Cass Business School
CITY UNIVERSITY LONDON

Working Paper:

Extreme Value Theory and mixed Canonical vine Copulas on modelling energy price risks

Authors:

Karimalis N. Emmanouil

Nomikos Nikos

London

25th of September, 2012

Abstract

In this paper we are concerned with modelling the joint distribution of spot and futures electricity portfolios employing a mixed canonical vine copula model. Our approach combines pseudo-maximum-likelihood fitting of GARCH models to estimate the conditional volatility, extreme value theory (EVT) for estimating the tails of innovations conditional distribution of time series models and pair copula construction (canonical vine copula) for modelling the dependency structure of the portfolio constituents. Our main findings are that VaR estimates based on normality assumption are satisfactory for higher quantiles but poor for extreme quantiles. We show that our procedure gives better 1-day VaR estimates for extreme quantiles than methods which ignore the heavy tails of innovations and the asymmetries of the joint return distribution. Moreover, we show that for inferences involving the tails, pair copula selection should not be based on likelihood-based criteria but rather on non-parametric dependence measures.

1 Introduction

The understanding of joint asset return distributions is an integral part for managing portfolio risks successfully. Modelling the joint return distribution of portfolio, however, is a non-trivial task. What makes this task non-trivial is the complex dynamics and specific characteristics of each particular asset in the portfolio on one hand, and the varying dependency structure between all portfolio constituents on the other. In this paper we are concerned with describing the joint return distribution of power portfolios and compute risk measures such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR).

Our modelling strategy comprises of two main stages. The first stage combines pseudo-maximum-likelihood fitting of time series models and extreme value theory for estimating both tails of the conditional innovations distribution of time series models. Within this framework we take into account the conditional volatility for each individual return series while assigning an explicit model to each tail of the conditional returns distribution. In the second stage, the dependency structure among portfolio return series is modelled employing a mixed canonical vine copula model. Hence, by combining a semi-parametric approach for the margins and a vine copula methodology for describing the dependency structure we aim to provide a flexible way of modelling the conditional distribution of asset returns while, at the same time, paying a particular attention to the tails of distribution, which is in practice the focus of all risk management applications.

The modelling of extreme events is the central concern of extreme value theory and the main objective of this theory is to provide asymptotic models allowing the modelling of tails of distribution. Extreme Value Theory (EVT) has found applications in many fields of modern science such as engineering, insurance, hydrology and many others (see e.g. Embrechts et al. (1999a), Reiss and Thomas (1997)). The last years, more and more research has been undertaken to analyze extreme events of financial series via EVT (Embrechts et al., 1999b; Danielsson and de Vries, 1997; McNeil and Frey, 2000 etc.). EVT-based methods are attractive for tail estimation of financial time series because provide better fit in extreme quantiles for heavy-tailed data. Furthermore, EVT-methods treat the tails of distribution separately, offering a distinct parametric form for each tail of distribution and allowing for asymmetry and extrapolation beyond the range of the data.

Therefore, EVT has found various risk management applications. For example, McNeil and Frey (2000) propose a method for estimating VaR and related risk measures by filtering return series with GARCH models and then applying threshold-based EVT techniques to residuals series. They found that a conditional approach that models the conditional distribution of returns is better suited for VaR estimation than the unconditional approach. Gençay and Selçuk (2004) apply EVT in daily stock market returns of emerging markets. They report that EVT-based model dominates other parametric models in terms of VaR forecasting in extreme quantiles. With respect to the energy markets, Byström (2005) focuses on the Nord Pool intra-daily price changes and calculate extreme quantiles by fitting traditional time-series models and an EVT-based model to empirical data. He found that both in-sample and out-of sample estimates of moderate and extreme tail quantiles are more accurate than corresponding estimates of time series models with normal or Student-t innovations. Chan and Gray (2006) propose an AR-EGARCH-EVT model to forecasting VaR using daily electricity prices from various power markets. In markets where the distribution of returns is characterized by high volatility, skewness and kurtosis the AR-EGARCH-EVT model dominates in terms of VaR performance.

Nevertheless, multivariate research on statistical properties of power portfolios is rather limited. Börger et al. (2007) analyze the joint return distribution of various energy futures series. With respect to power data, they focus on monthly and yearly Phelix futures series. According to their findings, the multivariate Normal hypothesis is strongly rejected within a commodity class as well as across commodities. They show that the class of Generalized Hyperbolic (GH) distributions is capable of fitting power future prices and clearly outperforms the Normal distribution. Moreover, they demonstrate how the multivariate fit of the distributions can be used to estimate risk measures and also state that the exact choice of the distribution has a major influence on risk measures.

To the best of our knowledge, however, there are no studies analyzing the dependence of power or other commodities portfolios using the concept of vine copula modelling. The use of copulas enables the separation of the dependence model from the marginal distributions. While the last years there is a growing literature on copulas, most of the research is still limited to the bivariate case where a rich variety of different copula families with distinct characteristics is available. However, the choice of multivariate copulas is rather limited compared to the bivariate case. Apart from the multivariate Gaussian copula that does not allow for tail dependence and Student-t copula than can only captures symmetric tail dependence, the exchangeable multivariate Archimedean copulas are extremely restrictive. Therefore, building higher-dimensional copulas is a natural next step.

The so-called pair copula construction (PCC) principals provide an alternative and flexible way for building multivariate distributions. The fist pairwise structure was proposed by Joe (1996) and further extended by Bedford and Cooke (2001, 2002) and Kurowicka and Cooke (2006). In particular, Bedford and Cooke realized that there are many possible ways to decompose a d-dimensional density using pairwise construction and hence they organized them in graphical way of sequentially designing trees where only products of bivariate copulas, the so-called pair-copulas, are involved. They called these distributions *regular vines*. Aas et al. (2009) were the first to recognized that bivariate copulas for any copula family as well as several copula families can be mixed in one pair-copula construction. They also developed algorithms that allow standard maximum likelihood estimation (MLE) and simulation for two special classes of regular vines; the canonical vine (C-vine) and D-vine, where each model provides a specific way of decomposing the density.

While the use of copulas in finance is popular nowadays, the literature on vine copula modelling is mainly concentrated on illustrative applications. For example, Aas and Berg (2009) fit two different classes of models for construction higher-dimensional dependence on a four-dimensional equity data set; the nested Archimedean construction (NAC) and the pair-copula construction (PCC). The goodness-of-fit tests employed, strongly reject the NCA while the PCC (D-vine) provides an adequate fit. They also show that the PCC does not overfit the data and works satisfactory for out-of-sample VaR calculations . Czado et al. (2011) developed a data driven sequential approach, based on Kendall's τ estimates, for C-vine specifications. They show through simulation studies the satisfactory performance of maximum likelihood estimation procedure for mixed/non-mixed C-vine models. They also consider an application involving US-exchange rates by fitting univariate time series models to the exchange rate return series and modelling the dependencies among the resulting standardized residuals through a mixed C-vine model.

There are few studies in finance literature recently that support the use of vine copulas for financial applications. Mendes et al. (2010) employ a D-vine model on a six-dimensional global portfolio with skewed-t distribution as the unconditional model for the marginals. They show how pair-copulas can be used on a daily basis for constructing efficient frontiers and computing VaR. Heinen and Valdesogo (2009) propose the canonical vine autoregressive (CAVA) model that can be viewed as a time-varying non-linear and non-Gaussian extension of the CAPM model. They focus on high-dimensional vine copula modelling. They show that neither the marginal distribution nor the dependency structure is Gaussian and the persistence in the dependence between all pairs is not the same, as implied by the DCC model. They also show that in terms of in-sample and out-of sample VaR, the CAVA model performs better than the DCC model in all portfolios examined. Brechmann and Czado (2011) also focus on high-dimensional vine copula modelling. They propose a factor model for analyzing the dependency structure among European stocks of Euro Stoxx 50 index, the so-called Regular Vine Market Sector (RVMS) model. The model employs a general R-vine copula construction, avoiding in this way to impose the independence assumptions of the CAVA model. The authors show that the RVMS model provide good fits of the data and accurate VaR forecasts. Moreover, they show that their proposed methodology can reduce the required risk capital in contrast to the DCC model with Gaussian innovations when employed for VaR forecasting.

We believe that the contribution of our study to the existing literature is threefold. Firstly, we introduce the concept of vine copula as an alternative and more flexible way to describe the joint return distribution of power portfolios. To our best knowledge, vine copula modelling has never been used before to describe the dependency structure of power related portfolios in the literature. Secondly, we propose an extension of extreme value theory in the context of vine copula modelling that focuses mainly on extreme quantiles. We strongly believe that our proposed methodology contributes and provides an alternative perspective to the study of multivariate extremes. Finally, we compute risk measures based on this framework and discuss the implications of our findings for portfolio risk management.

The rest of this paper is organized as follows. Section 2 summarizes the basic concepts of extreme value theory. Section 3 introduces the copula theory and presents the pair-copula construction principles and particularly the canonical vine copula modelling. Practical issues related to vines such as inference and model selection are also discussed. The model setup is explained in Section 4 whereas Section 5 presents the data. Section 6 reports the estimated model parameters. Section 7 reports VaR and CVaR forecasts and their corresponding backtest results and Section 8 concludes.

2 Extreme Value Theory

Let $X_1, \ldots, X_n$ be a sequence of n independent and identically distributed (i.i.d.) random variables with distribution function F. Univariate extreme value theory (EVT) centers on the distribution function F with a particular focus on extreme tail quantiles. In general, within the EVT context, there are two approaches of identifying extreme events in real data. The first approach, known as *block maxima*, is the direct modelling of the distribution of minimum or maximum realizations. The other approach, known as *peak over threshold*, concentrates on the realizations that exceed a prespecified threshold level. The peak over threshold (POT) method is employed in this particular study.

2.1 Peak over Threshold (POT) method

The peak over threshold method is concerned with the behaviour of large observations that exceed a high threshold. Let's consider an (unknown) distribution function F of a random sample $X_1, \ldots X_n$. The peak over threshold approach focuses on estimating the distribution function F_u of values of x above of a high threshold level u. The distribution function F_u is called the *conditional excess distribution function (cedf)* and is defined as

$$F_u(y) = P(X - u \leq y \mid X > u), \qquad 0 \leq y \leq x_F - u \tag{2.1}$$

where X is a random variable, u is a given threshold, $y = x - u$ are the excesses and $x_F \leq \infty$ is the right endpoint of F. This conditional probability may be written as

$$F_u(y) = \frac{\Pr\left(X - u \leq y, X > u\right)}{\Pr(X > u)} = \frac{F(u+y) - F(u)}{1 - F(u)} = \frac{F(x) - F(u)}{1 - F(u)} \tag{2.2}$$

The estimation of F_u is a difficult exercise as we have in general very little observations in this area. A theorem by Balkema and de Haan (1974) and Pickands (1975) is very useful at this point since it provides a powerful result about the *cedf*.

Theorem 2.1 (Pickands (1975), Balkema and de Haan (1974)) *For a large number of underlying distribution functions, $F_u(y)$ is well approximated by $G_{\xi,\sigma}$, the Generalized Pareto Distribution (GPD):*

$$F_u(y) \approx G_{\xi,\sigma}(y), \qquad u \to \infty,$$

where

$$G_{\xi,\sigma}(y) = \begin{cases} 1 - (1 + \frac{\xi}{\sigma} y)^{-1/\xi} & \text{if} \quad \xi \neq 0 \\ 1 - e^{-y/\sigma} & \text{if} \quad \xi = 0 \end{cases} \tag{2.3}$$

for $y \in [0, (x_F - u)]$ *if* $\xi \geq 0$ *and* $y \in \left[0, -\frac{\sigma}{\xi}\right]$ *if* $\xi < 0$. $G_{\xi,\sigma}$ *is the so-called generalized Pareto distribution (GPD),* $\xi = 1/\alpha$ *is a shape parameter and* α *is the tail index.*

The GPD includes a number of other distributions. For $\xi > 0$, it takes the form of the ordinary Pareto distribution, which is the most relevant for modelling the fat tails of financial time series. When also $\xi > 0$, $E[x^k]$ is infinite for $k > 1/\xi$. Moreover, when $\xi = 0$, the GPD corresponds to the exponential distribution whereas for $\xi < 0$ it corresponds to Pareto II type distribution. The GPD model can be estimated with the maximum likelihood method. More precisely, Hosking and Wallis (1987) showed that the maximum likelihood estimates are asymptotically normally distributed for $\xi > -0.5$.

3 Copula Theory

A copula is a multivariate distribution function $C(u_1 \dots u_d)$ defined on the unit cube $[0,1]^d$ with uniformly distributed marginals. It provides a way of isolating the dependency structure between d random variables while allowing for arbitrary marginal distributions. The copula concept was initially developed by Sklar (1959). The famous theorem of Sklar (1959) gives the connection between marginals and copulas to the joint distribution. In general, a copula function can be extended for an arbitrary dimension d, but since our mission is to develop multivariate copulas using only bivariate copulas as building blocks, we will only focus on the bivariate case $d = 2$.

Theorem 3.1 (Sklar (1959)) *Let* $F : \overline{\mathbb{R}}^2 \to [0,1]$ *with* $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$ *be a bivariate distribution with one-dimensional marginals* $F_1, F_2 : \overline{\mathbb{R}} \to [0,1]$. *Then there exists a two-dimensional copula* C, *such that for all* $(x_1, x_2) \in \overline{\mathbb{R}}^2$

$$F(x_1, x_2) = C\Big(F_1(x_1), F_2(x_2)\Big) \tag{3.1}$$

holds and vice versa

$$C(u_1, u_2) = F\Big(F_1^{-1}(u_1), F_2^{-1}(u_2)\Big), \tag{3.2}$$

where u_1 *and* $u_2 \in [0,1]$ *and* $F_1^{-1}(u_1)$ *and* $F_2^{-1}(u_2)$ *are the inverse distribution functions of the marginals.*

Let $f_X(x)$ and $f_Y(y)$ be marginal densities with joint density of $f_{XY}(x,y)$. It can be shown that the joint density can be decomposed as a product of marginal densities and copula density, $c(u,v)$, as follows

$$f_{XY}(X,Y) = \frac{\partial^2 F_{XY}(X,Y)}{\partial x \partial y} = \frac{\partial^2 C(F_X(x), F_Y(y))}{\partial x \partial y} = c(u,v) f_X(x) f_Y(y)$$

Tail dependence is another very useful copula-based measure of extreme co-movements. It is very important property for applications, such as risk management, that concern with the study of dependence of extreme values. Many empirical studies in finance (e.g. Longin and Solnik (1995, 2001), Ang and Chen (2002), Hong et al. (2007) among others) have indicated the presence of asymmetries in financial data, meaning that lower tail dependence can be stronger than upper tail dependence or vice versa. As such, standard symmetric multivariate distributions are inappropriate to address this feature. Moreover, tail dependence is one of the properties that distinguish between the different copula families. There are copula families that do not allow for tail dependence and others allow for either lower only or upper only tail dependence. There are also "reflection symmetric" copulas that imply same upper and lower tail dependence for any bivariate margin or "reflection asymmetric" copulas that allow for flexible upper and lower tail dependence.

3.1 Pair Copula Construction

Pairwise copula construction constitutes a very useful way for building flexible multivariate distributions. The modelling principle is based on a decomposition of a multivariate density into a cascade of bivariate copulas, which is applied on original variables and on their conditional and unconditional distributions. Aas et al. (2009) were the first that realized that this construction principle can be extended by allowing arbitrary pair-copula families as building blocks. As such, a multivariate distribution, which is decomposed using the pair-copula principle and allows for different copula families as building blocks, is called *mixed vine*. Regular vines include two main type of vines, C- and D-vines. Their main difference lies on the way they organize a multivariate density decomposition. C-vines utilize a star tree methodology whereas the D-vines employ a line tree methodology.

Let $\mathbf{X} = (X_1, \ldots, X_d)^t$ be a vector of random variables with a joint density $f(x_1, \ldots, x_d)$, marginal densities $f(x_1), \ldots, f(x_d)$ and marginal distributions $F_1(x_1), \ldots, F_d(x_d)$. This density can be decomposed as

$$f(x_1, \ldots, x_d) = f_d(x_d) \cdot f(x_{d-1}|x_d) \cdot f(x_{d-2}|x_{d-1}, x_d) \cdots f(x_1|x_2, \ldots, x_d), \tag{3.3}$$

and this factorization is unique up to a re-labeling of the variables. It can be shown that the joint distribution function $f(x_1, \ldots, x_d)$, for an absolutely continuous multivariate distribution F with strictly increasing, continuous marginal densities can be factorized as

$$f(x_1, \ldots, x_d) = c_{1\ldots d}\Big(F_1(x_1), \ldots, F_d(x_d)\Big) \cdot f_1(x_1) \cdots f_d(x_d) \tag{3.4}$$

where $c_{1\ldots d}$ is a uniquely identified d-variate copula density.

In the beginning of this section, we show that a bivariate joint distribution can be factorized as a product of a copula and marginal densities as follows

$$f(x_1, x_2) = c_{12}(F(x_1), F(x_2))f_1(x_1)f_2(x_2). \tag{3.5}$$

Moreover, the conditional density $f(x_1|x_2)$ can be expressed in terms of a copula as

$$f(x_1|x_2) = \frac{f(x_1, x_2)}{f_2(x_2)} = \frac{c_{12}(F(x_1), F(x_2))f_1(x_1)f_2(x_2)}{f_2(x_2)} = c_{12}(F(x_1), F(x_2))f_1(x_1) \tag{3.6}$$

For the d-dimensional case it holds that

$$f(x|\boldsymbol{\nu}) = c_{x,\nu_j|\boldsymbol{\nu}_{-j}}(F(x|\boldsymbol{\nu}_{-j}), F(\nu_j|\boldsymbol{\nu}_{-j})) \cdot f(x|\boldsymbol{\nu}_{-j}) \tag{3.7}$$

where $\boldsymbol{\nu}$ is a d-dimensional vector, ν_j is an arbitrary chosen component of $\boldsymbol{\nu}$ and $\boldsymbol{\nu}_{-j}$ represents the $\boldsymbol{\nu}$ vector, excluding this component. In general, under appropriate regularity conditions, a joint density can be decomposed as product of bivariate copulas, acting on several different conditional probability distributions. For example, one possible pair-copula decomposition for $d = 3$ is

$$\begin{aligned} f(x_1, x_2, x_3) &= f_3(x_3) \cdot f(x_2|x_3) \cdot f(x_1|x_2, x_3) \\ &= f_3(x_3) \cdot \underbrace{f_2(x_2) \cdot c_{23}(F_2(x_2), F_3(x_3))}_{f(x_2|x_3)} \cdot \\ &\quad \underbrace{c_{12|3}(F(x_1|x_3), F(x_2|x_3)) \cdot c_{12}(F_1(x_1), F_2(x_2)) \cdot f_1(x_1)}_{f(x_1|x_2,x_3)} \end{aligned} \tag{3.8}$$

It is also clear that given a specific decomposition, there are still many different re-parame - terizations. Thus we need to introduce rules that will enable us to decompose a joint density into a cascade of pair-copulas. However, before introducing these rules, we need to introduce another concept related with pair-copula construction. The pair-copula construction involves marginal conditional distributions of the form $F(x|\boldsymbol{\nu})$ and we need a way to evaluate such marginal conditional distribution functions. Joe (1996) showed that for every u_j in the vector $\boldsymbol{\nu}$, $F(x|\boldsymbol{\nu})$ can be written as

$$F(x|\boldsymbol{\nu}) = \frac{\partial C_{x,\nu_j|\boldsymbol{\nu}_{-j}}(F(x|\boldsymbol{\nu}_{-j}), F(\nu_j|\boldsymbol{\nu}_{-j}))}{\partial F(\nu_j|\boldsymbol{\nu}_{-j})} \tag{3.9}$$

where $C_{x,\nu_j|\boldsymbol{\nu}_{-j}}$ an arbitrary copula distribution function. Following the notation of Aas et al. (2009) we will use the function $h(x,\nu;\boldsymbol{\theta})$ to represent the conditional distribution function when x and ν are uniforms. For $x,\nu \sim U[0,1]$ it holds

$$h_{\boldsymbol{\theta}} = \frac{\partial C_{\theta}(F_x(x),F_{\nu}(\nu))}{\partial F_{\nu}(\nu)} = \frac{\partial C_{\boldsymbol{\theta}(x,\nu)}}{\partial \nu} \tag{3.10}$$

3.2 Canonical vines (C-vines)

There is a significant number of possible pair-copula decompositions for high-dimensional distributions. It is therefore crucial to introduce rules that enable reducing this complexity. Bedford and Cooke (2001, 2002) have introduced a graphical way of organizing pair-copula density decompositions, known as *regular vines*. However, the class of regular vines is still very extended. Canonical vine (C-vine) and D-vine (Kurowicka and Cooke, 2006) constitute two special cases of regular vines. In this study, we concentrate on C-vines because have not been extensively investigated in financial applications so far.

The joint density decomposition, as organized by a C-vine, is given in the form of a nested set of star trees. For the d-dimensional C-vine, the pairs in level 1 are $(1,i)$, for $i = 2 \ldots d$, and for level ℓ, $2 \leq \ell < d$, the (conditional) pairs are $(\ell, i|1,\ldots,\ell-1)$ for $i = \ell+1,\ldots,d$. Figure 3.1 shows the specification for a six-dimensional C-vine. It consists of five trees, T_j, $j = 1\ldots,d-1$. Tree T_j has $d+1-j$ nodes and $d-j$ edges. Each edge corresponds to a bivariate copula density whereas the edge labels correspond to the subscript of the bivariate copula density. For example, edge $34|12$ corresponds to the conditional copula density $c_{34|12}(\cdot)$, where copula density $c_{34|12}(\cdot)$ can be of any parametric form. In total, $d(d-1)/2$ pair-copula families should be chosen for the whole decomposition. The nodes in tree T_j are necessary for determining the labels of next tree T_{j+1}.

As can be seen from Figure 3.1, there is one node in each tree which is connected with the remaining nodes of this particular tree. In principle, the order starts with a variable that has the highest dependency with all remaining variables, the *"pilot"* variable. Conditioning all remaining variables on first pilot variable, we can obtain the second variable with the highest dependency with all other variables. This approach is completed once we select all pilot variables for every tree (Czado et al., 2011). The d-dimensional canonical decomposition is given by Aas et al. (2009) as follows

$$f(\mathbf{x}) = \prod_{k=1}^{d} f_k(x_k) \times \prod_{i=1}^{d-1}\prod_{j=1}^{d-i} c_{i,i+j|1:(i-1)}\Big(F(x_i|x_1,\ldots,x_{i-1}), F(x_{i+j}|x_1,\ldots,x_{i-1})\Big) \tag{3.11}$$

For example, the six-dimensional canonical vine structure of Figure 3.1 can be written as

$$
\begin{aligned}
f(x_1, x_2, x_3, x_4, x_5, x_6) = \\
& f_1(x_1) f_2(x_2) f_3(x_3) f_4(x_4) f_5(x_5) f_6(x_6) \cdot c_{1,2}(F(x_1), F(x_2)) \\
& \cdot c_{1,3}(F(x_1), F(x_3)) \cdot c_{1,4}(F(x_1), F(x_4)) \cdot c_{1,5}(F(x_1), F(x_5)) \\
& \cdot c_{1,6}(F(x_1), F(x_6)) \cdot c_{2,3|1}(F(x_{2|1}), F(x_{3|1})) \cdot c_{2,4|1}(F(x_{2|1}), F(x_{4|1})) \\
& \cdot c_{2,5|1}(F(x_{2|1}), F(x_{5|1})) \cdot c_{2,6|1}(F(x_{2|1}), F(x_{6|1})) \cdot c_{3,4|1,2}(F(x_{3|1,2}), F(x_{4|1,2})) \\
& \cdot c_{3,5|1,2}(F(x_{3|1,2}), F(x_{5|1,2})) \cdot c_{3,6|1,2}(F(x_{3|1,2}), F(x_{6|1,2})) \\
& \cdot c_{4,5|1,2,3}(F(x_{4|1,2,3}), F(x_{5|1,2,3})) \cdot c_{4,6|1,2,3}(F(x_{4|1,2,3}), F(x_{6|1,2,3})) \\
& \cdot c_{5,6|1,2,3,4}(F(x_{5|1,2,3,4}), F(x_{6|1,2,3,4}))
\end{aligned} \tag{3.12}
$$

It is clear that the construction is iterative by nature, and given a specific decomposition in (3.11) there are as many as $\frac{d!}{2}$ possible different C-vine structures. As already mentioned, for a mixed C-vine copula model we need also to choose a bivariate copula family for each of the $d(d-1)/2$ pair-copulas. It is therefore quite important to develop selection rules that allow selecting an appropriate copula family for each edge in the C-vine model. However, most of these selection rules are based on estimated C-vines. Thus, we now turn to the parameter estimation in C-vines.

3.3 Inference for a C-vine model

Assume a vector $x_i = (x_{i,1}, \ldots, x_{i,T})^T$, $i = 1, \ldots, d$ of random variables at T points in time. Further, assume that the T observations of each random variable $x_{i,t}$ are independent over time and uniformly distributed on [0,1]. The assumption of independence, as Aas et al. (2009) note, is not a limiting assumption. Most of financial time series are serial correlated and thus univariate time-series models can be fitted to the margins and filter the time dependencies. As a result, the analysis can be continued with the residuals. The log-likelihood of a C-vine model can be written, according to Aas et al. (2009), as

$$
\ell(x; \boldsymbol{\theta}) = \sum_{j=1}^{d-1} \sum_{i=1}^{d-j} \sum_{t=1}^{T} \log \Big[c_{j,j+i|1,\ldots,j-1}\Big(F(x_{j,t}|x_{1,t}, \ldots, x_{j-1,t}), F(x_{j+i,t}|x_{1,t}, \ldots, x_{j-1,t}) \Big) \Big] \tag{3.13}
$$

where $F(x_{j,t}|x_{1,t}, \ldots, x_{j-1,t})$ and $F(x_{j+i,t}|x_{1,t}, \ldots, x_{j-1,t})$ are conditional distributions and determined by (3.9) and the h-function definition of (3.10).

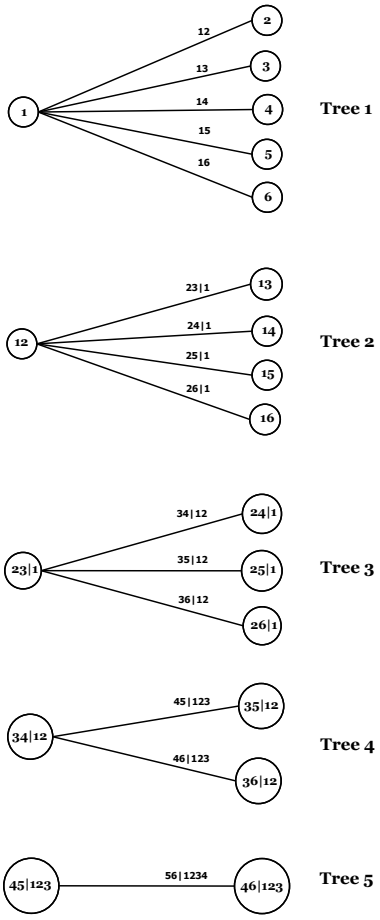

Figure 3.1: Tree representation of a canonical vine with 6 variables, 5 trees and 15 edges.

For each bivariate copula in the sum (3.13) there is at least one parameter to be determined. The number depends on the parametric assumption for each pair-copula in the C-vine model. For example, a Gaussian copula has one parameter whereas a Student-t copula has two parameters. If parametric margins are also estimated, i.e. $f_i(x_{i,t};\delta_i)$ with $i = 1,\dots,d$, the added contribution to (3.13) is

$$\sum_{t=1}^{T}\sum_{i=1}^{d} f_i(x_{i,t};\delta_i). \tag{3.14}$$

Under this setting, full MLE estimates can be obtained by maximizing (3.13) combined with (3.14) with respect to the parameters $(\boldsymbol{\theta},\delta_i,\dots,\delta_d)$. In general, the full MLE estimation would be the preferred choice of estimation due to the well-known optimality properties.

However, the Inference Functions for Margins (IFM) method is usually preferred to full MLE due to its computational tractability and comparable efficiency. The IFM method (see Joe and Hu (1996); Joe (1997)) is a multi-step optimization technique. It divides the parameter vector into separate parameters for each margin and parameters for the copula model. Therefore, one may break up the optimization problem into several smaller optimization steps with fewer parameters. For example, in the first step one may maximize the log-likelihood function of margins in (3.14) over the parameter vector $(\delta_i \ldots, \delta_d)$ and in the second step the log-likelihood of C-vine model in (3.13), given the estimated parameters of the margins from the first step, over the parameter vector $\boldsymbol{\theta}$. The IFM method has been found highly efficient compared to full MLE optimization for a number of multivariate models in a study by Joe (1997).

The IFM is a fully parametric method and thus misspecification of the marginal distributions may affect the performance of the estimator. In addition, marginal distributions are unknown almost always in practice. This fact increases the probability of marginal misspecification. The semi-parametric (SP) method, proposed by Genest et al. (1995), can tackle the marginal misspecification problem since treats the marginal distributions as unknown functions. The SP method is also known as *pseudo maximum likelihood* (PML) method. The PML method estimates each marginal distribution non-parametrically by the empirical distribution function (*edf*) without assuming any particular parametric distribution for the marginals. Once this is completed, the dependency structure between the marginals is estimated using a parametric multivariate copula family or in our particular study a C-vine model. Kim et al. (2007) have showed that MLE/IFM methods are non robust against marginals misspecification, and that the SP method performs better than the MLE and IFM methods, overall. Note that in our study the estimation of a C-vine model, corresponds to the IFM method when marginals are transformed to uniforms parametrically, or PML method when instead marginal transformations of the data are obtained non-parametrically.

We now turn to the estimation of C-vine models. C-vine parameters can be estimated using the sequential estimator (SE) or the maximum likelihood estimator (MLE). Following Czado et al. (2011) suppose i.i.d. data $\mathbf{u}_t = (u_{1,t}, \ldots, u_{d,t})^t$ for $t = 1, \ldots, T$ are available. For SE the parameters of unconditional bivariate copulas of tree 1 are first estimated. Then these estimates are used to estimate pair-copula parameters with single conditioning variable. The estimation proceeds tree by tree, since the conditional pairs in trees $2, \ldots, d-1$ depend on the specification of the previous trees via the h-function, defined in (3.10). Hence, C-vine models are estimated sequentially until all parameters are estimated. The estimation, under this context, can be carried out either using the inversion of Kendall's τ estimates for one parameter bivariate copulas or using MLE.

In particular, the parameter vector $\boldsymbol{\theta}_{1,j+1}$ of bivariate copula families $c_{1,j+1}$ in tree 1 are estimated using data $(u_{1,t}, u_{j+1,t})$, $t = 1, \ldots, T$ for $j = 1, \ldots, d-1$. Given the estimated parameter vector $\hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}$ of tree 1, we next want to estimate the parameter vector $\boldsymbol{\theta}_{2,j+1}$ corresponding to $c_{2,j+2|1}$ for $j = 1, \ldots, d-2$ in tree 2. Define

$$\hat{v}_{2|1,t} := F(u_{2,t}|u_{1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}_{1,1}) = h(u_{2,t}|u_{1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}{}_{1,1})$$

$$\hat{v}_{j+2|1,t} := F(u_{j+2,t}|u_{1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}_{1,j+1}) = h(u_{j+2,t}|u_{1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}{}_{1,j+1})$$

for $j = 1, \ldots, d-2$ and denote these estimates by $\hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}_{2,j}$. We can subsequently use data $(\hat{v}_{2|1,t}, \hat{v}_{j+2|1,t})$, $t = 1, \ldots, T$ to estimate $\boldsymbol{\theta}_{2,j}$ for $j = 1, \ldots, d-2$. For estimating the parameter vector $\theta_{3,j}$ corresponding to tree 3 and pair-copula families with double conditioning variables $c_{3,j+3|1,2}$ for $j = 1, \ldots, d-3$, we define

$$\hat{v}_{3|1,2,t} := h(\hat{v}_{2|1,t}|\hat{v}_{3|1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}_{2,1})$$

$$\hat{v}_{j+3|1,t} := h(\hat{v}_{j+3|1,t}|\hat{v}_{3|1,t}; \hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}_{2,j})$$

and estimate $\boldsymbol{\theta}_{3,j}$ using $(\hat{v}_{3|1,2,t}, \hat{v}_{j+3|1,2,t})$, $t = 1, \ldots, T$ for $j = 1, \ldots, d-3$.

Following the same reasoning, one can sequentially estimate the pair-copula parameters for each nested set of trees in C-vine structure until all unconditional and conditional bivariate copula parameters to be estimated. The sequential estimates have been recently found by Hobæk Haff (2010) asymptotically normal under some regularity conditions but their asymptotic covariance properties are intractable (Czado et al., 2011). To improve efficiently we can use MLE estimation. It is clear that MLE requires high-dimensional optimization of the log-likelihood and as such is much more time consuming compared to SE. Moreover, MLE requires good starting values of the parameters in the numerical maximization for quick convergence. Sequential estimates can be used as starting values for the optimization.

3.4 C-vine decomposition and Copula selection

In general, there are exactly $d!/2$ different C-vine structures available and the variation increases along with the dimension of the dataset. As a result, it is necessary to develop selection rules that uniquely decompose a C-vine model and provide an overall good fit to the data. So far, only empirical selection procedures exist for a C-vine decomposition (see Nikoloulopoulos et al. (2012); Czado et al. (2011)). The empirical selection rule of Czado et al. (2011) is followed in this study. The rule suggests a data driven sequential

approach to determine the pilot variable and the $d-1$ unconditional pair-copulas in level 1 of C-vine model. The rule works as follows. Estimate all possible pairwise Kendall's $\tau_{i,j}$ coefficients, noted $\hat{\tau}_{i,j}$, and find the variable i^* that maximizes

$$\hat{S}_i := \sum_{j=1}^{d} |\hat{\tau}_{i,j}| \tag{3.15}$$

over $i = 1, \ldots, d$. Once the most dependent with other variables, i^*, is selected for tree 1, we reorder the variables such as i^* to become the first variable. We can then link the pilot variable i^* with the remaining variables and select the unconditional pair-copulas $c_{1,j+1},\ j = 1, \ldots, d-1$. The selection of appropriate copula-families is based on multiple criteria such as graphical and analytical tools as well as goodness-of-fit tests. We will discuss this choice later and assume for the moment that we are able to select a pair-copula family for each unconditional bivariate copula $c_{1,j+1},\ j = 1, \ldots, d-1$. As in the sequential estimation procedure $d-1$ transformed variables are formed.

$$\hat{v}_{j+2|1,t} := h(u_{j+2,t}|u_{1,t}; \boldsymbol{\theta}^{\boldsymbol{SE}}_{1,j+1}) \quad j = 0, \ldots, d-2,\ t = 1, \ldots, T \tag{3.16}$$

Based on $d-1$ variables of size T all pairwise Kendall's τ coefficients are estimated and the pilot variable i^{**} of tree 2 that maximizes (3.15) is selected. Subsequently, we reorder the variables $i = 2, \ldots, d$ in such a way that i^{**} is variable 2. Having selected i^{**} as the pilot variable for tree 2, we can consequently select the copula families with single conditioning variable 1 $c_{2,j+2|1},\ j = 1, \ldots, d-2$. This procedure is continued until we determine the pilot variable for each tree as well as all corresponding pair-copulas and their associated sequential parameter estimates $\hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}$.

When modelling the dependency structure of random variables using copulas, the true copula is always unknown. Therefore, we need tools to specify an appropriate copula family to describe the observed dependence structure of the variables. For our sequential selection procedure we only require selection rules for bivariate copulas and hence our analysis will concentrate only on copula selection in two dimensions. The problem of selecting an appropriate (bivariate) copula family is a well studied problem in literature and many procedures have been proposed. In principle there are two different classes of tools that help in copula selection procedure, *graphical* and *analytical* tools. In our study, we use both sets of tool sets in copula selection procedure.

4 Model Setup

To compute risk estimates for given portfolios of power price series $1, \ldots, M$ on a daily basis, we choose a moving window size T for each series in the portfolio, a sample size $N = 1,000$ and portfolio weights $\omega_j, j = 1, \ldots, M$ with $\sum_{j=1}^{M} \omega_j = 1$ for equally weighted long only and $\sum_{j=1}^{M} \omega_j = -1$ for equally weighted short only portfolios, respectively.

(a) We fit an ARMA(p,q) - GARCH(m,r) model for each series in the portfolio. Therefore, for $j = 1, \ldots, M$ and $t = 1, \ldots, T$ we estimate:

$$r_{t,j} = \mu_j + \sum_{k=1}^{p} \phi_{k,j} r_{t-k,j} + \varepsilon_{t,j} + \sum_{k=1}^{q} \theta_{k,j} \varepsilon_{t-k,j}, \tag{4.1}$$

$$\sigma_{t,j}^2 = \omega_j + \sum_{k=1}^{m} \alpha_{k,j} \varepsilon_{t-k,j}^2 + \sum_{k=1}^{r} \beta_{k,j} \sigma_{t-k,j}^2, \tag{4.2}$$

$$\varepsilon_{t,j} = \sigma_{t,j} z_{t,j}, \tag{4.3}$$

where the error term $z_{t,j}$ in (4.3) is an i.i.d. sequence with zero mean, unit variance and distribution functions, denoted by $F_{\text{norm},j}$ and $F_{\text{t},j}$ (i.e., standard normal and standardized Student-t), respectively. Subsequently, we compute standardized residuals

$$\hat{z}_{t,j} = \frac{1}{\hat{\sigma}_{t,j}} \Big(r_{t,j} - \hat{\mu}_j - \sum_{k=1}^{p} \hat{\phi}_{k,j} r_{t-k,j} - \hat{\sigma}_{t,j} \hat{z}_{t,j} - \sum_{k=1}^{q} \hat{\theta}_{k,j} \hat{\sigma}_{t-k,j} \hat{z}_{t-k,j} \Big). \tag{4.4}$$

(b) For the C-vine-EVT and A-C-vine-EVT models we fix a high threshold level u for the upper and lower tails of residuals distribution and assume that excess residuals over this threshold follow the Generalized Pareto Distribution (GPD). The resulting piecewise semi-parametric distribution, denoted by $F_{\text{evt},j}$, encompasses the estimates of the parametric tails and the non-parametric kernel-smoothed interior.

(c) Transform the standardized residuals to copula data either parametrically (C-vine-Norm, C-vine-t) or non-parametrically (C-Vine-EVT, A-C-vine-EVT).

(d) Fit a canonical vine model to the set $\mathbf{u} = (\mathbf{u}_{t,1}, \ldots, \mathbf{u}_{t,M})$ obtained from previous step.

(e) For each $n = 1, \ldots, N$, we generate a sample of $\tilde{u}_1, \ldots, \tilde{u}_M$ uniforms from the estimated C-vine model, $\hat{C}_t$, of the previous step.

(f) Convert $\tilde{u}_1, \ldots, \tilde{u}_M$ to standardized residuals $(\hat{z}_{t+1,1}, \ldots, \hat{z}_{t+1,M})'$ using the inverse of the corresponding distribution function for each model, $\hat{F}_{\text{evt},j}^{-1}(\tilde{u}_j)$, $\hat{F}_{\text{norm},j}^{-1}(\tilde{u}_j)$ or $\hat{F}_{\text{t},j}^{-1}(\tilde{u}_j)$.

(g) Based on estimated conditional mean (4.1) and variance (4.2) of step (a), we compute

the ex-ante GARCH variance forecasts for $j = 1, \ldots, M$,

$$\sigma^2_{t+1,j} = \hat{\omega}_j + \sum_{k=1}^{m} \alpha_{k,j} \varepsilon^2_{t+1-k,j} + \sum_{k=1}^{r} \beta_{k,j} \sigma^2_{t+1-k,j}. \tag{4.5}$$

(h) The estimated ARMA parameters and the GARCH variance forecasts (4.5) are used to compute the ex-ante return forecast for $j = 1, \ldots, M$,

$$\hat{r}_{t+1,j} = \hat{\mu}_j + \sum_{k=1}^{p} \hat{\varphi}_{k,j} r_{t+1-k,j} + \hat{\sigma}_{t+1,j} \hat{z}_{t+1,j} + \sum_{k=1}^{q} \hat{\theta}_{k,j} \hat{\sigma}_{t+1-k,j} \hat{z}_{t+1-k,j}. \tag{4.6}$$

(i) The portfolio return forecast, $\hat{r}_{t+1,P}$, is given by

$$\hat{r}_{t+1,P} = \omega_1 \hat{r}_{t+1,1} + \ldots + \omega_M \hat{r}_{t+1,M}. \tag{4.7}$$

(j) Compute VaR and CVaR forecasts by taking α-quantiles of the portfolio return forecasts

$$\mathrm{VaR}_{t+1|t}(\alpha) = F^{-1}(q)\ \hat{r}_{t+1,P} \tag{4.8}$$

$$\mathrm{CVaR}_{t+1|t}(\alpha) = E(\hat{r}_{t+1,P} | \hat{r}_{t+1,P} > \mathrm{VaR}_{t+1|t}(\alpha)) \tag{4.9}$$

where $F^{-1}(q)$ is the qth quantile ($q = 1 - \alpha$) of the portfolio return forecasts, $\hat{r}_{t+1,P}$.

5 Data Description

This study focuses on power portfolios consisting of spot and futures Phelix power contracts traded at the European Energy Exchange (EEX). Phelix Futures are traded for the current week and the next four weeks (Phelix Week Futures), the current and the next nine months (Phelix Month Futures), the next eleven quarters (Phelix Quarter Futures) and the next six years (Phelix Year Futures). Year and quarter futures are fulfilled by cascading, i.e. futures contracts with longer delivery periods are replaced by equivalent futures contracts with shorter delivery periods on the last day of trading. Therefore, three exchange trading days before the beginning of delivery, year and quarter futures cascade into the respective quarter or month futures whereas the month Phelix futures remain tradeable during the delivery period and reach their expiry on the exchange trading day before the last delivery day. The power portfolios of the present analysis include time series of spot and futures contracts with delivery during the base (Phelix Baseload) and peak (Phelix Peakload) hours of each day. The choice of the portfolio constituents is based on their high level of trading interest and liquidity. The Baseload and Peakload portfolios include:

- One, two and three months ahead generic time series of daily electricity Phelix Baseload (F1BM, F2BM and F3BM) and Phelix Peakload (F1PM, F2PM and F3PM) futures prices.
- One and two quarters ahead generic time series of daily electricity Phelix Baseload (F1BQ and F2BQ) and Phelix Peakload (F1PQ and F2PQ) futures prices.
- One and two years ahead generic time series of daily electricity Phelix Baseload (F1BY and F2BY) and Phelix Peakload (F1PY and F2PY) futures prices.
- Historical time series of daily day-ahead electricity Phelix Baseload (SpotB) and Phelix Peakload (SpotP) prices.

Both datasets cover a time period of almost 8 years worth of data. The Baseload dataset covers the period from January 2, 2004 to February 7, 2012 whereas the Peakload dataset covers the period from January 2, 2004 to February 22, 2012. In total, there are 2,058 and 2,068 daily price observations for each series in Baseload and Peakload portfolios, respectively. However, the present study does not analyze the price levels of portfolio constituents directly, but instead focuses on log-returns of "generic time series".

Generic time series are artificially constructed time series, that represent the prices of futures with (approximately) same time to maturity. For example, a one-month-ahead generic Phelix Baseload futures is the price series that corresponds to the next-to-delivery contract. The employment of generic time series is useful in order to rule out the well-known Samuelson effect of futures contracts. In general, the variance of a future contract increases when the contract approaches maturity. This behavior is known as Samuelson effect and is observable both in price levels and log-returns of individual contracts. Since generic time series correspond always to futures contracts with same time before delivery, they will not exhibit the increasing volatility within one of the time series (Börger et al., 2007).

For generic return series, however, we transform the data further to account for the roll-over of the futures contracts and the large jumps in returns series that do not stem from price formation at exchange. For example, the price for an one-month-ahead generic on January 31, 2011 is the price of the Phelix Feb 2011 Month contract and the price of the same contract on the next trading day (February 01, 2011) is the the price of the Phelix Mar 2011 Month contract. When calculating the return between these two days may involve a significant jump in the return series due to the possibility of the products having different means. To overcome this problem, we apply an overlap of one day every time that a specific contract approaches its last trading day and a new contract comes into play. Another problem arises with the quarterly and yearly futures contracts because they cease trading

three business days before the delivery period. For example, the Phelix Apr 2011 Quarter contract is traded until March 29, 2011. The last return for this contract can be calculated using the prices on March 28 and 29, 2011. The return on March 29 is excluded since it is based on the price of Phelix Apr 2011 Quarter contract obtained on March 29, 2011 and that of Phelix Jul 2011 Quarter contract obtained on the next trading day, March 30, 2011. Therefore, as in the monthly contract case, we have an overlap of one day every time that a quarterly or yearly future contract reaches its last trading day and a new contract comes into play. Nevertheless, monthly and quarterly or yearly contracts, as explained above, expiry on different trading days within the same calendar month. Hence, when calculating returns of generic time series, the monthly return overlaps do not tally with that of quarterly or yearly return overlaps.

This problem is solved by excluding the monthly returns every time that correspond to the one day overlap of the quarterly or yearly contracts. Since we deal with multivariate time series and thus we are interested in the longest possible joint time series, we also delete trading days which are a statutory holiday or a day of the weekend. Phelix futures contracts are not traded during weekends or statutory holidays whereas Phelix day-ahead spot power contracts are traded every day of a single week during the entire year. Figures 5.1 display the actual and transformed generic return series of Baseload portfolio. It is clear from 5.1 the presence of price spikes due to the roll-over of the futures contracts. As such, if we do not take into account this fact and exclude them from our estimation sample, their presence will cause severe misspecification problems to our model.

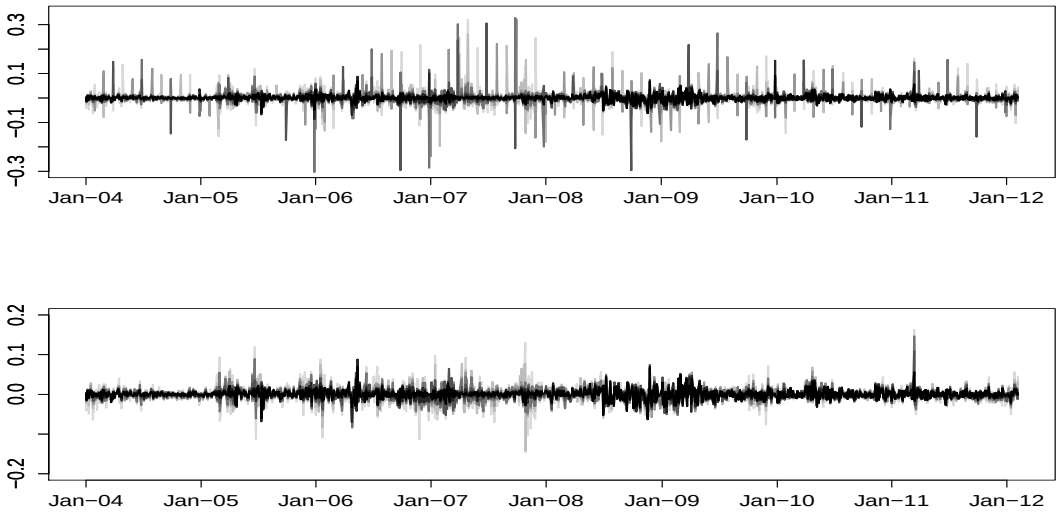

Figure 5.1: **Upper:** Log-returns of generic Phelix Baseload futures without adjustment for the roll-over of futures contracts. **Lower:** Log-returns of generic Phelix Baseload futures with adjustment for the roll-over of the futures contracts.

Tables 5.1 and 5.2 show the descriptive statistics of spot and generic futures returns, after excluding the artificial jumps from each futures contract. At first glance, we can spot that the number of observations in both samples has been reduced due to the adjustment in the futures contracts. The mean returns are close to zero and slightly negative for most of the series considered in both datasets. It is also evident the decreasing volatility behaviour as we move from spot series to futures contracts with longer time to maturity. In other words, the volatility term-structure can be read off across the different generics. Figure 5.2 display the one-month-ahead against the three-months-ahead generic Phelix Baseload futures returns (left) and Phelix Peakload futures returns (right). It can be shown that the one-month-ahead futures returns are more volatile compared to the three-months-ahead futures returns both in Baseload and Peakload portfolios. This behaviour can be also observed in Tables 5.1 and 5.2 by comparing the sample volatility estimates for each series in the two portfolios.

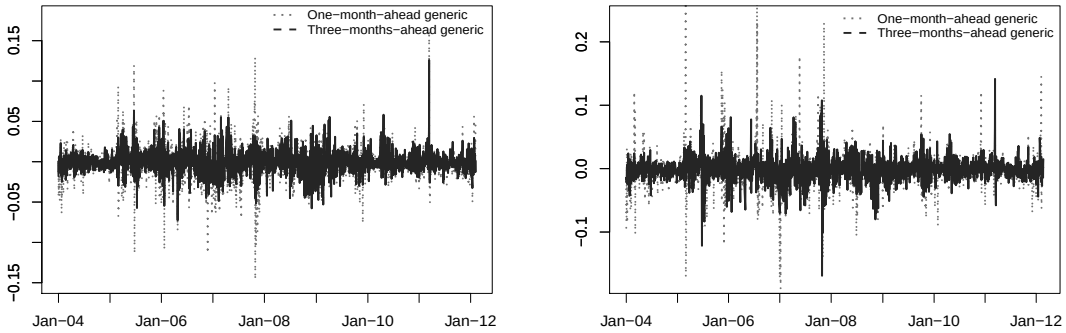

Figure 5.2: **Left:** Log-returns of generic Phelix Baseload one-month-ahead and three-months-ahead power futures. **Right:** Log-returns of generic Phelix Peakload one-month-ahead and three-months-ahead power futures.

The volatility of the spot series is the highest among the rest of the series in both portfolios whereas it declines gradually as the expiration of the futures contracts increases. The maximum and minimum returns of spot series reach extreme values in both portfolios, while the maximum and minimum returns of generic futures in the Peakload portfolio are higher compared to the values of the corresponding contracts in the Baseload portfolio. All return series are leptokurtotic with fat tails, as indicated by the positive values of (excess) kurtosis which is present in all series. The spot return series are characterized by negative skewness due to the extreme negative spikes that are present in the spot series whereas almost all generic return series in both portfolios have positive values of skewness. Furthermore, the high values of Jarque-Bera statistics and their respective zero p-values reject the assumption of normality for each series in the two portfolios.

	SpotB	F1BM	F2BM	F3BM	F1BQ	F2BQ	F1BY	F2BY
# of Obser.	1934	1934	1934	1934	1934	1934	1934	1934
Mean	0.0025	-0.0012	-0.0005	-0.0002	-0.0003	0.0000	0.0002	0.0002
Variance	0.0260	0.0005	0.0003	0.0002	0.0002	0.0002	0.0001	0.0001
Stand. Dev.	0.1613	0.0218	0.0173	0.0146	0.0144	0.0123	0.0112	0.0097
Min	-1.0764	-0.1461	-0.1423	-0.0727	-0.0836	-0.0615	-0.0705	-0.0634
Max	1.0964	0.1627	0.1489	0.1261	0.1463	0.1091	0.0884	0.0699
Exc. Kurtosis	6.8255	6.0707	7.4221	4.7930	7.9588	6.5901	6.3463	7.0465
Skewness	-0.0677	0.1914	0.2660	0.3097	0.5523	0.3647	0.0118	0.1531
JB	3766.7	29.7	4474.9	18.4	5217.3	3553.1	3255.4	4020.6
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 5.1: Descriptive statistics of spot and generic futures returns for Baseload portfolio

	SpotP	F1PM	F2PM	F3PM	F1PQ	F2PQ	F1PY	F2PY
# of Obser.	1945	1945	1945	1945	1945	1945	1945	1945
Mean	0.0031	-0.0019	-0.0019	-0.0010	-0.0004	-0.0007	-0.0003	-0.0001
Variance	0.0338	0.0008	0.0007	0.0004	0.0002	0.0002	0.0002	0.0001
Stand. Dev.	0.1838	0.0278	0.0272	0.0195	0.0158	0.0157	0.0123	0.0101
Min	-1.3757	-0.1898	-0.2263	-0.1691	-0.0918	-0.1457	-0.0886	-0.0615
Max	1.3233	0.3295	0.1944	0.1415	0.1549	0.1501	0.1254	0.0670
Exc. Kurtosis	8.8130	27.4211	9.5097	8.2933	9.3500	12.3454	10.4967	4.5930
Skewness	-0.1440	2.1953	-0.1580	0.0194	0.4320	0.2912	0.3861	0.0030
JB	6318.5	62640.4	7356.9	5589.6	7164.6	12410.5	9001.2	1715.4
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 5.2: Descriptive statistics of spot and generic futures returns for Peakload portfolio

6 Application to Phelix power portfolios[1]

In this section we implement the modelling framework, described in section 4, to forecast daily VaR and CVaR estimates. At first, we model the conditional mean and variance of each return series through time series models. An appropriate ARMA-GARCH model is selected for each individual series. For the C-vine-EVT models, a semi-parametric distribution is fitted to each series of standardized residuals. We transform the standardized residuals to uniforms either parametrically (C-vine-Norm and C-vine-t) or non-parametrically (C-vine-EVT and A-C-vine-EVT). The uniforms are used as input data for pair-copula construction. A canonical mixed vine copula structure is fitted to model the joint distribution of portfolio return series. Both sequential and full maximum likelihood estimates for the C-vine models are estimated. Based on the estimated C-vine structures and ARMA-GARCH models, next-period returns are simulated and VaR and CVaR estimates are computed.

[1]Due to space limitation we only present the results correspond to Phelix Baseload portfolio return series. Peakload portfolio results are available from the authors on request.

6.1 Univariate fitting of margins using ARMA-GARCH models

The modelling procedure for building an appropriate ARMA-GARCH model, especially when data are "messy", is not always straightforward. The stylized features of each time series analyzed in this study are very complex and hence the selection of an appropriate model is a non-trivial exercise. To select an appropriate model for the conditional mean and variance of each particular series we work as follows. Firstly, the AIC information criterion is applied to every return series of both portfolios and the models, as suggested by the criterion, are estimated. For the conditional variance equation a low-order GARCH(1,1) model is specified for each series and if cannot adequately model the volatility dynamics of the series an extra term is added each time until successfully removing all "ARCH-effects" from the squared residual series. If a model cannot remove the serial correlation from residual series due to significant higher order autocorrelation and partial autocorrelation lags then a suitable AR or MA term is added to the mean specification at the specific significant high lag orders, to address the seasonal behaviour, and the model is re-estimated and tested for its adequacy. The conditional mean equations for all series under consideration can be found in Appendix A. With respect to the conditional variance models, it seems that a GARCH(1,1) specification is adequate for all series. The estimated parameters of ARMA-GARCH models for almost all cases appear statistical significant at 5% level. Moreover, there are clear evidences that models' assumptions are satisfied[2]. Therefore, we assume that the selected time series models are adequate to model the conditional mean and variance of Phelix Baseload and Peakload portfolio return series, respectively.

6.2 Semi-parametric modelling of margins

The semi-parametric modelling of margins has the advantage that we have an explicit model for each tail. As such, we implement the peak over threshold (POT) method and fit the GPD to observations in the residual series that exceed a high threshold level u. The most important step in estimating the parameters of the GPD is the choice of the threshold u. Theorem 2.1 tells us that u should be high enough in order to approximate the conditional excess distribution by the GPD. On the other hand, the higher the threshold the less observations are left for the estimation of parameters and consequently the variance of the parameter estimates increases. So far, there is no formal way for optimal threshold level selection. In this study we follow McNeil and Frey (2000) and set the threshold level at the 8th and 92th percentiles of residuals distribution for the lower and upper tail, respectively.

[2]The estimation results and residual diagnostic tests are not presented due to space limitation but are available from the authors on request.

As a result of these threshold levels, 134 observations for Baseload portfolio residual series and 135 observations for Peakload portfolio residual series are used in the estimation process. Table 6.1 presents the percentage return threshold values u and maximum likelihood GPD estimates of the tail index ξ and the scale parameter σ with corresponding standard errors for upper and lower tails, respectively. The estimated tail index values for Phelix Baseload portfolio range between -0.124 (F2BQ, upper tail) and 0.180 (F1BY, lower tail) whereas for the Phelix Peakload portfolio range between -0.140 (F1PY, upper tail) and 0.138 (F2PM, upper tail). Recall that $\xi > 0$ corresponds to heavy-tailed distributions whose tails decay like power functions, such as the Pareto, Student's-t, Cauchy and Fréchet distributions. The case $\xi = 0$ corresponds to distributions whose tails decay exponentially. In this category belong the normal, exponential, gamma and lognormal distributions. Finally, the case $\xi < 0$ corresponds to short-tailed distributions with a finite right endpoint, such as the uniform and beta distributions (McNeil and Frey, 2000).

Baseload Portfolio										
	Lower tail					Upper tail				
Series	u (%)	$\hat{\xi}$	se(ξ)	$\hat{\sigma}$	se(σ)	u (%)	$\hat{\xi}$	se(ξ)	$\hat{\sigma}$	se(σ)
SpotB	-1.259	0.126	0.091	0.606	0.076	1.273	0.087	0.099	0.603	0.079
F1BM	-1.342	-0.037	0.078	0.549	0.064	1.340	-0.099	0.092	0.729	0.092
F2BM	-1.340	0.004	0.083	0.503	0.060	1.351	-0.018	0.091	0.650	0.081
F3BM	-1.337	-0.050	0.085	0.592	0.072	1.391	-0.072	0.070	0.610	0.068
F1BQ	-1.337	-0.060	0.095	0.580	0.075	1.391	0.000	0.090	0.577	0.072
F2BQ	-1.363	0.120	0.099	0.477	0.063	1.374	-0.124	0.086	0.612	0.074
F1BY	-1.279	0.180	0.096	0.530	0.068	1.368	-0.096	0.092	0.522	0.066
F2BY	-1.245	0.161	0.093	0.589	0.075	1.378	0.010	0.097	0.517	0.067
Peakload Portfolio										
	Lower tail					Upper tail				
Series	u	$\hat{\xi}$	se(ξ)	$\hat{\sigma}$	se(σ)	u	$\hat{\xi}$	se(ξ)	$\hat{\sigma}$	se(σ)
SpotP	-1.259	0.068	0.101	0.528	0.070	1.323	0.057	0.104	0.704	0.095
F1PM	-1.335	0.065	0.103	0.562	0.075	1.274	-0.013	0.106	0.727	0.099
F2PM	-1.365	-0.056	0.080	0.565	0.066	1.276	0.138	0.098	0.589	0.077
F3PM	-1.326	-0.087	0.091	0.672	0.084	1.285	0.096	0.108	0.612	0.084
F1PQ	-1.387	0.092	0.114	0.476	0.068	1.292	0.069	0.090	0.602	0.075
F2PQ	-1.323	-0.042	0.084	0.697	0.084	1.282	-0.008	0.073	0.613	0.069
F1PY	-1.278	0.110	0.092	0.568	0.071	1.333	-0.140	0.076	0.634	0.072
F2PY	-1.320	0.045	0.076	0.635	0.073	1.336	-0.014	0.106	0.552	0.075

Table 6.1: Maximum Likelihood Estimates (MLE) of the parameters of the Generalized Pareto Distribution (GPD) and Threshold percentage returns (u (%)) corresponding to 8% and 92% of empirical quantiles for Lower and Upper tail, respectively.

High values of the estimated tail index is an indication of extreme values since $\xi > 0$ reflects heavy-tailed distributions. According to tail index parameter estimates, most of the tail forms do not correspond to fat-tailed distributions. The highest positive tail index value equals to 0.180 and corresponds to F1BY Phelix Baseload series. This value, though, does

not imply a very fat-tail form. For example, an empirical study by Byström (2005), using extremely fat-tailed Nord Pool hourly electricity prices, found that a Fréchet distribution applies to the upper tail of standardized residuals. Similar results were also found by Chan and Gray (2006). The above evidences cannot strongly supported by our empirical GPD parameter estimates. The upper tail index GPD estimates for Spot Baseload and Peakload data are positive but their numerical values cannot reach the values of tail index parameters of the above studies. This is probably an indication that Phelix Spot market does not experience so many extreme price events compared to other electricity markets.

6.3 Selection and Estimation of C-vine models

Four different C-vine structures are estimated for both portfolios. Table 6.2 summarizes the specifications of the models investigated. The selection of copula families for each pair-copula in C-vine specification is mainly based on Akaike Information Criterion (AIC). There are two main reasons behind this choice. Firstly, it is practically impossible in large dimensions one to investigate every single unconditional and conditional pair-copula in the vine structure and define accordingly an appropriate copula family for each of these pairs. As a result, we use the AIC, which is the most frequently used criterion in copula selection literature. The range of all possible copula families employed by the criterion is defined in Appendix B. The second main reason that drives our copula selection strategy is related with the theoretical and empirical results of Joe et al. (2010) and Nikoloulopoulos et al. (2012) studies.

Joe et al. (2010) show that vine copulas can have a different upper and lower tail dependence for each bivariate margin when asymmetric bivariate copulas with upper/lower tail dependence are used in level 1 of the vine. In other words, in order for a vine copula to have tail dependence for all bivariate margins, it is necessary only for the bivariate copulas in level 1 to have tail dependence and not necessarily for the conditional bivariate copulas in levels $2, \ldots, d-1$ to have tail dependence. At levels 2 or higher, Independence or Gaussian copulas might be adequate to model the dependency structure. Moreover, Nikoloulopoulos et al. (2012) show that vine copulas with bivariate t linking copulas tend to be preffered by likelihood-based selection methods because they provide a better fit in the middle for the first level of the vine. They suggest that for inference involving the tails, the "best-fitting" copula should not be entirely likelihood-based but also depend on matching the non-parametric tail dependence measures and extreme quantiles. Taking into account these results, we also consider a hybrid to the C-vine-EVT model, where all t bivariate linking copulas selected by the AIC in level 1 are replaced by asymmetric copula families. If the empirical data present different degree of tail dependence, we expect to get more accurate risk measure estimates from vine models that allow asymmetries.

Model	Model type
C-vine-EVT	mixed C-vine model with semi-parametric margins
A-C-vine-EVT	same as C-vine-EVT model but all t-copulas in tree 1 are replaced by asymmetric copula families
C-vine-Norm	C-vine copula with all pair-copulas being Gaussian copulas and normal margins
C-vine-t	C-vine with all pair-copulas being t-copulas and Student-t margins

Table 6.2: Summary of models investigated for Phelix Baseload and Peakload portfolios.

6.3.1 Baseload selection and estimation results

After filtering the original return series with the appropriate ARMA-GARCH models, the resulting standardized residual series are transformed to uniform pseudo-observations. Figure 6.3 displays scatter plots and the estimated Kendall's τ coefficients of standardized residual series. The results highlight the positive dependence among generic residual series and the almost independence among the spot residual series and the rest of the series in the Phelix Baseload portfolio. Moreover, Figure 6.4 provide insights for the degree of tail dependence. In particular, Figure 6.4 displays Chi-plots of right upper and lower left quadrats of the series for detecting upper and lower tail dependence. Figure 6.4 illustrates a variety of tail dependence behaviour among the residual series. Most of the series do not show significant tail dependence or show symmetric tail dependence. However, there are few cases that display some degree of asymmetric tail dependence.

We apply now the sequential procedure of Czado et al. (2011) to select an appropriate C-vine copula model for the Phelix Baseload copula data. Table 6.3 reports the empirical Kendall's τ correlation matrix of the transformed residual series and the sum of their absolute values, denoted by $\hat{S}$ and defined in equation (3.14). According to Table 6.3, the i^* variable that maximizes the sum of absolute values, $\hat{S}$, is the F1BQ series and consequently it is placed as the pilot variable in level 1 of the vine structure. Table 6.4 presents the empirical Kendall's τ matrix of the series, conditioned on $i^* = \text{F1BQ}$, and the sum over the absolute entries of each row. The i^{**} variable that maximizes $\hat{S}$ is the F1BY series and is set as the pilot variable in level 2. Following the same identification procedure, the ordering of the variables for the mixed C-vine structure in Baseload portfolio is specified as follows

$$F1BQ - F1BY - F1BM - F2BY - F2BQ - SpotB - F3BM - F2BM$$

Table C.1 in Appendix C provides an overview of selected unconditional and conditional pairs for each level of the C-vine model. The above sequential procedure identifies not only an appropriate factorization for the mixed C-vine model but also the pair-copula families for each tree and provides sequential estimates $\hat{\boldsymbol{\theta}}^{SE}$.

	SpotB	F1BM	F2BM	F3BM	F1BQ	F2BQ	F1BY	F2BY	$\hat{S}$
SpotB	1.00	-0.02	-0.02	-0.03	-0.03	-0.03	-0.04	-0.05	1.20
F1BM	-0.02	1.00	0.66	0.58	0.65	0.48	0.45	0.39	4.22
F2BM	-0.02	0.66	1.00	0.68	0.75	0.57	0.54	0.48	4.70
F3BM	-0.03	0.58	0.68	1.00	0.79	0.61	0.58	0.50	4.76
F1BQ	-0.03	0.65	0.75	0.79	1.00	0.64	0.62	0.54	**5.02**
F2BQ	-0.03	0.48	0.57	0.61	0.64	1.00	0.72	0.63	4.67
F1BY	-0.04	0.45	0.54	0.58	0.62	0.72	1.00	0.75	4.70
F2BY	-0.05	0.39	0.48	0.50	0.54	0.63	0.75	1.00	4.35

Table 6.3: Empirical Kendall's τ matrix and the sum over the absolute entries of each row for the Phelix Baseload portfolio copula data.

Table 6.6 presents the resulting mixed C-vine-EVT model and sequential and maximum likelihood estimates. All pair-copula families are selected by the AIC without testing for independence whereas the sequential estimates are used as initial values to obtain maximum likelihood estimates. It can be seen that the sequential $\hat{\boldsymbol{\theta}}^{\boldsymbol{SE}}$ and maximum likelihood $\hat{\boldsymbol{\theta}}^{\boldsymbol{MLE}}$ estimates are pretty close for all estimated models. These findings support the employment of sequential estimation as the preffered optimization method. With respect to copula selection, 13 different copula types were selected for the 28 in total different pair-copulas in C-vine-EVT model. The majority of the selected copula families correspond to the t copula. The empirical results of our likelihood-based copula selection procedure seem to agree with the empirical findings of Nikoloulopoulos et al. (2012). In particular, 5 out of 8 selected copula families in level 1 belong to the t copula. Based on the above results, we specify the A-C-vine-EVT model by replacing the selected t copula families of level 1 with asymmetric copula families. For levels $2 \dots d-1$, the selection of the appropriate copula family is based on the AIC. We also test for independence in the C-vine model. The Independence copula is selected for pair-copulas that cannot reject the null hypothesis of independence.

	F3BM$\|i^*$	F1BY$\|i^*$	F2BM$\|i^*$	F2BQ$\|i^*$	F2BY$\|i^*$	F1BM$\|i^*$	SpotB$\|i^*$	$\hat{S}$
F3BM$\|i^*$	1.00	0.07	-0.01	0.08	0.02	-0.06	0.01	1.25
F1BY$\|i^*$	0.07	1.00	-0.00	0.48	0.59	-0.08	-0.02	**2.24**
F2BM$\|i^*$	-0.01	-0.00	1.00	0.01	-0.02	0.20	0.01	1.25
F2BQ$\|i^*$	0.08	0.48	0.01	1.00	0.38	-0.07	-0.01	2.03
F2BY$\|i^*$	0.02	0.59	-0.02	0.38	1.00	-0.10	-0.03	2.13
F1BM$\|i^*$	-0.06	-0.08	0.20	-0.07	-0.10	1.00	0.01	1.52
SpotB$\|i^*$	0.01	-0.02	0.01	-0.01	-0.03	0.01	1.00	1.09

Table 6.4: Empirical Kendall's τ matrix and the sum over the absolute entries of each row for the Phelix Baseload portfolio copula data.

The specification of appropriate asymmetric copula families in place of the selected t copula families in level 1 of the C-vine-EVT model is based on a set of analytical and graphical tools. In particular, we employ the Vuong (1989) and Clarke (2007) goodness-of-fit tests

as well as the empirical and theoretical contour plots and λ-function plots for assessing the fit of selected copula families to empirical data. Table 6.5 reports the Vuong and Clarke test results of the selected t pair-copulas in level 1 against copula families with different tail dependence. We employ copula families with upper only or lower only tail dependence and families that have upper tail dependence different from lower tail dependence. The goodness-of-fit results of Table 6.5 are in line with the AIC results of the C-vine-EVT model. Both tests provide the highest scores to t copula for each pair analyzed. The BB1 and SBB1 families obtain the second highest scores whereas the rest of copula families employed obtain negative scores for most of the pairs under consideration.

Test	Pairs	t	C	G	J	BB1	BB6	BB7	BB8	SBB1	SBB7
Vuong test	F1BQ-F1BY	8	-8	-1	-8	5	-3	-2	6	5	-2
	F1BQ-F1BM	9	-8	0	-8	6	-2	-1	-1	6	-1
	F1BQ-F2BQ	9	-8	0	-8	5	-2	-1	1	5	-1
	F1BQ-F3BM	9	-8	1	-8	6	-1	0	-5	6	0
	F1BQ-F2BM	9	-8	3	-8	6	1	-3	-3	6	-3
Clarke test	F1BQ-F1BY	9	-8	1	-8	5	-1	-4	5	5	-4
	F1BQ-F1BM	9	-8	3	-8	6	1	-4	-1	6	-4
	F1BQ-F2BQ	9	-7	2	-9	6	0	-4	1	6	-4
	F1BQ-F3BM	9	-8	3	-8	5	1	-2	-5	7	-2
	F1BQ-F2BM	9	-8	3	-8	6	1	-3	-3	6	-3

Table 6.5: Vuong and Clarke goodness-of-fit tests for copula family selection.

The above test results imply tail dependence in the pairs analyzed, since the copula families with the highest scores suggest symmetric (t copulas) or asymmetric (BB1 and SBB1) upper and lower tail dependence while copula families that have upper only or lower only tail dependence are disregarded. Therefore, it seems that the employment of asymmetric copula families are not necessary for our data set. As indicated by the Chi-plots of Figure 6.4 and the likelihood-based goodness-of-fit test results, it seems that there is a symmetric tail dependence and thus t copula is the appropriate copula family to model this behaviour.

However, we want to investigate whether the specification of asymmetric copula families in place of t selected copula families in level 1 of the vine structure can improve the model performance in VaR and CVaR forecasting. Thus, we replace the t copulas of F1BQ-F1BY, F1BQ-F1BM, F1BQ-F2BQ and F1BQ-F2BM pairs with the BB1 copula whereas for the F1BQ-F3BM pair we choose the SBB1 copula. Figures 6.1 and 6.2 compare the fit of t and BB1 copulas on the F1BQ-F1BM pair copula data through contour and λ-function plots[3].

[3]Similar graphs were also plotted for the rest of the pairs under investigation.

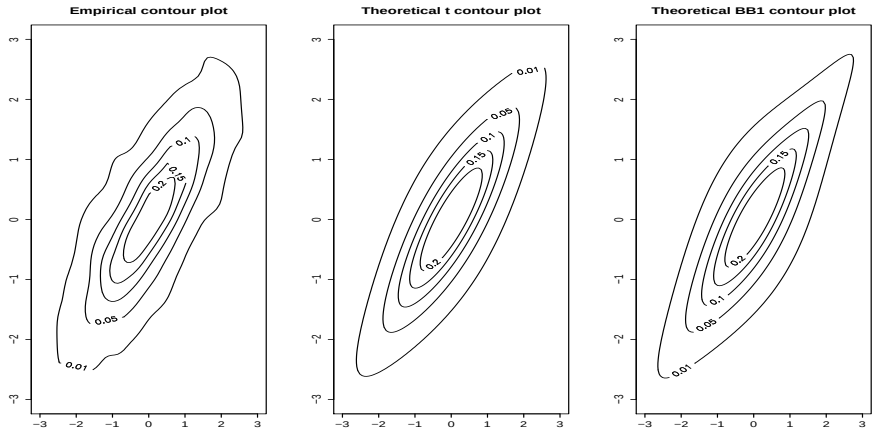

Figure 6.1: Empirical and Theoretical contour plots for assessing the fit of t and BB1 copulas on the transformed pair copula F1BQ-F1BY Phelix Baseload data.

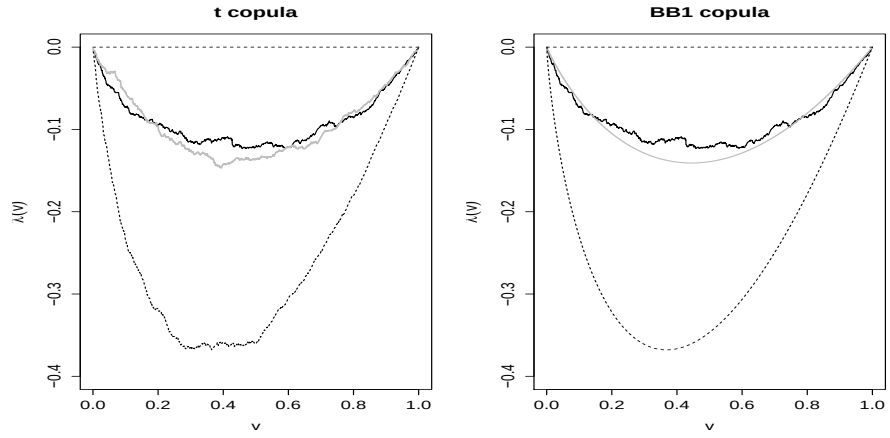

Figure 6.2: Plots of λ-functions for assessing the fit of t and BB1 copulas on the transformed pair copula F1BQ-F1BY Phelix Baseload data.

In general, the theoretical t contour plots for all pairs analyzed matching better the corresponding empirical contour plots than the theoretical BB1 contour plots. The same conclusions can be drawn from the λ-function plot comparisons. The simulated t copula λ-function plots seem to fit better the interior part of the bivariate distribution than the corresponding BB1 λ-function plots. However, the BB1 λ-functions appear adequate on fitting the tails of the distribution. Table 6.7 presents the resulting mixed A-C-vine-EVT model and sequential and maximum likelihood estimates.

Level	Block	Family	Param.	$\hat{\boldsymbol{\theta}}^{SE}$	$\hat{\boldsymbol{\theta}}^{MLE}$
1	$C_{1,2}$	t	$\hat{\rho}$	0.81	0.81
			$\hat{\nu}$	7.68	7.79
	$C_{1,3}$	t	$\hat{\rho}$	0.84	0.85
			$\hat{\nu}$	7.48	7.47
	$C_{1,4}$	F	$\hat{\theta}$	6.55	6.37
	$C_{1,5}$	t	$\hat{\rho}$	0.83	0.83
			$\hat{\nu}$	9.59	9.6
	$C_{1,6}$	BB8_90	$\hat{\theta}$	-1.08	-1.09
			$\hat{\delta}$	-0.97	-0.95
	$C_{1,7}$	t	$\hat{\rho}$	0.94	0.94
			$\hat{\nu}$	6.98	6.93
	$C_{1,8}$	t	$\hat{\rho}$	0.92	0.92
			$\hat{\nu}$	6.69	6.69
2	$C_{2,3\|1}$	N	$\hat{\rho}$	-0.12	-0.12
	$C_{2,4\|1}$	t	$\hat{\rho}$	0.79	0.8
			$\hat{\nu}$	6.31	6.33
	$C_{2,5\|1}$	t	$\hat{\rho}$	0.68	0.68
			$\hat{\nu}$	6.93	6.95
	$C_{2,6\|1}$	F	$\hat{\theta}$	-0.21	-0.18
	$C_{2,7\|1}$	t	$\hat{\rho}$	0.11	0.12
			$\hat{\nu}$	8.28	8.31
	$C_{2,8\|1}$	C 270	$\hat{\theta}$	-0.01	-0.01
3	$C_{3,4\|1,2}$	BB8_270	$\hat{\theta}$	-1.12	-1.13
			$\hat{\delta}$	-0.96	-0.94
	$C_{3,5\|1,2}$	F	$\hat{\theta}$	-0.17	-0.15
	$C_{3,6\|1,2}$	SC	$\hat{\theta}$	0.04	0.04
	$C_{3,7\|1,2}$	t	$\hat{\rho}$	-0.07	-0.09
			$\hat{\nu}$	9.34	9.38
	$C_{3,8\|1,2}$	SBB7	$\hat{\theta}$	1.29	1.29
			$\hat{\delta}$	0.26	0.24
4	$C_{4,5\|1,2,3}$	G	$\hat{\theta}$	1.04	1.04
	$C_{4,6\|1,2,3}$	t	$\hat{\rho}$	-0.04	-0.04
			$\hat{\nu}$	10.84	10.84
	$C_{4,7\|1,2,3}$	N	$\hat{\rho}$	-0.08	-0.08
	$C_{4,8\|1,2,3}$	t	$\hat{\rho}$	-0.03	-0.03
			$\hat{\nu}$	18.69	18.69
5	$C_{5,6\|1,2,3,4}$	C90	$\hat{\theta}$	-0.03	-0.03
	$C_{5,7\|1,2,3,4}$	t	$\hat{\rho}$	0.06	0.05
			$\hat{\nu}$	16.72	16.72
	$C_{5,8\|1,2,3,4}$	SJ	$\hat{\theta}$	1.05	1.05
6	$C_{6,7\|1,2,3,4,5}$	J	$\hat{\theta}$	1.02	1.01
	$C_{6,8\|1,2,3,4,5}$	SC	$\hat{\theta}$	0.01	0
7	$C_{7,8\|1,2,3,4,5,6}$	t	$\hat{\rho}$	0.01	0.01
			$\hat{\nu}$	6.1	6.09

Table 6.6: Estimated parameters of C-vine-EVT model for Phelix Baseload portfolio. $\hat{\boldsymbol{\theta}}^{SE}$ and $\hat{\boldsymbol{\theta}}^{MLE}$ correspond to sequential and maximum likelihood estimated parameters, respectively. Selected copula families are explained in Table B.1.

Level	Block	Family	Param.	$\hat{\boldsymbol{\theta}}^{SE}$	$\hat{\boldsymbol{\theta}}^{MLE}$
1	$C_{1,2}$	BB1	$\hat{\theta}$	0.49	0.49
			$\hat{\delta}$	1.91	1.91
	$C_{1,3}$	BB1	$\hat{\theta}$	0.54	0.54
			$\hat{\delta}$	2.05	2.05
	$C_{1,4}$	F	$\hat{\theta}$	6.55	6.55
	$C_{1,5}$	BB1	$\hat{\theta}$	0.58	0.58
			$\hat{\delta}$	1.95	1.95
	$C_{1,6}$	Ind.			
	$C_{1,7}$	SBB1	$\hat{\theta}$	0.29	0.29
			$\hat{\delta}$	3.76	3.76
	$C_{1,8}$	BB1	$\hat{\theta}$	0.56	0.56
			$\hat{\delta}$	2.87	2.87
2	$C_{2,3\|1}$	N	$\hat{\rho}$	-0.1	-0.1
	$C_{2,4\|1}$	t	$\hat{\rho}$	0.78	0.78
			$\hat{\nu}$	6.6	6.6
	$C_{2,5\|1}$	t	$\hat{\rho}$	0.68	0.68
			$\hat{\nu}$	6.97	6.97
	$C_{2,6\|1}$	Ind.			
	$C_{2,7\|1}$	t	$\hat{\rho}$	0.11	0.11
			$\hat{\nu}$	8.29	8.29
	$C_{2,8\|1}$	Ind.			
3	$C_{3,4\|1,2}$	BB8_270	$\hat{\theta}$	-1.15	-1.15
			$\hat{\delta}$	-0.97	-0.97
	$C_{3,5\|1,2}$	Ind.			
	$C_{3,6\|1,2}$	Ind.			
	$C_{3,7\|1,2}$	t	$\hat{\rho}$	-0.07	-0.07
			$\hat{\nu}$	10.08	10.08
	$C_{3,8\|1}$	SBB7	$\hat{\theta}$	1.29	1.29
			$\hat{\delta}$	0.26	0.26
4	$C_{4,5\|1,2,3}$	G	$\hat{\theta}$	1.04	1.04
	$C_{4,6\|1,2,3}$	Ind.			
	$C_{4,7\|1,2,3}$	N	$\hat{\rho}$	-0.09	-0.09
	$C_{4,8\|1,2,3}$	Ind.			
5	$C_{5,6\|1,2,3,4}$	Ind.			
	$C_{5,7\|1,2,3,4}$	t	$\hat{\rho}$	0.06	0.06
			$\hat{\nu}$	18.45	18.45
	$C_{5,8\|1,2,3,4}$	Ind.			
6	$C_{6,7\|1,2,3,4,5}$	Ind.			
	$C_{6,8\|1,2,3,4,5}$	Ind.			
7	$C_{7,8\|1,2,3,4,5,6}$	Ind.			

Table 6.7: Estimated parameters of A-C-vine-t model for Phelix Baseload portfolio. $\hat{\boldsymbol{\theta}}^{SE}$ and $\hat{\boldsymbol{\theta}}^{MLE}$ correspond to sequential and maximum likelihood estimated parameters, respectively. Selected copula families are explained in Table B.1.

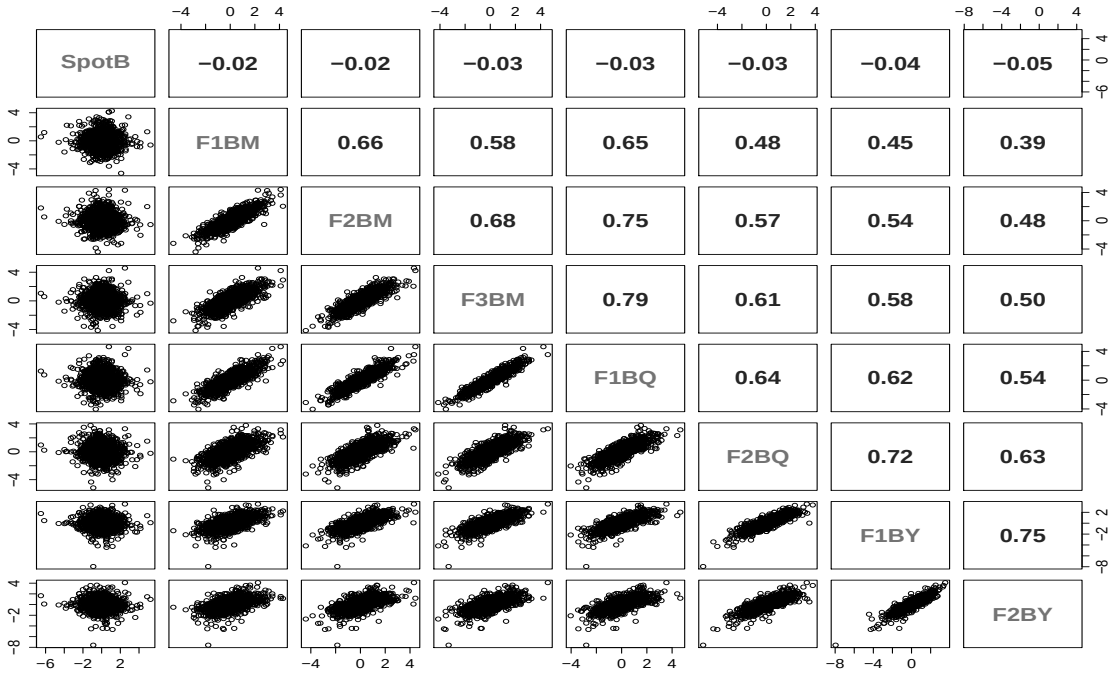

Figure 6.3: Dependence structure of standardized residuals in the Baseload portfolio. **Upper diagonal matrix:** Sample Kendall's τ correlation coefficients for each pair of standardized residual series. **Lower diagonal matrix:** Bivariate scatter plots for each pair of standardized residual series.

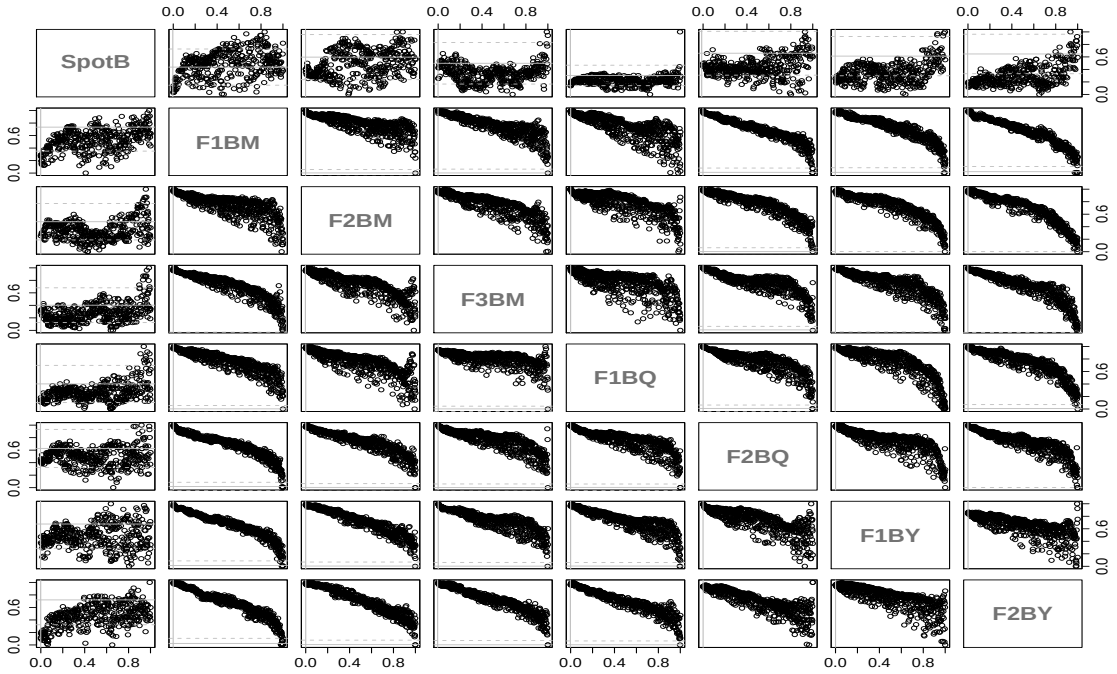

Figure 6.4: Tail dependence of transformed standardized residuals in the Baseload portfolio. **Upper diagonal matrix:** Chi-plots for each pair of transformed standardized residuals for right upper quadrat. **Lower diagonal matrix:** Chi-plots for each pair of transformed standardized residuals for left lower quadrat.

7 Empirical Results

We forecast one-day-ahead VaR and CVaR of equally weighted long and short only Phelix Baseload and Peakload portfolios. The competing models are specified in Table 6.2. We also employ the RiskMetrics (RM) model with smoothing constant $\lambda = 0.94$ as a naive benchmark. For model backtesting we use a period of 250 trading days that corresponds approximately to one year of trading. The forecasting period for the Baseload portfolio corresponds to the period from January 25, 2011 to February 07, 2012 and that of Peakload portfolio spans the period from February 10, 2011 to February 22, 2012. We evaluate all risk estimates at 1% and 5% confidence levels, since they constitute the most common levels used both in literature and financial markets for model evaluation. Figure 7.1 display the Baseload portfolio return series and the 95%-VaR forecasts for mixed C-vine-EVT and A-C-vine-EVT models. For a backtesting period of 250 observations and confidence levels 1% and 5% we expect 2.5 and 12.5 exceedances, respectively. According to Figure 7.1, both models seem to respond well in volatility changes and produce acceptable number of failures. However, the VaR performance for each single model is hard to be assessed visually and hence a two-stage selection procedure, similar to Sarma et al. (2003), is followed. In the first stage, all models are tested for statistical accuracy and, if survive rejection, a second stage filtering of the surviving models is employed using subjective loss functions.

7.1 Statistical Tests

The first-stage of the model selection procedure is useful in order to examine whether the VaR estimates coming out from alternative models satisfy the appropriate theoretical statistical properties. A well-specified VaR model should produce statistical meaningful VaR forecasts. As such, the proportion of exceedances should approximately equal the VaR confidence level (unconditional coverage) while the exceedances should not occur in clusters but instead independently. For example, a well-specified model should produce low VaR forecasts in times of low volatility and high VaR forecasts in times of high volatility, so that exceedances are spread over the entire sample period, and do not come in clusters. Therefore, a model which fails to capture the volatility dynamics of the underlying return distribution will suffer from clustering of failures, even if it may produce the correct unconditional coverage. The term "conditional coverage" include both properties. A wide range of tests have been proposed in the literature to test for these two properties. In principle, there are no universal guidelines on which tests to employ, since each test shares certain merits and demerits. As such, we employ the Kupiec (1995) unconditional coverage test, the conditional coverage test by Christoffersen (1998) and the duration-based Weibull test of independence by Christoffersen and Pelletier (2004) that appear very popular in the literature for testing the above two properties.

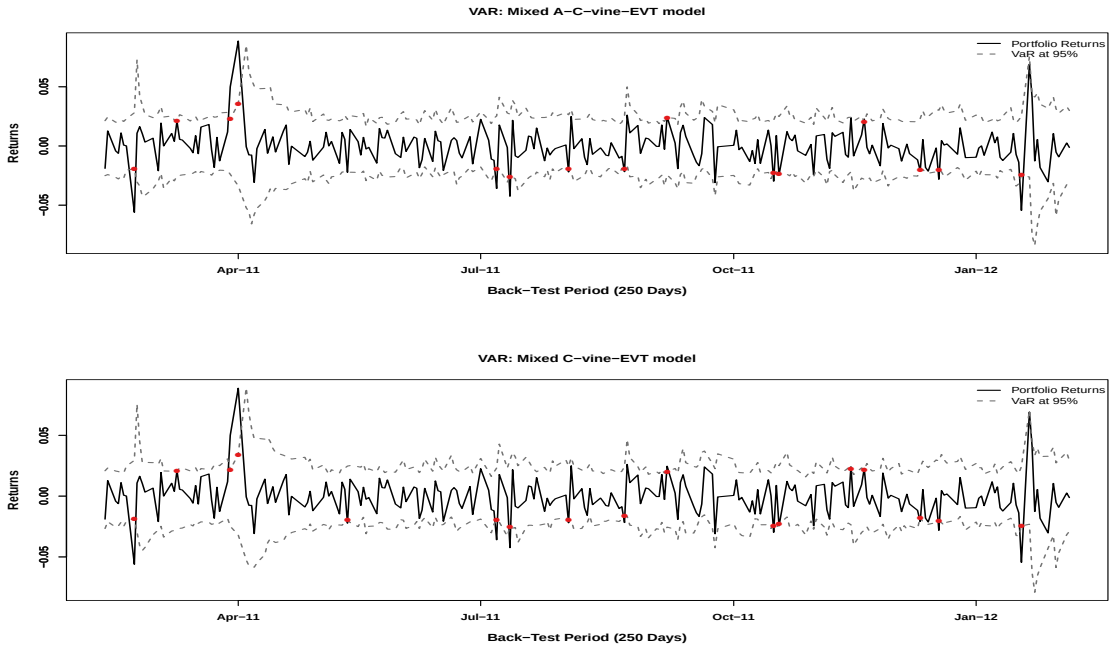

Figure 7.1: Baseload portfolio returns and 95%-VaR forecasts for the A-C-vine-EVT and C-vine-EVT models. VaR exceedances are marked by red circles

7.1.1 Kupiec (1995) test for unconditional coverage ($\mathbf{LR_{uc}}$)

The most well-know test based on failure rates has been proposed by Kupiec (1995). The test is a likelihood-ratio test and measures weather the number of exceedances is consistent with the confidence level. Under the null hypothesis that the model is "well-specified", the number of exceedances should follow the binomial distribution. Therefore, the idea is to examine whether the observed failure rate $\hat{\pi}$ is significantly different from α, the failure rate implied by the confidence level. The likelihood-ratio test statistic is given by:

$$LR_{uc} = -2\log\left[\frac{\alpha^{n_1}(1-\alpha)^{n_0}}{\hat{\pi}^{n_1}(1-\hat{\pi})^{n_0}}\right] \sim \chi^2_{(1)}, \tag{7.1}$$

where n_1 number of exceedances, n_0 number of non-exceedances, α is the confidence level at which VaR measures are estimated and $\hat{\pi} = n_1/(n_0 + n_1)$ is the MLE estimate of α. Under the null hypothesis, LR_{uc} is asymptotically χ^2 distributed with one degree of freedom. Nevertheless, Kupiec's test encounters two major shortcomings. First, the test is statistically weak with sample sizes consistent with the current regulatory framework (one year). Second, the test only considers the frequency of exceedances and do not take into account the time when they occur. This may fail to reject a model which suffers from clustered exceedances. Therefore, a backtesting should not rely entirely on unconditional coverage but clustered exceedances should be also taken into account.

7.1.2 Christoffersen (1998) test for conditional coverage ($\mathbf{LR_{cc}}$)

Christoffersen (1998) proposed a "conditional coverage" test that jointly examines whether the total number of exceedances is equal to the expected one, and the VaR failures are independently distributed. The test is carried out, given the realization of return series r_t and the *ex-ante* VaR for a $\alpha\%$ coverage rate by first defining an indicator function I_{t+1} that gets the value of 1 if a VaR violation occurs and 0 otherwise. If the model is "well-specified", then an exception today should not depend on whether or not an exception occurred on the previous day whereas the total number of exceedances should be consistent with the confidence level. Therefore, by combining the test statistic of independence (LR_{ind}) with Kupiec's test statistic (LR_{uc}), we obtain a conditional coverage test $LR_{cc} = LR_{uc} + LR_{ind}$ that jointly tests these two properties - i.e. the correct VaR failures and the independence of exceedances. The LR_{cc} test statistic for the correct conditional coverage is given by:

$$LR_{cc} = -2\log\frac{(1-\alpha)^{n_0}\alpha^{n_1}}{(1-\hat{\pi}_{01})^{n_{00}}\ \hat{\pi}_{01}^{n_{01}}\ (1-\hat{\pi}_{11})^{n_{10}}\ \hat{\pi}_{11}^{n_{11}}} \sim \chi^2_{(2)}, \tag{7.2}$$

where n_{ij} are the number of i values followed by a j value in the I_{t+1} series $(i, j = 0, 1)$, $\pi_{ij} = \Pr\{I_{t+1} = i | I_t = j\} \quad (i, j = 0, 1)$, $\hat{\pi}_{01} = n_{01}/(n_{00}+n_{01})$, $\hat{\pi}_{11} = n_{11}/(n_{00}+n_{01})$. LR_{cc} is also χ^2 distributed but in this case with two degrees of freedom. The above tests are reliable for detecting misspecified risk models when the temporal dependence in the sequence of VaR violations is a simple first-order Markov structure. However, we are interested in tests that have power against more general forms of dependence. The duration-based Weibull test of Christoffersen and Pelletier (2004) deals with this issue.

7.1.3 Christoffersen and Pelletier (2004) duration based test

The intuition behind the duration-based tests is that the clustering of exceedances will result in an excessive number of relatively short or relatively long no-hit durations (i.e., the duration of time (in days) between two VaR violations), corresponding to market turmoil and market calm, respectively. Following Christoffersen and Pelletier (2004), we denote by d_i the duration between two consecutive VaR violations (i.e., the no-hit duration):

$$d_i = t_i - t_{i-1},$$

where t_i denotes the day of the i^{th} exceedance. Under the conditional coverage hypothesis, the duration process d_i follows a geometric distribution with probability equal to α and a probability density function, given by:

$$f(d;\alpha) = \alpha(1-\alpha)^{d-1} \qquad d \in \mathbb{N}^*. \tag{7.3}$$

The general idea of the test is to specify a distribution that nests equation (7.3), so that the memoryless property can be tested via parameter restriction (Hurlin and Perignon, 2012). The only memory free (continuous) random distribution is the exponential, thus under the null hypothesis the distribution of the no-hit durations should be

$$f_{\text{exp}}(d; p) = p \, \exp(-pd). \tag{7.4}$$

In (7.4) we have $\mathbb{E}(d) = 1/\alpha$ and, since the conditional coverage hypothesis implies a mean of duration equals to $1/\alpha$, it also implies the condition $p = a$. With regards to statistical testing for independence, Christoffersen and Pelletier (2004), specify an alternative distribution that allows dependence. They consider the Weibull distribution where

$$f_{\text{w}}(d; p, b) = p^b b d^{b-1} \, \exp\Big((-pd_i)^b\Big). \tag{7.5}$$

This particular type of distribution is able to capture violation clustering. The Weibull will have a decreasing hazard fuction when $b < 1$ that corresponds to a high number of very short durations (very volatile period) and a high number of very long durations (low volatile period). This can be an evidence of misspecified volatility dynamics in the risk model (Christoffersen and Pelletier, 2004).

Since the exponential distribution corresponds to a Weibull with a flat hazard function, i.e., $b = 1$, the test for independence (Christoffersen and Pelletier, 2004) is :

$$H_{0,\text{ind}} : \quad b = 1. \tag{7.6}$$

Berkowitz et al. (2009) extended this approach to consider for the conditional coverage hypothesis, which is:

$$H_{0,\text{cc}} : \quad b = 1 \ \& \ p = \alpha. \tag{7.7}$$

7.2 Loss Functions

The idea of employing loss functions for assessing VaR performance was firstly introduced by Lopez (1998, 1999). The loss function evaluation method is not based on hypothesis-testing framework, but rather on assessing to VaR estimates a numerical score that reflects the evaluator's specific concerns. As such, it provides a measure of relative performance that can be utilized to assess the performance of VaR estimates. Under this approach, a model which minimizes the loss is preferred to other models (Lopez, 1998). The general form of a loss function for the i^{th} model is:

$$L_{i,t+1} = \begin{cases} f(r_{t+1}, \text{VaR}_{i,t+1|t}) & \text{if} \quad r_{t+1} < \text{VaR}_{i,t+1|t}, \\ g(r_{t+1}, \text{VaR}_{i,t+1|t}) & \text{if} \quad r_{t+1} \geq \text{VaR}_{i,t+1|t}, \end{cases} \tag{7.8}$$

where $f(x, y)$ and $g(x, y)$ are functions such that $f(x, y) \geq g(x, y)$ for a given y. Lower values of $L_{i,t+1}$ are preferred because exceedances are given higher scores than non exceedances. The total numerical score for each risk model i, is obtained by summing out the individual loss function estimates for the complete regulatory sample.

Under general conditions, accurate VaR estimates generates the lowest possible numerical scores. This approach for evaluating VaR performance is very flexible since the loss function can take different forms that represent the evaluator's specific concerns. However, this approach is vulnerable to misspecification of the loss function. Various loss functions have been proposed in the literature in order to address the regulatory specific concerns. In this study, we employ the magnitude loss function (MLF) of Lopez (1998), the regulatory loss function (RLF) of Sarma et al. (2003) and the predictive quantile loss function (PQLF) of Koenker and Bassett (1978) in order to evaluate the VaR performance of our risk models. We also evaluate the CVaR performance by employing the mean absolute error (MAE) and mean squared error (MSE) loss functions of Angelidis and Degiannakis (2007).

The magnitude loss function of Lopez (1998) incorporates two main regulatory concerns: the magnitude as well as the number of exceedances. The magnitude loss function has the following general representation:

$$QL_{i,t+1} = \begin{cases} 1 + (r_{i,t+1} - \text{VaR}_{i,t+1|t})^2 & \text{if} \quad r_{i,t+1} < \text{VaR}_{i,t+1|t}, \\ 0 & \text{if} \quad r_{i,t+1} \geq \text{VaR}_{i,t+1|t}. \end{cases} \tag{7.9}$$

A score of one is given when a violation occurs and an additional quadratic term, based on its magnitude, is also included. The numerical score increases with the magnitude of the exception and hence can provide additional information about the VaR relative performance. The second type of loss function we apply is the regulatory loss function of Sarma et al. (2003), which is similar to the magnitude loss function of Lopez (1998). This function does not penalizes for the number of exceedances but only for the magnitude of the failure. It gets the following functional form:

$$RLF_{i,t+1} = \begin{cases} (r_{i,t+1} - \text{VaR}_{i,t+1|t})^2 & \text{if} \quad r_{i,t+1} < \text{VaR}_{i,t+1|t}, \\ 0 & \text{if} \quad r_{i,t+1} \geq \text{VaR}_{i,t+1|t}. \end{cases} \tag{7.10}$$

As in loss function (7.9), the quadratic term in (7.10) ensures that the large VaR exceedances are penalized more than the small VaR exceedances. The predictive quantile loss function of Koenker and Bassett (1978) penalizes more heavily observations for which an exception occurs, and represents a measure of the predicted tail at a given confidence level. The loss function may taken the following functional form:

$$PQLF_{i,t+1} = \begin{cases} |r_{i,t+1} - \text{VaR}_{i,t+1|t}| \cdot (1-\alpha) & \text{if} \quad r_{i,t+1} < \text{VaR}_{i,t+1|t}, \\ |r_{i,t+1} - \text{VaR}_{i,t+1|t}| \cdot \alpha & \text{if} \quad r_{i,t+1} \geq \text{VaR}_{i,t+1|t}. \end{cases} \tag{7.11}$$

A backtest based on this loss functional form is carried on by calculating the sample average loss (T observations):

$$PQLF_i = \frac{1}{T} \sum_{j=1}^{T} PQLF_{i,t+j}. \tag{7.12}$$

The economic intuition behind the PQLF is that capital charges should also be considered in the evaluation of VaR performance. Therefore, the capital foregone for over-predicting the true VaR should not be neglected.

The above loss functions take into account only the magnitude and do not consider the size of expected loss, given a VaR violation. Since the size of expected loss is also an integral part in regulator's utility function, we backtest the risk models performance by using loss functions that take into account the Expected Shortfall (ES) or Conditional-VaR (CVaR), if a VaR violation occurs. As such, we employ the mean absolute error (MAE) and mean squared error (MSE) loss functions of Angelidis and Degiannakis (2007). The MAE loss function can be described by the following equation:

$$\Psi^1_{i,t+1} = \begin{cases} |r_{i,t+1} - \text{CVaR}_{i,t+1|t}| & \text{if} \quad r_{i,t+1} < \text{VaR}_{i,t+1|t}, \\ 0 & \text{if} \quad r_{i,t+1} \geq \text{VaR}_{i,t+1|t}, \end{cases} \tag{7.13}$$

whereas the MSE loss function can obtain the following functional form:

$$\Psi^2_{i,t+1} = \begin{cases} (r_{i,t+1} - \text{CVaR}_{i,t+1|t})^2 & \text{if} \quad r_{i,t+1} < \text{VaR}_{i,t+1|t}, \\ 0 & \text{if} \quad r_{i,t+1} \geq \text{VaR}_{i,t+1|t}. \end{cases} \tag{7.14}$$

A backtest based on loss functional forms of equations (7.13) and (7.14) is carried on by calculating the sample average expected loss (T observations) for MAE and MSE:

$$MAE_i = \frac{1}{T} \sum_{j=1}^{T} \Psi^1_{i,t+j}, \qquad MSE_i = \frac{1}{T} \sum_{j=1}^{T} \Psi^2_{i,t+j}.$$

7.3 Statistical Test Results

Table 7.1 presents the number of exceedances and p-values of the statistical tests for the Phelix Baseload and Peakload long and short only portfolios. It is evident from the p-values of the tests that almost all VaR forecasts at 1% and 5% confidence levels are sufficiently accurate. The null hypotheses of independence (Weibull), unconditional (uc) and conditional coverage (cc) cannot be rejected at 5% significance level for almost all risk models. The only exceptions are the rejection of the null hypothesis for the conditional coverage test at the 5% significance level (but not at 1%) for the Phelix Baseload short portfolio 99%-VaR forecasts for all C-vine models and the rejection of the null hypothesis of independence in Christoffersen and Pelletier test for C-vine-t model 99%-VaR forecasts for Phelix Peakload long portofolio.

The rejection of the conditional coverage null hypothesis, at 5% significance level, for all C-vine models 99%-VaR forecasts of Phelix Baseload short portfolios is due to the first-order Markov structure of independence test, which is embedded in the conditional coverage test. As explained, the unconditional and conditional coverage tests have low statistical power with the current regulatory backtest sample. This weakness is the reason for the rejection of the conditional coverage tests in our particular case. Nevertheless, the hypothesis of independence is pretty strong, since the Weibull independence test cannot be rejected at 5% level for any risk model except for the C-vine-t model 99%-VaR forecasts for Phelix Peakload long portofolio. This exception, however, is not of great concern since there is only 1 exceedance in this level and this affects the respective test statistic. The corresponding unconditional and conditional coverage tests cannot be rejected.

Moreover, the number of exceedances for all risk models are very close to the expected number of exceedances and therefore the hypothesis of unconditional coverage cannot be rejected. The p-values of the corresponding Kupiec tests are very high for all VaR forecast levels and all risk models. As expected, the benchmark RM model produces most of the times the highest number of exceedances without being rejected though by any statistical test at 5% level. It seems also that the C-vine-EVT and A-C-vine-EVT models obtain the highest p-values of unconditional coverage tests at 1% level whereas for the 5% VaR forecast level, no clear conclusion can be drawn on the basis of Kupiec test. In this stage, all VaR model forecasts can be considered as statistically adequate and hence can be subsequently used in the second stage of backtesting procedure, which entails evaluation through the employment of subjective loss functions introduced in section 7.2.

7.4 Loss Function Results

The second-stage of the selection procedure entails the employment of loss functions in order to assess the VaR and CVaR forecast performance of various risk models under consideration. Tables 7.2 and 7.3 report the average out-of-sample VaR and CVaR estimates, the numerical scores for the VaR-based regulatory (RLF), magnitude (MLF) and predictive quantile (PQLF) loss functions as well as the numerical scores of the CVaR-based MAE and MSE loss functions for Phelix Baseload and Peakload long and short only portfolios, respectively.

The numerical scores of VaR-based loss functions are pretty supportive for the A-C-vine-EVT model at the 1% confidence level. In particular, almost all numerical scores for every possible Phelix Baseload and Peakload long and short only portfolio combinations tend to favor the A-C-vine-EVT model over the rest of the models at 1% significance level. The A-C-vine-EVT model's numerical scores are minimized in this level in 11 out of 12 cases for all VaR-based loss functions employed. C-vine-EVT model also produces satisfactory results in this level. Moreover, the CVaR-based loss function resutls favour the A-C-Vine-EVT model. In total, 4 out of 8 times the corresponding MAE and MSE loss functions are minimized by the A-C-vine-EVT model.

Therefore, the statistical test results in the previous section and the loss function findings in this section, support our modelling approach for the extreme quantiles. It seems that a mixed A-C-vine structure can successfully describe the dependency structure of portfolio return series while the employment of extreme value theory can provide better fit in the the tails. Moreover, these findings support the theoretical and empirical findings of Joe et al. (2010) and Nikoloulopoulos et al. (2012). Based on loss function results at 1% level, it seems that VaR and CVaR forecasts produced by the A-C-vine-EVT model outperfom those of C-Vine-EVT model. Therefore, for inference involving the tails, the pair-copula selection procedure should not be entirely likelihood-based.

At 5% level the numerical score results seem to be mixed, since there is no strong preference towards a specific risk model. In general, the C-vine-Norm model estimates tend to minimize most of the times the VaR-based loss functions but there is no such a case where the C-vine-Norm model is preffered by all three loss functions simultaneously. The A-C-vine-EVT and C-vine-EVT models tend to minimize the VaR-based loss functions in few cases too, with the C-vine-t and RM being the least preferred models according to the VaR-based loss fuction numerical scores in this level. The CVaR-based loss function results at 5% level do not also provide more insights.

		Phelix Baseload Long Portfolio				Phelix Baseload Short Portfolio			
Model	**Level**	**No. of exceed.**	**Kupiec (uc)**	**Christoff. (cc)**	**Weibull (ind.)**	**No. of exceed.**	**Kupiec (uc)**	**Christoff. (cc)**	**Weibull (ind.)**
RM	99%	4	0.38048	0.63771	0.93515	3	0.75799	0.06329	1.00000
	95%	14	0.66907	0.39618	0.74427	13	0.88535	0.36735	0.57266
C-vine-EVT	99%	3	0.75799	0.91938	0.25047	2	0.74193	0.02235	1.00000
	95%	11	0.65706	0.54490	0.26182	8	0.16322	0.18966	0.19264
A-C-vine-EVT	99%	3	0.75799	0.91938	0.25047	2	0.74193	0.02235	1.00000
	95%	10	0.45291	0.49648	0.83140	7	0.08281	0.08830	0.30945
C-vine-Norm	99%	4	0.38048	0.63771	0.93515	2	0.74193	0.02235	1.00000
	95%	9	0.28602	0.40385	0.70539	7	0.08281	0.08830	1.00000
C-vine-t	99%	4	0.38048	0.63771	0.93515	2	0.74193	0.02235	1.00000
	95%	11	0.65706	0.54490	0.26182	9	0.28602	0.05618	1.00000

		Phelix Peakload Long Portfolio				Phelix Peakload Short Portfolio			
Model	**Level**	**No. of exceed.**	**Kupiec (uc)**	**Christoff. (cc)**	**Weibull (ind.)**	**No. of exceed.**	**Kupiec (uc)**	**Christoff. (cc)**	**Weibull (ind.)**
RM	99%	6	0.05935	0.14575	0.86344	6	0.05935	0.05032	0.56562
	95%	16	0.32937	0.62119	0.71425	15	0.48124	0.43564	0.18624
C-vine-EVT	99%	3	0.75799	0.91938	0.52624	4	0.38048	0.08733	1.00000
	95%	10	0.45291	0.49648	0.39991	10	0.45291	0.53023	0.82869
A-C-vine-EVT	99%	4	0.38048	0.63771	0.78387	3	0.75799	0.06329	1.00000
	95%	9	0.28602	0.40385	0.60371	11	0.65706	0.71727	0.65718
C-vine-Norm	99%	4	0.38048	0.63771	0.78387	5	0.16185	0.07766	1.00000
	95%	7	0.08281	0.18141	0.90648	10	0.45291	0.53023	0.48095
C-vine-t	99%	1	0.27807	0.55307	0.00686	4	0.38048	0.08733	1.00000
	95%	9	0.28602	0.40385	0.64173	11	0.65706	0.71727	0.31980

Table 7.1: Phelix Baseload and Peakload Long and Short portfolio p-values of statistical tests, described in Section 7.1, and number of VaR exceedances for each risk model at $\alpha = 0.01$ and $\alpha = 0.05$.

Phelix Baseload Long Portfolio								
Model	**Level**	**Av.VaR**	**RLF**	**MLF**	**PQLF**	**Av.CVaR**	**MAE**	**MSE**
RM	99%	-0.03515	0.0015646	4.0015646	0.00063489	-0.04027	0.0002127	0.00000418
	95%	-0.02485	0.0032817	14.0032817	**0.00181747**	-0.03117	**0.0004567**	0.00000871
C-vine-EVT	99%	-0.04351	0.0009507	3.0009507	0.00060692	-0.05586	0.0001885	0.00000273
	95%	-0.02706	0.0030201	11.0030201	0.00188056	-0.03779	0.0006105	0.00001073
A-C-vine-EVT	99%	-0.04447	**0.0004391**	**3.0004391**	**0.00057180**	-0.05717	**0.0001387**	**0.00000171**
	95%	-0.02780	0.0028869	10.0028869	0.00187775	-0.03869	0.0005913	**0.00000869**
C-vine-Norm	99%	-0.03965	0.0011447	4.0011447	0.00062836	-0.04553	0.0001531	0.00000272
	95%	-0.02828	**0.0027324**	**9.0027324**	0.00189274	-0.03541	0.0005316	0.00000921
C-vine-t	99%	-0.04275	0.0010770	4.0010770	0.00061632	-0.05578	0.0002377	0.00000403
	95%	-0.02650	0.0030967	11.0030967	0.00186608	-0.03723	0.0005666	0.00001030

Phelix Baseload Short Portfolio								
Model	**Level**	**Av.VaR**	**RLF**	**MLF**	**PQLF**	**Av.CVaR**	**MAE**	**MSE**
RM	99%	-0.03515	0.0033335	3.0033335	0.00072408	-0.04027	0.0003077	0.00000939
	95%	-0.02485	0.0060319	13.0060319	0.00188622	-0.03117	0.0005375	0.00001741
C-vine-EVT	99%	-0.04209	0.0015976	2.0015976	0.00062577	-0.05348	**0.0001290**	0.00000268
	95%	-0.02638	0.0039360	8.0039360	0.00174001	-0.03658	0.0003646	0.00001235
A-C-vine-EVT	99%	-0.04274	**0.0015196**	**2.0015196**	0.00062027	-0.05440	0.0001639	**0.00000261**
	95%	-0.02705	0.0037340	7.0037340	0.00175258	-0.03733	0.0003497	**0.00001051**
C-vine-Norm	99%	-0.03850	0.0019666	2.0019666	**0.00061009**	-0.04421	0.0001832	0.00000524
	95%	-0.02728	**0.0037337**	**7.0037337**	0.00176239	-0.03436	**0.0003213**	0.00001163
C-vine-t	99%	-0.04189	0.0017390	2.0017390	0.00062869	-0.05528	0.0001372	0.00000529
	95%	-0.02557	0.0042147	9.0042147	**0.00171525**	-0.03635	0.0005921	0.00003091

Table 7.2: Phelix Baseload Long and Short portfolio VaR and CVaR loss function results, described in Section 7.2; Av.VaR is the average forecasted Value-at-Risk for the out-of-sample period; Av.CVaR is the average forecasted Conditional Value-at-Risk (Expected Shortfall) for the out-of-sample period; Bold values indicate the minimum loss function score across different risk models and same confidence level.

Phelix Peakload Long Portfolio								
Model	**Level**	**Av.VaR**	**RLF**	**MLF**	**PQLF**	**Av.CVaR**	**MAE**	**MSE**
RM	99%	-0.03476	0.0002246	6.0002246	0.00046690	-0.03982	**0.0000514**	**0.00000015**
	95%	-0.02457	0.0018197	16.0018197	0.00170577	-0.03082	**0.0003675**	**0.00000328**
C-vine-EVT	99%	-0.04013	0.0000530	3.0000530	0.00044744	-0.04959	0.0005616	0.00002937
	95%	-0.02607	0.0017057	10.0017057	0.00167023	-0.03511	0.0008263	0.00001875
A-C-vine-EVT	99%	-0.04040	**0.0000103**	**4.0000103**	**0.00041746**	-0.04941	0.0004356	0.00001614
	95%	-0.02653	0.0015422	9.0015422	**0.00164987**	-0.03537	0.0008030	0.00001670
C-vine-Norm	99%	-0.03912	0.0000829	4.0000829	0.00043210	-0.04475	0.0003761	0.00001305
	95%	-0.02790	0.0013726	**7.0013726**	0.00168505	-0.03496	0.0007567	0.00001505
C-vine-t	99%	-0.04444	0.0000004	1.0000004	0.00044276	-0.05827	0.0008244	0.00004366
	95%	-0.02748	**0.0013035**	9.0013035	0.00168213	-0.03881	0.0008689	0.00002239

Phelix Peakload Short Portfolio								
Model	**Level**	**Av.VaR**	**RLF**	**MLF**	**PQLF**	**Av.CVaR**	**MAE**	**MSE**
RM	99%	-0.03476	0.0034548	6.0034548	0.00078469	-0.03982	0.0003358	0.00000990
	95%	-0.02457	0.0064371	15.0064371	0.00205432	-0.03082	0.0006236	0.00001797
C-vine-EVT	99%	-0.04402	0.0022995	4.0022995	0.00076175	-0.05596	0.0002175	0.00000536
	95%	-0.02721	0.0051004	**10.0051004**	**0.00199222**	-0.03807	0.0005311	0.00001767
A-C-vine-EVT	99%	-0.04482	**0.0010725**	**3.0010725**	**0.00064555**	-0.05705	0.0001116	0.00000300
	95%	-0.02794	**0.0048558**	11.0048558	0.00201820	-0.03898	0.0005384	0.00001737
C-vine-Norm	99%	-0.03852	0.0028474	5.0028474	0.00074071	-0.04417	0.0002642	0.00000686
	95%	-0.02714	0.0052233	10.0052233	0.00201169	-0.03427	**0.0004958**	0.00001775
C-vine-t	99%	-0.04364	0.0017716	4.0017716	0.00070101	-0.05916	**0.0000865**	**0.00000123**
	95%	-0.02678	0.0055913	11.0055913	0.00200585	-0.03841	0.0005459	**0.00001618**

Table 7.3: Phelix Peakload Long and Short portfolio VaR and CVaR loss function results, described in Section 7.2; Av.VaR is the average forecasted Value-at-Risk for the out-of-sample period; Av.CVaR is the average forecasted Conditional Value-at-Risk (Expected Shortfall) for the out-of-sample period; Bold values indicate the minimum loss function score across different risk models and same confidence level.

8 Conclusion

In this study we introduced the vine copula modelling as an alternative and more flexible way to describe the joint distribution of power portfolios. Our approach combines pseudo-maximum-likelihood fitting of time series models and extreme value theory to model the tails of resulting innovations conditional distribution. We modelled the dependency structure of portfolios using a canonical vine copula modelling approach. We also employed an alternative vine specification based on the theoretical and empirical results of Joe et al. (2010) and Nikoloulopoulos et al. (2012). In particular, we replaced the AIC selected t copula families in level 1 of the vine with copula families having asymmetric tail dependence whereas the Independence copula was employed for pair-copulas that could not reject the null hypothesis of independence. Baseload and Peakload portfolio VaR and CVaR forecasts were computed by five alternative models. The evaluation of their forecasting performance was conduced using standard statistical and subjective loss function techniques. In general, the statistical tests for all models show good unconditional and conditional coverage. However, the A-C-vine-EVT model shows superior performance according to VaR and CVaR-based loss function results at extreme quantiles ($\alpha = 0.01$) compared to the rest of the models employed.

These results provide new insights for risk management applications within the vine copula modelling framework and also support the findings of Joe et al. (2010) and Nikoloulopoulos et al. (2012). It seems that the A-C-vine-EVT model provides better fit in the tails of joint distribution and hence better risk estimates compared to the entirely AIC selected C-vine-EVT model. Even though both data sets do not show significant tail asymmetries, it seems that the asymmetric copulas instead of the likelihood-selected t copulas in level 1 of the vine structure as well as the employment of Independence copulas for the pairs that display independence improve the fit in the tails and as a consequence the risk measure estimates. We strongly believe that the employment of A-C-vine-EVT model for risk management applications in portfolios that exhibit more significant asymmetric tail dependences will further enhance the importance and predictive ability of the model in extreme quantiles.

The fit and performance of the model can be also improved in various ways. For example, the selection of asymmetric copula families in level 1 of the vine should not be entirely based on empirical and theoretical scatter and λ-function plots comparisons but also on matching the empirical tail dependence of the data with the theoretical tail dependence of corresponding asymmetric copula families. Moreover, the fit of semi-parametric marginals and canonical vine copula model can be improved by re-estimating the models in fixed periods of time, i.e., every month, when considering forecasting applications. All risk measure forecasts presented in this study were obtained by a static version of the model. To sum up, we believe that the proposed methodology have many potentials in modelling of large-scale portfolios and can substantially improve the existing portfolio risk management methods.

References

Aas, K., and D. Berg, 2009, Models for construction of multivariate dependence: A comparison study, *The European Journal of Finance* 15, 639–659.

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

Ang, A., and J. Chen, 2002, Asymmetric correlations of equity portfolios, *Journal of Financial Economics* 63, 443–494.

Angelidis, T., and S. Degiannakis, 2007, Backtesting VaR models: An expected shortfall approach, Working paper, University of Crete, Department of Economics.

Balkema, A. A., and L. de Haan, 1974, Residual life time at great age, *The Annals of Probability* 2, 792–804.

Bedford, T., and R. M. Cooke, 2001, Probability density decomposition for conditionally dependent random variables modeled by vines, *Annals of Mathematics and Artificial Intelligence* 32, 245–268.

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

Berkowitz, J., P. Christoffersen, and D. Pelletier, 2009, Evaluating value-at-risk models with desk-level data, *Management Science* forthcoming.

Börger, R. H., À. Cartea, R. Kiesel, and G. Schindlmayr, 2007, A multivariate commodity analysis and applications to risk management, Birkbeck working papers in economics and finance, Birkbeck, Department of Economics, Mathematics & Statistics.

Brechmann, E. C., and C. Czado, 2011, Risk management with high-dimensional vine copulas: An analysis of the euro stoxx 50, Working paper, Technische Universität München, http://mediatum.ub.tum.de/doc/1079276/1079276.pdf.

Byström, H. N.E., 2005, Extreme value theory and extremely large electricity price changes, *International Review of Economics & Finance* 14, 41–55.

Chan, K.F., and P. Gray, 2006, Using extreme value theory to measure value-at-risk for daily electricity spot prices, *International Journal of Forecasting* 22, 283–300.

Christoffersen, P., 1998, Evaluating interval forecasts, *International Economic Review* 39, 841–62.

Christoffersen, P., and D. Pelletier, 2004, Backtesting value-at-risk: A duration-based approach, *Journal of Financial Econometrics* 2, 84–108.

Clarke, Kevin A., 2007, A simple distribution-free test for nonnested model selection, *Political Analysis* 15, 347–363.

Czado, C., U. Schepsmeier, and A. Min, 2011, Maximum likelihood estimation of mixed C-vines with application to exchange rates, *To appear in Statistical Modelling* .

Danielsson, J., and C. G. de Vries, 1997, Tail index and quantile estimation with very high frequency data, *Journal of Empirical Finance* 241–257.

Embrechts, P., C. Kluppelberg, and T. Mikosch, 1999a, *Modelling Extremal Events: for Insurance and Finance*, second edition (Springer, Berlin).

Embrechts, Paul, Sidney I. Resnick, and Gennady Samorodnitsky, 1999b, Extreme value theory as a risk management tool, *North American Actuarial Journal* 30–41.

Gençay, R., and F. Selçuk, 2004, Extreme value theory and value-at-risk: Relative performance in emerging markets, *International Journal of Forecasting* 20, 287–303.

Genest, C., K. Ghoudi, and L.-P. Rivest, 1995, A semiparametric estimation procedure of dependence parameters in multivariate families of distributions, *Biometrika* 82, 543–552.

Heinen, A., and A. Valdesogo, 2009, Asymmetric CAPM dependence for large dimensions: The canonical vine autoregressive copula model, Core discussion paper 2009069, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Hobæk Haff, I., 2010, Parameter estimation for pair-copula constructions, *preprint* .

Hong, Y., J. Tu, and G. Zhou, 2007, Asymmetries in stock returns: Statistical tests and economic evaluation, *Review of Financial Studies* 20, 1547–1581.

Hosking, J. R. M., and J. R. Wallis, 1987, Parameter and quantile estimation for the generalized pareto distribution, *Technometrics* 29, 339–349.

Hurlin, C., and C. Perignon, 2012, Margin backtesting, *SSRN Electron. Journal* .

Joe, H., 1996, Families of m-variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters, in L. Rüschendorf, B. Schweizer, and M. D. Taylor, eds., *Distributions with fixed marginals and related topics*, 120–141 (Hayward: Institute of Mathematical Statistics, Hayward, CA).

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

Joe, H., and T. Hu, 1996, Multivariate distributions from mixtures of max-infinitely divisible distributions, *Journal of Multivariate Analysis* 57, 240–265.

Joe, H., H. Li, and A. K. Nikoloulopoulos, 2010, Tail dependence functions and vine copulas, *Journal of Multivariate Analysis* 101, 252–270.

Kim, G., M. J. Silvapulle, and P. Silvapulle, 2007, Comparison of semiparametric and parametric methods for estimating copulas, *Computational Statistics & Data Analysis* 51, 2836–2850.

Koenker, R., and G. Bassett, 1978, Regression quantiles, *Econometrica* 46, 33–50.

Kupiec, P. H., 1995, Techniques for verifying the accuracy of risk measurement models, *The Journal of Derivatives* 3, 73–84.

Kurowicka, D., and R. M. Cooke, 2006, *Uncertainty Analysis with High Dimensional Dependence Modelling*, first edition (Wiley).

Longin, F., and B. Solnik, 1995, Is the correlation in international equity returns constant: 1960-1990 ?, *Journal of International Money and Finance* 14, 3–26.

Longin, F., and B. Solnik, 2001, Extreme correlation of international equity markets, *The Journal of Finance* 56, 649–676.

Lopez, J. A., 1999, Methods for evaluating value-at-risk estimates, Federal reserve bank of san francisco, Economic Review, 23–17.

Lopez, J.A., 1998, Regulatory evaluation of value-at-risk models, Federal reserve bank of new york.

McNeil, A.J., and R. Frey, 2000, Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach, *Journal of Empirical Finance* 7, 271–300.

Mendes, V. d. M. B., S. M. Mendes, and P. C. R. Leal, 2010, Pair-copulas modeling in finance, *Financial Markets and Portfolio Management* 24, 193–213.

Nikoloulopoulos, A. K., H. Joe, and H. Li, 2012, Vine copulas with asymmetric tail dependence and applications to financial return data, *Computational Statistics & Data Analysis* 56, 3659–3673.

Pickands, J., 1975, Statistical inference using extreme order statistics, *Annals of Statistics* 3, 119–131.

Reiss, R.-D., and M. Thomas, 1997, *Statistical Analysis of Extreme Values with Applications to Insurance, Finance, Hydrology and Other Fields* (Birkhäuser Verlag).

Sarma, M., S. Thomas, and A. Shah, 2003, Selection of value-at-risk models, *Journal of Forecasting* 22, 337–358.

Sklar, A., 1959, Fonctions de répartition à n dimensions et leurs marges, *Publications del'Institut de Statistique de L'Université de Paris,* 8, 229–231.

Vuong, Q. H., 1989, Likelihood ratio tests for model selection and non-nested hypotheses, *Econometrica* 57, 307–33.

Appendices

A Time series models for the conditional mean

The conditional mean time series models for the Phelix Baseload portfolio return series are:

SpotB-ARMA(1,1): $r_t = \varphi_1 r_{t-1} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1BM - ARMA(15,1): $r_t = \varphi_8 r_{t-8} + \varphi_{15} r_{t-15} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2BM - ARMA(8,1): $r_t = \varphi_8 r_{t-8} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F3BM - ARMA(8,1): $r_t = \varphi_8 r_{t-8} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1BQ - ARMA(8,1): $r_t = \varphi_8 r_{t-8} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2BQ - ARMA(4,1): $r_t = \varphi_4 r_{t-4} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1BY - MA(1): $r_t = \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2BY - ARMA(4,1): $r_t = \varphi_4 r_{t-4} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

The conditional mean time series models for the Phelix Peakload portfolio return series are:

SpotP-ARMA(1,1): $r_t = \varphi_1 r_{t-1} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1PM - ARMA(15,1): $r_t = \varphi_8 r_{t-8} + \varphi_9 r_{t-9} + \varphi_{15} r_{t-15} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2PM - ARMA(18,1): $r_t = \varphi_8 r_{t-8} + \varphi_{18} r_{t-18} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F3PM - ARMA(15,1): $r_t = \varphi_8 r_{t-8} + \varphi_{15} r_{t-15} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1PQ - ARMA(18,1): $r_t = \varphi_8 r_{t-8} + \varphi_{15} r_{t-15} + \varphi_{18} r_{t-18} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2PQ - ARMA(8,1): $r_t = \varphi_3 r_{t-3} + \varphi_4 r_{t-4} + \varphi_8 r_{t-8} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F1PY - ARMA(9,1): $r_t = \varphi_4 r_{t-4} + \varphi_8 r_{t-8} + \varphi_9 r_{t-9} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

F2PY - ARMA(4,1): $r_t = \varphi_4 r_{t-4} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$

B Copula family abbreviations

No.	Short name	Long name
0	I	Independence
1	N	Gaussian
2	t	t
3	C	Clayton
4	G	Gumbel
5	F	Frank
6	J	Joe
7	BB1	Clayton-Gumbel
8	BB6	Joe-Gumbel
9	BB7	Joe-Clayton
10	BB8	Frank-Joe
13	SC	Survival Clayton
14	SG	Survival Gumbel
16	SJ	Survival Joe
17	SBB1	Survival Clayton-Gumbel
18	SBB6	Survival Joe-Gumbel
19	SBB7	Survival Joe-Clayton
20	SBB8	Survival Joe-Frank
23	C90	Rotated Clayton 90 degrees
24	G90	Rotated Gumbel 90 degrees
26	J90	Rotated Joe 90 degrees
27	BB1_90	Rotated Clayton-Gumbel 90 degrees
28	BB6_90	Rotated Joe-Gumbel 90 degrees
29	BB7_90	Rotated Joe-Clayton 90 degrees
30	BB8_90	Rotated Frank-Joe 90 degrees
33	C270	Rotated Clayton 270 degrees
34	G270	Rotated Gumbel 270 degrees
36	J270	Rotated Joe 270 degrees
37	BB1_270	Rotated Clayton-Gumbel 270 degrees
38	BB6_270	Rotated Joe-Gumbel 270 degrees
39	BB7_270	Rotated Joe-Clayton 270 degrees
40	BB8_270	Rotated Frank-Joe 270 degrees

Table B.1: No., Short and Long names of copula families in **CDVine** package, R statistical software.

C Pair-copulas for Phelix Baseload portfolio

Level	Numerical representation	Unconditional & Conditional pair-copulas
1	1,2	F1BQ,F1BY
	1,3	F1BQ,F1BM
	1,4	F1BQ,F2BY
	1,5	F1BQ,F2BQ
	1,6	F1BQ,SpotB
	1,7	F1BQ,F3BM
	1,8	F1BQ,F2BM
2	2,3|1	F1BY,F1BM|F1BQ
	2,4|1	F1BY,F2BY|F1BQ
	2,5|1	F1BY,F2BQ|F1BQ
	2,6|1	F1BY,SpotB|F1BQ
	2,7|1	F1BY,F3BM|F1BQ
	2,8|1	F1BY,F2BM|F1BQ
3	3,4|12	F1BM,F2BY|F1BQ F1BY
	3,5|1,2	F1BM,F2BQ|F1BQ,F1BY
	3,6|1,2	F1BM,SpotB|F1BQ,F1BY
	3,7|1,2	F1BM,F3BM|F1BQ,F1BY
	3,8|1,2	F1BM,F2BM|F1BQ,F1BY
4	4,5|123	F2BY,F2BQ|F1BQ,F1BY,F1BM
	4,6|1,2,3	F2BY,SpotB|F1BQ,F1BY,F1BM
	4,7|1,2,3	F2BY,F3BM|F1BQ,F1BY,F1BM
	4,8|1,2,3	F2BY,F2BM|F1BQ,F1BY,F1BM
5	5,6|1234	F2BQ,SpotB|F1BQ,F1BY,F1BM,F2BY
	5,7|1,2,3,4	F2BQ,F3BM|F1BQ,F1BY,F1BM,F2BY
	5,8|1,2,3,4	F2BQ,F2BM|F1BQ,F1BY,F1BM,F2BY
6	6,7|12345	SpotB,F3BM|F1BQ,F1BY,F1BM,F2BY,F2BQ
	6,8|1,2,3,4,5	SpotB,F2BM|F1BQ,F1BY,F1BM,F2BY,F2BQ
7	7,8|1,2,3,4,5,6	F3BM,F2BM|F1BQ,F1BY,F1BM,F2BY,F2BQ,SpotB

Table C.1: Determination of unconditional and conditional pair-copulas for all levels of canonical vine structure selected by the empirical rule of Czado et al. (2011).