Pair-copula constructions for non-Gaussian DAG models

Alexander Bauer, Claudia Czado, and Thomas Klein
Technische Universität München, Zentrum Mathematik
Boltzmannstr. 3
85748 Garching, Germany
E-mail: a.bauer@ma.tum.de, cczado@ma.tum.de, tpklein@ma.tum.de

1 INTRODUCTION

Graphical models are multivariate statistical models in which the corresponding joint distribution of a family of random variables is restricted by a list of conditional independence assumptions. This list is conveniently summarised in a graph whose vertices represent the variables and whose edges represent interrelations of these variables. Lauritzen (1996) and Cowell et al. (2003) are standard references on the theory of graphical models but are mainly limited to the assumption of joint normality as far as continuous variables are concerned. We propose a new type of statistical model that permits non-Gaussian distributions as well as the inclusion of conditional independence assumptions induced by a directed acyclic graph (DAG). This combination of features is achieved by using so-called pair-copula constructions (PCCs) in which a multivariate distribution is decomposed into bivariate, potentially conditional distributions based on iterated applications of Sklar's theorem on copulas. The basic idea of applying PCCs to distributions with certain conditional independence properties goes back to Hanea et al. (2006) and Kurowicka and Cooke (2006). While these authors restricted their considerations to Gaussian margins and copulas we utilise PCCs to specifically construct non-Gaussian distributions in order to capture features such as tail behaviour and non-linear, asymmetric dependence as observed, for instance, in financial data.

2 PAIR-COPULA CONSTRUCTIONS AND REGULAR VINES

A copula is a multivariate cumulative distribution function (cdf) $C\colon [0,1]^d \to [0,1]$, $d \in \mathbb{N}$, such that all univariate marginals are uniform on the interval $[0,1]$. By Sklar's theorem (Sklar, 1959) every cdf $F\colon \mathbb{R}^d \to [0,1]$ may be written as $F(x_1,\dots,x_d) = C\big(F_1(x_1),\dots,F_d(x_d)\big)$ for some suitable copula C, where $F_1,\dots,F_d$ are the univariate marginals of F. If F is absolutely continuous and $F_1,\dots,F_d$ are strictly increasing we can pass to its probability density function (pdf) and write $f(x_1,\dots,x_d) = c\big(F_1(x_1),\dots,F_d(x_d)\big)\prod_{i=1}^d f_i(x_i)$, where the copula pdf c is uniquely determined. Various examples of copulas together with the underlying theory are presented in Joe (1997) and Nelsen (2006).

While there is a plethora of literature on bivariate copula families (also called *pair-copula families*), the range of higher-variate copula families is rather limited. Based on work of Joe (1996), Bedford and Cooke (2001, 2002) therefore proposed a flexible way of constructing multivariate copulas that uses (conditional) pair copulas as building blocks only. The core of their approach is a graphical representation called a *regular vine* that consists of a sequence of trees, each edge of which is associated with a certain pair copula. Figure 1 shows an example of a *D-vine*, a popular type of regular vine. According to Aas et al. (2009) the pdf of a D-vine pair-copula construction (PCC) on $\mathbb{R}^d$ is given by

$$f(\boldsymbol{x}) = \prod_{i=1}^{d-1}\prod_{j=1}^{d-i} c_{j,j+i|(j,j+i)}\big(F_{j|(j,j+i)}(x_j \mid \boldsymbol{x}_{(j,j+i)}), F_{j+i|(j,j+i)}(x_{j+i} \mid \boldsymbol{x}_{(j,j+i)}) \,\big|\, \boldsymbol{x}_{(j,j+i)}\big) \prod_{i=1}^{d} f_i(x_i).$$

Here we have written $\boldsymbol{x}_{(j,j+i)} := (x_{j+1},\dots,x_{j+i-1})$. Also, $F_{j|(j,j+i)}(\,\cdot \mid \boldsymbol{x}_{(j,j+i)})$ denotes the conditional cdf of X_j given $\boldsymbol{X}_{(j,j+i)} = \boldsymbol{x}_{(j,j+i)}$, where $\boldsymbol{X} = (X_1,\dots,X_d)$ is distributed as f. The values of $F_{j|(j,j+i)}$

and $F_{j+i|(j,j+i)}$ in tree $i \leq d-1$ can be computed iteratively using a recursive formula given in Joe (1996). The only copulas needed in this computation are the ones specified in the preceding trees.

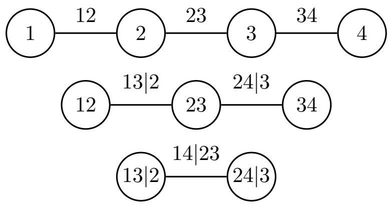

Figure 1: A four-variate D-vine specifying the pair copulas C_{12}, C_{23}, C_{34}, $C_{13|2}$, $C_{24|3}$, and $C_{14|23}$.

We can construct a multitude of multivariate *vine copulas* by selecting a number of pair copulas and by setting all univariate marginals to be uniform distributions on $[0,1]$. The construction of a d-variate vine copula requires the specification of $\binom{d}{2}$ pair copulas, a number increasing quadratically in d. The actual number of decisions to make in practical applications might, however, be lower if the analysed data exhibit conditional independences. In that case the corresponding pair copulas are nothing but product copulas with pdf equal to one. Instead of starting one's analysis with a set of regular vines it may therefore be more fruitful to look for conditional independences first. One class of models tailored for this task are Bayesian networks. By transferring the pair-copula concept to graphical models, Hanea et al. (2006) provided an opportunity to exploit the advantages of both worlds. However, their analysis is restricted to Gaussian pair copulas. We will review PCCs for directed graphs in the next section. A more detailed exposition of regular vines can be found in Kurowicka and Cooke (2006, Section 4.4) and Kurowicka and Joe (2011).

3 PAIR-COPULA CONSTRUCTIONS AND DIRECTED ACYCLIC GRAPHS

Let $V \neq \emptyset$ be a finite set and let E be a subset of $\{(v,w) \in V^2 \,|\, v \neq w\}$ such that $(w,v) \notin E$ whenever $(v,w) \in E$. Then $D = (V,E)$ is a directed graph with vertex set V and edge set E. We denote a pair $(v,w) \in E$ by an arrow $v \to w$. A path in D is a sequence $v_1, \ldots, v_n \in V$, $n \geq 2$, such that D contains all arrows $v_i \to v_{i+1}$, $i \leq n-1$. In the special case $v_1 = v_n$ the path $v_1, \ldots, v_n$ is a cycle. If there are no cycles in D, then D is called a *directed acyclic graph* (DAG). For all $v \in V$ we further set $\mathrm{pa}(v) := \{w \in V \,|\, D \text{ contains } w \to v\}$ (set of *parents* of v), $\mathrm{de}(v) := \{w \in V \,|\, D \text{ contains a path from } v \text{ to } w\}$ (set of *descendants* of v), and $\mathrm{nd}(v) := V \setminus (\mathrm{de}(v) \cup \{v\})$ (set of *non-descendants* of v).

Now let $D = (V,E)$ be a DAG, P a probability measure on $\mathbb{R}^{|V|}$, and $\boldsymbol{X}$ an $\mathbb{R}^{|V|}$-valued random variable distributed as P. We write $\boldsymbol{X}_I \perp\!\!\!\perp \boldsymbol{X}_J \mid \boldsymbol{X}_K$ whenever $\boldsymbol{X}_I := (X_i)_{i \in I}$ and $\boldsymbol{X}_J$ are conditionally independent given $\boldsymbol{X}_K$ for pairwise disjoint sets $I, J, K \subseteq V$. P is said to be *D-Markovian* or, equivalently, to possess the *local D-Markov property* if $X_v \perp\!\!\!\perp \boldsymbol{X}_{\mathrm{nd}(v) \setminus \mathrm{pa}(v)} \mid \boldsymbol{X}_{\mathrm{pa}(v)}$ for all $v \in V$. In case P has a Lebesgue-density f the D-Markov property is equivalent to f admitting a *D-recursive factorisation* of the form $f(\boldsymbol{x}) = \prod_{v \in V} f_{v\,|\,\mathrm{pa}(v)}\left(x_v \,\middle|\, \boldsymbol{x}_{\mathrm{pa}(v)}\right)$ for all $\boldsymbol{x} \in \mathbb{R}^{|V|}$, where $f_{v\,|\,\mathrm{pa}(v)}$ is the pdf of $P_{v\,|\,\mathrm{pa}(v)}$. A comprehensive introduction to graphical models is found in Lauritzen (1996).

As an example consider the DAG from Figure 2 and an absolutely continuous probability measure P on $\mathbb{R}^4$ possessing the respective local D-Markov property. Straightforward application of above definition yields the restrictions $X_2 \perp\!\!\!\perp X_3 \mid X_1$ and $X_1 \perp\!\!\!\perp X_4 \mid \boldsymbol{X}_{23}$. There are no other implicit conditional independence properties in this example. The pdf of P admits the D-recursive factorisation $f(\boldsymbol{x}) = f_1(x_1)\; f_{2|1}(x_2 \,|\, x_1)\; f_{3|1}(x_3 \,|\, x_1)\; f_{4|23}(x_4 \,|\, \boldsymbol{x}_{23})$ for all $\boldsymbol{x} \in \mathbb{R}^4$.

Our aim is to derive a pair-copula decomposition for the pdf f of a D-Markovian probability measure P on $\mathbb{R}^{|V|}$. For every $v \in V$ we therefore order the elements of $\mathrm{pa}(v)$ increasingly (with respect to

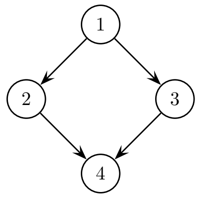

Figure 2: A DAG $D = (V, E)$ inducing the conditional independence properties $X_2 \perp\!\!\!\perp X_3 \mid X_1$ and $X_1 \perp\!\!\!\perp X_4 \mid \boldsymbol{X}_{23}$.

some strict total order $<$ on V) and set $\mathrm{pa}(v; w) := \{u \in \mathrm{pa}(v) \mid u < w\}$, $w \in \mathrm{pa}(v)$. A proof of the following theorem is given in Bauer et al. (2011).

Theorem 1. *Let $D = (V, E)$ be a DAG and let P be an absolutely continuous D-Markovian probability measure whose univariate marginal cdfs are strictly increasing. Then P is uniquely determined by its univariate margins P_v, $v \in V$, and its (conditional) pair copulas $C_{vw \mid \mathrm{pa}(v;w)}$, $v \in V$, $w \in \mathrm{pa}(v)$.*

By Theorem 1 we can decompose f into

$$(1) \quad f(\boldsymbol{x}) = \prod_{v \in V} \prod_{w \in \mathrm{pa}(v)} f_v(x_v)\; c_{vw \mid \mathrm{pa}(v;w)}\big(F_{v \mid \mathrm{pa}(v;w)}(x_v \mid \boldsymbol{x}_{\mathrm{pa}(v;w)}), F_{w \mid \mathrm{pa}(v;w)}(x_w \mid \boldsymbol{x}_{\mathrm{pa}(v;w)}) \,\big|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}\big).$$

The remaining question is whether DAGs actually extend the set of pair-copula decompositions beyond regular vines. This question is closely related to the computation of the conditional cdfs $F_{v \mid \mathrm{pa}(v;w)}$ and $F_{w \mid \mathrm{pa}(v;w)}$ in equation (1). In contrast to vines, DAGs allow for the specification of conditional cdfs that cannot be computed by simply plugging in results from preceding trees. Hence, the values of these functions have to be computed via integration of other margins. In view of their application for statistical inference, however, DAG-PCC models may be more parsimonious than vine-PCC models.

The lowest dimensional DAG models exhibiting this characteristic are the ones having the same Markov properties as the DAG in Figure 2. In this case, equation (1) yields

$$\begin{aligned} f(\boldsymbol{x}) &= f_1(x_1) \cdots f_4(x_4)\, c_{12}\big(F_1(x_1), F_2(x_2)\big)\, c_{13}\big(F_1(x_1), F_3(x_3)\big)\, c_{24}\big(F_2(x_2), F_4(x_4)\big) \\ &\quad \cdot c_{34|2}\big(F_{3|2}(x_3 \mid x_2), F_{4|2}(x_4 \mid x_2) \,\big|\, x_2\big), \quad \boldsymbol{x} \in \mathbb{R}^4, \end{aligned}$$

where $F_{4|2}(x_4 \mid x_2) = \frac{\partial C_{24}\left(F_2(x_2), F_4(x_4)\right)}{\partial F_2(x_2)}$, see Joe (1996). Many popular copulas exhibit a closed form expression for this partial derivative, see Aas et al. (2009). Since the copula C_{23} is not available in the decomposition of f we exploit the conditional independence property $X_2 \perp\!\!\!\perp X_3 \mid X_1$ to get

$$F_{3|2}(x_3 \mid x_2) = \int_0^1 c_{12}\big(u_1, F_2(x_2)\big)\; \frac{\partial C_{13}\big(u_1, F_3(x_3)\big)}{\partial u_1}\, \mathrm{d}u_1.$$

A detailed derivation of this formula is given in Bauer et al. (2011). There is no general closed-form solution for the integral.

4 LIKELIHOOD INFERENCE: A SIMULATION STUDY

With reference to Sklar's theorem we restrict our considerations to *DAG copulas* in this section, that is, DAG distributions with uniform [0,1] univariate margins. Let $D = (V, E)$ be a DAG and let $\{P_{\boldsymbol{\theta}} \mid \boldsymbol{\theta} \in \Theta\}$ be a family of D-Markovian probability measures on $[0,1]^{|V|}$ satisfying the usual assumptions. Given a realisation $\boldsymbol{u} = (\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n) \in \big([0,1]^{|V|}\big)^n$, $n \in \mathbb{N}$, of a sample of i.i.d. random variables $\boldsymbol{U}^1, \ldots, \boldsymbol{U}^n$, equation (1) yields the log-likelihood function

$$(2) \quad L(\boldsymbol{\theta}; \boldsymbol{u}) = \sum_{k=1}^{n} \sum_{v \in V} \sum_{w \in \mathrm{pa}(v)} \log c_{vw \mid \mathrm{pa}(v;w)}\Big(F_{v \mid \mathrm{pa}(v;w)}\big(u_v^k \mid \boldsymbol{u}^k_{\mathrm{pa}(v;w)}; \boldsymbol{\theta}\big), F_{w \mid \mathrm{pa}(v;w)}\big(u_w^k \mid \boldsymbol{u}^k_{\mathrm{pa}(v;w)}; \boldsymbol{\theta}\big); \boldsymbol{\theta}\Big)$$

for all $\boldsymbol{\theta} = (\boldsymbol{\theta}_{vw\,|\,\mathrm{pa}(v;w)})_{v\in V, w\in\mathrm{pa}(v)} \in \Theta$. On the right hand side we have omitted parameter subscripts as well as the values of the conditioning variables in the pair-copula pdfs $c_{vw\,|\,\mathrm{pa}(v;w)}$. The latter omission is based on the assumption of constant copula parameters $\boldsymbol{\theta}_{vw\,|\,\mathrm{pa}(v;w)}$ for all observations $\boldsymbol{u}^k$, $k \le n$. By that assumption the copula pdfs involved are treated as if they were unconditional. The cdfs $F_{v\,|\,\mathrm{pa}(v;w)}$ and $F_{w\,|\,\mathrm{pa}(v;w)}$, however, remain conditional. This assumption reduces model complexity while still encompassing a rich class of DAG copulas and has become common practice in likelihood inference for PCCs, see Aas et al. (2009) and Hobæk Haff et al. (2010).

We examined the practical performance of the DAG-copula model by a simulation study. To this end we drew samples from two families of DAG copulas with conditional independence properties given by the DAG D in Figure 2. More precisely, these families of DAG copulas emerge from two different choices of pair-copula families involved. In either setting we considered six different parameter configurations, resulting in the twelve simulation scenarios described in Table 1. The copula families used (Clayton, Gumbel, Gaussian, and Student's t) exhibit notable differences in tail behaviour. Within each copula family a range of values of Kendall's τ as a dependence measure may be obtained by suitable parameter choices. See Aas (2004) for the relations between copula parameters, tail dependence, and Kendall's τ.

Scenario			1	2	3	4	5	6
C_{12}	Clayton	δ	0.67 (0.25)	6.00 (0.75)	6.00 (0.75)	0.67 (0.25)	0.67 (0.25)	0.67 (0.25)
C_{13}	Gumbel	δ	1.33 (0.25)	4.00 (0.75)	1.33 (0.25)	4.00 (0.75)	1.33 (0.25)	1.33 (0.25)
C_{24}	Student	ρ, ν	0.38, 5 (0.25)	0.92, 5 (0.75)	0.38, 5 (0.25)	0.38, 5 (0.25)	0.92, 5 (0.75)	0.38, 5 (0.25)
$C_{34\|2}$	Gauss	ρ	0.38 (0.25)	0.92 (0.75)	0.38 (0.25)	0.38 (0.25)	0.38 (0.25)	0.92 (0.75)
Scenario			7	8	9	10	11	12
C_{12}	Clayton	δ	0.67 (0.25)	6.00 (0.75)	6.00 (0.75)	0.67 (0.25)	0.67 (0.25)	0.67 (0.25)
C_{13}	Clayton	δ	0.67 (0.25)	6.00 (0.75)	0.67 (0.25)	6.00 (0.75)	0.67 (0.25)	0.67 (0.25)
C_{24}	Clayton	δ	0.67 (0.25)	6.00 (0.75)	0.67 (0.25)	0.67 (0.25)	6.00 (0.75)	0.67 (0.25)
$C_{34\|2}$	Clayton	δ	0.67 (0.25)	6.00 (0.75)	0.67 (0.25)	0.67 (0.25)	0.67 (0.25)	6.00 (0.75)

Table 1: Twelve simulation scenarios given by various choices of copula families and parameters for C_{12}, C_{13}, C_{24}, $C_{34|2}$, with notation as in Aas (2004). Parameters were chosen with regard to Kendall's τ as a measure of dependence. The values of τ for each scenario are given in parentheses.

In each simulation run we generated $n = 1000$ i.i.d. observations, and in each scenario we performed $N = 100$ such runs. For each of the 1200 runs we then numerically calculated maximum-likelihood (ML) estimates of the parameters of the underlying DAG-copula model (as specified in Table 1) and compared them to the respective true parameters. Details on the sampling method and the routine for finding ML estimates as well as an overview of the results in terms of empirical bias and mean squared error (MSE) for this *non-Gaussian DAG-copula model* are given in Bauer et al. (2011). Figure 3 shows kernel density estimates of the maximised log-likelihoods for each scenario.

DAG-PCC models may be viewed as generalisations of Gaussian graphical models as presented in Lauritzen (1996, Chapter 4). ML estimation in these models reduces to an estimation of the correlation matrix, observing the restrictions captured in the DAG specifying the model, see Cox and Wermuth (1996, Chapter 3). By choosing all univariate margins as well as all pair-copula families in a DAG-PCC model based on a DAG D to be Gaussian we obtain the Gaussian graphical model based on D. The Gaussian DAG copula corresponding to the DAG of our simulation study, for instance, is a four-variate Gaussian copula, the shape of the correlation matrix of which can be found in Bauer et al. (2011). Kernel density estimates of the maximised log-likelihoods for this *Gaussian DAG-copula model* are given in Figure 3.

As stated in Section 3 the DAG-copula model corresponding to the DAG D of our simulation study cannot be represented by a regular vine. The D-vine with first tree (3)—(1)—(2)—(4) features the same unconditional pair copulas C_{12}, C_{13}, and C_{24} as the DAG copula. The second and third tree comprise the conditional pair copulas $C_{23|1}$, $C_{14|2}$, and $C_{34|12}$. We performed ML estimation in the corresponding D-vine-copula model for the 1200 runs of our simulation study and compared the maximised log-likelihoods to those obtained in the Gaussian and the non-Gaussian DAG-copula model. Details on the selection procedure for the copula families for $C_{14|2}$ and $C_{34|12}$ can be found in Bauer et al. (2011). There is no actual choice to be made for $C_{23|1}$ since we have $U_2 \perp\!\!\!\perp U_3 \mid U_1$ by construction. Hence $C_{23|1}$ is the product copula.

Figure 3 shows that maximised log-likelihoods in the non-Gaussian DAG-copula models are roughly 50% higher than those obtained in their Gaussian counterparts. Even though the Gaussian DAG-copula model in scenarios 1 through 6 has one parameter less than its non-Gaussian competitor this difference in log-likelihood clearly shows the latter model's superiority. Figure 3 also shows the performance of the approximating D-vine-copula model. As this model and the non-Gaussian DAG-copula model share the same unconditional pair copulas, it is not surprising that the maximised log-likelihoods in these models differ mainly in those scenarios with high values of $\boldsymbol{\theta}_{34|2}$, namely, scenarios 2, 6, 8, and 12. Note, however, that the vine-copula model involves a higher number of parameters than the non-Gaussian DAG-copula model and that its maximised log-likelihood is always slightly inferior to that of the latter model.

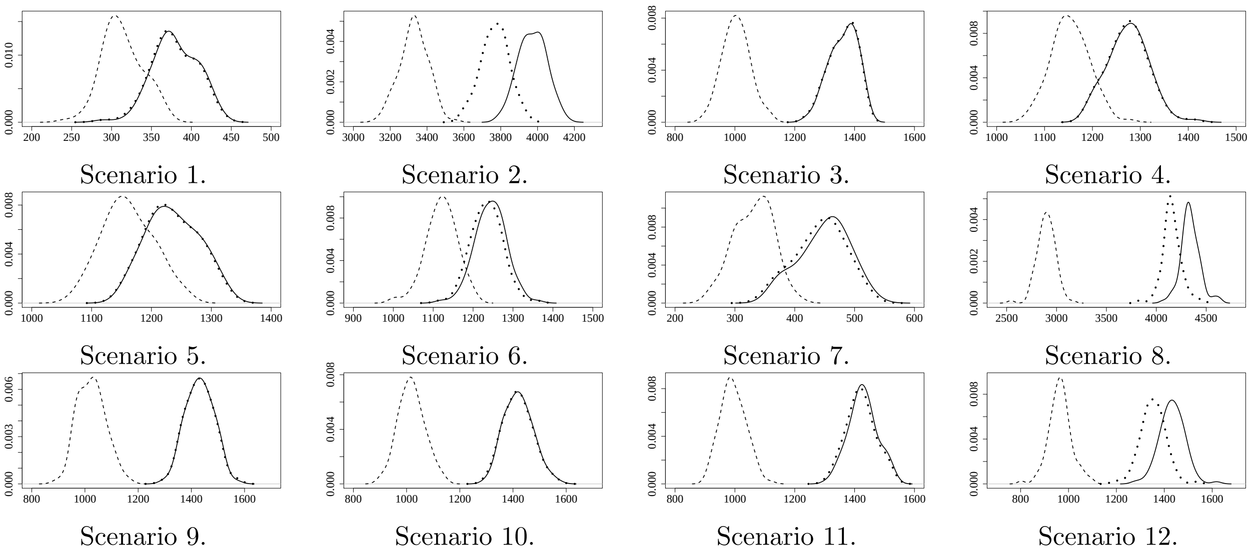

Figure 3: Kernel density estimates of the maximised log-likelihoods in 100 runs for each scenario (see Table 1), based on the Gaussian DAG (dashed line), the non-Gaussian DAG (solid) and the D-vine-copula model (dotted).

5 CONCLUSION

We have defined non-Gaussian DAG-copula models based on density factorisations that are driven by two ideas: First, the D-recursive factorisation known from the literature on DAG models and second, pair-copula constructions (PCCs) for the conditional densities resulting from that factorisation. Such PCCs have been widely used in regular-vine models. Our approach in Section 3, however, is tailored to constructing models with pre-specified Markov properties (as given by a DAG), a feature that cannot be precisely implemented in vine models. The theoretical justification for that approach is given in Theorem 1, which also states that univariate margins and pair copulas in these models can be freely chosen and are not limited to the Gaussian case. As is demonstrated in Section 4 our non-Gaussian DAG-copula models are suitable for likelihood inference and outperform both Gaussian DAG and

D-vine-copula models in certain settings. Gaussian DAG-copula models exhibit particularly poor performance in presence of non-normal tail behaviour.

We shall mention that our approach, which extends an idea from Hanea et al. (2006), involves considerable computational effort. First, the pair-copula constructions used in our expansion of the likelihood function give rise to conditional cdfs whose evaluation requires numerical integration. Second, the routine for finding ML estimates faces the usual difficulties of non-linear high-dimensional optimisation problems. Both issues were addressed in more detail in Bauer et al. (2011). Moreover, selecting a suitable model from the class of non-Gaussian DAG-PCC models is still an open problem and involves both the question for the assumed Markov structure and the question for suitable pair copulas. In a similar vein, there is yet no clear answer to the question of whether regular-vine models or non-Gaussian DAG-PCC models are preferable for modeling a given data set.

REFERENCES

K. Aas. Modelling the dependence structure of financial assets: A survey of four copulas. Research report SAMBA/22/04, Norwegian Computing Centre, 2004.

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

A. Bauer, C. Czado, and T. Klein. Pair-copula constructions for non-Gaussian DAG models. Submitted for publication, 2011.

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

T. Bedford and R. M. Cooke. Vines—A New Graphical Model for Dependent Random Variables. *The Annals of Statistics*, 30(4):1031–1068, 2002.

R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer, New York, second edition, 2003.

D. R. Cox and N. Wermuth. *Multivariate Dependencies*. Chapman & Hall, London, 1996.

A. M. Hanea, D. Kurowicka, and R. M. Cooke. Hybrid Method for Quantifying and Analyzing Bayesian Belief Nets. *Quality and Reliability Engineering International*, 22:709–729, 2006.

I. Hobæk Haff, K. Aas, and A. Frigessi. On the simplified pair-copula construction—Simply useful or too simplistic? *Journal of Multivariate Analysis*, 101(5):1296–1310, 2010.

H. Joe. 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*, volume 28 of *Lecture Notes—Monograph Series*, pages 120–141. Institute of Mathematical Statistics, Hayward, 1996.

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

D. Kurowicka and R. Cooke. *Uncertainty Analysis with High Dimensional Dependence Modelling*. John Wiley & Sons, Chichester, 2006.

D. Kurowicka and H. Joe (Eds.). *Dependence Modeling*. World Scientific, Singapore, 2011.

S. L. Lauritzen. *Graphical Models*. Oxford University Press, Oxford, 1996.

R. B. Nelsen. *An Introduction to Copulas*. Springer, New York, second edition, 2006.

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