

X-VINE MODELS FOR MULTIVARIATE EXTREMES

Anna Kiriliouk, Jeongjin Lee, Johan Segers

REPRINT | 2025 / 13

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/isba/publication>

X-vine models for multivariate extremes

Anna Kiriliouk¹ , Jeongjin Lee¹  and Johan Segers² 

¹Namur Institute for Complex Systems, University of Namur, Rue Grafé 2, Namur 5000, Belgium

²LIDAM/ISBA, UCLouvain, Voie du Roman Pays 20, Louvain-la-Neuve 1348, Belgium

Address for correspondence: Johan Segers, LIDAM/ISBA, UCLouvain, Voie du Roman Pays 20, 1348 Louvain-la-Neuve, Belgium. Email: johan.segers@uclouvain.be

Abstract

Regular vine sequences permit the organization of variables in a random vector along a sequence of trees. Vine-based dependence models have become greatly popular as a way to combine arbitrary bivariate copulas into higher-dimensional ones, offering flexibility, parsimony, and tractability. In this project, we use regular vine sequences to decompose and construct the exponent measure density of a multivariate extreme value distribution, or, equivalently, the tail copula density. Although these densities pose theoretical challenges due to their infinite mass, their homogeneity property offers simplifications. The theory sheds new light on existing parametric families and facilitates the construction of new ones, called X-vines. Computations proceed via recursive formulas in terms of bivariate model components. We develop simulation algorithms for X-vine multivariate Pareto distributions as well as methods for parameter estimation and model selection on the basis of threshold exceedances. The methods are illustrated by Monte Carlo experiments and a case study on US flight delay data.

Keywords: exponent measure, graphical model, multivariate Pareto distribution, pair copula construction, regular vine, tail copula

1 Introduction

For multivariate extremes, margin-free tail dependence models based on max-stable distributions arise from classical limit theory for sample extremes (de Haan & Resnick, 1977). A question of high interest is the construction of such models that are flexible, parsimonious, and computationally tractable, and scale well as the dimension grows (Engelke & Ivanovs, 2021). To do so, we propose a novel approach based on regular vine tree sequences (Bedford & Cooke, 2001, 2002), called X-vines. The models can easily be built in arbitrary dimension by combining bivariate components only. The latter can be chosen independently from one another, giving great flexibility. The pairs are grouped in trees, the number of which determines the complexity of the model. Computations proceed by recursive algorithms.

For copula-based dependence modelling, outside the context of extreme value analysis, vine constructions have grown into a versatile and widely applied approach (Czado, 2019; Czado & Nagler, 2022; Nagler & Vatter, 2023). Our contribution is to make vine-based methods operational, in theory and practice, for densities of exponent measures of multivariate extreme value distributions. The challenge to overcome is that these densities do not integrate to one but to infinity. A change of margin transforms these into tail copula densities, which still have infinite mass, but with structural properties that resemble copula densities more closely.

Figure 1 shows a regular vine sequence $\mathcal{V}$ in dimension $d = 5$. The sequence consists of four trees $T_1, \dots, T_4$, in which the edges of one tree become nodes in the next one. Each of the $d(d-1)/2 = 10$ pairs of variables appears as the leading pair before the semicolon on exactly one edge. The numbers behind the semicolons refer to conditioning variables. The first tree represents a Markov tree, while the subsequent trees add higher-order dependence relations.

Received: December 22, 2023. Revised: June 27, 2024. Accepted: October 9, 2024

© The Royal Statistical Society 2024. All rights reserved. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.




Figure 1. A regular vine sequence $\mathcal{V}$ in dimension $d = 5$ consisting of trees $T_1, \dots, T_4$, where T_j has $d - j + 1$ nodes and $d - j$ edges. The nodes of T_1 are $\{1, \dots, 5\}$, while the edges of T_j become the nodes of T_{j+1} .

An X-vine specification consists of a regular vine sequence $\mathcal{V}$, together with, for each edge in the first tree, a bivariate exponent measure or tail copula density, and, for each edge of the subsequent trees, a bivariate copula density, not necessarily stemming from extreme value theory. For the example in Figure 1, four bivariate exponent measure densities (T_1) and six bivariate copula densities (T_2, T_3, T_4) would be required. These bivariate components can be chosen independently from another, without any constraints. Our main result shows how to combine bivariate building blocks into a single multivariate exponent measure or tail copula density. Further, we show that commonly used parametric models in extreme value analysis are actually examples of such X-vine models. We leverage the recursive nature of regular vine sequences both in theory and computations.

The models constructed in this way turn out to be the tail limits of regular vine copulas. For the special case of D-vine copulas, such tail limits were already computed in Joe (1996) and Joe et al. (2010), focusing on tail dependence functions rather than densities. Li and Wu (2013) define tail densities and compute those of D-vine copulas in dimensions three and four. In Simpson et al. (2021), tails of D-vine and C-vine copulas were investigated from the perspective of the limiting shapes of sample clouds.

X-vine constructions are related to but different from recently proposed graphical models for extremes, either for directed or undirected graphs (Engelke & Hitz, 2020; Gissibl & Klüppelberg, 2018). Markov trees (Hu et al., 2024; Segers, 2020) are a special case in the intersection of graphical and X-vine models.

After reviewing background on multivariate extremes and regular vine sequences in Section 2, we state a version of Sklar's theorem for tail copula densities in Section 3. Some parametric examples are worked out in Section 4, with a focus on the simplifying assumption that the conditional copula densities only depend on the indices of the conditioning variables, but not their actual values. Section 5 contains the paper's main results, showing on the one hand how a general tail copula density can be decomposed along any regular vine sequence, and, on the other hand, how to construct a tail copula density from a regular vine sequence and bivariate building blocks. The so-called X-vine models arising in this way are put to work in the subsequent sections, covering simulation algorithms (Section 6), parameter estimation and model selection (Section 7), simulation studies (Section 8), and a case study on US flight delay data (Section 9) taken from Hentschel et al. (2024). Section 10 concludes. The online supplementary material contains a detailed example to illustrate our theory, the proofs of the paper's results, expressions based on exponent measure densities, and additional numerical results. The methods are implemented in the R (R Core Team, 2023) package *Xvine*,¹ relying in particular on packages *graphicalExtremes* (Engelke et al., 2022), and *VineCopula* (Nagler et al., 2023).

2 Background

We write $[d] = \{1, \dots, d\}$ for the index set of the variables. Bold symbols refer to multivariate quantities. For a point $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and a subset $J \subseteq [d]$, write subvectors as

¹ Available from <https://github.com/JeongjinLee88/Xvine>.

$\mathbf{x}_J = (x_j)_{j \in J}$ and $\mathbf{x}_{J^c} = (x_j)_{j \in [d] \setminus J}$. Mathematical operations on vectors such as addition, multiplication and comparison are considered component-wise.

2.1 Multivariate extreme value theory: tail copulas and their densities

Tail copulas. Classical extreme value theory starts from the assumption that the distribution function F of a random vector $\mathbf{X} = (X_1, \dots, X_d)$ is in the max-domain of attraction of a multivariate extreme value distribution (Beirlant et al., 2004; de Haan & Ferreira, 2007). This assumption concerns the tails of the univariate marginal distribution functions $F_1, \dots, F_d$ and the tail dependence structure of $\mathbf{X}$. Our focus is on the latter aspect, that is, on probabilities of high values occurring jointly among the variables X_j . Let C be the survival copula of $\mathbf{X}$: under the standing assumption that $F_1, \dots, F_d$ are continuous, C is the distribution function of the random vector $\mathbf{U} = (U_1, \dots, U_d)$ with uniformly distributed components $U_j = 1 - F_j(X_j)$ for $j \in [d]$. Our interest is in high values of X_j and thus in low outcomes of U_j .

Definition 2.1 Tail copula and tail copula measure. The (lower) tail copula R of C is the function on $\mathbb{E} = (0, \infty]^d \setminus \{\infty\}$ defined by

$$R(\mathbf{x}) = \lim_{t \searrow 0} t^{-1} C(t\mathbf{x}), \quad \mathbf{x} \in \mathbb{E}, \quad (2.1)$$

provided the limit exists. If it does, then the *tail copula measure*, denoted by the same symbol R , is the Borel measure on $\mathbb{E}$ determined by $R((0, \mathbf{x}]) = R(\mathbf{x})$ for $\mathbf{x} \in \mathbb{E}$.

The term ‘tail copula’ stems from Schmidt and Stadtmüller (2006) in case $d = 2$. The tail copula measure already appears in Einmahl et al. (2001) and is closely linked to the exponent measure A in equation (2.9), introduced in de Haan and Resnick (1977). The limit (2.1) appears in Jaworski (2006) and in Joe et al. (2010), who call it ‘tail dependence function’.

The total mass of the tail copula measure R equals infinity but $R(B)$ is finite for Borel sets B contained in $\{\mathbf{x} \in \mathbb{E} : \min \mathbf{x} \leq M\}$ for some $M > 0$. Its univariate margins are equal to the one-dimensional Lebesgue measure:

$$\forall j \in [d], \forall x_j \in (0, \infty), \quad R(\{\mathbf{y} \in \mathbb{E} : y_j \leq x_j\}) = x_j. \quad (2.2)$$

Equation (2.1) implies that R is homogeneous of order one, both as a function and as a measure: for $\mathbf{x} \in \mathbb{E}$, Borel sets $B \subseteq \mathbb{E}$, and $s \in (0, \infty)$,

$$R(s\mathbf{x}) = s R(\mathbf{x}) \quad \text{and} \quad R(sB) = s R(B). \quad (2.3)$$

Any Borel measure R on $\mathbb{E}$ satisfying (2.2) and (2.3) is a tail copula measure, i.e. is the limit in (2.1) for some copula C ; see Lemma B.1 in the online supplementary material.

Tail copula densities. Throughout, we assume that the tail copula measure R has no mass on the hyperplanes through infinity, $R(\{\mathbf{x} \in \mathbb{E} : x_j = \infty\}) = 0$ for all $j \in [d]$, so R is supported on $(0, \infty)^d$. Further, we assume that R is absolutely continuous with respect to the d -dimensional Lebesgue measure with continuous density $r : (0, \infty)^d \rightarrow [0, \infty)$, that is, $R(B) = \int_B r(\mathbf{x}) d\mathbf{x}$ for all Borel sets $B \subseteq (0, \infty)^d$. Choosing $B = (0, \mathbf{x}]$ for $\mathbf{x} \in (0, \infty)^d$, we recover r from R by $r(\mathbf{x}) = \frac{\partial^d}{\partial x_1 \dots \partial x_d} R(\mathbf{x})$ for all $\mathbf{x} \in (0, \infty)^d$. The marginal constraint (2.2) implies

$$\forall j \in [d], \forall x_j \in (0, \infty), \quad \int_{(0, \infty)^{d-1}} r(\mathbf{x}) d\mathbf{x}_{\setminus j} = 1. \quad (2.4)$$

In view of (2.3), r is homogeneous of order $1 - d$:

$$\forall s \in (0, \infty), \forall \mathbf{x} \in (0, \infty)^d, \quad r(s\mathbf{x}) = s^{1-d} r(\mathbf{x}). \quad (2.5)$$

Properties (2.4) and (2.5) characterize the set of d -variate *tail copula densities*; see Section 4 for some parametric families.

By Scheffé's lemma, if the copula C has density c and the tail copula measure R has density r , then (2.1) is implied by $\lim_{t \searrow 0} t^{1-d} c(t\mathbf{x}) = r(\mathbf{x})$ for all $\mathbf{x} \in (0, \infty)^d$. Li and Wu (2013) call r the (lower) tail density function of C . A word of caution: for a given copula density c , the limit $\lim_{t \searrow 0} t^{1-d} c(t\mathbf{x})$ may exist for all $\mathbf{x} \in (0, \infty)^d$ but not be a tail copula density, as the marginal constraint (2.4) may fail. A case in point is the independence copula, with density $c \equiv 1$, in which case the said limit is zero.

Margins of tail copulas. For nonempty $J \subseteq [d]$, let $\pi_J : (0, \infty)^d \rightarrow (0, \infty)^J$ denote the coordinate projection $\mathbf{x} \mapsto \pi_J(\mathbf{x}) = \mathbf{x}_J$ and let $R_J = R \circ \pi_J^{-1}$ denote the J th marginal measure

$$R_J(B) = R(\pi_J^{-1}(B)) = R(\{\mathbf{x} \in (0, \infty)^d : \mathbf{x}_J \in B\}), \quad (2.6)$$

for Borel sets $B \subseteq (0, \infty)^J$. The choice $B = \prod_{j \in J} (0, x_j]$ for $\mathbf{x}_j \in (0, \infty)$ shows that R_J is the tail copula measure of the copula C_J of $U_J = \pi_J(U)$, as in Definition 2.1. For $\mathbf{x}_J \in (0, \infty]^J \setminus \{\infty\}$, we have $R_J(\mathbf{x}_J) = R(\tilde{\mathbf{x}}_J)$ with $\tilde{x}_j = x_j$ if $j \in J$ and $\tilde{x}_j = \infty$ if $j \in [d] \setminus J$. For $\mathbf{x}_J \in (0, \infty)^J$, the density r_J of the measure R_J is

$$r_J(\mathbf{x}_J) = \frac{\partial^{|J|}}{\prod_{j \in J} \partial x_j} R_J\left(\prod_{j \in J} (0, x_j]\right) = \int_{\mathbf{x}_{[d] \setminus J} \in (0, \infty)^{d-|J|}} r(\mathbf{x}) d\mathbf{x}_{[d] \setminus J}, \quad (2.7)$$

and, provided $J \neq [d]$, is obtained from r by integrating out the variables x_j with indices $j \notin J$.

Relations between tail copulas and other concepts in multivariate extremes. We review a number of related terms from multivariate extreme value theory. This paragraph can be skipped at first reading, as the contributions in Sections 3 to 5 do not depend on it.

Equation (2.1) is equivalent to the assumption that the random vector $\mathbf{V} = (V_1, \dots, V_d)$ with $V_j = 1/U_j$ for $j \in [d]$ satisfies

$$\lim_{n \rightarrow \infty} P(\mathbf{V} \leq n\mathbf{z})^n = G(\mathbf{z}) = \exp\{-\ell(1/\mathbf{z})\}, \quad \mathbf{z} \in (0, \infty)^d, \quad (2.8)$$

that is, $\mathbf{V}$ is in the max-domain of attraction of a *multivariate extreme value* (or *max-stable*) distribution G defined in terms of the *stable tail dependence function* $\ell(\mathbf{x}) = R(\mathbb{E} \setminus [\mathbf{x}, \infty])$ for $\mathbf{x} \in [0, \infty)^d$. The margins of $\mathbf{V}$ are unit-Pareto, $P(V_j > x) = 1/x$ for $x \geq 1$, while those of G are unit-Fréchet, $G_j(x) = \exp(-1/x)$ for $x > 0$. Equivalently, the distribution of $\mathbf{V}$ is *multivariate regularly varying* (de Haan & Resnick, 1977; Resnick, 2007) with limit measure

$$\Lambda(\cdot) = R(\{\mathbf{x} \in \mathbb{E} : (1/x_1, \dots, 1/x_d) \in \cdot\}) \quad \text{on } [0, \infty)^d \setminus \{0\}, \quad (2.9)$$

that is, $tP(\mathbf{V} \in tB) \rightarrow \Lambda(B)$ as $t \rightarrow \infty$ for Borel sets B contained in $[0, \infty)^d \setminus [0, \varepsilon]^d$ for some $\varepsilon > 0$ and satisfying $\Lambda(\partial B) = 0$, with ∂B the boundary of B . The measure Λ is called *exponent measure* because of the identity $G(\mathbf{z}) = \exp\{-\Lambda([0, \infty) \setminus [0, \mathbf{z}])\}$ for $\mathbf{z} \in (0, \infty)^d$.

Another statement equivalent to (2.1) is that the conditional distribution of $\mathbf{V}/t$ given $\max \mathbf{V} > t$ is asymptotically *multivariate Pareto*, that is, we have the weak convergence

$$(\mathbf{V}/t \mid \max \mathbf{V} > t) \rightsquigarrow \mathbf{Y}, \quad t \rightarrow \infty, \quad (2.10)$$

where $\mathbf{Y}$ is a random vector supported on $\mathbb{L}_{>1} := \{\mathbf{y} \in [0, \infty)^d : \max \mathbf{y} > 1\}$ with distribution equal to the restriction of Λ in (2.9) to $\mathbb{L}_{>1}$ and normalized to a probability measure,

$$P(\mathbf{Y} \in B) = \Lambda(B \cap \mathbb{L}_{>1}) / \Lambda(\mathbb{L}_{>1}) \quad (2.11)$$

for Borel sets $B \subseteq \mathbb{R}^d$ and normalizing constant $\Lambda(\mathbb{L}_{>1}) = \ell(\mathbf{1})$. For $\mathbf{y} \in \mathbb{L}_{>1}$, we also have $P(\mathbf{Y} \geq \mathbf{y}) = R(\mathbf{1}/\mathbf{y})/\ell(\mathbf{1})$.

Up to a location shift, multivariate Pareto distributions are a special case of *multivariate generalized Pareto distributions* introduced in Rootzén and Tajvidi (2006). They play a prominent role in the theory of graphical models for extremes (Engelke & Hitz, 2020).

In practice, it may be convenient to reduce the information contained in the tail copula measure to a dependence coefficient. In this paper, we will encounter the *tail dependence coefficients* $\chi_J = R(\{x \in (0, \infty]^d : \max_{j \in J} x_j < 1\}) = R_J(1, \dots, 1)$ for index sets $J \subseteq [d]$ with two or more elements.

If the tail copula measure R is concentrated on $(0, \infty)^d$ and has density r , the exponent measure and the multivariate Pareto distribution have densities too. The exponent measure λ is concentrated on $(0, \infty)^d$ and has density

$$\lambda(\mathbf{y}) = r(\mathbf{1}/\mathbf{y}) \prod_{j=1}^d y_j^{-2}, \quad \mathbf{y} \in (0, \infty)^d, \quad (2.12)$$

while the multivariate Pareto vector $\mathbf{Y}$ has probability density $\lambda(\mathbf{y})/\ell(\mathbf{1})$ for $\mathbf{y} \in \mathbb{L}_{>1}$.

2.2 Regular vine sequences and the vine telescoping product formula

Informally, a regular vine sequence on d elements consists of a linked set of $d - 1$ trees, where the edges in tree $j - 1$ become nodes in tree j , joined by an edge only if they share a common node as edges in tree $j - 1$, for $j \in \{2, \dots, d - 1\}$ (Bedford & Cooke, 2002). *Regular vine copulas*, also called *pair copula constructions*, allow for flexible high-dimensional modelling using parametric families of bivariate copulas as building blocks for the edges of each tree. For recent overviews, see, for example, Czado (2019) or Czado and Nagler (2022). Our objective is to do the same for tail copula densities. All concepts introduced below are illustrated via a five-dimensional example in [Section A of the online supplementary material](#).

A *tree* $T = (N, E)$ is a connected, acyclic graph comprising a finite node set N and an edge set $E \subseteq \{\{x, y\} : x, y \in N, x \neq y\}$. Any two distinct nodes in a tree are connected by a unique path.

Definition 2.2 Regular vine (tree) sequence. An ordered set of trees $\mathcal{V} = (T_1, \dots, T_{d-1})$ is a *regular vine tree sequence* on $d \geq 3$ elements if $T_1 = (N_1, E_1)$ is a tree with node set $N_1 = [d]$, while for $j \in \{2, \dots, d - 1\}$, the tree $T_j = (N_j, E_j)$ has node set $N_j = E_{j-1}$ and the *proximity condition* holds: for any $\{a, b\} \in E_j$, we have $|a \cap b| = 1$, that is, two nodes in T_j can be connected only if, as edges in T_{j-1} , they share a common node.

Definition 2.3 Complete union, conditioning set, conditioned set. Let $\mathcal{V} = (T_1, \dots, T_{d-1})$ be a regular vine sequence, with $T_j = (N_j, E_j)$. For any edge $e \in E_j$, the *complete union* of e is $A_e = e$ if $j = 1$ and $A_e = A_a \cup A_b$ if $e = \{a, b\}$ and $j \in \{2, \dots, d - 1\}$. Informally, A_e is the subset of nodes in $[d]$ reachable from e by the membership relation. The *conditioning set* of an edge $e = \{a, b\} \in E_j$ with $j \geq 2$ is $D_e = A_a \cap A_b$ and the *conditioned set* of e is $C_e = C_{e,a} \cup C_{e,b}$ with $C_{e,a} = A_a \setminus D_e = A_a \setminus A_b$ and $C_{e,b} = A_b \setminus D_e = A_b \setminus A_a$. If $e = \{a, b\} \in E_1$, then $C_{e,a} = \{a\}$, $C_{e,b} = \{b\}$, $C_e = e$ and $D_e = \emptyset$.

For an edge $e = \{a, b\} \in E_j$, the sets $C_{e,a}$, $C_{e,b}$ and D_e are disjoint and their union equals A_e , a node set with $j + 1$ elements. The sets $C_{e,a}$ and $C_{e,b}$ are singletons while D_e has $j - 1$ elements. Clearly, $D_e = A_a \setminus C_e = A_b \setminus C_e$. The edge e is also written as $e = (C_{e,a}, C_{e,b}; D_e)$ and, since $C_{e,a}$ and $C_{e,b}$ are singletons, $\{a_e\}$ and $\{b_e\}$, say, we abbreviate $e = (a_e, b_e; D_e)$, with $a_e, b_e \in [d]$ and the convention that $a_e < b_e$, that is, $a_e = \min C_e$ and $b_e = \max C_e$. Similarly, for $e = \{a, b\} \in E_1$, the convention is that $a < b$. In [Figure 1](#), for instance, the edge e labelled $(14; 23)$ in T_3 has $a_e = 1, b_e = 4$, and $D_e = \{2, 3\}$. See [Section B.1 in the online supplementary material](#) for some additional properties.

For positive scalars $\gamma_1, \dots, \gamma_d$, the identity $\prod_{j=2}^d (\gamma_j/\gamma_{j-1}) = \gamma_d/\gamma_1$ is a *telescoping product*. Regular vine sequences enjoy a similar property that will be key to Theorem 5.1 and which, to the best of our knowledge, has not yet been stated in the literature.

Lemma 2.4 Vine telescoping product. Let $d \geq 3$ be an integer and let $\gamma_j \in (0, \infty)$ for every nonempty $J \subseteq [d]$, with $\gamma_j := \gamma_{\{j\}} = 1$ for every $j \in [d]$. Given a regular vine sequence $\mathcal{V} = (T_j = (N_j, E_j))_{j=1}^{d-1}$ on $[d]$, we have

$$\gamma_{[d]} = \prod_{e \in E_1} \gamma_e \cdot \prod_{j=2}^{d-1} \prod_{e=\{a,b\} \in E_j} \frac{\gamma_{D_e} \cdot \gamma_{A_e}}{\gamma_{A_a} \cdot \gamma_{A_b}}. \quad (2.13)$$

With the notation $\gamma_{I|J} = \gamma_{I \cup J}/\gamma_J$, the last factor in (2.13) is

$$\frac{\gamma_{D_e} \cdot \gamma_{A_e}}{\gamma_{A_a} \cdot \gamma_{A_b}} = \frac{\gamma_{\{a_e, b_e\} | D_e}}{\gamma_{a_e | D_e} \cdot \gamma_{b_e | D_e}}. \quad (2.14)$$

Remark B.4 in Appendix B.1 in the online supplementary material generalizes Lemma 2.4 to factorizations of γ_j for certain nonempty subsets $J \subseteq [d]$, while **Remark B.5** of the online supplementary material provides a generalization of (2.13) without the restriction that $\gamma_j = 1$ for all $j \in [d]$.

3 Sklar's theorem for tail copula densities

Let r be a d -variate tail copula density, so properties (2.4) and (2.5) hold. Recall its multivariate margins in (2.7). For nonempty, disjoint $I, J \subseteq [d]$ and for $\mathbf{x}_I \in (0, \infty)^I$ and $\mathbf{x}_J \in (0, \infty)^J$ such that $r_J(\mathbf{x}_J) > 0$, define $r_{I|J}(\mathbf{x}_I | \mathbf{x}_J) := r_{I \cup J}(\mathbf{x}_{I \cup J})/r_J(\mathbf{x}_J)$ for all $\mathbf{x}_{I \cup J} \in (0, \infty)^{I \cup J}$. Viewing the 'conditional' tail copula density $r_{I|J}$ as an $|I|$ -variate probability density, we can decompose it into $|I|$ univariate probability densities and a copula density.

Proposition 3.1 Sklar's theorem: direct part. Let r be a d -variate tail copula density and let $I, J \subseteq [d]$ be nonempty and disjoint. For $\mathbf{x}_J \in (0, \infty)^J$ such that $r_J(\mathbf{x}_J) > 0$, the function $r_{I|J}(\cdot | \mathbf{x}_J) : (0, \infty)^I \rightarrow [0, \infty)$, defined via $\mathbf{x}_I \mapsto r_{I|J}(\mathbf{x}_I | \mathbf{x}_J)$, is a probability density on $(0, \infty)^I$. For $i \in I$, it has marginal probability density $x_i \mapsto r_{i|J}(x_i | \mathbf{x}_J)$ and cumulative distribution function $x_i \mapsto R_{i|J}(x_i | \mathbf{x}_J) := \int_0^{x_i} r_{i|J}(t | \mathbf{x}_J) dt$, with quantile function $u_i \mapsto R_{i|J}^{-1}(u_i | \mathbf{x}_J)$. If $|I| \geq 2$, then the copula density of $r_{I|J}(\cdot | \mathbf{x}_J)$ is

$$\forall \mathbf{u}_I \in (0, \infty)^I, \quad c_{I|J}(\mathbf{u}_I; \mathbf{x}_J) := \frac{r_{I|J}(\mathbf{x}_I | \mathbf{x}_J)}{\prod_{i \in I} r_{i|J}(x_i | \mathbf{x}_J)} \quad \text{with} \quad x_i = R_{i|J}^{-1}(u_i | \mathbf{x}_J). \quad (3.1)$$

The homogeneity property (2.5) induces certain scaling properties.

Lemma 3.2 Scale equi- and invariance. In Proposition 3.1, for $\mathbf{x}_J \in (0, \infty)^J$ such that $r_J(\mathbf{x}_J) > 0$, the family $\{r_{I|J}(\cdot | t\mathbf{x}_J) : t \in (0, \infty)\}$ of probability densities is a scale family: if the random vector ξ_I has probability density $r_{I|J}(\cdot | \mathbf{x}_J)$, then the random vector $t\xi_I$ has probability density $r_{I|J}(\cdot | t\mathbf{x}_J)$. As a consequence, for all $t \in (0, \infty)$, $\mathbf{x}_I \in (0, \infty)^I$, $\mathbf{u}_I \in (0, 1)^I$ and $i \in I$, we have $r_{i|J}(tx_i | t\mathbf{x}_J) = r_{i|J}(x_i | \mathbf{x}_J)$, $R_{i|J}^{-1}(u_i | t\mathbf{x}_J) = t \cdot R_{i|J}^{-1}(u_i | \mathbf{x}_J)$, and

$$c_{I|J}(\mathbf{u}_I; t\mathbf{x}_J) = c_{I|J}(\mathbf{u}_I; \mathbf{x}_J). \quad (3.2)$$

In particular, if $J = \{j\}$ for some $j \in [d]$, then $c_{I;j}(\cdot; \mathbf{x}_j) = c_{I;j}(\cdot; 1)$ does not depend on the value of $\mathbf{x}_j \in (0, \infty)$.

Remark 3.3 Conditional independence. Let $\mathbf{Y}$ be the multivariate Pareto vector in (2.11), suppose $d \geq 3$, and let i, j be distinct elements in $[d]$, with $J = [d] \setminus \{i, j\}$. Provided r is positive and continuous, Y_i and Y_j are conditionally independent given $\mathbf{Y}_J$ in the sense of Definition 1 in Engelke and Hitz (2020) if and only if $c_{i,j;J}(\mathbf{u}_i, \mathbf{u}_j; \mathbf{x}_J) = 1$ for all $(\mathbf{u}_i, \mathbf{u}_j) \in (0, 1)^2$ and $\mathbf{x}_J \in (0, \infty)^J$, the density of the independence copula. This statement follows by combining the factorization property in Proposition 1 in Engelke and Hitz (2020) with equations (2.12) and (3.1); see also (C.1) in Appendix C in the online supplementary material for an expression of $c_{I;J}$ in terms of the exponent measure density λ .

Definition 3.4 Simplifying assumption. The family of copula densities $c_{I;J}(\cdot; \mathbf{x}_J)$ in (3.1) satisfies the *simplifying assumption* if the copula densities do not depend on $\mathbf{x}_J \in (0, \infty)^J$, that is, there exists a single copula density $c_{I;J}$ such that $c_{I;J}(\mathbf{u}_I; \mathbf{x}_J) = c_{I;J}(\mathbf{u}_I)$ for all $\mathbf{u}_I \in (0, 1)^I$ and $\mathbf{x}_J \in (0, \infty)^J$.

If J is a singleton, $J = \{j\}$, the scale invariance (3.2) implies that $\{c_{I;j}(\cdot; \mathbf{x}_j) : \mathbf{x}_j \in (0, \infty)\}$ always satisfies the simplifying assumption. In Section 4, we will see that the copula families induced by many commonly used parametric exponent measure models satisfy the simplifying assumption for all I and J .

As a converse to Proposition 3.1, we can combine several tail copula densities along a scale-invariant family of copula densities to form a tail copula density in higher dimensions. With slight abuse of notation, write $i \cup J := \{i\} \cup J$ for $i \in I$. A family of k -variate copula densities $\{c(\cdot; \theta) : \theta \in \Theta\}$ indexed by a Borel set $\Theta \subseteq \mathbb{R}^m$ is *jointly measurable* if the map $(\mathbf{u}, \theta) \mapsto c(\mathbf{u}; \theta)$ is Borel measurable on $(0, 1)^k \times \Theta$.

Proposition 3.5 Sklar's theorem: converse part. Let $I \cup J$ be a partition of $[d]$ into two disjoint, nonempty sets, with $d \geq 3$ and $|I| \geq 2$. Let $(r_{i \cup J})_{i \in I}$ be an $|I|$ -tuple of $(|J| + 1)$ -variate tail copula densities with common margin r_J . Let $\{c_{I;J}(\cdot; \mathbf{x}_J) : \mathbf{x}_J \in (0, \infty)^J\}$ be a jointly measurable family of $|I|$ -variate copula densities such that (3.2) holds. Then the function $r : (0, \infty)^d \rightarrow [0, \infty)$ defined by

$$r(\mathbf{x}) := r_J(\mathbf{x}_J) \cdot \prod_{i \in I} r_{i \cup J}(\mathbf{x}_i | \mathbf{x}_J) \cdot c_{I;J}(\mathbf{R}_{i \cup J}(\mathbf{x}_i | \mathbf{x}_J), i \in I; \mathbf{x}_J), \quad (3.3)$$

for $\mathbf{x} \in (0, \infty)^d$ such that $r_J(\mathbf{x}_J) > 0$ and zero otherwise, is a d -variate tail copula density with margins $r_{i \cup J}$ for $i \in I$.

4 Parametric families of tail copula densities

We will compute the copula densities $c_{I;J}(\cdot; \mathbf{x}_J)$ in (3.1) for tail copula densities r arising from two parametric families: the scaled extremal Dirichlet model (Section 4.1) due to Belzile and Nešlehová (2017), encompassing the logistic, negative logistic and Dirichlet max-stable models, and the Hüsler–Reiss model (Section 4.2). For both families, the simplifying assumption (Definition 3.4) is satisfied for all choices of I and J , and the copula densities $c_{I;J}$ take on known parametric forms. This section can be skipped at first reading.

4.1 Scaled extremal Dirichlet model

In dimension $d \geq 2$, consider parameters $\alpha_1, \dots, \alpha_d > 0$ and $\rho > -\min(\alpha_1, \dots, \alpha_d)$ with $\rho \neq 0$. The angular measure density of the *scaled extremal Dirichlet* model is introduced in Belzile and Nešlehová (2017, Proposition 5). By equation (2.12) and Lemma B.2 in the online

supplementary material, the corresponding tail copula density is

$$r(\mathbf{x}) = \frac{\Gamma(\bar{\alpha} + \rho)}{d|\rho|^{d-1} \prod_{j=1}^d \Gamma(\alpha_j)} \cdot \left\{ \sum_{j=1}^d c(\alpha_j, \rho)^{1/\rho} x_j^{-1/\rho} \right\}^{-\rho-\bar{\alpha}} \cdot \prod_{j=1}^d c(\alpha_j, \rho)^{\alpha_j/\rho} x_j^{-\alpha_j/\rho-1}, \quad (4.1)$$

for $\mathbf{x} \in (0, \infty)^d$, where $\bar{\alpha} = \alpha_1 + \dots + \alpha_d$ and $c(\alpha, \rho) = \Gamma(\alpha + \rho)/\Gamma(\alpha)$. The model is closed under marginalization: for $J \subseteq [d]$ with $|J| \geq 2$, the marginal tail copula density r_J in (2.7) is of the same form in dimension $|J|$ and with parameters $(\alpha_j)_{j \in J}$ and ρ . The model unites and extends several well-known parametric families in multivariate extreme value analysis:

- if $\alpha_1 = \dots = \alpha_d = 1$ and $\rho > 0$, we get the negative logistic model (Joe, 1990);
- if $\alpha_1 = \dots = \alpha_d = 1$ and $-1 < \rho < 0$, we obtain the logistic model (Gumbel, 1960);
- if $\rho = 1$, we get the Coles–Tawn extremal Dirichlet model (Coles & Tawn, 1991).

In McNeil and Nešlehová (2010, Definition 3), the family of *Liouville* copulas is defined as the collection of survival copulas of random vectors of the form $\mathbf{Y} = S\mathbf{W}$, where $\mathbf{W}$ has a Dirichlet distribution on the unit simplex and is independent of the positive random variable S . The special case where $\mathbf{W}$ is uniformly distributed on the unit simplex yields the family of *Archimedean* copulas (McNeil & Nešlehová, 2009).

Proposition 4.1 Let r be the scaled extremal Dirichlet tail copula density in (4.1) in dimension $d \geq 3$. For every pair of disjoint, nonempty subsets $I, J \subset [d]$ with $|I| \geq 2$, the copula densities $c_{I,J}(\cdot; \mathbf{x}_J)$ in (3.1) satisfy the simplifying assumption (Definition 3.4). If $\rho > 0$, then $c_{I,J}$ is equal to the density of the $|I|$ -variate Liouville copula with Dirichlet parameters $(\alpha_i)_{i \in I}$ and with radial density on $(0, \infty)$ proportional to $s \mapsto s^{\sum_{i \in I} \alpha_i - 1} \cdot (s+1)^{-\rho - \sum_{k \in I \cup J} \alpha_k}$. If $\rho < 0$, then $c_{I,J}$ is equal to the density of the survival copula of the above Liouville copula.

In special cases, the copula densities $c_{I,J}$ can be identified more explicitly (calculations in Appendix B.3 in the online supplementary material):

- if $\alpha_1 = \dots = \alpha_d = 1$ and $\rho > 0$, then $c_{I,J}$ is the density of the $|I|$ -variate Clayton copula with parameter $\theta/(1 + |J|\theta)$ with $\theta = 1/\rho > 0$;
- if $\alpha_1 = \dots = \alpha_d = 1$ and $-1 < \rho < 0$, then $c_{I,J}$ is the density of the $|I|$ -variate Clayton survival copula with parameter $\theta/(|J|\theta - 1)$ with $\theta = -1/\rho > 1$.

In both cases, the dependence *decreases* as the number $|J|$ of conditioning variables increases.

4.2 Hüsler–Reiss model

Let $\mathcal{D}_d$ be the set of symmetric, strictly conditionally negative definite matrices, that is, all matrices $\Gamma = (\Gamma_{ij})_{i,j=1}^d \in \mathbb{R}^{d \times d}$ of the form $\Gamma_{ij} = \text{Var}(A_i - A_j)$ for a d -variate random vector $\mathbf{A} = (A_1, \dots, A_d)$ with positive definite covariance matrix, i.e. Γ is a variogram matrix. Fix $k \in [d]$ and consider the $(d-1) \times (d-1)$ -dimensional positive definite covariance matrix

$$\Sigma^{(k)} = \frac{1}{2} (\Gamma_{ik} + \Gamma_{jk} - \Gamma_{ij})_{i,j \neq k}; \quad (4.2)$$

for Γ as in the previous sentence, we have $\Sigma_{ij}^{(k)} = \text{Cov}(A_i - A_k, A_j - A_k)$. The d -variate *Hüsler–Reiss* tail copula density is

$$r(\mathbf{x}; \Sigma^{(k)}) = \left(\prod_{i: i \neq k} x_i^{-1} \right) \phi_{d-1}(\tilde{\mathbf{x}}_{\setminus k}; \Sigma^{(k)}), \quad \mathbf{x} \in (0, \infty)^d, \quad (4.3)$$

where ϕ_{d-1} is the centred $(d-1)$ -dimensional Gaussian density with the stated covariance matrix and $\bar{x}_{\setminus k} = (\log(x_i/x_k) - \frac{1}{2}\Gamma_{ik})_{i:i \neq k}$. The exponent measure density λ associated with r via (2.12) appears in Engelke et al. (2015), where it is shown that it does not depend on the choice of $k \in [d]$. The corresponding max-stable distribution G in (2.8) goes back to Hüsler and Reiss (1989), who studied maxima of triangular arrays of Gaussian random vectors.

For a matrix M and for index sets K and L , write $M_{KL} = (M_{ij})_{i \in K, j \in L}$. Let Φ denote the standard normal distribution function and recall that the m -variate Gaussian copula density with $(m \times m)$ -dimensional positive definite correlation matrix R evaluated in $\mathbf{u} \in (0, 1)^m$ is

$$c_R(\mathbf{u}) = |\mathbf{R}|^{-1/2} \cdot \exp\left\{-\frac{1}{2}\mathbf{z}^\top(\mathbf{R}^{-1} - \mathbf{I}_p)\mathbf{z}\right\} \quad \text{where} \quad \mathbf{z}_i = \Phi^{-1}(u_i). \quad (4.4)$$

Proposition 4.2 Let Γ be a $d \times d$ variogram with $d \geq 3$ and let r be the Hüsler–Reiss tail copula density in (4.3). For every pair of disjoint, nonempty subsets $I, J \subset [d]$ with $|I| \geq 2$, the copula densities $c_{I|J}(\cdot; \mathbf{x}_J)$ in (3.1) satisfy the simplifying assumption (Definition 3.4). The copula density $c_{I|J}$ is equal to the $|I|$ -variate Gaussian copula density (4.4) with correlation matrix $\mathbf{R}_{I|J} = (\Delta_{I|J}^{(k)})^{-1/2} \Sigma_{I|J}^{(k)} (\Delta_{I|J}^{(k)})^{-1/2}$ for any $k \in J$, where $\Delta_{I|J}^{(k)}$ is the diagonal matrix with the same diagonal as $\Sigma_{I|J}^{(k)} = \Sigma_{II}^{(k)} - \Sigma_{IJ}^{(k)} (\Sigma_{JJ}^{(k)})^{-1} \Sigma_{JI}^{(k)}$.

The matrix $\mathbf{R}_{I|J}$ in Proposition 4.2 is the correlation matrix of the conditional Gaussian distribution with covariance matrix $\Sigma_{I|J}^{(k)}$, the corresponding Schur complement. Its expression depends on the choice of $k \in J$, although the actual matrix does not. It remains to be investigated how to express it in terms of the Hüsler–Reiss precision matrix Θ introduced in Hentschel et al. (2024).

5 Vine decompositions of tail copula densities

5.1 X-vines: density decomposition and construction

In this section, we show that any tail copula density r in dimension $d \geq 3$ can be decomposed along any regular vine sequence (Definition 2.2) on d elements into $d-1$ bivariate tail copula densities and $\sum_{i=1}^{d-2} i = d-12$ bivariate copula densities. Moreover, copula densities with only a single conditioning variable satisfy the simplifying assumption (Definition 3.4). Conversely, starting from any regular vine sequence on $d \geq 3$ elements, any collection of $d-1$ bivariate tail copula densities and any collection of $d-12$ bivariate copula densities, a d -variate tail copula density can be assembled. We coin tail copula densities constructed in this way as *X-vines*. Appendix C in the online supplementary material states the main results in terms of exponent measure densities.

Theorem 5.1 Tail copula density decomposition along a regular vine. Let r be a d -variate ($d \geq 3$) tail copula density and let $\mathcal{V} = (T_j)_{j=1}^{d-1}$ with $T_j = (N_j, E_j)$ be a regular vine sequence on $[d]$. For $\mathbf{x} \in (0, \infty)^d$, we have

$$r(\mathbf{x}) = \prod_{e \in E_1} r_{a_e, b_e}(x_{a_e}, x_{b_e}) \cdot \prod_{j=2}^{d-1} \prod_{e \in E_j} c_{a_e, b_e; D_e}(R_{a_e|D_e}(x_{a_e} | \mathbf{x}_{D_e}), R_{b_e|D_e}(x_{b_e} | \mathbf{x}_{D_e}); \mathbf{x}_{D_e}), \quad (5.1)$$

where, for $e = (a_e, b_e; D_e) \in E_2 \cup \dots \cup E_{d-1}$, the pair-copula density $c_{a_e, b_e; D_e}$ is

$$c_{a_e, b_e; D_e}(u_{a_e}, u_{b_e}; \mathbf{x}_{D_e}) = \frac{r_{a_e, b_e|D_e}(x_{a_e}, x_{b_e} | \mathbf{x}_{D_e})}{r_{a_e|D_e}(x_{a_e} | \mathbf{x}_{D_e}) \cdot r_{b_e|D_e}(x_{b_e} | \mathbf{x}_{D_e})}, \quad (5.2)$$

with $x_{a_e} = R_{a_e|D_e}^{-1}(u_{a_e} | x_{D_e})$ and $x_{b_e} = R_{b_e|D_e}^{-1}(u_{b_e} | x_{D_e})$ for $u_{a_e}, u_{b_e} \in (0, 1)$.

The decomposition (5.1) also applies to marginal tail copula densities r_f for index sets $J = A_f$ for some edge f in the vine; see [Remark B.7 in the online supplementary material](#). For the pair copula $C_{a_e, b_e; D_e}(\cdot, \cdot; x_{D_e})$ associated with the edge $e = (a_e, b_e; D_e)$ in the regular vine sequence, consider the first-order partial derivatives

$$\begin{aligned} C_{a_e|b_e; D_e}(u_{a_e} | u_{b_e}; x_{D_e}) &:= \frac{\partial}{\partial u_{b_e}} C_{a_e, b_e; D_e}(u_{a_e}, u_{b_e}; x_{D_e}) = \int_{v=0}^{u_{b_e}} c_{a_e, b_e; D_e}(v, u_{b_e}; x_{D_e}) dv, \\ C_{b_e|a_e; D_e}(u_{b_e} | u_{a_e}; x_{D_e}) &:= \frac{\partial}{\partial u_{a_e}} C_{a_e, b_e; D_e}(u_{a_e}, u_{b_e}; x_{D_e}) = \int_{v=0}^{u_{a_e}} c_{a_e, b_e; D_e}(u_{a_e}, v; x_{D_e}) dv. \end{aligned} \quad (5.3)$$

Theorem 5.2 Recursion and uniqueness. In the setting of Theorem 5.1, for any $e = (a_e, b_e; D_e) \in E_2 \cup \dots \cup E_{d-1}$, we have

$$\begin{aligned} R_{a_e|D_e \cup b_e}(x_{a_e} | x_{D_e \cup b_e}) &= C_{a_e|b_e; D_e}(R_{a_e|D_e}(x_{a_e} | x_{D_e}) | R_{b_e|D_e}(x_{b_e} | x_{D_e}); x_{D_e}), \\ R_{b_e|D_e \cup a_e}(x_{b_e} | x_{D_e \cup a_e}) &= C_{b_e|a_e; D_e}(R_{b_e|D_e}(x_{b_e} | x_{D_e}) | R_{a_e|D_e}(x_{a_e} | x_{D_e}); x_{D_e}). \end{aligned} \quad (5.4)$$

As a consequence, r is determined uniquely by the bivariate tail copula densities $r_{a,b}$ for $e = \{a, b\} \in E_1$ and the bivariate copula densities $c_{a_e, b_e; D_e}(\cdot; x_{D_e})$ for $e \in E_2 \cup \dots \cup E_{d-1}$ and $x_{D_e} \in (0, \infty)^{D_e}$.

Equation (5.4) is effectively a recursive relation, allowing to reduce the number of conditioning variables until there is only one conditioning variable left. The reason is that each of the indices $a_e | D_e$ and $b_e | D_e$ in (5.4) is itself of the form $a_f | D_f \cup b_f$ or $b_f | D_f \cup a_f$ for an edge $f \in e$ in E_{j-1} , i.e. one level lower than e .

Definition 5.3 X-vine tail copula density. A d -variate tail copula density r is an *X-vine* along a regular vine sequence $\mathcal{V}$ if for each edge $e \in E_2 \cup \dots \cup E_{d-1}$, the pair copula densities $c_{a_e, b_e; D_e}(\cdot, \cdot; x_{D_e})$ do not depend on the value of $x_{D_e} \in (0, \infty)^{D_e}$.

Example 5.4 Trivariate case. By Lemma 3.2, a trivariate tail copula density r is always an X-vine, and this along any of the three possible regular vine sequences on $\{1, 2, 3\}$. For the vine determined by $E_1 = \{\{1, 2\}, \{2, 3\}\}$, for instance, we have

$$\begin{aligned} r(x_1, x_2, x_3) &= r_{12}(x_1, x_2) r_{23}(x_2, x_3) \cdot c_{13;2}(R_{1|2}(x_1 | x_2), R_{3|2}(x_3 | x_2)), \\ \text{where } c_{13;2}(u_1, u_3) &= \frac{r(x_1, 1, x_3)}{r_{12}(x_1, 1) r_{23}(1, x_3)} \quad \text{with } x_j = R_{j|2}^{-1}(u_j | 1), \quad j \in \{1, 3\}. \end{aligned} \quad (5.5)$$

The function r is thus completely specified by the two bivariate tail copula densities r_{12} and r_{23} and one bivariate copula density $c_{13;2}$. The form (5.5) was already discovered for tail copula densities of D-vine copulas in [Li and Wu \(2013\)](#).

By Propositions 4.1 and 4.2, the scaled extremal Dirichlet model (including the logistic, negative logistic and extremal Dirichlet models) and the Hüsler–Reiss family have conditional copula densities $c_{i,j}(\cdot; x_j)$ that always satisfy the simplifying assumption. As a consequence, they are examples of X-vine tail copula densities too, and this along any regular vine sequence.




Figure 2. (a) An X-vine specification $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ on $d = 5$ variables. The bivariate tail copula densities r_{a_e, b_e} in the first tree T_1 are selected from the Hüsler–Reiss (HR), negative logistic (NL), logistic (L), and Dirichlet (Diri) parametric families, each with specified tail dependence coefficients χ . The bivariate copula densities $c_{a_e, b_e; D_e}$ in the subsequent trees T_2, T_3, T_4 are chosen from the Clayton (Clay), Gumbel (Gum), and Gaussian (Ga) copula families, each with specified Kendall's tau τ . (b) Box-plots (first four for T_1 , next three for T_2 , next two for T_3 and last one for T_4) of dependence measure estimates from bivariate parametric families selected sequentially from the data for the X-vine specification in (a). The lines in the middle of the boxes indicate the specified parameter values.

Definition 5.5 X-vine specification. The triplet $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ is an *X-vine specification* on d elements ($d \geq 3$) if:

1. $\mathcal{V} = (T_j)_{j=1}^{d-1}$ with $T_j = (N_j, E_j)$ is a regular vine sequence on $[d]$;
2. $\mathcal{R} = \{r_{a_e, b_e} : e = \{a_e, b_e\} \in E_1\}$ is a family of bivariate tail copula densities;
3. $\mathcal{C} = \{c_{a_e, b_e; D_e} : e = (a_e, b_e; D_e) \in \bigcup_{j \geq 2} E_j\}$ is a family of bivariate copula densities.

Figure 2a shows an example of an X-vine specification involving a mix of parametric models for the bivariate (tail) copula densities.

Theorem 5.6 X-vine tail copula density construction. Let $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ be an X-vine specification on $d \geq 3$ elements. Then the function r defined by

$$r(\mathbf{x}) = \prod_{e \in E_1} r_{a_e, b_e}(x_{a_e}, x_{b_e}) \cdot \prod_{j=2}^{d-1} \prod_{e \in E_j} c_{a_e, b_e; D_e}(R_{a_e | D_e}(x_{a_e} | \mathbf{x}_{D_e}), R_{b_e | D_e}(x_{b_e} | \mathbf{x}_{D_e})) \quad (5.6)$$

with the functions $R_{\cdot | \cdot}$ defined recursively by

$$\left. \begin{aligned} R_{a_e | D_e \cup b_e}(x_{a_e} | \mathbf{x}_{D_e \cup b_e}) &= C_{a_e | b_e; D_e}(R_{a_e | D_e}(x_{a_e} | \mathbf{x}_{D_e}) | R_{b_e | D_e}(x_{b_e} | \mathbf{x}_{D_e})) \\ R_{b_e | D_e \cup a_e}(x_{b_e} | \mathbf{x}_{D_e \cup a_e}) &= C_{b_e | a_e; D_e}(R_{b_e | D_e}(x_{b_e} | \mathbf{x}_{D_e}) | R_{a_e | D_e}(x_{a_e} | \mathbf{x}_{D_e})) \end{aligned} \right\} \quad (5.7)$$

is a d -variate tail copula density. For $e \in E_1$, the bivariate margin of r is equal to $r_e \in \mathcal{R}$, while for $e = (a_e, b_e; D_e) \in E_2 \cup \dots \cup E_{d-1}$, the pair copula density $c_{a_e, b_e; D_e}(\cdot, \cdot; \mathbf{x}_{D_e})$ in (5.2) is equal to $c_{a_e, b_e; D_e} \in \mathcal{C}$. In particular, r is an X-vine.

In Example 5.4 with $d = 3$, the vine $\mathcal{V} = (T_1, T_2)$ is determined by $E_1 = \{\{1, 2\}, \{2, 3\}\}$, while $\mathcal{R} = \{r_{12}, r_{23}\}$ and $\mathcal{C} = \{c_{13;2}\}$, and (5.6) reduces to (5.5).

5.2 X-vines as limits of regular vine copula densities

A natural question is whether X-vine tail copula densities arise as the lower tail dependence limits of regular vine copula densities, as introduced in Bedford and Cooke (2002). Below, we show that this is indeed the case, provided the pair copula densities at the edges of the first tree have the corresponding bivariate tail copula densities as lower tail dependence limits. In the passage to the limit, the regular vine sequence is preserved and so are the pair copulas at all trees starting from the second one.

Let $\mathcal{V} = (T_j)_{j=1}^{d-1}$ with $T_j = (N_j, E_j)$ be a regular vine sequence on d elements, for $d \geq 3$. For every edge $e = (a_e, b_e; D_e) \in \bigcup_{j=1}^{d-1} E_j$, let $c_{a_e, b_e; D_e}$ be a bivariate copula density, with copula $C_{a_e, b_e; D_e}$ and conditional distribution functions $C_{a_e | b_e; D_e}$ and $C_{b_e | a_e; D_e}$. For $e \in E_1$, the conditioning set D_e is empty and we simply write c_{a_e, b_e} and so on. Then there is a d -variate regular vine copula density c given by

$$c(\mathbf{u}) = \prod_{e \in E_1} c_{a_e, b_e}(u_{a_e}, u_{b_e}) \cdot \prod_{j=2}^{d-1} \prod_{e \in E_j} c_{a_e, b_e; D_e}(C_{a_e | D_e}(u_{a_e} | \mathbf{u}_{D_e}), C_{b_e | D_e}(u_{b_e} | \mathbf{u}_{D_e})), \quad (5.8)$$

where, for a random vector U with density c , the conditional distribution function of (U_{a_e}, U_{b_e}) given $U_{D_e} = \mathbf{u}_{D_e}$ is

$$C_{a_e, b_e | D_e}(u_{a_e}, u_{b_e} | \mathbf{u}_{D_e}) = C_{a_e, b_e; D_e}(C_{a_e | D_e}(u_{a_e} | \mathbf{u}_{D_e}), C_{b_e | D_e}(u_{b_e} | \mathbf{u}_{D_e})). \quad (5.9)$$

Assumption 5.7 Let c be a d -variate regular vine copula density as in (5.8), with $d \geq 3$.

- (i) For every edge $e \in E_1$, there exists a bivariate tail copula density r_{a_e, b_e} such that $\lim_{t \searrow 0} t \cdot c_{a_e, b_e}(tx, ty) = r_{a_e, b_e}(x, y)$, for $(x, y) \in (0, \infty)^2$.
- (ii) For every edge $e \in \bigcup_{j=2}^{d-1} E_j$, the pair copula density $c_{a_e, b_e; D_e}$ is continuous.

Proposition 5.8 Under Assumption 5.7, we have $\lim_{t \searrow 0} t^{d-1} \cdot c(tx) = r(x)$ for $x \in (0, \infty)^d$, where the X-vine tail copula density r is generated by the triple $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ as in Theorem 5.6 with

- $\mathcal{V}$ the same regular vine sequence as in (5.8),
- $\mathcal{R} = \{r_{a_e, b_e} : e \in E_1\}$ for r_{a_e, b_e} in Assumption 5.7(i), and
- $\mathcal{C}$ the same bivariate copula densities for edges $e \in E_j$ with $j \geq 2$ as in (5.8).

In particular, the copula C of c has tail copula R with tail copula density r .

Proposition 5.8 is foreshadowed by Theorem 3.4 and Example 3.5 in Li and Wu (2013), who consider D-vine copulas and who require stronger convergence properties.

5.3 Truncated X-vines

If the bivariate copula density $c_{a_e, b_e; D_e}$ associated with an edge $e \in E_2 \cup \dots \cup E_{d-1}$ is equal to the independence one, $c_{a_e, b_e; D_e} \equiv 1$, the corresponding factor drops out in (5.1) and the recursive formulas in (5.4) and (5.7) simplify to $R_{a_e | D_e \cup b_e}(x_{a_e} | \mathbf{x}_{D_e \cup b_e}) = R_{a_e | D_e}(x_{a_e} | \mathbf{x}_{D_e})$ and $R_{b_e | D_e \cup a_e}(x_{b_e} | \mathbf{x}_{D_e \cup a_e}) = R_{b_e | D_e}(x_{b_e} | \mathbf{x}_{D_e})$. Sparse X-vine specifications arise when $c_{a_e, b_e; D_e} \equiv 1$ for many edges $e \in E_2 \cup \dots \cup E_{d-1}$. Model selection of sparse vine copulas in high dimensions has been investigated in Müller and Czado (2019) and Nagler et al. (2019). The case where all pair copulas are equal to the independence one for all edges starting from a given tree is of practical importance (Brechmann et al., 2012; Brechmann & Joe, 2015).

Definition 5.9 A *truncated regular vine tree sequence* $\mathcal{V} = (T_1, \dots, T_q)$ on $d \geq 3$ elements with truncation level $q \in \{1, \dots, d-2\}$ is an ordered set of trees for which there exists a regular vine tree sequence $\mathcal{V}' = (T'_1, \dots, T'_{d-1})$ on d elements such that $T_j = T'_j$ for all $j \in \{1, \dots, q\}$.

Definition 5.10 The triplet $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ is a *truncated X-vine specification* on d elements ($d \geq 3$) with truncation level $q \in \{1, \dots, d-2\}$ if:

1. $\mathcal{V} = (T_j)_{j=1}^q$ is a truncated regular vine tree sequence on $[d]$;
2. $\mathcal{R} = \{r_{a,b} : e = \{a, b\} \in E_1\}$ is a family of bivariate tail copula densities;
3. $\mathcal{C} = \{c_{a_e, b_e; D_e} : e = (a_e, b_e; D_e) \in \bigcup_{j=2}^q E_j\}$ is a family of bivariate copula densities; in case $q = 1$, we have $\mathcal{C} = \emptyset$.

By definition, any truncated X-vine specification can be completed to a full X-vine specification by completing the truncated regular vine tree sequence $\mathcal{V}$ to a full one $\mathcal{V}'$ as in Definition 5.9 and by setting $c_{a_e, b_e; D_e} \equiv 1$ for all $e \in E'_{q+1} \cup \dots \cup E'_{d-1}$. The resulting X-vine copula density r does not depend on the way in which the truncated regular vine tree sequence $\mathcal{V}$ is completed, since the factorization in (5.1) simplifies anyway to

$$r(\mathbf{x}) = \prod_{e \in E_1} r_{a_e, b_e}(x_{a_e}, x_{b_e}) \cdot \prod_{j=2}^q \prod_{e \in E_j} c_{a_e, b_e; D_e}(R_{a_e | D_e}(x_{a_e} | \mathbf{x}_{D_e}), R_{b_e | D_e}(x_{b_e} | \mathbf{x}_{D_e})), \quad (5.10)$$

The right-hand side of (5.10) involves $\sum_{j=2}^q (d-j)$ bivariate copula densities. The truncation level q allows to tune the trade-off between sparsity and flexibility. If $q = 1$, the second product is empty and the model is a Markov tree as in Engelke and Hitz (2020), Segers (2020) and Engelke and Volgushev (2022). Increasing q and adding trees ('layers') yields more complex dependence models. Truncated X-vine specifications will be shown at work in Section 9.

6 Sampling from X-vine Pareto distributions

6.1 Inverted multivariate Pareto distributions

In Section 2, we introduced the multivariate Pareto distribution (2.11) as a limit model for high threshold excesses in (2.10). In the context of tail copula measures, it is more convenient to work with their reciprocals.

Definition 6.1 The d -variate random vector $\mathbf{Z}$ has an *inverted multivariate Pareto distribution* if there exists a d -variate tail copula measure R such that $P(\mathbf{Z} \in B) = R(B \cap \mathbb{L})/R(\mathbb{L})$, for Borel sets B , with $\mathbb{L} = \{\mathbf{x} \in (0, \infty]^d : \min \mathbf{x} < 1\}$; equivalently, if there exists a multivariate Pareto random vector $\mathbf{Y}$ such that $\mathbf{Z} = 1/\mathbf{Y}$.

If R is concentrated on $(0, \infty)^d$ and has density r , then $\mathbf{Z}$ has probability density $\mathbf{x} \mapsto r(\mathbf{x})1(\mathbf{x} \in \mathbb{L})/R(\mathbb{L})$. In general, for $j \in [d]$, the conditional distribution of Z_j given $Z_j < 1$ is uniform on $(0, 1)$; this follows from the marginal constraint (2.2). For Borel sets B , we have

$$P(\mathbf{Z} \in B \mid Z_j < 1) = R(B \cap \mathbb{L}^{(j)}), \quad (6.1)$$

where the set $\mathbb{L}^{(j)} = \{\mathbf{x} \in (0, \infty)^d : x_j < 1\}$ has R -measure one. Hence, on $\mathbb{L}^{(j)}$, the probability density $\mathbf{x} \mapsto r(\mathbf{x})1(x_j < 1)$ of $(\mathbf{Z} \mid Z_j < 1)$ coincides with the tail copula density r .

For the random vector U in Definition 2.1, the inverted multivariate Pareto vector Z is the weak limit in

$$(tU \mid \min U < 1/t)Z, \quad t \rightarrow \infty. \quad (6.2)$$

For nonempty $J \subseteq [d]$, we have $(tU_J \mid \min U_J < 1/t) \rightsquigarrow Z_{|J|} \stackrel{d}{=} (Z_J \mid \min Z_J < 1)$, as $t \rightarrow \infty$. The conditional marginal $Z_{|J|}$ has a $|J|$ -variate inverted multivariate Pareto distribution associated with the marginal tail copula measure R_J in (2.6). In contrast, Z_J does *not* necessarily have an inverted multivariate Pareto distribution, as $\min Z_J < 1$ is not guaranteed.

6.2 Sampling from (inverted) multivariate Pareto distributions

Through equations (2.10) and (6.2), the (inverted) multivariate Pareto distribution serves as a model for a random vector conditionally on the event that at least one variable takes a value far in the tail of its respective marginal distribution. The L-shaped support of the (inverted) multivariate Pareto distribution makes direct random sampling from it a little awkward. Lemma 2 in Engelke and Hitz (2020) provides an ingenious algorithm that reduces the task of sampling a multivariate Pareto random vector Y to sampling the conditional distributions $(Y \mid Y_j > 1)$ for every $j \in [d]$. Below, we study the equivalent problem of simulating from the conditional distribution of $(Z \mid Z_j < 1)$ for X-vine tail copula densities.

We will do so by inverting the Rosenblatt transformation (Rosenblatt, 1952), applying conditional quantile functions with an increasing number of conditioning variables successively to independent uniform random variables. A judicious choice of the ordering of the variables permits to compute the required conditional quantile functions recursively in terms of the bivariate ingredients of the X-vine specification. This *sampling order* (Cooke et al., 2015) is encoded by the permutation constructed in the next lemma.

Lemma 6.2 Let $\mathcal{V}$ be a regular vine sequence on d elements. For all $j \in [d]$, there exists a permutation σ_j of $\{1, \dots, d\}$ such that (i) $\sigma_j(1) = j$ and (ii) there exist edges $e_{j,1} \in E_1, \dots, e_{j,d-1} \in E_{d-1}$ such that $\sigma_j(k) \in C_{e_{j,k-1}}$ and $\{\sigma_j(i) : i = 1, \dots, k\} = A_{e_{j,k-1}}$ for all $k \in \{2, \dots, d\}$.

Let r be an X-vine tail copula density as in Definition 5.3. For each edge $e \in E_2 \cup \dots \cup E_{d-1}$ and for fixed $0 < u_{b_e} < 1$, let $u_{a_e} \mapsto C_{a_e|b_e;D_e}^{-1}(u_{a_e} \mid u_{b_e})$ be the inverse of the distribution function $u_{a_e} \mapsto C_{a_e|b_e;D_e}(u_{a_e} \mid u_{b_e})$, with $C_{a_e|b_e;D_e}$ defined in (5.3). Similarly for $C_{b_e|a_e;D_e}^{-1}$. Recall from Proposition 3.1 that for $i \in [d]$, nonempty $J \subseteq [d] \setminus \{i\}$, and $\mathbf{x}_J \in (0, \infty)^d$ such that $r_J(\mathbf{x}_J) > 0$, the quantile function $u_i \mapsto R_{i|J}^{-1}(u_i \mid \mathbf{x}_J)$ is the inverse of the distribution function $x_i \mapsto R_{i|J}(x_i \mid \mathbf{x}_J)$. Inverting (5.4) yields the recursive relations

$$\begin{aligned} R_{a_e|D_e \cup b_e}^{-1}(u_{a_e} \mid \mathbf{x}_{D_e \cup b_e}) &= R_{a_e|D_e}^{-1}\left(C_{a_e|b_e;D_e}^{-1}(u_{a_e} \mid R_{b_e|D_e}(\mathbf{x}_{b_e} \mid \mathbf{x}_{D_e})) \mid \mathbf{x}_{D_e}\right), \\ R_{b_e|D_e \cup a_e}^{-1}(u_{b_e} \mid \mathbf{x}_{D_e \cup a_e}) &= R_{b_e|D_e}^{-1}\left(C_{b_e|a_e;D_e}^{-1}(u_{b_e} \mid R_{a_e|D_e}(\mathbf{x}_{a_e} \mid \mathbf{x}_{D_e})) \mid \mathbf{x}_{D_e}\right). \end{aligned} \quad (6.3)$$

Let $W_1, \dots, W_d$ be independent random variables, all uniformly distributed on $(0, 1)$. For $j \in [d]$, let σ_j be a permutation of $[d]$ satisfying the two requirements in Lemma 6.2. Define a random vector $\mathbf{Z}^{(j)} = (Z_1^{(j)}, \dots, Z_d^{(j)})$ recursively as follows:

$$Z_j^{(j)} = W_j \quad \text{and} \quad Z_{\sigma_j(k)}^{(j)} = R_{\sigma_j(k)|\sigma_j(\{1, \dots, k-1\})}^{-1}\left(W_{\sigma_j(k)} \mid \mathbf{Z}_{\sigma_j(\{1, \dots, k-1\})}^{(j)}\right), \quad k \in \{2, \dots, d\}. \quad (6.4)$$

Proposition 6.3 Let Z be an inverted multivariate Pareto random vector associated with the X-vine tail copula density r . For $j \in [d]$, the distribution of $\mathbf{Z}^{(j)}$ in (6.4) is equal to the one of Z conditionally on $Z_j < 1$. For every $k \in \{2, \dots, d\}$, the conditional quantile function $R_{\sigma_j(k)|\sigma_j(\{1, \dots, k-1\})}^{-1}$ is of the form $R_{a_e|D_e \cup b_e}^{-1}$

or $R_{b_e|D_e \cup a_e}^{-1}$ for some edge $e = e_{j,k-1} \in E_{k-1}$ and can thus be computed recursively via (6.3).

The simulation algorithm based on Proposition 6.3 relies on structure matrices encoding regular vine sequences as explained in [Appendix A in the online supplementary material](#). In Section 8, we apply the algorithm to assess the estimation methods from Section 7 through Monte Carlo experiments.

7 Estimation and model selection

Let $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,d})$ for $i \in \{1, \dots, n\}$ be an independent random sample from a distribution function F with continuous but unspecified margins $F_1, \dots, F_d$ and whose survival copula C has lower tail copula R (Definition 2.1). Suppose that the tail copula density r is an X-vine with specification $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ (Definitions 5.3 and 5.5 and Theorem 5.6). We propose a procedure to estimate r from the excesses over a high multivariate threshold.

The regular vine sequence $\mathcal{V} = (T_j)_{j=1}^{d-1}$ with trees $T_j = (N_j, E_j)$ may be known or not. The bivariate tail copulas $r_{a,b}$ for edges $e = \{a, b\} \in E_1$ in the first tree and the bivariate copula densities $c_{a_e, b_e; D_e}$ for edges $e = (a_e, b_e; D_e) \in E_j$ in trees $j \geq 2$ are assumed to belong to prespecified (lists of) parametric families. The X-vine specification may be truncated (Definition 5.10), leading to a simpler model.

The basis of the method is a link between the conditional copula densities $c_{i,j}$ in Sklar's theorem (Proposition 3.1) on the one hand and the inverted multivariate Pareto distribution (Definition 6.1) on the other hand (Section 7.1). In Section 7.2, we propose parameter estimates, supposing that $\mathcal{V}$ is given and that parametric families of (tail) copula densities have been specified for all edges. In Section 7.3, finally, we treat model selection, which comprises the selection of the parametric families of the bivariate model components, the selection of the regular vine sequence $\mathcal{V}$, and the selection of the truncation level q .

7.1 Copulas and inverted multivariate Pareto distributions

Let r be a d -variate tail copula density, not necessarily an X-vine. In Sklar's theorem (Proposition 3.1), suppose that the copula density $c_{i,j}$ satisfies the simplifying assumption (Definition 3.4). The following proposition shows how to transform an inverted multivariate Pareto random vector $\mathbf{Z}$ associated with r into a random vector with density $c_{i,j}$.

Proposition 7.1 Let r be a d -variate tail copula density and let $\mathbf{Z}$ be an inverted multivariate Pareto random vector associated with r . Let $I, J \subset [d]$ be nonempty and disjoint.

- (i) For $\mathbf{z}_J \in (0, \infty)^J$ such that $\min \mathbf{z}_J < 1$ and $r_J(\mathbf{z}_J) > 0$, the conditional density of $\mathbf{Z}_I$ given $\mathbf{Z}_J = \mathbf{z}_J$ is $r_{I|J}(\cdot | \mathbf{z}_J)$.
- (ii) Suppose $|I| \geq 2$. If $c_{I,J}(\cdot; \cdot)$ satisfies the simplifying assumption (Definition 3.4), then, conditionally on the event $\min \mathbf{Z}_J < 1$, the random vector $(R_{i|J}(Z_i | \mathbf{Z}_J))_{i \in I}$ is independent of $\mathbf{Z}_J$ and its density is $c_{I,J}$.

By statement (ii), the density of $(R_{i|J}(Z_i | \mathbf{Z}_J))_{i \in I}$ given $\mathbf{Z}_J \in A$ is equal to $c_{I,J}$ for any nonempty set $A \subseteq \{\mathbf{z} \in (0, \infty)^J : \min \mathbf{z} < 1\}$. For estimation, we will use this property for $A = (0, 1)^J$, requiring in effect that *all* variables (rather than *at least one*) in J exceed a high threshold. We do so in order to avoid a potential bias stemming from including too many nonextreme values in the procedure. An alternative would be to opt for a censored-likelihood approach (Ledford & Tawn, 1996).

7.2 Sequential maximum-likelihood estimation of X-vines

Let the d -variate tail copula density r be an X-vine specified by $(\mathcal{V}, \mathcal{R}, \mathcal{C})$ as in Theorem 5.6. Assume that the bivariate tail copula densities ($\mathcal{R}$) and bivariate copula densities ($\mathcal{C}$) belong to specified parametric families. Let the parameter vector be denoted by $\boldsymbol{\theta} = (\boldsymbol{\theta}_R, \boldsymbol{\theta}_C)$: here, $\boldsymbol{\theta}_R = (\theta_e)_{e \in E_1}$ contains the parameters (or parameter vectors) $\theta_e \in \Theta_e$ associated with each pairwise tail copula density r_{a_e, b_e} for $e \in E_1$, while $\boldsymbol{\theta}_C = (\theta_{C,j})_{j=2}^{d-1}$ for $\theta_{C,j} = (\theta_e)_{e \in E_j}$ denotes the parameters (or parameter

vectors) $\theta_e \in \Theta_e$ associated with each bivariate copula density $c_{a_e, b_e; D_e}$ for $e \in \bigcup_{j=2}^{d-1} E_j$. While it is possible to derive the full likelihood of an X-vine (inverted) multivariate Pareto distribution, performing parameter estimation with the full model in high dimensions is challenging. Instead, using the X-vine decomposition into *bivariate* components and recursively defined quantities (Theorem 5.6), we outline a *sequential* procedure for parameter estimation, tree by tree. This approach is inspired by the one for regular vine copulas (see, e.g. Czado, 2019), but with suitable adaptations to the extreme value context.

(1) **Standardizing the margins and selecting sub-samples.** Recall that $X_1, \dots, X_n$ is an independent random sample from F , with survival copula C , tail copula R and tail copula density r . For $i = 1, \dots, n$ and $j \in [d]$, let $\widehat{U}_{i,j} = 1 - \widehat{F}_j(X_{i,j})$, where $\widehat{F}_j$ denotes any estimator of the marginal distribution function $F_j(x) = P(X_{i,j} \leq x)$. One possibility is the empirical distribution function, and to avoid boundary effects, we set

$$\widehat{U}_{i,j} = 1 - (\text{rk}_{i,j} - 0.5)/n, \quad (7.1)$$

where $\text{rk}_{i,j} = \sum_{s=1}^n 1(X_{s,j} \leq X_{i,j})$ is the (maximal) rank of $X_{i,j}$ among $X_{1,j}, \dots, X_{n,j}$. We view the points $\widehat{U}_i = (\widehat{U}_{i,1}, \dots, \widehat{U}_{i,d})$ as pseudo-observations from the survival copula C .

By equation (6.2), for large $t > 0$, the rescaled points $t\widehat{U}_i$ for $i \in \{1, \dots, n\}$ such that $\min \widehat{U}_i < 1/t$ constitute pseudo-observations from a distribution that approximates the inverted multivariate Pareto distribution associated with r . We set $t = n/k$ where, in an asymptotic setting, $k = k_n \in \{1, \dots, n\}$ satisfies $k \rightarrow \infty$ and $k/n \rightarrow 0$. For $j \in [d]$ and $\widehat{U}_{i,j}$ as in (7.1), let

$$K_j = \{i = 1, \dots, n : \widehat{U}_{i,j} < k/n\} \quad (7.2)$$

be the set of indices i corresponding to the k largest observations for the j th component. Further, put $K_j = \bigcap_{i \in J} K_j$ and $K = \bigcup_{j \in [d]} K_j$, which are, respectively, the set of indices i corresponding to large observations in *all* variables in $J \subseteq [d]$ *simultaneously* and the set of indices with large observations in *at least one* variable. Write $\widehat{Z}_i = (\widehat{Z}_{i,j})_{j \in [d]}$ where $\widehat{Z}_{i,j} = (n/k)\widehat{U}_{i,j}$. In view of equation (6.2), we treat $\{\widehat{Z}_i\}_{i \in K}$ as a sample of $|K|$ pseudo-observations of the inverted multivariate Pareto distribution associated with r . The sample size $|K|$ is random, and from (7.1), we have $k \leq |K| \leq dk$. For nonempty $J \subseteq [d]$ and $i \in K$, we write $\widehat{Z}_{i,J} = (\widehat{Z}_{i,j})_{j \in J}$.

(2) **Estimating the tail copula parameters θ_R in T_1 .** For each edge $e = \{a_e, b_e\} \in E_1$, we estimate the parameter (vector) θ_e associated with the bivariate tail copula density $r_{a_e, b_e}(\cdot; \theta_e)$. By equation (6.1), for $j \in \{a_e, b_e\}$, the conditional density of (Z_{a_e}, Z_{b_e}) given $Z_j < 1$ is $r_{a_e, b_e}(z_{a_e}, z_{b_e}; \theta_e)1(z_j < 1)$. We use maximum pseudo-likelihood estimation to fit this density to the sub-samples $(\widehat{Z}_{i,a_e}, \widehat{Z}_{i,b_e})$ for both $i \in K_{a_e}$ and $i \in K_{b_e}$, with K_j as in equation (7.2). More precisely, we maximize each of the two pseudo-likelihoods

$$\mathcal{L}_{R,e}(\theta_e^{(j)}; (\widehat{Z}_{i,a_e}, \widehat{Z}_{i,b_e}), i \in K_j) = \prod_{i \in K_j} r_{a_e, b_e}(\widehat{Z}_{i,a_e}, \widehat{Z}_{i,b_e}; \theta_e^{(j)}), \quad j \in \{a_e, b_e\}, \quad (7.3)$$

over $\theta_e^{(j)} \in \Theta_e$, yielding estimates $\hat{\theta}_e^{(a_e)}$ and $\hat{\theta}_e^{(b_e)}$, respectively. The final estimate is $\hat{\theta}_{a_e, b_e} = \{\hat{\theta}_e^{(a_e)} + \hat{\theta}_e^{(b_e)}\}/2$. The idea of averaging maximum pseudo-likelihood estimators on product spaces has already been proposed for the Hüsler–Reiss model (Engelke et al., 2022, 2015). Other estimation approaches include censored likelihoods (de Haan et al., 2008; Ledford & Tawn, 1996) or empirical stable tail dependence functions (Einmahl et al., 2018, 2008).

(3) **Estimating the copula parameters θ_C in $T_2, \dots, T_{d-1}$.** Estimation of the parameters associated with edges $e = (a_e, b_e; D_e)$ in E_j for $j \in \{2, \dots, d-1\}$ is based on a similar procedure as the one employed in regular vine copulas (Czado, 2019, Section 7.2). The main difference concerns the definition of the pseudo-observations: if $\hat{\theta}(E_{1:(j-1)})$ denotes the parameter estimates

associated with the edges in $E_{1:(j-1)} = E_1 \cup \dots \cup E_{j-1}$, by Proposition 7.1(i), for $i \in K_{D_e}$, pseudo-observations $(\widehat{U}_{i,a_e;D_e}, \widehat{U}_{i,b_e;D_e})$ from $c_{a_e,b_e;D_e}$ can be defined by

$$\widehat{U}_{i,a_e;D_e} = R_{a_e|D_e}(\widehat{Z}_{i,a_e} | \widehat{Z}_{i,D_e}; \hat{\theta}(E_{1:j-1})), \quad \widehat{U}_{i,b_e;D_e} = R_{b_e|D_e}(\widehat{Z}_{i,b_e} | \widehat{Z}_{i,D_e}; \hat{\theta}(E_{1:j-1})). \quad (7.4)$$

The full procedure is detailed in [Section D of the online supplementary material](#).

7.3 Model selection for X-vines

Model selection for X-vines based on a random sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ from F involves:

1. selecting a regular vine sequence $\mathcal{V} = (T_1, \dots, T_{d-1})$;
2. given the regular vine sequence obtained in (1), choosing adequate bivariate parametric (tail) copula families, $\mathcal{R}$ and $\mathcal{C}$.

In fact, the two procedures are executed together sequentially, progressing from one tree to the next tree. First, we describe step (2) given the regular vine sequence $\mathcal{V}$.

Selecting parametric (tail) copula families given a truncated regular vine sequence. We first consider the specification of $\mathcal{R}$ in Definition 5.5 in tree T_1 . Let $\mathcal{B}_{\mathcal{R},1:T} = (\mathcal{B}_{\mathcal{R},1}, \dots, \mathcal{B}_{\mathcal{R},T})$ be the list of candidate bivariate tail copula families. For each edge $e = (e_a, e_b) \in E_1$, specifying $\mathcal{R}$ involves choosing the bivariate tail copula family among $\mathcal{B}_{\mathcal{R},1:T}$. Similar to the idea of averaged maximum pseudo-likelihood estimates in Section 7.2, we use the averaged-AIC value for selecting bivariate tail copula families. For each edge e and each $\mathcal{B}_{\mathcal{R},t}$, $t = 1, \dots, T$, we obtain two maximum pseudo-likelihood estimates $\hat{\theta}_e^{(t,a_e)}$ and $\hat{\theta}_e^{(t,b_e)}$, derived through the maximization of log-likelihood functions on product spaces; see (7.3). The averaged-AIC value is $\text{AIC}_e(\mathcal{B}_{\mathcal{R},t}) = 2v^{(t)} - \frac{1}{2} \{ \log \mathcal{L}_{\mathcal{R},e}(\hat{\theta}_e^{(t,a_e)}) + \log \mathcal{L}_{\mathcal{R},e}(\hat{\theta}_e^{(t,b_e)}) \}$, where $v^{(t)}$ is the number of parameters in $\mathcal{B}_{\mathcal{R},t}$. We select the bivariate tail copula family with the lowest averaged-AIC among $\text{AIC}_e(\mathcal{B}_{\mathcal{R},1}), \dots, \text{AIC}_e(\mathcal{B}_{\mathcal{R},T})$.

Similarly, to specify $\mathcal{C}$ and select the bivariate parametric copula family among a list of candidates for each edge in $E_2 \cup \dots \cup E_{d-1}$, we follow common practice in vine copula modelling and choose the family with the lowest AIC; see [Brechmann \(2010\)](#) and [Czado \(2019, Section 8.1\)](#).

It is worthwhile to note that when implementing the model selection step for an edge e in E_j , only the trees $T_1, \dots, T_j$ need to have been selected, but not the trees $T_{j+1}, \dots, T_{d-1}$. This sequential approach aligns well with the vine learning procedure in the next paragraph.

Selecting the regular vine sequence. [Morales-Napoles \(2010\)](#) showed that the number of regular

vine sequences on d elements is equal to $d! 2^{\binom{d-2}{2}-1}$, making it impossible to go through all possible vine sequences. We adopt a model selection approach similar to the one in [Dissmann et al. \(2013\)](#), choosing trees sequentially from T_1 to T_{d-1} .

To select the first tree, T_1 , on the node set $N_1 = [d]$, we follow a procedure as in [Engelke and Volgushev \(2022\)](#) and [Hu et al. \(2024\)](#). For every pair $\{a, b\}$ of distinct elements in $[d]$, let $w_{a,b}$ be a nonnegative weight derived from the data. We use an empirical version of the tail dependence coefficient $\chi_{a,b} = R_{a,b}(1, 1)$, setting $w_{a,b}$ equal to $\widehat{\chi}_{a,b} = \frac{1}{k} \sum_{i=1}^n 1\{\widehat{U}_{i,a} \leq k/n, \widehat{U}_{i,b} \leq k/n\}$, where $\widehat{U}_{i,a}$ and $\widehat{U}_{i,b}$ are defined as in (7.1). Another possible edge weight could be the empirical extremal variogram as in [Engelke and Volgushev \(2022\)](#). For subsequent trees $T_2, \dots, T_{d-1}$, the edge weight is chosen to be the absolute value of the empirical Kendall's tau. All trees are selected as maximum spanning trees; see, for example, [Czado \(2019, Section 8.3\)](#).

Selecting truncated regular vine tree sequences. The model selection procedure described above may set all pair-copulas to the independence copula in the subsequent trees from T_{q+1} to T_{d-1} for a truncation level $q \in \{1, \dots, d-2\}$. In this case, the resulting model corresponds to the truncated X-vine model in Definition 5.10. Besides an information criterion, we consider two additional

criteria for selecting the independence copula at an edge e : when the effective sample size $|K_{D_e}| = n_{D_e}$ falls below a certain low value, or when the absolute value of the empirical Kendall's tau is close to zero. Sparsity induced by a smaller effective sample size, is more likely when the total sample size n is relatively small with respect to the dimension d .

Even when the above criteria are not met, when d is large, it is natural to limit the number of model parameters by considering truncated X-vines, since Dissmann's algorithm captures as much dependence as possible in the first few trees. We use a modified Bayesian information (mBIC) to determine the truncation level, inspired by the one for regular vine copulas in Nagler et al. (2019). This modified version adjusts the prior probability in the BIC to penalize dependence copulas more severely in trees at higher levels. More specifically, assuming that for any edge $e \in E_j$ for $j \geq 2$, the parametric family has a single parameter θ_e and that a value of $\theta_e = 0$ corresponds to the independence copula, the mBIC includes independent Bernoulli variables $1(\theta_e \neq 0)$ with mean $\psi_e = \psi_0^{j-1}$ for $e \in E_j$ and a hyperparameter $\psi_0 \in (0, 1)$: set $\text{mBIC}(1) = 0$ and, for $q \in \{2, \dots, d-1\}$, put

$$\text{mBIC}(q) = \sum_{j=2}^q \sum_{e \in E_j} \left[1(\theta_e \neq 0) \left\{ \log n_{D_e} - 2 \log \left(\frac{\psi_e}{1 - \psi_e} \right) \right\} - 2 \log \left\{ \frac{\mathcal{L}_{C,e}(\hat{\theta}_e)}{(1 - \psi_e)^{-1}} \right\} \right] \quad (7.5)$$

with $\mathcal{L}_{C,e}$ as in (D.1). The selected truncation level q^* is the value of q in $\{1, \dots, d-1\}$ that minimizes $\text{mBIC}(q)$. In practice, we will set $\psi_0 = 0.9$ as in Nagler et al. (2019).

8 Simulation study

We conduct three simulation studies, evaluating the proposed procedures for parameter estimation, selection of bivariate parametric (tail) copula families, and vine truncation.

8.1 Parameter estimation

We first consider the 5-dimensional X-vine model r in Figure 2a. The bivariate tail copula densities r_{a_e, b_e} in the first tree $T_1 = (N_1, E_1)$ are chosen from the Hüsler–Reiss model, the negative logistic, logistic, and Dirichlet families (Section 4), while the bivariate copula densities $c_{a_e, b_e; D_e}$ in the subsequent trees $T_j = (N_j, E_j)$ for $j \in \{2, 3, 4\}$ are taken from the Clayton, Gumbel, and Gaussian copula families. Figure 2a shows the values specified for the tail dependence coefficient $\chi = R(1, 1)$ or for Kendall's τ for each edge. The formulas connecting the parameters of the families of tail copula densities with χ are given in Appendix E.1 in the online supplementary material.

Implementing simulation algorithms as detailed in Section 6 and Appendix A in the online supplementary material, we generate multivariate inverted Pareto random samples $Z_1, \dots, Z_n$ associated with r . As in equation (7.1), we transform to $\hat{U}_1, \dots, \hat{U}_n$, with $\hat{U}_{i,j}$ based on the rank of $Z_{i,j}$ among $Z_{1,j}, \dots, Z_{n,j}$. We then take the sub-sample $\{\hat{Z}_i\}_{i \in K}$ with $K = \bigcup_{j \in [d]} K_j$ and K_j as in (7.2).

We set $(n, k) = (4\,000, 200)$ and perform sequential parameter estimation as in Section 7.2 over 200 repetitions. We obtain maximum pseudo-likelihood estimates and corresponding dependence measures, $\hat{\chi}$ and $\hat{\tau}$, using the relations in Appendix E.1 of the online supplementary material, for each of the 10 edges in the vine. In Figure 2b, box-plots present dependence measure estimates of parametric families based on the X-vine specification in Figure 2a. The four left-most box-plots (light sky blue) show the tail dependence coefficient estimates $\hat{\chi}$ for the four edges $e \in E_1$, and the remaining plots display Kendall's tau estimates $\hat{\tau}$ for the six edges $e \in E_2 \cup E_3 \cup E_4$. The plot supports the validity of the sequential method for estimating tail dependence measures. We see that estimation uncertainty becomes larger at higher tree levels. Sections F.1.1–F.1.3 in the online supplementary material include additional box-plots: first, of dependence measures as in Figure 2b, but supposing that the bivariate parametric families are known, second, of dependence measures with varying sample sizes and threshold exceedances, and third, of maximum-likelihood estimates of tail copula densities for a specific edge $e \in E_1$ with varying (n, k) .

8.2 Selecting parametric (tail) copula families

We assess algorithm effectiveness in selecting bivariate parametric families for each edge in each tree, using the X-vine specification in Figure 2a. While the flexibility of X-vine models allows us

T_1	99%	40%	60%	49%
T_2		58%	93%	82%
T_3			42%	68%
T_4				11%

to consider any bivariate parametric family, we simplify the process by considering a catalogue of four candidate tail copula models for T_1 : the Hüsler–Reiss, logistic, negative logistic, and Dirichlet models, along with a catalogue of nine candidate pair-copula families for T_2, T_3, T_4 : Independence, Gaussian, Clayton, Survival Clayton, Gumbel, Survival Gumbel, Frank, Joe, and Survival Joe copulas, as implemented in the R package *VineCopula* (Nagler et al., 2023).

Using again 200 repetitions with $(n, k) = (4\,000, 200)$, we assess the accuracy of the (averaged) AIC in Section 7.3 in selecting bivariate parametric families. The overall proportion of correctly selected families across all trees is 60%. For individual trees, the proportions are 62% for T_1 , 78% for T_2 , 55% for T_3 , and 10% for T_4 . The specific proportions for each edge in each tree are

The lower proportions observed for e_{23} and e_{45} in E_1 result from the challenge in distinguishing between the logistic and negative logistic models. Overall, the proportion of accurately selected families decreases with declining effective sample size and Kendall's tau estimate across tree levels. We observe the lowest proportion for the deepest edge $e_{15;234}$ where the corresponding pair-copula exhibits weak dependence. For this edge, the algorithm selects the independence copula 110 times out of 200.

In Section F.1.4 of the online supplementary material, we investigate the effective sample sizes for each tree.

8.3 Tree selection and truncation

We consider a higher dimensional X-vine model where residual dependence weakens with increasing tree level. Specifically, for $\mathcal{V}$ we consider a 50-dimensional C-vine, that is, in each tree T_j , there is a single node $a_j \in [d]$ such that a_j belongs to the conditioned set $\mathcal{C}_e$ of all edges $e \in E_j$. The first tree includes the Hüsler–Reiss models and negative logistic models with randomly assigned parameter values $\theta_e \in [1, 2]$ for $e \in E_1$. Subsequent trees contain bivariate Gaussian copulas with partial correlations $\rho_e = 1.1 - 0.1j$ for $e \in E_j$ and $j \in \{2, \dots, 9\}$, and $\rho_e = 0.1$ for $e \in E_j$ with $j \geq 10$. This X-vine specification allows us to explore truncated X-vine models by setting pair-copulas with weak dependence to independence copulas.

We use a single inverted multivariate Pareto sample of size $n = 1\,000$ from the X-vine model directly, i.e. without rank transformation and thresholding. As in Sections 7.2 and 7.3, we sequentially select trees $T_1, \dots, T_{d-1}$ using $\hat{\chi}_e$ for $e \in E_1$ and $\hat{\tau}_e$ for $e \in E_j, j \geq 2$, as edge weights, with $\hat{\chi}_e$ as in equation (E.1) in the online supplementary material. For each selected tree, we choose the bivariate parametric families with the lowest (averaged) AIC for each edge and estimate the associated parameters. Additionally, independence copulas are chosen in subsequent trees if either $|\hat{\tau}_e| < 0.05$ or $n_{D_e} < 10$ for each edge.

To explore truncated X-vine models, we use the mBIC in (7.5) with $\psi_0 = 0.9$ as in Nagler et al. (2019). The estimated mBIC-optimal truncation level is $q^* = 19$ (dotted line in Figure 3a). We assess the goodness of fit by comparing pairwise χ -values from the true X-vine model with those from the fitted 50-dimensional X-vine model via Monte Carlo simulations (as explained in Section E.2 in the online supplementary material) in Figure 3b. For the truncated X-vine model, the χ -plot in Figure 3c resembles that of the full model but exhibits more variability, particularly for lower χ -values.

9 Application: US flight delay data

We apply X-vine models to investigate extremal dependence among large flight delays in the US flight network. The raw data set is accessible through the US Bureau of Transportation Statistics.² These flight delay data were analysed by Hentschel et al. (2024) who first selected

² <https://www.bts.dot.gov>.




Figure 3. Simulation study: (a) The modified Bayesian information (mBIC) plotted across tree levels with a dotted line indicating the selected mBIC-optimal truncation level of $q^* = 19$. (b) χ -plot comparing pairwise χ values from the true X-vine model to those obtained from the fitted 50-dimensional X-vine model via Monte Carlo simulations. (c) Similar to (b) but for the truncated fitted X-vine model.

airports with a minimum of 1,000 flights per year and applied a k -medoids clustering algorithm to identify homogeneous clusters in terms of extremal dependence. This clustering approach not only makes the analysis suitable for modelling extremal dependence but also reveals shared frequent flight connections between airports and similar geographical characteristics in each cluster. Focusing on the Hüsler–Reiss family, Hentschel et al. (2024) fitted an extremal graphical model to large flight delays of airports in the Texas cluster to investigate conditional independence.

For the purpose of model comparison, we also analyse daily total delays (in minutes) in the Texas cluster. The cluster comprises $d = 29$ airports and counts $n = 3,603$ days from 2005 to 2020, during which all airports have recorded measurements. These preprocessed data are available through the R-package `graphicalExtremes` (Engelke et al., 2022). A graphical representation of the actual flight connections between airports is presented in Figure 4a.

Let $\hat{x}_i$, for $i = 1, \dots, n$, denote the measurements. Following the standardization process with the rank transformation (7.1) as in Section 7.2, we take the sub-sample $\hat{z}_i = (n/k)\hat{u}_i$ for $i \in K = \bigcup_{j \in [d]} K_j$ as in (7.2). We choose $k = 0.13n$ in order to have a large enough effective sample size with respect to the number of variables.

Assigning $\hat{\chi}_e$ for $e \in E_1$ and $\hat{\tau}_e$ for $e \in E_j, j \geq 2$, as edge weights, we select the trees of the regular vine sequentially as in Section 7.3. Figure F.5a in the online supplementary material shows the first maximum spanning tree. The corresponding estimated tail dependence coefficients vary between 0.42 and 0.65, with a mean of 0.55. Out of the $d - 1 = 28$ edges in T_1 , selected bivariate tail copula families include the Hüsler–Reiss (8 edges), negative logistic (3 edges) and Dirichlet models (17 edges). Subsequent trees with a total of 378 edges are then chosen with the following pair-copula families (number of edges in parentheses): Independence (201), Gaussian (31), Clayton (8), Gumbel (26), Frank (34), Joe (36), Survival Clayton (31), Survival Gumbel (9), and Survival Joe (2). As the tree level rises, we observe a decline in residual dependence and an increase in the number of independence copulas. The X-vine algorithm sets all pair copulas to the independence copula from T_{21} , that is, the vine is truncated at level $q = 20$.

To evaluate the X-vine model's efficacy in capturing extremal dependence for large flight delays between airports, we compare bivariate and trivariate empirical tail dependence coefficients with those from the fitted X-vine models in Section F.2.2 in the online supplementary material. Fig. F.6 in the online supplementary material indicates a satisfactory fit, both for the full 29-dimensional model and for the truncated model.

To further explore truncated X-vine models with a lower truncation level, we use the mBIC in equation (7.5). Figure F.7a in the online supplementary material illustrates the mBIC-optimal truncation level of $q^* = 7$, which corresponds to 113 dependence copulas [Gaussian (17), Clayton (3), Gumbel (23), Frank (16), Joe (25), Survival Clayton (22), Survival Gumbel (6), and Survival Joe (1)] out of the 147 total edges from T_2 to T_7 . The total number of bivariate model components is thus $28 + 113 = 141$. Additionally, we investigate how the truncation level changes over the threshold range in Figure F.7b in the online supplementary material. It appears that the mBIC-optimal truncation levels are not overly sensitive to the choice of the threshold. The




Figure 4. (a) Flight graph illustrating connection flights between airports joined by edges in the Texas cluster. (b) The X-vine graph with the first seven trees superimposed, showing for trees T_2 to T_7 only the 113 out of 147 edges for which a copula other than the independence copula is selected, for 141 edges in total. (c) The estimated extremal Hüsler–Reiss graph structure using the regularized method with a tuning parameter value of $\rho^* = 0.1$, totalling 148 edges. Corresponding χ -plots comparing empirical pairwise tail dependence coefficients with those from the fitted graphical models: (d) the flight graph, (e) the truncated X-vine graph, and (f) the extremal Hüsler–Reiss graph using the EGlearn algorithm.

superimposed graph of the first seven vine trees is shown in Figure 4b, showing for T_2 to T_7 only the 113 out of 147 edges for which a copula other than the independence one is selected. The vine tree sequence from T_1 to T_7 is presented in Figure F.5 of the online supplementary material.

Returning to model comparison, we consider extremal Hüsler–Reiss graphical models (Engelke et al., 2021; Hentschel et al., 2024). Hentschel et al. (2024) obtained a sparse Hüsler–Reiss graphical model using the EGlearn algorithm (Engelke et al., 2021). The tuning parameter $\rho \geq 0$ controls sparsity. The empirical extremal variogram matrix $\hat{\Gamma}$ is used as edge weight for the *minimum* spanning tree with smaller elements in $\hat{\Gamma}$ indicating stronger extremal dependence. Following Engelke et al. (2022), the data set is split in half: a training set for model fitting and a test set for tuning parameter selection. We determine the optimal tuning parameter, $\rho^* = 0.1$, by evaluating the Hüsler–Reiss log-likelihood across tuning parameter values ρ from the test set. The resulting sparse extremal graph with 148 edges and $\rho^* = 0.1$ is shown in Figure 4c.

We assess the goodness of fit using the entire sub-sample and compare empirical tail dependence coefficients to those obtained from the fitted graphical models as explained in Appendix E.2 in the online supplementary material: the Hüsler–Reiss extremal graphical model for the flight graph in Figure 4d, the truncated X-vine model with $q^* = 7$ in Figure 4e, and the extremal Hüsler–Reiss graphical model with $\rho^* = 0.1$ in Figure 4f, respectively.

The flexibility of a vine-based dependence model allows the selection of various (tail) copula families with the lowest AIC. Recall from Section 4.2 that the Hüsler–Reiss model arising from the X-vine specification is constructed using bivariate Hüsler–Reiss tail copula densities in the first tree and bivariate Gaussian copulas in subsequent trees. In the truncated X-vine model, 8 out of 28 edges prefer Hüsler–Reiss models in T_1 , while in trees T_j for $j \geq 2$, a total of 96 out of 113 edges favour copula families other than the Gaussian one. However, uncertainty in sequential parameter

estimation tends to accumulate across tree levels. Consequently, the χ -plot of the truncated X-vine model shows a better fit but more variability. In contrast, the χ -plot of the extremal Hüsler–Reiss graphical model has less variability but seems biased towards higher tail dependence. [Figure F.8 in the online supplementary material](#) presents a model comparison similar to [Figure 4d–f](#), but focusing exclusively on the Hüsler–Reiss model and using empirical tail dependence coefficients derived from the empirical extremal variogram matrix.

10 Discussion

We have opened the door to the construction and the use of extremal dependence models based on regular vine tree sequences. For (ordinary) copulas, the methodology has been extensively developed in the literature since its inception more than two decades ago. Our contribution is to deliver the theoretical and methodological advances required to apply vine machinery in the extreme value context too. The key consists of a version of Sklar’s theorem applied to tail copula densities, together with a telescoping product formula for regular vine sequences.

While the proposed methodology is fully operational, it is clearly open to improvement, while many open questions remain, just as for ordinary regular vine copulas. We name just a few. The vine sequence learning method inspired by [Dissmann et al. \(2013\)](#) aims to capture dependence by the first several trees in the sequence, but there is no guarantee that it retrieves the true structure. The regular vine atlas of [Morales-Nápoles et al. \(2023\)](#) can serve as a test bed for evaluating vine learning approaches. While the parameter estimators were shown to perform well in simulations, their large-sample theory remains to be developed, in line with [Hobæk Haff \(2013\)](#). For the Hüsler–Reiss model, the precise connection between the variogram matrix and the correlation parameters of the bivariate Gaussian copulas remains to be elucidated and the effect of truncation on the Hüsler–Reiss precision matrix ([Hentschel et al., 2024](#)) to be uncovered. The relation between X-vines and graphical models for extremes as in [Engelke and Hitz \(2020\)](#) deserves further investigation, perhaps leveraging results of [Zhu and Kurowicka \(2022\)](#). Finally, as our approach is limited to dependence structures generated from max-stable distributions, a completely open question is whether it can be extended to other settings in multivariate extreme value analysis, such as the conditional extremal dependence model ([Heffernan & Tawn, 2004](#)) or the geometric sample-cloud approach ([Nolde, 2014](#); [Simpson et al., 2021](#)).

Acknowledgments

The authors thank Manuel Hentschel for providing the US flight delay data and gratefully acknowledge two reviewers and an associate editor for valuable comments.

Conflicts of interest: None declared.

Funding

This work is supported financially by the Fonds de la Recherche Scientifique—FNRS, Belgium (grant T.0203.21).

Author contributions

Anna Kiriliouk, Jeongjin Lee, and Johan Segers contributed equally to this work.

Data availability

The R package *Xvine*, available from <https://github.com/JeongjinLee88/Xvine>, permits to replicate all results, including tables and figures, in this article. The flight data used in Section 9 is also included in the package.

Supplementary material

[Supplementary material](#) is available online at *Journal of the Royal Statistical Society: Series B*.

References

- Bedford T., & Cooke R. M. (2001). Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence*, 32(1/4), 245–268. <https://doi.org/10.1023/A:1016725902970>
- Bedford T., & Cooke R. M. (2002). Vines—a new graphical model for dependent random variables. *The Annals of Statistics*, 30(4), 1031–1068. <https://doi.org/10.1214/aos/1031689016>
- Beirlant J., Goegebeur Y., Segers J., & Teugels J. L. (2004). *Statistics of extremes: Theory and applications* (Vol. 558). John Wiley & Sons.
- Belzile L. R., & Nešlehová J. G. (2017). Extremal attractors of Liouville copulas. *Journal of Multivariate Analysis*, 160, 68–92. <https://doi.org/10.1016/j.jmva.2017.05.008>
- Brechmann E. (2010). *Truncated and simplified regular vines and their applications* [Master's thesis]. Technische Universität München.
- Brechmann E. C., Czado C., & Aas K. (2012). Truncated regular vines in high dimensions with application to financial data. *Canadian Journal of Statistics*, 40(1), 68–85. <https://doi.org/10.1002/cjs.v40.1>
- Brechmann E. C., & Joe H. (2015). Truncation of vine copulas using fit indices. *Journal of Multivariate Analysis*, 138, 19–33. <https://doi.org/10.1016/j.jmva.2015.02.012>
- Coles S. G., & Tawn J. A. (1991). Modelling extreme multivariate events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(2), 377–392. <https://doi.org/10.1111/j.2517-6161.1991.tb01830.x>
- Cooke R. M., Kurowicka D., & Wilson K. (2015). Sampling, conditionalizing, counting, merging, searching regular vines. *Journal of Multivariate Analysis*, 138, 4–18. <https://doi.org/10.1016/j.jmva.2015.02.001>
- Czado C. (2019). Analyzing dependent data with vine copulas. In *A practical guide with R: Lecture Notes in Statistics* (Vol. 222). Springer.
- Czado C., & Nagler T. (2022). Vine copula based modeling. *Annual Review of Statistics and Its Application*, 9(1), 453–477. <https://doi.org/10.1146/statistics.2022.9.issue-1>
- de Haan L., & Ferreira A. (2007). *Extreme value theory: An introduction*. Springer Science & Business Media.
- de Haan L., Neves C., & Peng L. (2008). Parametric tail copula estimation and model testing. *Journal of Multivariate Analysis*, 99(6), 1260–1275. <https://doi.org/10.1016/j.jmva.2007.08.003>
- de Haan L., & Resnick S. I. (1977). Limit theory for multivariate sample extremes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 40(4), 317–337. <https://doi.org/10.1007/BF00533086>
- Dissmann J., Brechmann E. C., Czado C., & Kurowicka D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, 59, 52–69. <https://doi.org/10.1016/j.csda.2012.08.010>
- Einmahl J. H., Kirilouk A., & Segers J. (2018). A continuous updating weighted least squares estimator of tail dependence in high dimensions. *Extremes*, 21, 205–233. <https://doi.org/10.1007/s10687-017-0303-7>
- Einmahl J. H., Krajina A., & Segers J. (2008). A method of moments estimator of tail dependence. *Bernoulli*, 14(4), 1003–1026. <https://doi.org/10.3150/08-BEJ130>
- Einmahl J. H. J., Piterbarg V. I., & de Haan L. (2001). Nonparametric estimation of the spectral measure of an extreme value distribution. *The Annals of Statistics*, 29(5), 1401–1423. <https://doi.org/10.1214/aos/1013203459>
- Engelke S., Hitz A., Gnecco N., & Hentschel M. (2022). graphicalExtremes: Statistical methodology for graphical extreme value models. R package version 0.2.0.
- Engelke S., & Hitz A. S. (2020). Graphical models for extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4), 871–932. <https://doi.org/10.1111/rssb.12355>
- Engelke S., & Ivanovs J. (2021). Sparse structures for multivariate extremes. *Annual Review of Statistics and Its Application*, 8, 241–270. <https://doi.org/10.1146/statistics.2021.8.issue-1>
- Engelke S., Lalancette M., & Volgushev S. (2021). ‘Learning extremal graphical structures in high dimensions’, arXiv:2111.00840, preprint: not peer reviewed. <https://doi.org/10.48550/arXiv.2111.00840>
- Engelke S., Malinowski A., Kabluchko Z., & Schlather M. (2015). Estimation of Hüsler–Reiss distributions and Brown–Resnick processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 77(1), 239–265. <https://doi.org/10.1111/rssb.12074>
- Engelke S., & Volgushev S. (2022). Structure learning for extremal tree models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(5), 2055–2087. <https://doi.org/10.1111/rssb.12556>
- Gissibl N., & Klüppelberg C. (2018). Max-linear models on directed acyclic graphs. *Bernoulli*, 24(4A), 2693–2720. <https://doi.org/10.3150/17-BEJ941>
- Gumbel E. J. (1960). Distributions des valeurs extrêmes en plusieurs dimensions. *Publications de l'Institut de statistique de l'Université de Paris*, 9, 171–173.
- Heffernan J. E., & Tawn J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(3), 497–546. <https://doi.org/10.1111/j.1467-9868.2004.02050.x>

- Hentschel M., Engelke S., & Segers J. (2024). Statistical inference for Hüsler–Reiss graphical models through matrix completions. *Journal of the American Statistical Association*. <https://doi.org/10.1080/01621459.2024.2371978>
- Hobæk Haff I. (2013). Parameter estimation for pair-copula constructions. *Bernoulli*, 19(2), 462–491. <https://doi.org/10.3150/12-BEJ413>
- Hu S., Peng Z., & Segers J. (2024). Modelling multivariate extreme value distributions via Markov trees. *Scandinavian Journal of Statistics*, 51(2), 760–800. <https://doi.org/10.1111/sjos.v51.2>
- Hüsler J., & Reiss R.-D. (1989). Maxima of normal random vectors: Between independence and complete dependence. *Statistics & Probability Letters*, 7(4), 283–286. [https://doi.org/10.1016/0167-7152\(89\)90106-5](https://doi.org/10.1016/0167-7152(89)90106-5)
- Jaworski P. (2006). On uniform tail expansions of multivariate copulas and wide convergence of measures. *Applicationes Mathematicae*, 33(2), 159–184. <https://doi.org/10.4064/am33-2-3>
- Joe H. (1990). Families of min-stable multivariate exponential and multivariate extreme value distributions. *Statistics & Probability Letters*, 9(1), 75–81. [https://doi.org/10.1016/0167-7152\(90\)90098-R](https://doi.org/10.1016/0167-7152(90)90098-R)
- Joe H. (1996). Families of m -variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. In *Distributions with fixed marginals and related topics* (Seattle, WA, 1993) of IMS Lecture Notes Monogr. Ser. (Vol. 28 pp. 120–141). Institute of Mathematical Statistics.
- Joe H., Li H., & Nikoloulopoulos A. K. (2010). Tail dependence functions and vine copulas. *Journal of Multivariate Analysis*, 101(1), 252–270. <https://doi.org/10.1016/j.jmva.2009.08.002>
- Ledford A. W., & Tawn J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1), 169–187. <https://doi.org/10.1093/biomet/83.1.169>
- Li H., & Wu P. (2013). Extremal dependence of copulas: A tail density approach. *Journal of Multivariate Analysis*, 114, 99–111. <https://doi.org/10.1016/j.jmva.2012.07.005>
- McNeil A. J., & Nešlehová J. (2009). Multivariate Archimedean copulas, d -monotone functions and ℓ_1 -norm symmetric distributions. *The Annals of Statistics*, 37(5B), 3059–3097. <https://doi.org/10.1214/07-AOS556>
- McNeil A. J., & Nešlehová J. (2010). From archimedean to liouville copulas. *Journal of Multivariate Analysis*, 101(8), 1772–1790. <https://doi.org/10.1016/j.jmva.2010.03.015>
- Morales-Napoles O. (2010). Counting vines. In D. Kurowicka & H. Joe (Eds.), *Dependence modeling: Vine copula handbook* (pp. 189–218). World Scientific. https://doi.org/10.1142/9789814299886_0009
- Morales-Nápoles O., Rajabi-Bahaabadi M., Torres-Alves G. A., & 't Hart C. M. P. (2023). Chimera: An atlas of regular vines on up to 8 nodes. *Scientific Data*, 10(1), 337. <https://doi.org/10.1038/s41597-023-02252-6>
- Müller D., & Czado C. (2019). Selection of sparse vine copulas in high dimensions with the lasso. *Statistics and Computing*, 29(2), 269–287. <https://doi.org/10.1007/s11222-018-9807-5>
- Nagler T., Bumann C., & Czado C. (2019). Model selection in sparse high-dimensional vine copula models with an application to portfolio risk. *Journal of Multivariate Analysis*, 172, 180–192. <https://doi.org/10.1016/j.jmva.2019.03.004>
- Nagler T., Schepsmeier U., Stoeber J., Brechmann E. C., Graeler B., & Erhardt T. (2023). *VineCopula: Statistical Inference of Vine Copulas*. R package version 2.4.5.
- Nagler T., & Vatter T. (2023). *rvinecopulib: High Performance Algorithms for Vine Copula Modeling*. R package version 0.6.3.1.1.
- Nolde N. (2014). Geometric interpretation of the residual dependence coefficient. *Journal of Multivariate Analysis*, 123, 85–95. <https://doi.org/10.1016/j.jmva.2013.08.018>
- R Core Team (2023). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Resnick S. I. (2007). *Heavy-tail phenomena: Probabilistic and statistical modeling*. Springer Science & Business Media.
- Rootzén H., & Tajvidi N. (2006). Multivariate generalized Pareto distributions. *Bernoulli*, 12(5), 917–930. <https://doi.org/10.3150/bj/1161614952>
- Rosenblatt M. (1952). Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3), 470–472. <https://doi.org/10.1214/aoms/1177729394>
- Schmidt R., & Stadtmüller U. (2006). Non-parametric estimation of tail dependence. *Scandinavian Journal of Statistics*, 33(2), 307–335. <https://doi.org/10.1111/sjos.2006.33.issue-2>
- Segers J. (2020). One-versus multi-component regular variation and extremes of Markov trees. *Advances in Applied Probability*, 52(3), 855–878. <https://doi.org/10.1017/apr.2020.22>
- Simpson E. S., Wadsworth J. L., & Tawn J. A. (2021). A geometric investigation into the tail dependence of vine copulas. *Journal of Multivariate Analysis*, 184, 104736. <https://doi.org/10.1016/j.jmva.2021.104736>
- Zhu K., & Kurowicka D. (2022). Regular vines with strongly chordal pattern of (conditional) independence. *Computational Statistics & Data Analysis*, 172, 107461. <https://doi.org/10.1016/j.csda.2022.107461>