{i,t}$$

Step 11. Update the rolling estimation window by removing the oldest observation and including the realized returns $r_{i,t}$.

After running the rolling window procedure across the entire out-of-sample period, all three copula-based portfolios are constructed for all chosen risk aversion levels.

6 Out-of-Sample Performance Analysis

In this chapter the out-of-sample portfolio performance is analysed is in order to evaluate how the different copula-based portfolio strategies, constructed in Section 5.3, perform in a frictionless market setting using historical data. The analysis is divided into three sections. In first section, the copula-based strategies are compared to an equally weighted benchmark portfolio under a fixed risk aversion parameter in order to assess their performance relative to a naive portfolio construction. In the second section, the copula-based portfolios are evaluated individually across different levels of risk aversion within the mean-CVaR framework to examine how changes in risk preferences affect their performance. Finally, in the last section, the different copula-based portfolios are compared in order to evaluate how alternative dependence structures affect portfolio performance

6.1 Comparison with Benchmark Portfolio

In order to compare the performance of the three copula-based portfolios with the naive benchmark portfolio, an equally weighted portfolio is constructed over the out-of-sample period. The equally weighted portfolio uses the weights

$$\omega_i = \frac{1}{10},$$

over the entire period, such that its portfolio return, R_t^{EW}, at trading day t is defined as

$$R_t^{\text{EW}} = \sum_{i=1}^{10} \omega_i r_i,$$

This portfolio is referred to as *Equal-weight* for the remainder of the thesis, while the copula-based portfolios are referred to as *Gaussian*, *Student's t* and *R-vine*.

For the copula-based portfolios, the risk aversion parameter is set to $\lambda = 0.5$, representing a balanced trade-off between the expected return and downside risk within the mean–CVaR framework. This level of risk aversion is used throughout the remainder of this section, as it provides a baseline for the analysis, where extreme preferences for either maximisation of the expected return or minimisation of downside risk is avoided. However, to ensure robustness across different levels of risk aversion, the cumulative returns for $\lambda = 0.2$ and $\lambda = 0.8$ are plotted in Appendix A.8.

To compare the performance of the three copula-based portfolios with Equal-weight, the cumulative returns[3] are plotted in Figure 6.1 over the entire out-of-sample period.

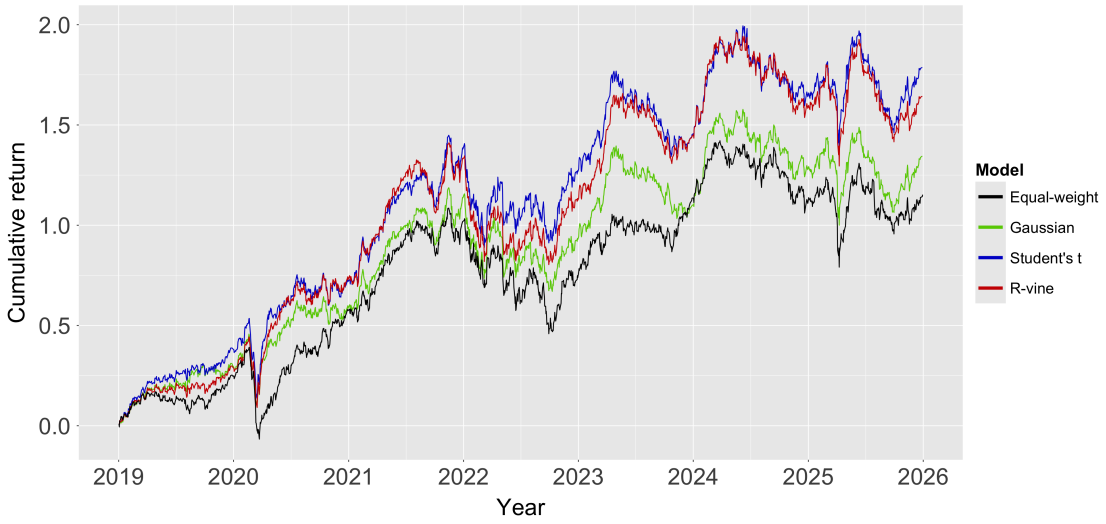

Figure 6.1: Plot of the cumulative return for each portfolio under a risk aversion parameter of $\lambda = 0.5$ over the out-of-sample-period.

From Figure 6.1, it is evident that the three copula-based portfolios outperform Equal-weight throughout most of the out-of-sample period. Equal-weight consistently exhibits the lowest cumulative return, only temporarily reaching a performance level similar to the copula-based portfolios during the early parts of the sample and in connection with the stock market crash in early 2020 due to COVID-19.

Among the copula-based portfolios, Gaussian shows the closest resemblance of Equal-weight in terms of cumulative performance. This is particularly evident during the periods where their two graphs almost align, such as from early 2021 to mid-2022, around the turn of 2024 and during the market crash in April 2025 following the introduction of new tariff policies by the United States. However, outside these periods, a more persistent gab between Gaussian and Equal-weight is observed with Gaussian outperforming Equal-weight.

The cumulative returns of the three copula-based portfolios appear to follow the same overall pattern, just on different levels. Since they all share the same underlying marginal distributions, these differences must stem from their copula specifications, which determine the dependence structure between the assets and thereby influence the joint return distribution used in the portfolio optimization. Looking at Figure 6.1, Student's t and R-vine tend to follow each other closely over the entire period, alternating between who achieves the highest cumulative return, although Student's t appears to have a small advantage. In contrast, Gaussian consistently underperforms, only shortly approaching the same level of

[3]The cumulative returns are computed using simple returns.

cumulative return as the other two copula-based portfolios around mid-2022, before falling behind for the rest of the period.

These differences in their cumulative performance are also reflected by their performance statistics in Table 6.1, where Equal-weight achieves both the lowest total and annualized returns. Over the out-of-sample period Equal-weight increases its initial capital by approximately 115.25%, with an average annual growth rate of 11.6%. In comparison, the copula-based portfolios achieve substantially higher total returns, ranging from 134.50% for Gaussian to 178.70% for Student's, with R-vine positioned in between at 164.35%.

Model	**Total Return** (%)	**Annualized Return** (%)	**Sortino Ratio**	$\mathbf{VaR_{0.95}}$ (%)	$\mathbf{CVaR_{0.95}}$ (%)
Equal Weight	115.25	11.57	0.0609	1.91	2.66
Gaussian	134.50	12.95	0.0725	1.68	2.47
Student's t	178.70	15.78	0.0860	1.65	2.49
R-vine	164.35	14.91	0.0818	1.64	2.48

Table 6.1: Performance statistics for each portfolio under a risk aversion parameter of $\lambda = 0.5$ over the out-of-sample-period.

The risk measures in Table 6.1 confirm the ranking of their overall performances, while also providing a more nuanced view of their underlying risk characteristics. These measures include Sortino ratio, VaR and CVaR , which are all defined in Appendix A.7. When calculating the Sortino ratio value, the *Minimum Acceptable Return (MAR)* is set to zero, meaning that it measures the excess return generated per unit of downside risk, where downside risk is defined as the variability of returns below zero. This value is lowest for Equal-weight, which only generates 0.0609 units of return per unit of downside risk. As for the tail-risk measures, Equal-weight also exhibits the most adverse VaR and CVaR values, implying a 5% probability that the daily losses exceed 1.91%, and an expected loss of 2.66% in these extreme cases.

Among the copula-based portfolios, Student's t achieves the highest Sortino ratio at 0.0860, followed closely by R-vine at 0.0818 and Gaussian at 0.0725. This suggests that Student's t delivers the most efficient trade-off between return and downside risk among all portfolios, and a similar ranking is also observed in the tail-risk measures, although the differences in their values are relatively small.

Overall, the results from both the cumulative returns in Figure 6.1 and the performance statistics in Table 6.1 indicate that the copula-based portfolios consistently outperform Equal-weight in terms of both return and risk-adjusted performance. This is reflected by both higher annualized and total returns, where the copula-based portfolios achieve between 19.25 and 63.45 percentage points higher total return than Equal-weight. In

addition, they exhibit improved Sortino ratio values, which is further supported by more favourable tail-risk outcomes, as reflected by their consistently lower CVaR values across all copula specifications.

6.2 Comparison Across Different Levels of Risk Aversion

In this section the performance of the three copula-based portfolios across different levels of risk aversion is examined. This is done for $\lambda \in \{0.2, 0.5, 0.8\}$ which are selected in order to represent low, moderate and high risk aversion. In contrast to the previous section, where the copula-based portfolios were compared to Equal-weight at a fixed risk aversion level, each copula is now evaluated separately in order to assess how changes in risk preferences affect portfolio performance.

6.2.1 Gaussian

The first copula-based portfolio considered is Gaussian, where the cumulative return dynamics across different levels of risk aversion are examined using the plot in Figure 6.2.

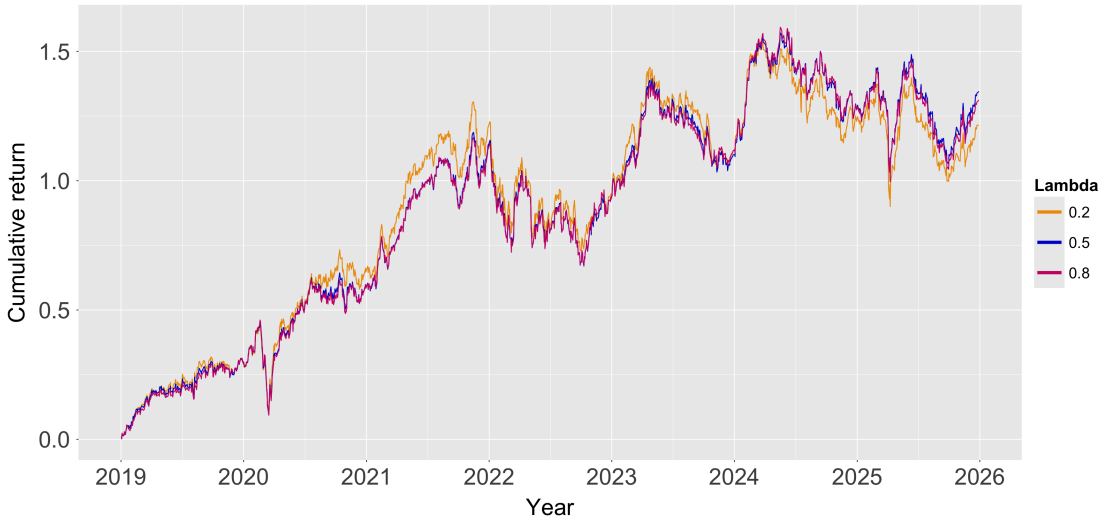

Figure 6.2: Plot of the cumulative return for Gaussian under different levels of risk aversion over the out-of-sample-period.

Overall, the three strategies exhibit very similar patterns, with $\lambda = 0.5$ and $\lambda = 0.8$ following each other so closely that they are often difficult to distinguish visually. These two strategies frequently alternate between who delivers the highest cumulative return, although $\lambda = 0.5$ ultimately achieves the highest return, which is also reflected by the performance statistics in Table 6.2, where it gets the highest total return of 134.50%. In contrast, the $\lambda = 0.2$ strategy generally attains higher cumulative returns over most of the

period. However, from early 2024 onwards this pattern reverses, with performance deteriorating and ultimately resulting in the lowest total return of 121.54%, despite exhibiting the highest cumulative return over the first five years.

λ	**Total-Return** (%)	**Annualized Return** (%)	**Sortino Ratio**	$\mathbf{VaR_{0.95}}$ (%)	$\mathbf{CVaR_{0.95}}$ (%)
0.2	121.54	12.40	0.0661	1.79	2.58
0.5	134.50	12.96	0.0725	1.68	2.47
0.8	131.18	12.73	0.0714	1.61	2.47

Table 6.2: Performance statistics for Gaussian under different levels of risk aversion over the out-of-sample-period.

When comparing the tail-risk measures, the $\lambda = 0.2$ strategy exhibits slightly higher VaR and CVaR values compared to the other strategies, implying a 5% probability of daily losses exceeding 1.79% and an expected loss of 2.58% in these extreme cases. This is also reflected by its Sortino ratio value, which is 0.0661, meaning that it generates relatively low excess return per unit of downside risk compared to the other strategies. The strategies with $\lambda = 0.5$ and $\lambda = 0.8$ are more similar in terms of risk characteristics, as they have identical CVaR values, while their VaR and Sortino ratio values only show minor differences, with $\lambda = 0.5$ performing marginally better. However, given that CVaR is identical for $\lambda = 0.5$ and $\lambda = 0.8$, the lower Sortino ratio value for $\lambda = 0.8$ appears to be driven by lower realized returns rather than improved tail-risk management.

Overall, Gaussian does not exhibit a consistent pattern in performance across different levels of risk aversion. In the mean-CVaR framework, an increase in λ would typically be expected to lead to higher CVaR and lower returns, as higher level of risk aversion places a higher emphasis on minimizing downside risk relative to expected return. This is, however, not the case for Gaussian, which instead shows non-monotonic behaviour in terms of downside risk and return measures across different values of λ.

In particular, for the strategy with $\lambda = 0.2$, where a stronger emphasis is placed on the maximisation of the expected return, Gaussian unexpectedly delivers the lowest total and annualized returns among the three strategies. Conversely, when increasing the risk aversion to $\lambda = 0.5$, the downside risk is reduced as expected, however, this is accompanied by an increase in both total and annualized returns, which deviates from the expectation.

6.2.2 Student's t

The second copula-based portfolio considered is Student's t, where the cumulative return across the different levels of risk aversion is plotted in Figure 6.3 for the entire out-of-sample period.

In contrast to Gaussian, the cumulative return for Student's t exhibit slightly more dispersion across different risk aversion levels. While the strategies appear to follow broadly similar patterns, the differences in magnitude over time make it easier to visually distinguish them. It therefore becomes clear that Student's t displays a more stable ranking of the strategies, with fewer crossovers, suggesting a more consistent trade-off between downside risk and return across the different levels of risk aversion.

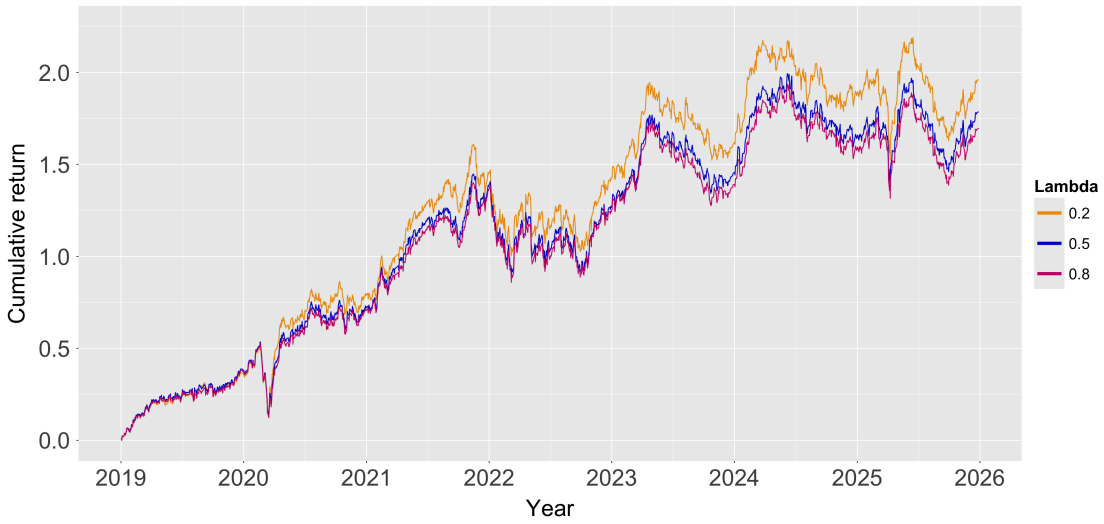

Figure 6.3: Plot of the cumulative return for Student's t under different levels of risk aversion over the out-of-sample-period.

Looking at Figure 6.3 it is evident that the strategy with $\lambda = 0.2$ delivers the highest cumulative return over the entire period, with the gap to the other strategies gradually increasing over time. In comparison, the $\lambda = 0.5$ and $\lambda = 0.8$ strategies appear more similar, although their differences become slightly more apparent over time, with $\lambda = 0.5$ achieving a total return that is approximately 9%-point higher. These observations are also confirmed by the performance statistics summarized in Table 6.3, where the $\lambda = 0.2$ strategy achieves both the highest total and annualized returns at 196.05% and 16.78% respectively. Generally, both total and annualized returns decline as λ increases, reflecting the reduced weight placed on expected return in mean-CVaR framework.

λ	**Total-Return** (%)	**Annualized Return** (%)	**Sortino Ratio**	$\mathbf{VaR_{0.95}}$ (%)	$\mathbf{CVaR_{0.95}}$ (%)
0.2	196.05	16.78	0.0893	1.70	2.50
0.5	178.70	15.78	0.0860	1.65	2.49
0.8	169.67	15.23	0.0835	1.66	2.49

Table 6.3: Performance statistics for Student's t under different levels of risk aversion over the out-of-sample-period.

Looking at the tail-risk measures in Table 6.3, the VaR for $\lambda = 0.2$ is slightly higher than for the other strategies, which is expected as $\lambda = 0.2$ places less weight on downside risk minimisation. However, interestingly enough the tail-risk measures for $\lambda = 0.5$ and $\lambda = 0.8$ are almost identical, suggesting that increasing risk aversion beyond $\lambda = 0.5$ does not necessary reduce realized tail risk, but instead appears to lower the expected return.

Given that the values of the risk measures are almost identical across these two strategies, the small deterioration in risk-adjusted performance, as captured by the Sortino ratio, which decreases slightly from 0.0860 to 0.0835 between $\lambda = 0.5$ and $\lambda = 0.8$, is therefore most likely driven by lower realized returns for $\lambda = 0.8$, rather than improved tail-risk management. Student's t therefore only partly confirms the expected behaviour within the mean-CVaR framework, as higher levels of risk aversion do not consistently lead to lower realized tail risk.

6.2.3 Vine

Lastly, R-vine is examined to see how a more flexible and complex dependence structure performs across different risk aversion levels.

Looking at the plot of it cumulative returns over the out-of-sample period in Figure 6.4, it is evident that the three strategies exhibit a high degree of similarity across the different levels of risk aversion, with only small deviations between the strategies over time. In particular, the $\lambda = 0.5$ and $\lambda = 0.8$ strategies evolve almost identically over time, with their paths overlapping throughout most of the period. This is also reflected in the performance statistics in Table 6.4, where the difference in their total and annualized returns is only 3.76%-point and 0.24%-point, respectively. The strategy with $\lambda = 0.2$ shows a bit more variation, with both periods of underperforming and outperforming the other strategies. In particular, it appears to generate a lower cumulative return from late 2021 to mid-2023, whereafter it recovers and converges towards the other strategies before ultimately taking the lead. It therefore still achieves the highest total and annualized returns with values of 181.16% and 15.92% respectively.

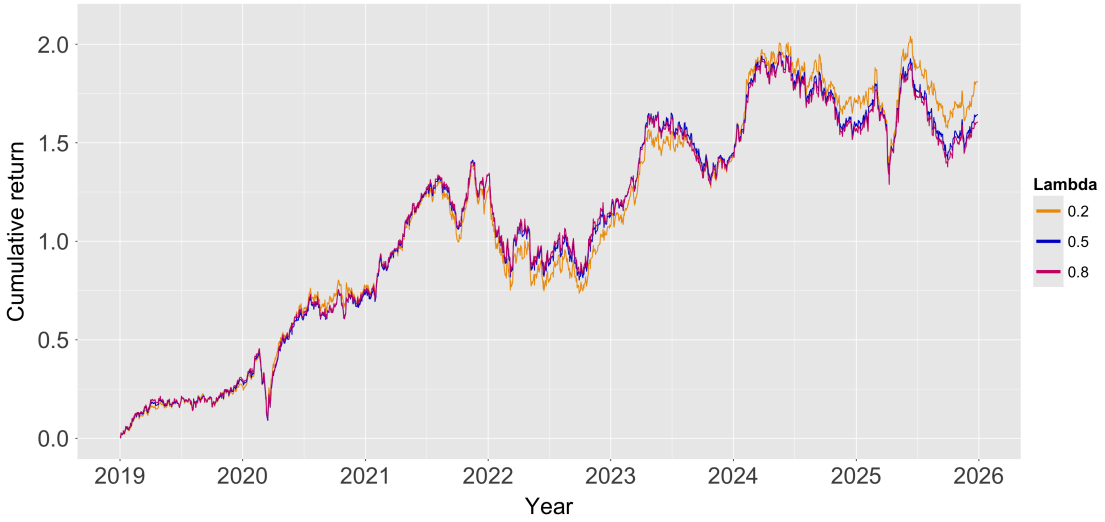

Figure 6.4: Plot of the cumulative return for R-vine under different levels of risk aversion over the out-of-sample-period.

Looking at the risk measures in Table 6.4, the strategy with $\lambda = 0.2$ exhibits the highest VaR and CVaR values, indicating a greater exposure to downside risk, which is consistent with its relatively low emphasis on downside risk minimisation in the mean–CVaR framework. When increasing the level of risk aversion to $\lambda = 0.5$ both tail-risk measures decline, reflecting a reduction in downside risk exposure. However, this pattern becomes less consistent when looking at $\lambda = 0.8$, where the CVaR remains unchanged while the VaR increases slightly relative to $\lambda = 0.5$. This therefore suggests that placing a stronger emphasis on CVaR minimisation, beyond $\lambda = 0.5$, does not necessarily lead to further improvements in the realised tail risk.

λ	**Total-Return** (%)	**Annualized Return** (%)	**Sortino Ratio**	$\mathbf{VaR_{0.95}}$ (%)	$\mathbf{CVaR_{0.95}}$ (%)
0.2	181.16	15.92	0.0854	1.76	2.51
0.5	164.35	14.91	0.0818	1.64	2.48
0.8	160.59	14.67	0.0806	1.66	2.48

Table 6.4: Performance statistics for R-vine under different levels of risk aversion over the out-of-sample-period.

Overall, the results for R-vine do not reveal a fully consistent pattern reflecting that a increased emphasis on maximizing the expected return in the mean-CVaR framework actually leads to a higher realized return. Although the strategy with $\lambda = 0.2$ ultimately achieves the highest total return, the cumulative return dynamics in Figure 6.4 show extended periods, in which it underperforms relative to the other strategies. At the same

time, the performance statistics suggest that increasing the emphasis on CVaR minimisation beyond $\lambda = 0.5$ does not lead to any substantial improvement in downside risk management.

So even though the Sortino ratio decreases slightly form 0.0818 to 0.0806 for $\lambda = 0.5$ to $\lambda = 0.8$, this appears to be driven primarily by lower realized returns rather than any meaningful improvements in downside risk. Hence, for R-vine a higher degree of risk aversion does not necessarily translate into lower tail risk, at least not beyond $\lambda = 0.5$.

6.2.4 Summary of results

In summary, the results from the analyses of the three copula-based portfolios across different levels of risk aversion reveal some differences in how they respond to changes in risk aversion within the mean-CVaR framework. In particular, Gaussian did not fulfil the expected behaviour regarding how realised performance would respond to different levels of risk aversion. Specifically, CVaR decreased only slightly as λ increased, while realised returns exhibited a non-monotonic pattern across the different levels of risk aversion.

In contrast, both Student's t and R-vine displayed a more coherent response to changes in risk aversion. For Student's t, total and annualized returns declined more consistently as λ increases, while CVaR remained relatively stable. R-vine showed a similar pattern, with total and annualized return as decreasing as λ increased, while CVaR only decreased slightly. So, although both Student's t and R-vines results are more consistent with the expectations within the mean-CVaR framework, no further improvements in downside risk were seen beyond $\lambda = 0.5$.

Overall, these results suggest that while more flexible dependence structures, such as Student's t and R-vine, improve the consistency of the relationship between risk aversion and realised performance, this does not necessarily translate into improved downside risk management beyond $\lambda = 0.5$.

6.3 Comparison of Dependence Structures

In this section, a more comprehensive comparison of the three copula-based portfolios is conducted. This is done by analysing Table 6.5, which summarises their performance statistics from the previous section across the three levels of risk aversion, thereby allowing for a direct comparison of how the different dependence structures affect portfolio performance within the mean-CVaR framework.

	$\lambda = 0.2$			$\lambda = 0.5$			$\lambda = 0.8$		
	Gaussian	**Student's t**	**R-vine**	**Gaussian**	**Student's t**	**R-vine**	**Gaussian**	**Student's t**	**R-vine**
Total Return	121.54	196.05	181.16	134.50	178.70	164.35	131.18	169.67	160.59
Annualized Return	12.40	16.78	15.92	12.96	15.78	14.91	12.73	15.23	14.67
Sortino Ratio	0.0661	0.0893	0.0854	0.0725	0.0860	0.0818	0.0714	0.0835	0.0806
$VaR_{0.95}$(%)	1.79	1.70	1.76	1.68	1.65	1.64	1.61	1.66	1.66
$CVaR_{0.95}$(%)	2.58	2.50	2.51	2.47	2.49	2.48	2.47	2.49	2.48

Table 6.5: Performance statistics for the copula-based portfolios under different levels of risk aversion over the out-of-sample-period, where the best performing value within each risk aversion level is highlighted in blue.

From Table 6.5 it is quite clear, that Student's t delivers the best overall performance among the copula-based portfolio strategies, achieving the highest total and annualized returns as well as the highest Sortino ratio across all levels of risk aversion. However, when considering the tail-risk measures, the ranking becomes less distinct, with Gaussian generally performing relatively well, achieving the lowest VaR and CVaR for $\lambda = 0.8$, as well as the lowest CVaR for $\lambda = 0.5$, which is somewhat unexpected, as it is unable to capture tail dependence. In contrast, R-vine achieves the lowest VaR for $\lambda = 0.5$, while Student's t achieves the lowest VaR and CVaR for $\lambda = 0.2$. However, overall these differences in their tail-risk measures remain relatively small and do not indicate any consistent dominance across the three copula specifications.

This therefore suggests that the differences in realized portfolio performance may be attributed to differences in the simulated return distributions generated for each trading day within the rolling window framework and then used to determine the portfolio weights through the mean-CVaR optimization problem. To investigate this further, the CVaR values obtained during the rolling window estimation are plotted over time for each copula-based portfolio under $\lambda = 0.5$ in Figure 6.5. These CVaR values are referred to as *model-implied CVaR*, as they are obtained from the simulated portfolio return distributions generated by the fitted GARCH-copula models and then used within the Mean-CVaR opti-

mization problem to determine the portfolio weights. They therefore reflect each model's expectation for the downside risk based on the simulated return distribution. For comparison, the corresponding plots for $\lambda = 0.2$ and $\lambda = 0.8$ are provided in Appendix A.9.

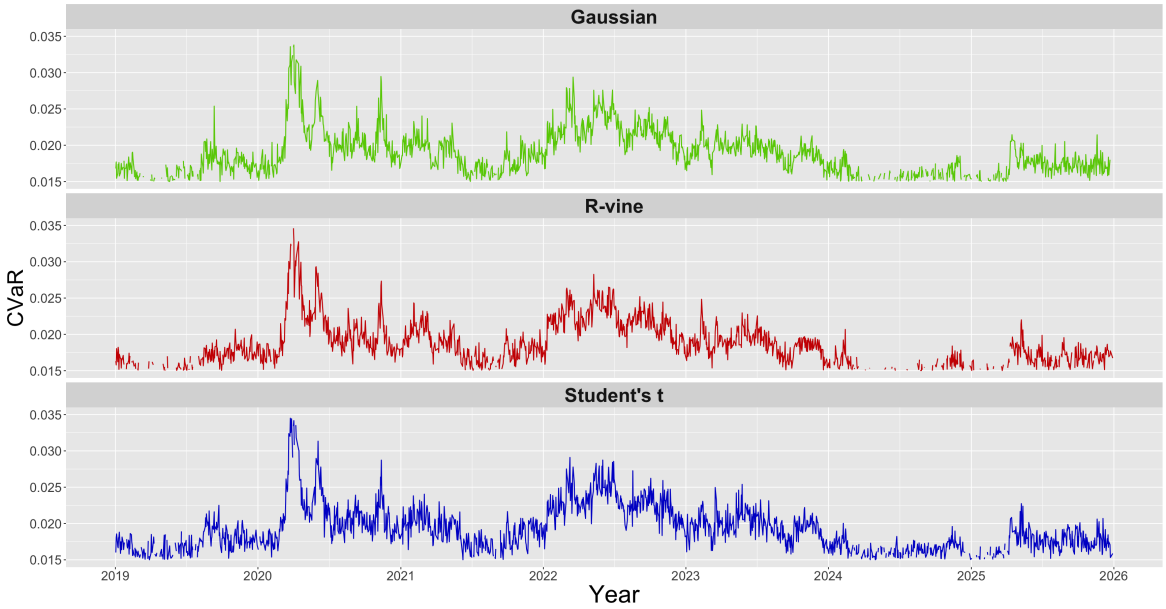

Figure 6.5: Plot of model-implied CVaR for the copula-based portfolios with $\lambda = 0.5$ over the out-of-sample-period from January 1st 2019 to December 31st 2025.

Overall, the three copula-based portfolios exhibit a very similar pattern, responding almost simultaneously to periods of increased market stress, which is especially evident around 2020 during COVID-19 as well as during the Russian invasion of Ukraine in 2022. During these periods, the model-implied CVaR increases notably for all portfolios, reflecting a common response to market stress in terms of increased estimated downside risk across all dependence structures. These similar responses to market stress are also reflected in Table 6.6, reporting the annual mean and standard deviations of the model-implied CVaR for each portfolio.

While the model-implied CVaR follows a similar pattern across all three copula-based portfolios, small but consistent differences can still be observed. In particular, Student's t generally produces slightly higher CVaR values, while R-vine tends to yield slightly lower values, with Gaussian positioned in between. However, these differences remain relatively stable and are small compared to how much CVaR changes over time. In particular, Gaussian exhibits annual mean CVaR values ranging from 1.58% to 2.22%, while Student's t ranges from 1.63% to 2.28% and R-vine from 1.52% to 2.20%, which confirms that all three models operate within a very similar risk range throughout the entire sample period. In terms of the standard deviations of the model-implied CVaR, the three portfolios are also very similar, with volatility peaking in 2022 and reaching its lowest level in 2024.

	Gaussian		Student's t		R-vine	
Year	**Mean**	**SD**	**Mean**	**SD**	**Mean**	**SD**
2019	1.66	0.167	1.71	0.164	1.60	0.157
2020	2.11	0.395	2.15	0.408	2.08	0.409
2021	1.82	0.200	1.85	0.206	1.80	0.207
2022	2.22	0.222	2.28	0.226	2.20	0.219
2023	1.93	0.152	1.97	0.174	1.87	0.157
2024	1.58	0.113	1.63	0.107	1.52	0.115
2025	1.69	0.156	1.72	0.164	1.64	0.146

Table 6.6: Annual mean and standard deviation of model-implied CVaR (in %) for the copula-based portfolio strategies under a risk aversion parameter of $\lambda = 0.5$ over the out-of-sample-period.

Overall, these results indicate that the different copulas play a more important role in shaping the realized portfolio returns than in determining the dynamics of the model-implied CVaR. While the different copula specifications lead to noticeable differences in their returns, the model-implied CVaR remains highly similar across all specifications in terms of both pattern and level.

This therefore suggests that, within the considered GARCH–copula framework, differences in the modelling of the dependence structure primarily affect the portfolio allocation and the realized returns, whereas the estimation of downside risk appears to be largely driven by common market conditions and identical marginal distributions. Consequently, the observed differences in portfolio performance cannot be attributed to differences in tail-risk estimation, but rather to how the simulated return distributions implied by the different copula specifications are translated into optimal portfolio weights within the mean–CVaR framework.

7 Discussion

In this thesis, multiple portfolios have been constructed based on different assumptions about dependence structures in order to examine how they compare to a more naive equally weighted benchmark portfolio, as well as how they perform under different levels of risk aversion. This chapter therefore aims to discuss and highlight some of the limitations and choices made when constructing the portfolios.

The first part of the discussion focuses on the choice of the sample period and how it affects the selection of the marginal models. The sample period was chosen as the period from January 1st 2017 to December 31st 2025, which includes several events affecting the financial market, such as the COVID-19 pandemic, the Russian invasion of Ukraine and multiple periods of heightened uncertainty due to Donald Trump's tariff policies. Consequently, choosing to perform the marginal model selection using data from 2017 to 2018 may be seen as a limitation, since this period is relatively stable compared to the subsequent period. As a result, the selected marginal models may not fully capture changes in the volatility dynamics during the periods of market stress in the out-of-sample period, and may therefore provide a less accurate representation of the returns over time. Choosing the period from 2017 to 2025 was however intentional, as it reflects a reality where sudden global events and political decisions affects the financial markets. Including periods affected by major global events therefore provides a more realistic setting for evaluating how GARCH-copula models perform under different market conditions, particularly in terms of capturing changes in volatility and dependence structures during periods of market stress, rather then solely focusing on more stable periods.

Building on the discussion about the selection of the marginal models, these were selected based on the lowest BIC, which is a relatively conservative information criterion, imposing a strong penalty on model complexity. This is also reflected in the selected model specifications, where most stocks are modelled using a GARCH(1,1) specification combined with the student's t distribution. While this specification contributes to model stability within the rolling estimation, it may be too restrictive to fully capture more complex or time-varying volatility dynamics over time. However, this trade-off was considered necessary, as the marginal models were re-estimated daily over a period of seven years, which requires a specification that is stable over time.

In an addition to the standard GARCH specification, other more flexible specifications, such as EGARCH and GJR-GARCH, were also considered, as they are able to capture asymmetric volatility, where negative returns tend to increase volatility more than positive returns with the same magnitude. However, implementing these models within the rolling estimation framework with daily re-estimation proved challenging, as the behaviour of the stocks changed over time, leading unstable model estimation. This problem could

potentially have been addressed by re-selecting the marginal model specifications throughout the sample period in order to ensure that they we able to capture the current market conditions. However, this would introduce additional model uncertainty and reduce comparability over time, as changes in the portfolio performance may then come from changes in the marginal specification rather then differences in the copula specifications themselves.

An important aspect to consider next is the choice of portfolio optimization problem. In this thesis the mean–CVaR optimization problem was chosen even though more traditional methods such as mean–variance optimization could have been used. Mean–variance is a widely used method within portfolio optimization, known for its simplicity and its ability to account for co-movements between portfolio stocks through the covariance matrix. However, it relies on the assumption that risk can be adequately captured by the variance of the return distribution, which is a rather restrictive assumption within financial markets, as return distributions are typically non-normal, exhibiting skewness and heavy tails. In contrast, mean–CVaR focuses explicitly on downside risk by measuring the expected loss in the tail of the return distribution beyond a given confidence level. This means that, instead of minimizing overall variance, it concentrates solely on reducing extreme losses, making it particularly well suited for financial markets where extreme negative returns tend to co-occur more frequently than extreme positive returns. It also aligns more closely with investor preferences, as the primary concern is typically to avoid large losses rather than large fluctuations in both directions.

When using CVaR within the mean-CVaR framework, the accuracy of the estimated return distribution is important, as CVaR is directly driven by the tail of this distribution. In this thesis, the distribution is obtained through simulated returns from the copula models, meaning that the accuracy of the CVaR estimates depends both on the number of simulations as well as the chosen copulas ability to capture the dependence in the tails. To balance the computational efficiency and the estimation accuracy, 5.000 simulations were used in the portfolio allocation. While this number is considered sufficient in practice to obtain reasonably stable estimates, a higher number of simulations may have improved the precision of the CVaR estimates and thereby may have led to more pronounced performance differences across the copula-based portfolios. In regards to the choice of copulas, each specification has different properties in terms of how dependence is captured, particularly in the tails of the distribution. For instance, the Gaussian copula does not allow for tail dependence, implying that extreme co-movements between stocks are systematically underestimated, which may explain why Gaussian had a overall worse performance compared to R-vine and Student's t. These two copulas are both better suited for the mean–CVaR framework, where the portfolio allocation is directly driven by the minimization of downside risk, making an accurate modelling of joint tail dependence particularly important.

A further critique of the mean–CVaR framework in this thesis is that it is only implemented for three fixed values of λ, whereas in a standard mean–CVaR setting, performance is typically evaluated along an efficient frontier, which allows for a more complete assessment

of the trade-off between expected return and downside risk across different levels of risk aversion. In this thesis, it can therefore only be concluded that increasing the level of risk aversion beyond $\lambda = 0.5$ does not reduce CVaR substantially, but rather just reduces expected returns. Had the efficient frontier been computed, it may have provided a more detailed picture of how this trade-off evolves across different risk aversion levels. However, estimating the full efficient frontier within the rolling estimation framework would be very computationally demanding, which is why three different levels of risk aversion, representing low, moderate and high risk aversion, were considered sufficient to capture the variation in portfolio performance.

Lastly, in terms of the market assumptions, the portfolios where constructed under a frictionless market setting, where short-selling was allowed, and frequently used as the portfolio weights were updated daily within the rolling estimation procedure. This was done to increase flexibility and improve the ability to manage risk exposure, however, it also makes the portfolio more sensitive to estimation errors, leading to larger and more frequent adjustments in portfolio weights than in a market that prohibits short-selling. Short-selling is therefore seen as a useful modelling assumption supporting the mean–CVaR framework. Within the frictionless market setting, no transaction coasts, taxes or other trading fees were included, meaning that the daily rebalancing of the portfolios could be carried out without any cost. While this simplifies the analysis and isolates the effect of the different copulas, more dynamic strategies may appear more attractive than they would in a realistic market setting, where frequent trading would be costly. Overall, these assumptions should therefore be considered when interpreting the results, as they reflect a simplified and frictionless market setting.

8 Conclusion

For this thesis the main goal was to assess how different dependence structures, modelled using different copulas, affect portfolio performance within a mean–CVaR framework. This was done by constructing portfolios of ten stocks from the OMX C24 index over the period from January 1st 2017 to December 31st 2025, where the joint return distribution of the stocks were modelled using a GARCH-copula framework based on Gaussian, Student's t and R-vine copulas.

The results showed that all three copula-based portfolios consistently outperform the naive equally weighted benchmark portfolio in terms of both return and downside risk. In particular, the copula-based portfolios generate substantially higher cumulative, total and annualized returns over the out-of-sample period, while also exhibiting lower downside risk as measured by VaR and CVaR.

When analysing each copula-based portfolio across different levels of risk aversion, the results indicate that the relationship between risk aversion and realised portfolio performance is not fully monotonic for all models. While Student's t and R-vine generally exhibit more stable and theoretically consistent behaviour across different values of λ, Gaussian shows a less consistent response, particularly in terms of its return dynamics. In addition, improvements in downside risk appear to be limited beyond $\lambda = 0.5$, as further increases in the risk aversion level mainly reduce the realised returns while CVaR remains unchanged.

Finally, when comparing the three copula-based portfolio strategies, the results show that Student's t delivers the strongest overall performance, followed by R-vine, while the Gaussian copula generally underperforms. However, despite the differences in their realised returns, the model-implied CVaR dynamics are highly similar across all specifications, both in terms of their level and pattern, with all responding almost identically to periods of market stress. This therefore suggests that the estimation of downside risk is primarily driven by the underlying market conditions and identical marginal distributions, rather than the specific copula's ability to model tail dependence.

Overall, these results suggest that the chosen copula specification mainly affects portfolio allocation through the simulated joint return distribution rather than the assessment of tail risk. As a result, differences in performance between the copula-based portfolios are primarily driven by how their dependence structures are used to find the optimal weights within the mean–CVaR framework, rather than by differences in estimated downside risk. This thesis therefore highlights that, while more flexible dependence structures such as Student's t and R-vine may improve overall performance, these improvements are mainly reflected in their realised returns rather than in improved downside risk. This indicates that, within the considered modelling framework, the choice of copula is more relevant for portfolio optimisation and allocation decisions than for the risk estimation itself.

9 | Bibliography

David Allen and Elisa Luciano. *Risk Analysis and Portfolio Modelling*. MDPI, 2019. ISBN 978-3-03921-625-3.

Tim Bedford and Roger M. Cooke. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence*, 32(1-4):245–268, 2001. doi: 10.1023/A:1016725902970.

Tim Bedford and Roger M. Cooke. Vines: A new graphical model for dependent random variables. *Annals of Statistics*, 30:1031–1068, 2002. ISSN 0090-5364. URL `https://www.jstor.org/stable/1558694`.

Patrick Billingsley. *Probability and measure*. Wiley Probability and Statistics. John Wiley Sons, 3 edition, 1995. ISBN 0-471-00710-2.

Eric Bouye, Valdo Durlleman, Ashkan Nikeghbali, Gaël Riboulet, and Thierry Roncalli. Copulas for finance, 2000. URL `https://www.proquest.com/working-papers/copulas-finance/docview/1699255593/se-2`.

Umberto Cherubini, Elisa Luciano, and Walter Vecchiato. *Copula methods in finance*. Wiley Finance Series. John Wiley Sons Inc, 1 edition, 2004. ISBN 978-0-470-86345-9.

Claudia Czado. *Analyzing Dependent Data with Vine Copulas: A Practical Guide with R*. Lecture Notes in Statistics. Springer Cham, 1 edition, 2019. ISBN 978-3-030-13785-4. doi: 10.1007/978-3-030-13785-4.

Jeffrey Dißmann, Eike C. Brechmann, Claudia Czado, and Dorota Kurowicka. Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics Data Analysis*, 59:52–69, 2013. doi: https://doi.org/10.1016/j.csda.2012.08.010. URL `https://www.sciencedirect.com/science/article/pii/S0167947312003131`.

Fabrizio Durante and Carlo Sempi. *Principles of Copula Theory*. CRC Press, 2016. ISBN 978-1-4398-8442-3.

Jürgen Franke, Wolfgang Karl Härdle, and Christian Matthias Hafner. *Statistics of Financial Markets: An Introduction*. Universitext. Springer, 5 edition, 2019. ISBN 978-3-030-13751-9.

Alexios Galanos. Introduction to the rugarch package, 2026. URL `https://cran.r-project.org/web/packages/rugarch/vignettes/Introduction_to_the_rugarch_package.pdf`.

Kaddour Hadri and William Mikhail. *Econometric Methods and Their Applications in Finance, Macro and Related Fields*. World Scientific Publishing Co Pte Ltd, 2014. ISBN 978-981-4513-48-7.

Piotr Jaworski, Fabrizio Durante, Wolfgang Karl Härdle, and Tomasz Rychlik, editors. *Copula Theory and Its Applications*. Lecture Notes in Statistics. Springer Berlin, Heidelberg, 1 edition, 2010. ISBN 978-3-642-12465-5. doi: 10.1007/978-3-642-12465-5.

Fei Luan, Weiguo Zhang, and Yongjun Liu. Robust international portfolio optimization with worst-case mean-cvar. *European Journal of Operational Research*, 303:877–890, 2022. ISSN 0377-2217. doi: 10.1016/j.ejor.2022.03.011.

Chizoba Morah. Bond vs. stock markets: Key differences explained. `https://www.investopedia.com/ask/answers/09/difference-between-bond-stock-market.asp`, 2026. Accessed: 2026-05-10.

Oswaldo Morales-Nápoles, Roger M. Cooke, and Dorota Kurowicka. About the number of vines and regular vines on n nodes. Submitted for publication, 2010.

Thomas Nagler et al. Statistical inference of vine copulas, 2026. URL `https://cran.r-project.org/web/packages/VineCopula/VineCopula.pdf`.

Nasdaq. Omx copenhagen 25 index (omxc25). https://www.nasdaq.com/da/european-market-activity/indexes/omxc25?id=IX14474.

Roger B. Nelsen. *An Introduction to Copulas*. Springer Series in Statistics. Springer New York, NY, 2 edition, 2006. ISBN 978-0-387-28678-5. doi: 10.1007/0-387-28678-0.

Brian G. Peterson Peterson et al. *Econometric Tools for Performance and Risk Analysis*, 2026. URL `https://cran.r-project.org/web/packages/PerformanceAnalytics/PerformanceAnalytics.pdf`. Documentation on R package.

David Ruppert and David S. Matteson. *Statistics and Data Analysis for Financial Engineering with R examples*. Springer Texts in Statistics. Springer New York, NY, 2 edition, 2015. ISBN 978-1-4939-2614-5.

Ralph Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk, 1999. URL `https://sites.math.washington.edu//~rtr/papers/rtr179-CVaR1.pdf`.

Yahoo Finance. What is the adjusted close? https://help.yahoo.com/kb/SLN28256.html, a.

Yahoo Finance. Yahoo! finance. https://finance.yahoo.com, b. The daily adjusted closing prices were retrieved using the R package quantmod.

A | Appendix

This chapter provides a selection of fundamental definitions and theorems from probability theory, included to support the theoretical framework presented in Chapter 2 (Copula Theory).

First is the definition of a d-dimensional distribution function.

Definition A.0.1. d-Dimensional Distribution Function
Let $d \in \mathbb{N}$ and let $\mathbf{X} = (X_1, \ldots, X_d)$ be a random vector on the probability space $(\Omega, \mathscr{F}, \mathbb{P})$. The *d-dimensional distribution function*, $F : \mathbb{R}^d \to [0,1]$, of $\mathbf{X}$ is then defined as

$$F_{\mathbf{X}}(x_1, \ldots, x_d) = \mathbb{P}\left(X_1 \leq x_1, \ldots, X_d \leq x_d\right),$$

for all $x_1, \ldots, x_d \in \mathbb{R}^d$.

By Definition A.0.1, a d-dimensional distribution function is defined as the joint probability that all components of the random vector $\mathbf{X} = (X_1, \ldots, X_d)$ are less than or equal to some specified values. Although this definition provides a formal description of a d-dimensional distribution function, a theorem specifying the necessary properties is required to determine whether an arbitrary function represents a valid d-dimensional distribution function.

Theorem A.0.2. Properties of d-Dimensional Distribution Functions
Let $d \in \mathbb{N}$. A function $F : \mathbb{R}^d \to [0,1]$ is a distribution function of a random vector $\mathbf{X} = (X_1, \ldots, X_d)$ be a random vector on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$, if and only if

- F is right continuous in each coordinate,
- F is d-increasing,
- F satisfies the lower boundary condition, meaning

$$F(x_1, \ldots, x_d) \to 0, \quad \text{if at least one coordinate } x_i \to -\infty,$$

- F satisfies the upper boundary condition, meaning

$$F(x_1, \ldots, x_d) \to 1, \quad \text{as } \min(x_1, \ldots, x_d) \to +\infty.$$

Proof. This proof is omitted, but it can be found in [Billingsley, 1995, p. 259 – 260]. □

Having defined the necessary properties for a function to qualify as a valid d-dimensional distribution function, it is useful to consider the one-dimensional case, as it provides the

foundation for mapping a random variable to a standard uniform distribution.

Theorem A.0.3. Probability Integral Transform
Let X be a random variable on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$ with the distribution function $F : \mathbb{R} \to [0, 1]$. Then

- if F is continuous, then $F(X) \sim \text{Uniform}(0, 1)$.
- if $U \sim \text{Uniform}(0, 1)$, then $F^{-1}(U) \sim F$.

Proof. This proof is omitted, but can be found in [Durante and Sempi, 2016, p. 5 – 6]. □

The probability integral transform states that applying the distribution function F to a continuous random variable X produces a new random variable $F(X)$, which is uniformly distributed on $[0, 1]$. Conversely, a uniform random variable on $[0, 1]$ can be transformed back to the original distribution using the inverse of F, which is defined in the following definition.

Definition A.0.4. Quasi-Inverse
Let $F : \mathbb{R} \to [0, 1]$ be a distribution function. The *quasi-inverse*, $F^{(-1)} : [0, 1] \to \mathbb{R}$, of F is defined such that

- if $t \notin \text{Ran}(F)$, then

$$F^{(-1)}(t) = \begin{cases} \inf\{x \in \mathbb{R} : F(x) \geq t\}, & t \in (0, 1], \\ \inf\{x \in \mathbb{R} : F(x) > t\}, & t = 0. \end{cases}$$

- If $t \in \text{Ran}(F)$, then

$$F(F^{(-1)}(t)) = t.$$

Remark. *Note, that when F is continuous and strictly increasing the quasi-inverse, $F^{(-1)}$, coincides with the standard inverse, F^{-1}.*

So basically, the quasi-inverse provides a way of mapping values from the codomain $[0, 1]$ of a distribution function back to the corresponding values in its domain. In the special case where $t \in \text{Ran}(F)$, meaning that t is a value that F actually attains, the quasi-inverse satisfies $F(F^{-1}(t)) = t$. Durante and Sempi [2016]

A.1 C- and D-vines

In this section, Canonical vines (C-vines) and Drawable vines (D-vines) are introduced as special cases of R-vines, obtained by imposing additional constrains on the tree sequence of R-vines defined in Definition 3.2.5 (Regular Vine (R-Vine)). More specifically, they introduce restrictions on how the nodes are connected in each tree, as specified in the definition below.

Definition A.1.1. Canonical Vine (C-Vine) and Drawable Vine (D-Vine)
Let $d \geq 2$. A regular vine $\mathcal{V} = (T_1, \ldots, T_{d-1})$ is called a

- *C-vine*, if in each tree T_i, there exists one *root node*, meaning a node $n \in N_i$, such that $|\{e \in E_i \mid n \in e\}| = d - i$,
- *D-vine*, if each node $n \in N_i$ satisfies $|\{e \in E_i \mid n \in e\}| \leq 2$,

for $i = 1, \ldots, d-1$.

So, in addition to the conditions required for a general R-vine, C-vines and D-vines impose additional restrictions on the way the nodes are connected in each tree.

For a C-vine, each tree must contain one root node, which is connected to all other nodes in the tree. An visual example on this is seen in Figure A.1 (left), where the first tree T_1 is depicted for a C-vine with four nodes. Here the root node is $n_1 = 1$, as it is connected to all other nodes in the tree through an edge.

In contrast, a D-vine has a more straight path-like structure, where each node is connected to two nodes at the most. An example of this is seen in Figure A.1 (right), where the first tree for a D-vine is depicted. Here, nodes 1 and 4 are connected to only one other node, while nodes 2 and 3 have two connections. This structure, therefore creates a linear ordering of the variables, where the dependence structure is modelled sequentially along the path. (See Czado [2019])

The joint density associated with a C-vine is given as

$$f(x_1, \ldots, x_d) = \prod_{k=1}^{d} f_k(x_k) \prod_{j=1}^{d-1} \prod_{i=1}^{d-j} c_{j,j+i;1,\ldots,j-1}.$$

Similarly, for a D-vine it is given by

$$f(x_1, \ldots, x_d) = \prod_{k=1}^{d} f_k(x_k) \prod_{j=1}^{d-1} \prod_{i=1}^{d-j} c_{i,(i+j);(i+1)\ldots,(i+j-1)}.$$

These expressions show that both their joint densities can be decomposed into bivariate copulas.

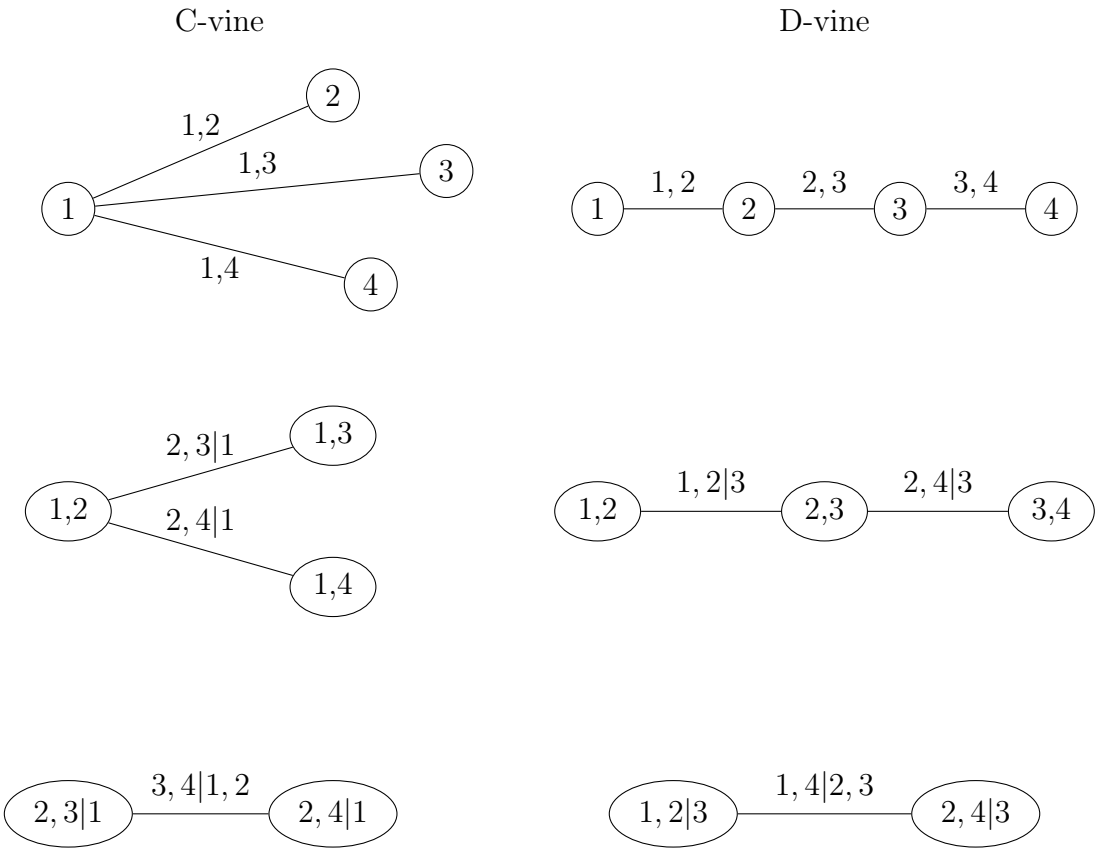

Figure A.1: Examples of a C-vine (left) and a D-vine (right)

A.2 The Dißmann Algorithm

In this section the Dißmann algorithm for sequential estimation of an R-vine copula is provided.

Algorithm 1 Dißmann's algorithm (using Kendall's tau)

Compute the empirical Kendall's tau $\hat{\tau}_{i,j}$ for all pairs $1 \leq i < j \leq d$
Construct the first tree as

$$T_1 = \underset{T=(N,E) \text{ spanning tree}}{\arg\max} \sum_{e=(\mathcal{C}_{e,a},\mathcal{C}_{e,b})\in E} |\hat{\tau}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}}|$$

for each edge $e \in E_1$ **do**
    Estimated parameter(s) for each copula family $\hat{\theta}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}}$ by ML
    Select a copula $C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}}$ by AIC
    for $k = 1, \ldots, n$ **do**

$$\hat{u}_{k,\mathcal{C}_{e,a}|\mathcal{C}_{e,b}} = F_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b}}\left(u_{k,\mathcal{C}_{e,a}} \mid u_{k,\mathcal{C}_{e,b}}; \hat{\theta}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}}\right)$$
$$\hat{u}_{k,\mathcal{C}_{e,b}|\mathcal{C}_{e,a}} = F_{\mathcal{C}_{e,b}|\mathcal{C}_{e,a}}\left(u_{k,\mathcal{C}_{e,b}} \mid u_{k,\mathcal{C}_{e,a}}; \hat{\theta}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}}\right)$$

    end for
end for
for $i = 2, \ldots, d-1$ **do**
    Determine $\tau_{\mathcal{C}_{e,a},\mathcal{C}_{e,b};D_e}$ for all admissible edges $e = \left(\mathcal{C}_{e,a}, \mathcal{C}_{e,b}; D_e\right)$ in the set $E_{p,i}$
    Select the maximum spanning tree among the edges $e \in T_i$ as

$$T_i = \underset{T=(N,E) \text{ spanning tree with } E\subset E_{P,i}}{\arg\max} \sum_{e=(\mathcal{C}_{e,a},\mathcal{C}_{e,b};D_e)\in E} |\hat{\tau}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b};D_e}|$$

    for each edge $e =\in E_i$ **do**
        Estimated parameter(s) for each copula family $\hat{\theta}_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$ by ML.
        Select a copula $C_{\mathcal{C}_{e,a},\mathcal{C}_{e,b}|D_e}$ by AIC.
        for $k = 1, \ldots, n$ **do**

$$\hat{u}_{k,\mathcal{C}_{e,a}|D_e} = F_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b}\cup D_e}\left(u_{k,\mathcal{C}_{e,a}} \mid u_{\mathcal{C}_{e,b}}, u_{k,D_e}; \hat{\theta}_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b}\cup D_e}\right)$$

$$\hat{u}_{k,\mathcal{C}_{e,b}|D_e} = F_{\mathcal{C}_{e,b}|\mathcal{C}_{e,a}\cup D_e}\left(u_{k,\mathcal{C}_{e,b}} \mid u_{\mathcal{C}_{e,a}}, u_{k,D_e}; \hat{\theta}_{\mathcal{C}_{e,a}|\mathcal{C}_{e,b}\cup D_e}\right)$$

        end for
    end for
end for
return The model estimate $\left(\hat{\mathcal{V}}, \hat{\mathcal{B}}, \hat{\Theta}\right)$

A.3 ACF plots for log returns

In this section the remaining ACF plots are provided.

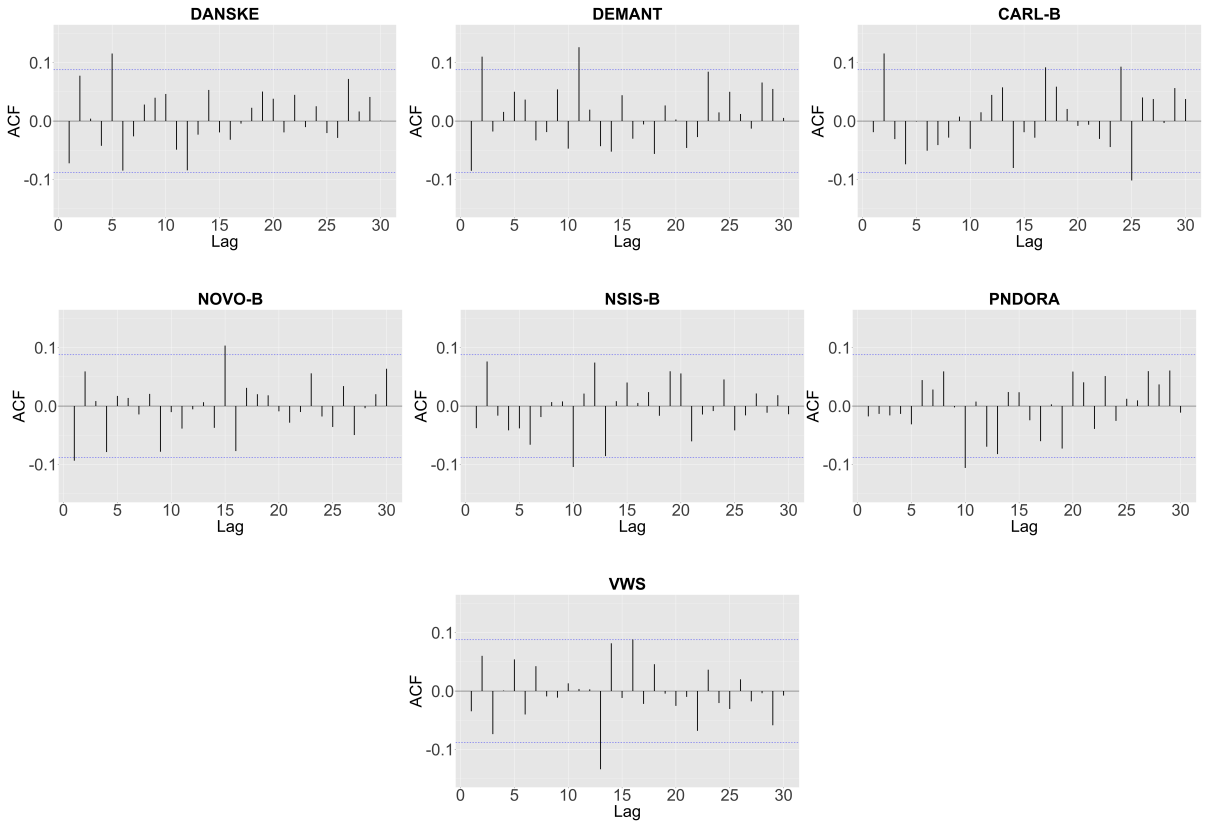

Figure A.2: ACF plots over 30 lags for the log returns based on the in-sample period from January 1st 2017 to December 31st 2018.

A.4 ACF plots for squared log returns

In this section the remaining ACF plots for the squared log returns is provided.

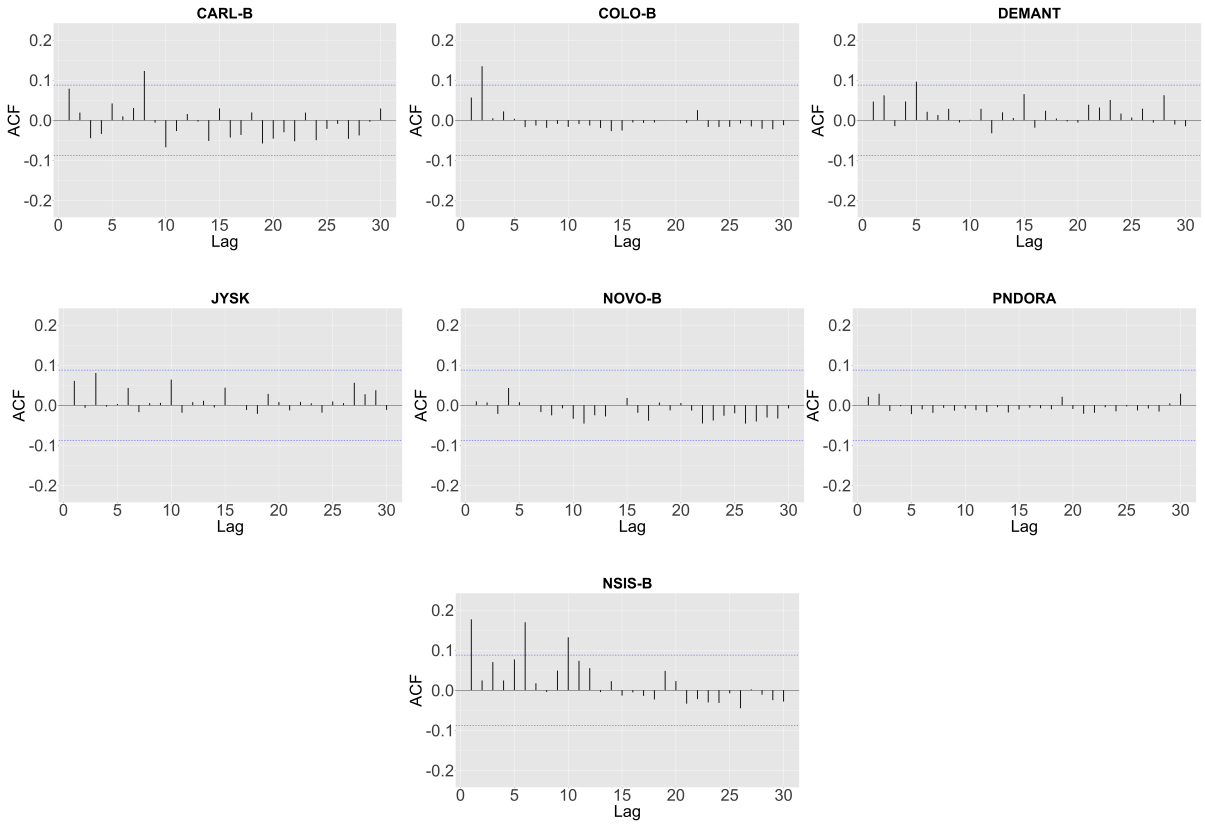

Figure A.3: ACF plots over 30 lags for the squared log returns based on the in-sample period from January 1st 2017 to December 31st 2018.

A.5 QQ-plots plots for the standardized residuals

In this section the remaining QQ-plots for the standardized residuals is provided.

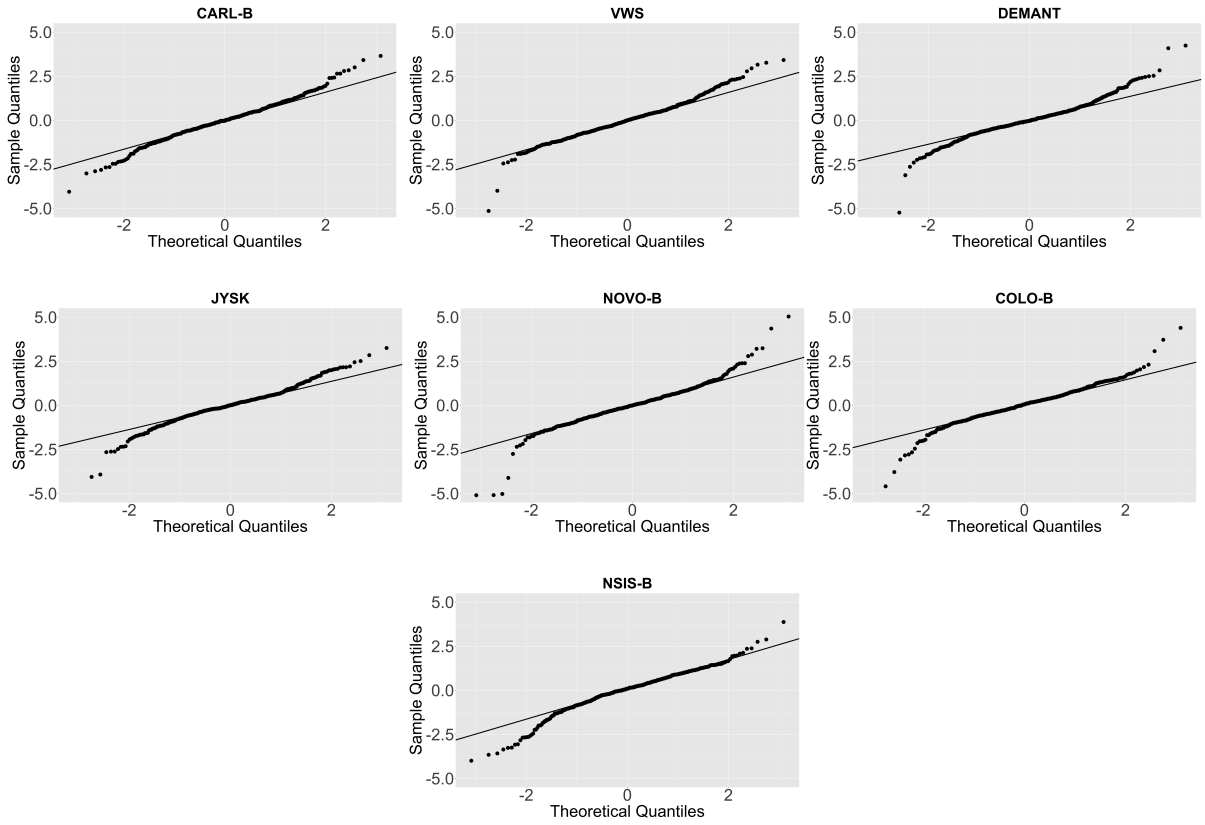

Figure A.4: QQ-plots of the sample quantiles of the standardized residuals against the theoretical quantiles of the standard normal distribution for the remaining stocks based on the in-sample period from January 1st 2017 to December 31st 2018.

A.6 Distributions

In this section the distributions used to model the standardized residuals in Section 5.2 are introduced.

A random variable $X \in \mathbb{R}$ is said to follow a *student's t distribution* with $\nu > 2$ degrees of freedom, if it density function is given by

$$f(x;\nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{(\nu-2)\pi}\Gamma\left(\frac{\nu}{2}\right)}\left(1+\frac{x^2}{\nu-2}\right)^{-\left(\frac{\nu+1}{2}\right)},$$

where Γ denotes the gamma function. (See Galanos [2026]) The distribution is then standardized when

$$\mathbb{E}[X] = 0, \quad \mathrm{Var}(X) = 1.$$

If the random variable is said to follow a *Normal Inverse Gaussian (NIG)* distribution, its density function is given by

$$f(x;\alpha,\beta,\mu,\delta) = \frac{\alpha\delta}{\pi}\exp\left(\delta\sqrt{\alpha^2+\beta^2}+\beta(x-\mu)\right)\frac{K_1\left(\alpha\sqrt{\delta^2+(x-\mu)^2}\right)}{\sqrt{\delta^2+(x-\mu)^2}},$$

where the kurtosis parameter α and the skewness parameter β satisfy the condition $0 \leq |\beta| < \alpha$, $\delta > 0$ is the scale parameter and $\mu \in \mathbb{R}$ is the location parameter. K_1 denotes the modified Bessel function of the second kind. (See Hadri and Mikhail [2014])

The mean and variance is given by

$$E[X] = \mu + \frac{\delta\beta}{\sqrt{\alpha^2-\beta^2}}, \quad \mathrm{Var}(X) = \frac{\delta\alpha^2}{(\alpha^2-\beta^2)^{\frac{3}{2}}},$$

where the distribution is standardized if

$$\mathbb{E}[X] = 0, \quad \mathrm{Var}(X) = 1.$$

A.7 Risk and Performance measures for Portfolio Analysis

In this section the risk and performance measures used to evaluate the portfolio performance in Chapter 6 is defined.

First is the performance measure the cumulative return, which measures the total wealth growth over the chosen period. This is defined as

$$R_{\text{cum}} = \prod_{t=1}^{T}(1 + R_t) - 1,$$

where R_t is the portfolio return at time $t = 1, \ldots, T$. To ensure comparability across time the annualized return is also included, providing the average growth or loss of the portfolio per year. This is defined as

$$R_{\text{annual}} = (1 + R_{\text{cum}})^{\frac{250}{T}} - 1,$$

where $T \in \mathbb{N}$ it the total number of observations, and 250 represents the number of trading days per year.

The risk-adjusted performance is measured using Sortino ratio, which is defined as

$$\text{Sortino Ratio} = \frac{\mathbb{E}[R_t] - \text{MAR}}{\delta_{\text{MAR}}},$$

where MAR is the *Minimum Acceptable Return* and δ_{MAR} is the downside deviation, which is given by

$$\delta_{\text{MAR}} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\min(R_t - \text{MAR}, 0)^2}.$$

To capture tail risk, the empirical VaR and CVaR are calculated based on the return distribution. The VaR at confidence level β is defined as

$$\text{VaR}_\beta = \inf\left\{\alpha \in \mathbb{R} : \mathbb{P}(R_t \leq \alpha) \geq \beta\right\},$$

where $\alpha \in \mathbb{R}$ is the threshold, where the probability of observing returns below this level is at least β. The CVaR is then defined as the expected loss conditioned on exceeding this VaR threshold and is computed as

$$\text{CVaR}_\beta = \mathbb{E}\left[R_t \mid R_t \leq VaR_\beta\right].$$

Finally, the downside risk over time is captured using the maximum drawdown, which measures the largest peak-to-trough decline in the portfolio wealth over the sample period.

The maximum drawdown is defined as

$$\min_t \left(\frac{W_t}{\max_{s \leq t} W_s} - 1 \right),$$

where W_t denotes the portfolio wealth at time t, calculated as

$$W_t = \prod_{i=1}^{t} (1 + R_i).$$

The maximum drawdown therefore captures the portfolio's worst observed loss from a previous high over the sample period. (See Peterson et al. [2026])

A.8 Cumulative return plots

In this section the plots for the cumulative returns for the copula-based portfolios and Equal-weight for the remaining risk aversion parameter values $\lambda = 02$ and $\lambda = 0.8$ are provided.

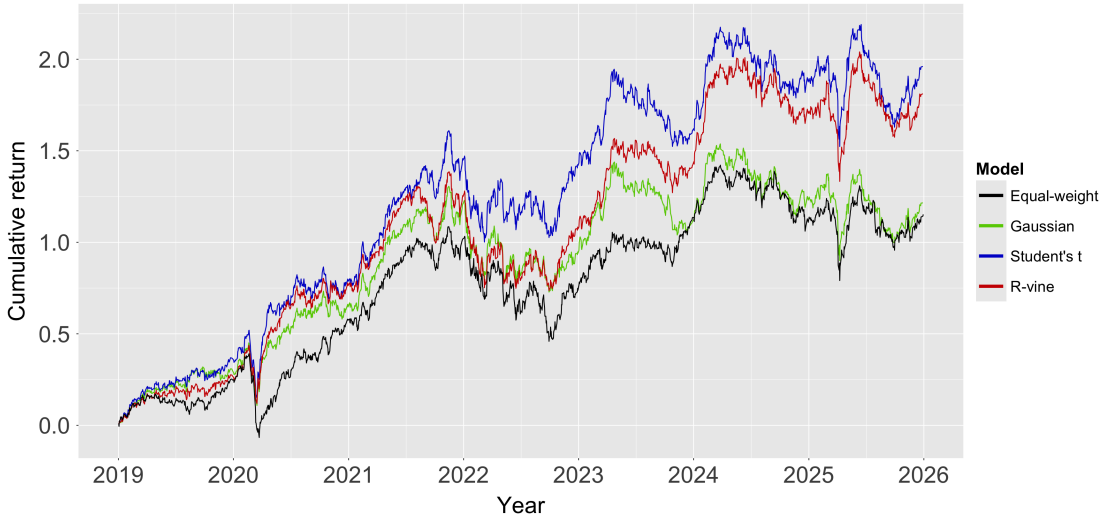

Figure A.5: Plot of the cumulative return for each portfolio under a risk aversion parameter of $\lambda = 0.2$ over the out-of-sample-period from January 1st 2019 to December 31st 2025.

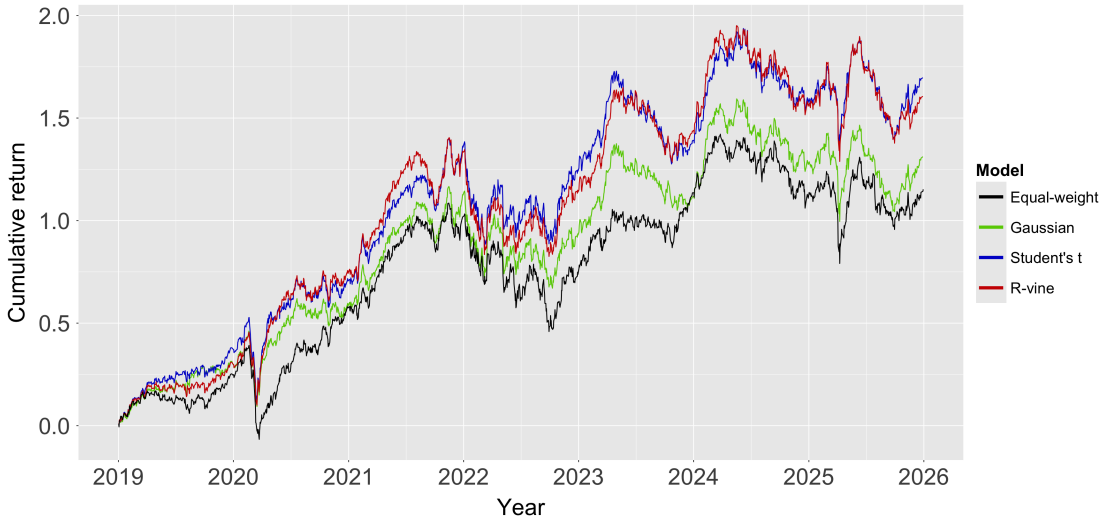

Figure A.6: Plot of the cumulative return for each portfolio under a risk aversion parameter of $\lambda = 0.8$ over the out-of-sample-period from January 1st 2019 to December 31st 2025.

A.9 Model-implied CVaR plots

In this section the plots for the model-implied CVaR for the remaining risk aversion parameter values $\lambda = 02$ and $\lambda = 0.8$ are provided.

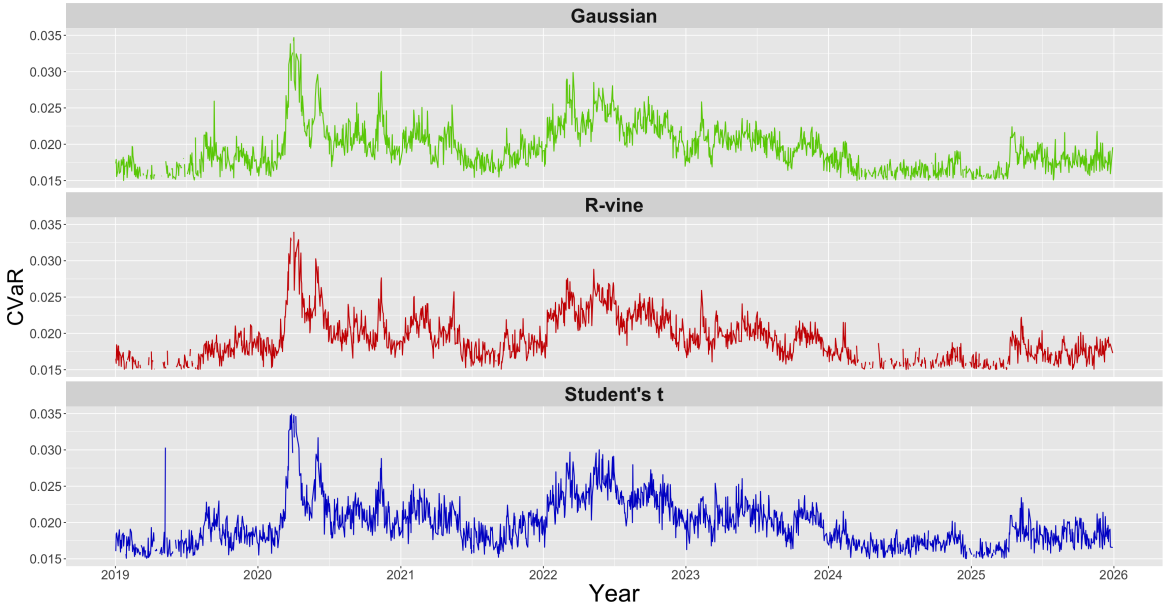

Figure A.7: Plot of model-implied CVaR for the copula-based portfolios with $\lambda = 0.2$ over the out-of-sample-period from January 1st 2019 to December 31st 2025.

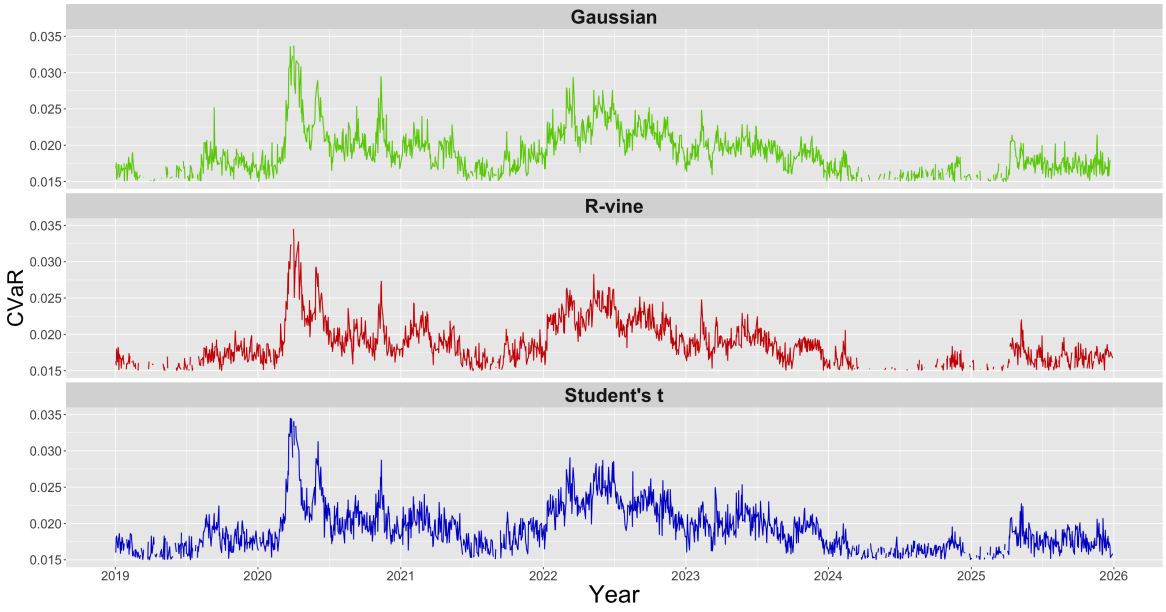

Figure A.8: Plot of model-implied CVaR for the copula-based portfolios with $\lambda = 0.8$ over the out-of-sample-period from January 1st 2019 to December 31st 2025.