Research Article Open Access

Nazih Benoumechiara, Nicolas Bousquet*, Bertrand Michel, and Philippe Saint-Pierre

Detecting and modeling critical dependence structures between random inputs of computer models

https://doi.org/10.1515/demo-2020-0016 Received March 3, 2020; accepted August 31, 2020

Abstract: Uncertain information on input parameters of computer models is usually modeled by considering these parameters as random, and described by marginal distributions and a dependence structure of these variables. In numerous real-world applications, while information is mainly provided by marginal distributions, typically from samples, little is really known on the dependence structure itself. Faced with this problem of incomplete or missing information, risk studies that make use of these computer models are often conducted by considering independence of input variables, at the risk of including irrelevant situations. This approach is especially used when reliability functions are considered as black-box models. Such analyses remain weakened in absence of in-depth model exploration, at the possible price of a strong risk misestimation. Considering the frequent case where the reliability output is a quantile, this article provides a methodology to improve risk assessment, by exploring a set of pessimistic dependencies using a copula-based strategy. In dimension greater than two, a greedy algorithm is provided to build input regular vine copulas reaching a minimum quantile to which a reliability admissible limit value can be compared, by selecting pairwise components of sensitive influence on the result. The strategy is tested over toy models and a real industrial case-study. The results highlight that current approaches can provide non-conservative results.

Keywords: computer model, output quantile, structural reliability, copula, regular vine

MSC: 62G07, 62G32, 62P30, 62H20

1 Introduction

Many industrial companies, like energy producers or vehicle and aircraft manufacturers, have to ensure a high level of safety for their facilities or products. In each case, the structural reliability of certain so-called critical components plays an essential role in overall safety. For reasons related to the fact that these critical components are highly reliable, and that real robustness tests can be very expensive or even hardly feasible, structural reliability studies generally use simulation tools [18, 43]. The physical phenomenon of interest being reproduced by a numerical model η (roughly speaking, a *computer model*), such studies are based on the calculation of a reliability indicator based on the comparison of $y = \eta(\mathbf{x})$ and a safety margin, where \mathbf{x} corresponds to a set of input parameters. In the framework of this article, such models are considered as black-boxes and can be explored only by simulation means.

Nazih Benoumechiara: Sorbonne Université, LPSM & EDF R&D, E-mail: nazih.benoumechiara@sorbonne-universite.fr

*Corresponding Author: Nicolas Bousquet: Sorbonne Université, LPSM & EDF R&D,

E-mail: nicolas.bousquet@sorbonne-universite.fr

Bertrand Michel: École Centrale Nantes, E-mail: bertrand.michel@ec-nantes.fr

Philippe Saint-Pierre: Institut de Mathématiques de Toulouse, Université Toulouse III,

E-mail: philippe.saint-pierre@math.univ-toulouse.fr

Open Access. CC BY © 2020 Nazih Benoumechiara et al., published by De Gruyter. This work is licensed under the Creative Commons Attribution alone 4.0 License.

While the problems of checking the validity of η and selecting inputs $\mathbf{x} \in \chi \subseteq \mathbb{R}^d$ are addressed by an increasing methodological corpus [3, 15], a perennial issue is the modeling of \mathbf{x} . Differing from the specification of η itself, this input vector is known with uncertainty, either because the number of experiments to estimate is limited, or because some inputs reflect intrinsically variable phenomena [59]. In most cases, these epistemic and aleatory uncertainties are jointly modeled by probability distributions [34]. Consecutively, the reliability indicator is often defined as the probability that y be lower than a threshold (*failure probability*), or a limit quantile for y. This article focuses on this last indicator, which provides an upper or lower bound of the mean effect of the output variable uncertainty.

Therefore the modeling of \mathbf{x} stands on the assessment of a joint probability distribution with support χ , divided between marginal and dependencies features. Though information on each dimension of \mathbf{x} can often be accessible experimentally or using physical or expert knowledge [5], the dependence structure between the component of \mathbf{x} remains generally unknown. Typically, statistical data are only available per dimension, but not available for two or more dimensions simultaneously. For this reason, most of robustness studies are conducted by sampling within independent marginal distributions. Doing so, reliability engineers try to capture input situations that minimize the reliability indicator.

However, the assumption of independence between inputs has been severely criticized since the works by [30] and [62], who showed that output failure probabilities of industrial systems can significantly vary and be underestimated if the input dependencies are neglected. More generally, [60, 61] showed that tail dependencies between inputs can have major expected effects on the uncertainty analysis results. A noticeable contribution to this analysis was made by [64], who highlighted, using vine copula models, the strong impact of input dependence structures on reliability indicators computed by well-known methods of uncertainty quantification.

In this paper, we consider a general probabilistic and statistical approach at modelling dependence with a supervised target. In particular, we study the specific problem of structural reliability as an application of such method. In decision-making problems, [40] proposed a general definition of the worst penalizing distribution as the minimizer of an excepted cost among a set of possible distributions. More recently, [2] extended this approach to account for incomplete dependence information. These theoretical works, that propose selection rules over the infinite set of all possible joint distributions, remain hard to apply in practice. Recent applied works made use of copulas [50] to model dependencies between stochastic inputs [60, 61], following other researchers confronted to similar problems in various fields: finance [13], structural safety [29], environmental sciences [57] or medicine [4]. These studies mainly consider bivariate copulas, which makes theses analysis effective only when two random variables are correlated. Cases where a greater number of variables is involved were explored by [37] then [64], who used vine copulas to approach complex multidimensional correlation problems in structural reliability, and by [63] in a specific context of regression by polynomial chaos expansion. A vine copula is a graphical representation of the pair-copula construction (PCC), proposed by [39], which defines a multidimensional dependence structure using conditional bivariate copulas. Various class of vines exist (see [17] for a review), and among them the regular vines (R-vines) introduced by [6, 7] are known for their appealing computational properties, while inference on PCC is usually demanding [20, 33].

R-vine parametric copulas seem promising to improve the search for a worst-case dependence between stochastic inputs, while keeping the benefits of a small number of parameters, as favoring inference and conducting simple sensitivity analyses *a posteriori*. To our knowledge, however, no practical methodology has been yet proposed to this end for which the notion of worst case is defined by the minimization of an output quantile. This is the subject of this article. More precisely, the aim of this research is to determine a parametric copula over \mathbf{x} , close to the worst case dependence structure, which is associated to a minimum value of the quantile of the distribution of y. Given a vine structure defined by a parameter vector, the optimization problem involves to conduct empirical quantile estimations for each value of this vector in a finite set of interest (chosen as a grid). The proposed methodology stands on an encompassing greedy algorithm exploring copula structures, which integrates several sub-algorithms of increasing complexity and is based on some simplifying assumptions. These algorithms are made available in the Python library dep-impact [8].

The article is therefore organized as follows. Section 2 introduces the framework and studies the consistency of a statistical estimation of the minimum quantile, given an input copula family and a growing

sequence of grids. The wider problem of selecting copulas in high-dimensional settings using a sequential algorithm of quantile minimization is considered in Section 3.

While the choice of R-vines is defended, similarly to [64] a sparsity hypothesis is made to diminish the computational burden, according to which only a limited number of pairwise dependencies is influent on the result. A greedy algorithm is proposed to carry out the complete procedure of optimization and modeling. This heuristic is tested in Section 4 over toy examples, using simulation, and a real industrial case-study. The results highlight that worst-case scenarios produced by this algorithm are often bivariate copulas reaching the Fréchet-Hoeffding bounds [27, 35] (describing perfect dependence between variables), as it could be expected in monotonic frameworks, but that other nontrivial copulas can be exhibited in alternative situations. Results and avenues for future research are extensively discussed in the last section of this article. We also refer to Appendix A and B for supplementary material on consistency proofs, on R-vine copulas and on R-vine iterative construction.

2 Minimization of the quantile of the output distribution

This section introduces a general framework for the calculation of the minimum quantile of the output distribution of a computational model, when the input distribution can be taken from a large family of distributions, each one corresponding to a particular choice of dependencies between the input variables.

2.1 A general framework for the computation of the minimum quantile

To be general, let us consider a computer code which takes a vector $\mathbf{x} \in \chi \subseteq \mathbb{R}^d$ as an input and produces a real quantity y in output. This code is represented by a deterministic function $\eta : \mathbb{R}^d \to \mathbb{R}$ such that $\eta(\mathbf{x}) = y$. The sets \mathbb{R} and \mathbb{R}^d are endowed with their Borel sigma algebras and we assume that η is measurable. The general expression of the function η is unknown but for some vector $\mathbf{x} \in \mathbb{R}^d$ it is assumed that the quantity $\eta(\mathbf{x})$ can always be computed. In particular, the derivatives of η , when they exist, are never assumed to be known. Let P_1, \ldots, P_d be a fixed family of d distributions, all supported on \mathbb{R} . We introduce the set $\mathbb{D}(P_1, \ldots, P_d)$ of all multivariate distributions P on \mathbb{R}^d such that the marginal distributions of P are all equal to the $(P_j)_{j=1...d}$. Henceforth, we use the shorter notation \mathbb{D} for $\mathbb{D}(P_1, \ldots, P_d)$.

For some $P \in \mathcal{D}$, let G be the cumulative distribution function of the model output. In other terms dG is the push-forward measure of P by η . For $\alpha \in (0, 1)$, let G^{-1} be the α -quantile of the output distribution:

$$G^{-1}(\alpha) := \inf\{y \in \mathbb{R} : G(y) \ge \alpha\}. \tag{1}$$

For the rest of this document, we denote as *output quantile* the α -quantile of the output distribution.

In many real situations, the function η corresponds to a known physical phenomenon. The input variables \mathbf{x} of the model are subject to uncertainties and are quantified by the distribution P. The propagation of these uncertainties leads to the calculation of the output quantile. Due to the difficulties to gather information, it is common to have this distribution incompletely defined and only known through its marginal distributions. Therefore, the set $\mathbb D$ corresponds to all the possible distributions that are only known through their marginal distributions $(P_j)_{j=1...d}$. Detecting very unfavorable correlations between inputs is similar to detecting correlations leading to low quantile values. We define as the *worst quantile*, the minimum value of the quantile by considering all the possible input distributions $P \in \mathbb D$. This conservative approach consists in minimizing $G^{-1}(\alpha)$ over the family $\mathbb D$ such as

$$G^{-1^{\star}}(\alpha) := \min_{P \in \mathcal{D}} G^{-1}(\alpha). \tag{2}$$

Since the function η has no closed form in general, it is not possible to give a simple expression of $G^{-1}(\alpha)$ in function of the distribution P, and consequently the minimum $G^{-1}(\alpha)$ does not have a simple expression

Figure 1: Illustration of the link between the dependence parameter $\boldsymbol{\theta}$ and the quantile function $G_{\boldsymbol{\theta}}^{-1}$. The joint CDF $F_{\boldsymbol{\theta}}$ is obtained using (3) from a copula $C_{\boldsymbol{\theta}}$ and marginal CDF's $(F_j)_{j=1}^d$. The push-forward of $F_{\boldsymbol{\theta}}$ through the model η leads to the CDF $G_{\boldsymbol{\theta}}$ and quantile function $G_{\boldsymbol{\theta}}^{-1}$ of the output distribution.

too. In this paper we propose to study a simpler problem than (2), by minimizing $G^{-1}(\alpha)$ over a subset of \mathcal{D} . This subset is a family of distributions $(P_{\theta})_{\theta \in \Theta}$ associated to a parametric family of copula $(C_{\theta})_{\theta \in \Theta}$, where Θ is a compact set of \mathbb{R}^p and p is the number of copula parameters.

2.2 Copula-based approach

We introduce the real-values random vector $\mathbf{X} = (X_1, \dots, X_d) \in \mathbb{R}^d$ associated to the distribution $P_{\boldsymbol{\theta}}$. Each component X_j , for $j = 1, \dots, d$, is a real-value random variable with distribution P_j . A copula describes the dependence structure between a group of random variables. Formally, a copula is a multidimensional continuous cumulative distribution function (CDF) linking the margins of \mathbf{X} to its joint distribution. Sklar's Theorem [58] states that every joint distribution $F_{\boldsymbol{\theta}}$ associated to the measure $P_{\boldsymbol{\theta}}$ can be written as

$$F_{\boldsymbol{\theta}}(\mathbf{x}) = C_{\boldsymbol{\theta}} \left(F_1(x_1), \dots, F_d(x_d) \right), \tag{3}$$

with some appropriate d-dimensional copula C_{θ} with parameter $\theta \in \Theta$ and the marginal CDF's $F_j(x_j) = \mathbb{P}[X_j \le x_j]$. If all marginal distributions are continuous functions, then there exists a unique copula satisfying

$$C_{\theta}(u_1,\ldots,u_d) = F_{\theta}(F_1^{-1}(u_1),\ldots,F_d^{-1}(u_d))$$

where $u_j = F_j(x_j)$. For F_{θ} absolutely continuous with strictly increasing marginal distributions, one can derive (3) to obtain the joint density of **X**:

$$f_{\boldsymbol{\theta}}(\mathbf{x}) = c_{\boldsymbol{\theta}} \left(F_1(x_1), \dots, F_d(x_d) \right) \prod_{i=1}^d f_i(x_i), \tag{4}$$

where c_{θ} denotes the copula density function of C_{θ} and $f_j(x_j)$ are the marginal densities of **X**. Numerous parametric copula families are available and are based on different dependence structures. Most of these families have bidimensional dependencies, but some can be extended to higher dimensions. However, these extensions have a lack of flexibility and cannot describe all types of dependencies [50]. To overcome these difficulties, tools like vine copulas [38] (described in Section 3) combine bivariate copulas, from different families, to create a multidimensional copula.

Let G_{θ} and G_{θ}^{-1} be respectively the CDF and quantile function of the push-forward distribution of P_{θ} by η (see Figure 1). For a given parametric family of copula $(C_{\theta})_{\theta \in \Theta}$ and a given $\alpha \in (0, 1)$, the minimum output quantile for a given copula is defined by

$$G_C^{-1^*}(\alpha) := \inf_{\boldsymbol{\theta} \in \Theta} G_{\boldsymbol{\theta}}^{-1}(\alpha)$$
 (5)

and if it exists, we consider a minimum

$$\boldsymbol{\theta}_{\mathcal{C}}^{\star} \in \operatorname*{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} G_{\boldsymbol{\theta}}^{-1}(\alpha).$$
 (6)

We call this quantity the minimum quantile parameter or worst dependence structure.

Note that there is no reason for $G_{\theta}^{-1}(\alpha)$ to be a convex function of θ . The use of gradient descent algorithms is thus not straightforward in this context. Moreover, the gradient of $\theta \to G_{\theta}^{-1}$ is unknown and only zero-order optimization methods can be applied to solve (6). For this reason, in the following of this section, we

analyze the basic approach which consists in estimating θ_C^* by approximating θ with a finite regular grid θ_N of cardinality N. Therefore, for a given parametric copula $(C_{\theta})_{\theta \in \theta}$ and a given $\alpha \in (0, 1)$, we restrict the problem (6) to

$$\boldsymbol{\theta}_{N}^{\star} \in \operatorname*{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{N}} G_{\boldsymbol{\theta}}^{-1}(\alpha). \tag{7}$$

2.3 Estimation with a grid search strategy

In the restricted problem (7), the greater N, the closer $\boldsymbol{\theta}_N^*$ to the minimum $\boldsymbol{\theta}^*$ of $\boldsymbol{\Theta}$; obviously the convergence rate should depend on the regularity of the function η and on the regularity of the quantile function $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}^{-1}(\alpha)$. Because η has no closed form, the quantile function $G_{\boldsymbol{\theta}}^{-1}(\alpha)$ has no explicit expression. The minimizer $\boldsymbol{\theta}_N^*$ can be estimated by coupling the simulation of independent and identically distributed (i.i.d) data (Y_1, \ldots, Y_n) , defined as realizations of the model output random variable $Y := \eta(\mathbf{X})$ with distribution $dG_{\boldsymbol{\theta}}$, with a minimization of the empirical quantile over $\boldsymbol{\Theta}_N$. Note that this problem does not deal with fitting vines on a set of observations: here we generate data according to various vine distributions and we estimate quantiles under these distributions.

For θ taking a value over the grid Θ_N , the empirical CDF of Y is defined for any $y \in \mathbb{R}$ by

$$\widehat{G}_{\boldsymbol{\theta}}(y) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{Y_i \le y}.$$
 (8)

The corresponding empirical quantile function $\widehat{G}_{\theta}^{-1}(\alpha)$ is defined as in (1) by replacing G with its empirical estimate. For a given α , the worst quantile on the fixed grid Θ_N is given by

$$\min_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_N}G_{\boldsymbol{\theta}}^{-1}(\alpha).$$

and can be estimated by replacing the quantile function with its empirical function:

$$\min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_N} \widehat{G}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\alpha}). \tag{9}$$

Finally the estimation of the minimum quantile parameter over the grid Θ_N is denoted by

$$\widehat{\boldsymbol{\theta}}_{N} = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{N}}{\operatorname{argmin}} \widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha). \tag{10}$$

The construction of the grid Θ_N can be difficult because Θ can be unbounded (e.g. $\Theta = [1, \infty]$ for a Gumbel copula). To tackle this issue, we chose to construct Θ_N among a normalized space using a concordance measure, which is bounded in [-1, 1] and does not rely on the marginal distributions. We chose the commonly used Kendall rank correlation coefficient (or Kendall's tau) [41] as a concordance measure to create this transitory space. This non-linear coefficient $\tau \in [-1, 1]$ is related to the copula function as follows:

$$\tau = 4 \int_{-1}^{1} \int_{-1}^{1} C_{\theta}(u_1, u_2) dC(u_1, u_2) - 1.$$

For many copula families, this relation is much more explicit (see for instance [28]). Therefore, the finite grid is created among $[-1, 1]^p$ and each element of this grid is converted to the copula parameter θ . Moreover, the use of concordance measures gives a normalized expression of the strength of dependencies for all pairs of variables, independently of the used copula families.

The consistency of estimators (9) and (10) is studied in next section, under general regularity and geometric assumptions on η and the functional $\theta \mapsto P_{\theta}$.

2.4 Consistency of worst quantile-related estimators

In this section, we give consistency results of the estimators $\min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_N} \widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)$ and $\widehat{\boldsymbol{\theta}}_N$, for a growing sequence of grids on the domain $\boldsymbol{\theta}$. For easier reading, we skip some definitions needed for our assumptions. Section A in Appendix provides a more complete presentation, including the formal definition of the modulus of increase of the quantile function.

Let α be a fixed value in (0, 1). To approximate Θ , we consider a sequence of finite discrete grids $(\Theta_N)_{N\geq 1}$ on Θ where N is the cardinal of Θ_N and such that

$$\sup_{\boldsymbol{\theta} \in \boldsymbol{\Theta}, \; \boldsymbol{\theta}' \in \boldsymbol{\Theta}_N} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2 \to 0 \quad \text{ as } N \text{ tends to infinity.}$$
 (11)

We first introduce technical hypotheses required for the consistency result which are commented further in the text.

Assumption A. For all $\theta \in \Theta$, the distribution P_{θ} admits a density f_{θ} for the Lebesgue measure and the copula C_{θ} admits a density c_{θ} for the Lebesgue measure on $[0,1]^d$ such that

$$m{ heta} imes [0,1]^d \longrightarrow \mathbb{R}$$
 $m{ heta} imes (x_1, \dots, x_d) \longrightarrow c_{m{ heta}}(x_1, \dots, x_d)$

is a continuous function.

Assumption B. For all $\theta \in \Theta$, G_{θ} is a continuous function.

Assumption C. For all $\theta \in \Theta$, G_{θ} is strictly increasing and the modulus of increase of G_{θ} at $G_{\theta}^{-1}(\alpha)$ is lower bounded by a positive function $\underline{\epsilon}_{\Theta}$.

Assumption D. There exists a unique $\theta^* \in \Theta$ minimizing $\theta \mapsto G_{\theta}^{-1}(\alpha)$.

Let $(N_n)_{n\geq 1}$ be a sequence of integers such that $N_n \lesssim n^{\beta}$ for some $\beta > 0$. For every $n \geq 1$ we consider the grid Θ_{N_n} and for every $\theta \in \Theta_{N_n}$ we compute the empirical quantile $\widehat{G}_{\theta}^{-1}(\alpha)$ from a sample of n i.i.d variables Y_1, \ldots, Y_n with $Y_i = \eta(\mathbf{X}_i)$, where the $\mathbf{X}_i's$ are i.i.d. random vectors with distribution P_{θ} . We then introduce the extremum estimator

$$\hat{\boldsymbol{\theta}} := \widehat{\boldsymbol{\theta}}_{N_n}. \tag{12}$$

Theorem 1. Under Assumptions **A**, **B** and **C**, for all $\varepsilon > 0$ we have

$$P\left(\left|\widehat{G}_{\widehat{\boldsymbol{\theta}}}^{-1}(\alpha) - G_{C}^{-1}^{*}(\alpha)\right| > \varepsilon\right) \xrightarrow{n \to \infty} 0. \tag{13}$$

Moreover, if Assumption **D** is also satisfied, then for all h > 0 we have

$$\mathbb{P}[|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_C^{\star}| > h] \xrightarrow{n \to \infty} 0$$

(proof given in Appendix A).

It would be possible to provide rates of convergence for this extremum quantile and for $\boldsymbol{\theta}^*$ at the price of more technical proofs, by considering also the dimension metric of the domain $\boldsymbol{\theta}$ and the modulus of increase of the function $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}(\alpha)$ (see for instance the proofs of Theorems 1 and 2 in [12] for an illustration of such computations). It would be also possible to derive similar results for alternative extremum quantities. One first example, useful in many applications, would be to estimate some "limit probability" by determining an extremum $\inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} G_{\boldsymbol{\theta}}(y)$ of the CDF for a fixed y.

This consistency result could also be extended for regular functional of G_{θ} or G_{θ}^{-1} , such that

$$\inf_{\boldsymbol{\theta}\in\Theta}\int_{y\geq y_0}G_{\boldsymbol{\theta}}(y)\mathrm{d}y\quad\text{or}\quad\inf_{\boldsymbol{\theta}\in\Theta}\int_{\alpha\geq\alpha_0}G_{\boldsymbol{\theta}}^{-1}(\alpha)\mathrm{d}y,$$

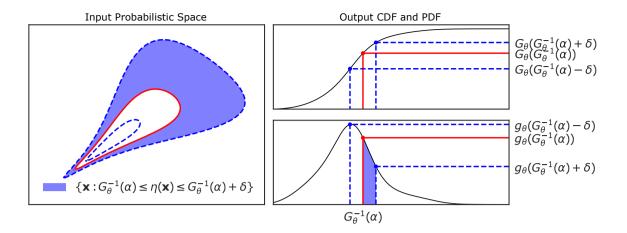


Figure 2: Pre-image (left) and image (right) of a modulus of increase of G_{θ} at the point $G_{\theta}^{-1}(\alpha)$ for a deviation $\pm \delta$.

for some fixed values y_0 and α_0 . Extending our results for such quantities is possible essentially because the Dvoretsky-Kiefer-Wolfowitz (DKW) inequality [21], used in the proof, gives an uniform control on the estimation of the CDF and the quantile function.

We now discuss the three first assumptions and provide some geometric and probabilistic interpretations of them. Assumption **A** requires some regularity of the input distribution with respect to $\boldsymbol{\theta}$. This is indeed necessary to locate the minimum of the quantile. Assumption **B** and **C** ensure that the output quantile function $G_{\boldsymbol{\theta}}^{-1}$ has a regular behavior in a neighborhood of the computed quantile $G_{\boldsymbol{\theta}}^{-1}(\alpha)$. Assumption **B** ensures that the output distribution $dG_{\boldsymbol{\theta}}$ has no Dirac masses whereas Assumption **C** ensures that there is no area of null mass inside the domain of $dG_{\boldsymbol{\theta}}$.

Figure 2 illustrates Assumption **B** with a possible configuration of the input distribution. For $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ an $\delta > 0$, we consider a small neighborhood $[G_{\boldsymbol{\theta}}^{-1}(\alpha) - \delta, G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta]$ of $G_{\boldsymbol{\theta}}^{-1}(\alpha)$, and the pre-image of this neighborhood. The two right figures are the CDF $G_{\boldsymbol{\theta}}$ (top) and PDF $g_{\boldsymbol{\theta}}$ (bottom) of the output variable Y for a given $\boldsymbol{\theta}$. The figure at the left hand represents the contours of the pre-image in the input space. The red plain line is the level set $\eta^{-1}(G_{\boldsymbol{\theta}}^{-1}(\alpha))$ and the dot blue line is the perturbed level set $\eta^{-1}(G_{\boldsymbol{\theta}}^{-1}(\alpha) \pm \delta)$). The blue area in the right figure corresponds to $[G_{\boldsymbol{\theta}}^{-1}(\alpha) - \delta, G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta]$ and the pre-image of this neighborhood is the blue area in the left figure. Assumption **B** requires that the mass of the blue domain is lower bounded by a positive function $\underline{\varepsilon}_{\boldsymbol{\theta}}(\delta)$ that does not depend on $\boldsymbol{\theta}$.

It is possible to give sufficient conditions on the input distribution F_{θ} and on the geometry of η to obtain Assumptions **B** and **C**. Using the definition of the modulus of continuity from Equation (A.1) in Appendix, it comes

$$\begin{split} \epsilon_{G_{\boldsymbol{\theta}}}(\delta, G_{\boldsymbol{\theta}}^{-1}(\alpha)) &= \max \left[\int\limits_{\{G_{\boldsymbol{\theta}}^{-1}(\alpha) \leq g \leq G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta\}} f_{\boldsymbol{\theta}}(\mathbf{x}) \mathrm{d}\lambda(\mathbf{x}); \int\limits_{\{G_{\boldsymbol{\theta}}^{-1}(\alpha) - \delta \leq g \leq G_{\boldsymbol{\theta}}^{-1}(\alpha)\}} f_{\boldsymbol{\theta}}(\mathbf{x}) \mathrm{d}\lambda(\mathbf{x}) \right] \\ &\geq \int\limits_{\{G_{\boldsymbol{\theta}}^{-1}(\alpha) \leq g \leq G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta\}} f_{\boldsymbol{\theta}}(\mathbf{x}) \mathrm{d}\lambda(\mathbf{x}) \end{split}$$

Assume that the code η is a Lipschitz and differentiable function with no null derivatives almost everywhere in the neighborhood of $G_{\theta}^{-1}(\alpha)$. Then, using the *coarea formula* (see for instance [26], Section 3.4.4, Proposition 3), we find that

$$\epsilon_{G_{\boldsymbol{\theta}}}(\delta, G_{\boldsymbol{\theta}}^{-1}(\alpha)) \geq \int\limits_{G_{\boldsymbol{\theta}}^{-1}(\alpha)}^{G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta} \left[\int\limits_{\eta^{-1}\{u\}} \frac{f_{\boldsymbol{\theta}}}{\|\nabla \eta\|} \mathrm{d}\mathcal{H}^{d-1} \right] \mathrm{d}u,$$

where \mathcal{H}^{d-1} is the d-1 dimensional Hausdorff measure (see for instance Chapter 2 in [26]). If the copula and the code are such that there exists a constant I such that for any $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and any u in the support of $dG_{\boldsymbol{\theta}}$

$$\int_{\eta^{-1}\{u\}} f_{\boldsymbol{\theta}} \, \mathrm{d} \mathcal{H}^{d-1} \leq I,$$

then we find that

$$\epsilon_{G_{\boldsymbol{\theta}}}(\delta, G_{\boldsymbol{\theta}}^{-1}(\alpha)) \geq \delta \frac{I}{\|\nabla \eta\|_{\infty}}.$$

Note that $\|\nabla \eta\|_{\infty} < \infty$ since η is assumed to be Lipschitz. We have proved that Assumption \mathbf{C} is satisfied in this context. Finally, by rewriting again the co-area formula for $G_{\theta}(y)$, we find that Assumption \mathbf{B} is satisfied as soon as the set of stationary points $(\|\nabla \eta(x)\| = 0)$ of all level set $\eta^{-1}\{u\}$ has null mass for the Hausdorff measure.

In conclusion, we see that for smooth copulas, Assumptions **C** and **B** mainly depend on the regularity of the function η , by requiring on one side that η does not oscillate to much and on the other side that the set of stationary points does not have a positive mass on the level sets of η .

3 Quantile minimization and choice of penalized correlation structure

A preliminary study of the influence of the dependence structure, specific to quantile minimization, is conducted in Appendix C. It also illustrates that the worst quantile is not always reached for a trivial copula (i.e., reaching the Fréchet-Hoeffding bounds).

The current section first provides a rationale for choosing the so-called R-vine structure as a preferential copula structure for modeling the variety of correlations between inputs. Then, the search for a minimum quantile is presented in two times. Subsection 3.2 proposes an exhaustive grid-search algorithm for estimating this quantile when the R-vine copula structure is fixed with a given pair-copula families and indexed by the parameter vector $\boldsymbol{\theta}$. Subsection 3.3 extends this rigid framework by permitting the search of particular subcopula pairwise structures, such that the minimization be more significant. In each situation, examples are provided to demonstrate the feasibility of the approach.

3.1 A rationale for R-vine copula structures

Representing dependence structures in high dimensional settings is a challenging problem. For the following definition, we simplify the expressions by omitting the use of θ : $f = f_{\theta}$, $F = F_{\theta}$ and $c = c_{\theta}$. By recursive conditioning, the joint density can be written as a product of conditioning distributions such as

$$f(x_1, \dots, x_d) = f_1(x_1) \cdot f_{2|1}(x_2|x_1) \cdot f_{3|1,2}(x_3|x_1, x_2) \cdots \cdots f_{d|1,2,\dots,d-1}(x_d|x_1, x_2, \dots, x_{d-1}).$$

$$(14)$$

For clarity reason, we now simplify the expression with $f_{3|1,2} = f_{3|1,2}(x_3|x_1,x_2)$ and so on for other orders. From (4), the conditioning densities of (14) can be rewritten as products of conditioning copula and marginal densities. For example, in a case of three variables and using (14), one possible decomposition of the the joint density can be written as

$$f(x_1, x_2, x_3) = f_1 \cdot f_{2|1} \cdot f_{3|1,2}. \tag{15}$$

Using (4), the reformulation of $f_{3|1,2}$ leads to

$$f_{3|1,2} = \frac{f_{1,3|2}}{f_{1|2}} = \frac{c_{1,3|2} \cdot f_{1|2} \cdot f_{3|2}}{f_{1|2}} = c_{1,3|2} \cdot f_{3|2}$$
 (16)

where $c_{1,3|2} = c_{1,3|2}(F_{1|2}(x_1|x_2), F_{3|2}(x_3|x_2))$. By developing $f_{3|1,2}$ in the same way, we find that

$$f_{3|1,2} = c_{1,3|2} \cdot c_{2,3} \cdot f_3. \tag{17}$$

Thus, by replacing the expression of $f_{3|1,2}$ in (15) and doing the same procedure for $f_{2|1}$, the joint density can be written as

$$f(x_1, x_2, x_3) = f_1 \cdot f_2 \cdot f_3 \cdot c_{1,2} \cdot c_{2,3} \cdot c_{1,3|2}. \tag{18}$$

This final representation of the joint density based on pair-copulas has been developed in [39] and is called the pair-copula construction (PCC). The resulting copula represented by the product of conditional copulas in (18) offers a very flexible way to construct high-dimensional copulas. However, it is not unique; indeed, (14) has numerous decomposition forms and it increases with the dimension.

To describe all such possible constructions in an efficient way, [6, 7] introduced the vine models. This graphical tool, based on a sequence of trees, gives a specific way to decompose the multivariate probability distribution. Basically, a vine model is defined by

- a structure of trees which can be represented by a matrix [48],
- a copula family for each pair of the structure,
- a parameter for each pair-copula.

A R-vine is the general construction of a vine model, but particular cases exists such as the D-vines and C-vines, described in Appendix B. Vine models were deeply studied in terms of density estimation and model selection using maximum likelihood [1], sequential estimation [20, 42], truncation [11] and Bayesian techniques [31]. Inspired by recent studies like [64], their popularity and well-known flexibility led us to use R-vines in this article, despite the fact that in our context we are looking for a conservative form and not to select the most appropriate form with given data, in absence of correlated observations.

3.2 Estimating a minimum quantile from a given R-vine

3.2.1 Grid-search algorithm

Let $\Omega = \{(i,j): 1 \le i,j \le d\}$ be the set of all the possible pairs of \mathbf{X} , in a d-dimensional problem. The number of pairs p is associated to the size of Ω such as $p = |\Omega| = {d \choose 2} = d(d-1)/2$. We define $\mathcal V$ as the vine structure and we consider fixed copula families for each pair. In this article, we only consider single parameter pair-copulas, such that the parameter $\boldsymbol \theta$ is a p-dimensional parameter vector with a definition space $\boldsymbol \Theta := \prod_{(i,j)\in\Omega} \theta_{i,j}$ where $\theta_{i,j}$ is the parameter space for the pair-copula of the pair (i,j). However, the methodology can easily be extended to multi-parameter pair-copulas. Note that a pair-copula can be conditioned to other variables, depending on its position in the vine structure $\mathcal V$. Thus, the input distribution $\mathrm{d} F_{\boldsymbol \theta}(\mathcal V)$ is defined by the vine structure $\mathcal V$, the copula families and the parameter $\boldsymbol \theta$. Also note that the copula parameter $\boldsymbol \theta$ is associated to the R-vine structure $\mathcal V$ (i.e., $\boldsymbol \theta = \boldsymbol \theta_{\mathcal V}$), see Section 3.2.2. For the sake of clarity, we simplify the notation to $\boldsymbol \theta$ only.

The most direct approach to estimate the minimum quantile is the Exhaustive Grid-Search algorithm, described by Algorithm 1.

Algorithm 1: Exhaustive Grid-Search algorithm to minimize the output quantile.

Data: A vine structure \mathcal{V} , a fixed grid $\boldsymbol{\Theta}_N$, a sample size n

1 for $\boldsymbol{\theta} \in \boldsymbol{\Theta}_N$ do

3

- **1.** Simulate a sample $\{X_i\}_{i=1}^n$ according to $dF_{\theta}(\mathcal{V})$;
 - **2.** Evaluate $\{Y_i = \eta(\mathbf{X}_i)\}_{i=1}^n$;
- **3.** Compute $\widehat{G}_{\theta}^{-1}(\alpha)$: empirical quantile of $\{Y_i\}_{i=1}^n$;

Result: $\min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_N} \widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)$

For a given vine structure \mathcal{V} , copula families, a grid $\boldsymbol{\Theta}_N$ and a sample size n, three steps are needed for each $\boldsymbol{\theta} \in \boldsymbol{\Theta}_N$. The first step simulates an input sample $\{\mathbf{X}_i\}_{i=1}^n$ according to the distribution $\mathrm{d}F_{\boldsymbol{\theta}}(\mathcal{V})$ for a given sample size n. The second evaluates the sample through the model η . The third estimates the output quantile from the resulting sample $\{Y_i = \eta(\mathbf{X}_i)\}_{i=1}^n$. The minimum quantile is took among the results of each loop.

3.2.2 Influence of the vine structure

Using R-vines, the dependence parameter θ is associated to the vine structure \mathcal{V} . Due to the hierarchy of the vine structure, some pair-copulas are conditioned to other variables and thus for their parameters. As an illustration, let us consider two vine structures with the two following copula densities, with the same simplified expressions as for (18):

$$c_{\mathcal{V}_1}(x_1, x_2, x_3, x_4) = c_{\theta_{1,3}} \cdot c_{\theta_{1,2}} \cdot c_{\theta_{2,4}} \cdot c_{\theta_{2,3|1}} \cdot c_{\theta_{1,4|2}} \cdot c_{\theta_{3,4|1,2}}$$
(19)

$$c_{\mathcal{V}_2}(x_1, x_2, x_3, x_4) = c_{\theta_{1,3}} \cdot c_{\theta_{3,4}} \cdot c_{\theta_{2,4}} \cdot c_{\theta_{1,4|3}} \cdot c_{\theta_{2,3|4}} \cdot c_{\theta_{1,2|3,4}}. \tag{20}$$

The difference between these densities is the conditioning of some pairs, the dependence parameters of theses vines are $\boldsymbol{\theta}_{\mathcal{V}_1} = [\theta_{1,2}, \theta_{1,3}, \theta_{1,4|2}, \theta_{2,3|1}, \theta_{2,4}, \theta_{3,4|1,2}]$ and $\boldsymbol{\theta}_{\mathcal{V}_2} = [\theta_{1,2|3,4}, \theta_{1,3}, \theta_{1,4|3}, \theta_{2,3|4}, \theta_{2,4}, \theta_{3,4}]$. Applying the same grid for these two vines may give different results due to the conditioning order from the vine structure. For example, if the pair X_3 - X_4 is very influential on minimizing the output quantile, it would be more difficult to find a minimum with \mathcal{V}_1 than \mathcal{V}_2 due to the conditioning of the pair with X_1 and X_2 in \mathcal{V}_1 . However, if the grid is thin enough, the minimum from these two vines should be equivalent.

To counter this difficulty, one possible option consists in randomly permuting the indexes of the variables and repeating the algorithm several times to visit different vines structures.

3.2.3 Computational cost

For one given R-vine structure and one fixed copula family at each pair, the overall cost of the method is equal to nN. However, as explained in § 2.4, the finite grid Θ_N , should be thin enough to reasonably explore Θ . Therefore, N should increase with the number of dimensions d and more specifically with the number of pairs $p = \binom{d}{2}$. A natural form for N would be to write it as $N = \gamma^p$, where $\gamma \in \mathbb{R}^+$. Thus, the overall cost of the exhaustive grid-search would be equal to $n\gamma^{\binom{d}{2}}$. The cost is in $\mathbb{O}(\gamma^{d^2})$ which makes the method hardly scalable when the dimension d increases. Ideas about alternative optimizations methods are discussed in Section 5.

3.3 Iterative search for a penalizing R-vine structure: a greedy heuristic based on pairwise copula

3.3.1 Going further in quantile minimization

With Algorithm 1, the previous subsection proposes an exhaustive grid-search strategy to determine a R-vine copula $C_{\tilde{\theta}}$ such that the associated output quantile $G_{\tilde{\theta}}^{-1}(\alpha)$ be the smallest (and also the most conservative in a structural reliability context). This approach remains however limited in practice since $C_{\tilde{\theta}}$ for fixed pair-copula families (e.g., Archimedean or max-stable copulas) and \mathcal{V} which is a member of the set \mathcal{F}_d of all the possible d-dimensional R-vine structure. Intuitively, a more reliable approach to quantile minimization should be based on mixing this estimation method with a selection among all members of the finite set \mathcal{F}_d , as well for the copula families. It is indeed likely that searching within an associative class of copulas like Archimedean ones, allowing modeling dependence in arbitrarily high dimensions, be a too rigid choice for estimating the minimum $G_{\tilde{\theta}}^{-1}(\alpha)$.

A minimum quantile can probably be found using a R-vine structure defined by conditional pairwise subcopulas (according to (18)) that are not part of the same rigid structure. However, a brute force exploration of \mathcal{F}_d would be conducted at an exponential cost increasing with d [49]. If we also consider the large computational cost of an exhaustive grid-search for a large number of dependent variables (as explained in § 3.2), this approach is not feasible in practice for high dimensions.

For this reason, it is proposed to extend Algorithm 1 by a greedy *heuristic* that dynamically selects the most influential correlations between variables while limiting the search to pairwise correlations. Doing so, minimizing the output quantile can be conducted in a reasonable computational time. Therefore the selected d-dimensional vine structure would be filled with independent pair-copulas except for the pairs that are influential on the minimization.

This working limitation, interpreted as a sparsity constraint, is based on the following assumption: it is hypothesized that only few pairs of variables have real influences on the minimization. It is close in spirit to the main assumption of global sensitivity analysis applied to computer models, according to which only a limited number of random variables has a major impact on the output [36, 56].

3.3.2 General principle

The method basically relies on an iterative algorithm exploring pairwise correlations between the uniform random variables $U_j = F_j^{-1}(X_j)$ and progressively building a non-trivial R-vine structure, adding one pair of variable to the structure at each iteration. Starting at step k = 0 from the simple independent copula

$$C_{\boldsymbol{\theta}^{(0)}}(u_1,\ldots,u_d)=\prod_{j=1}^d u_j,$$

the algorithm finally stops at a given step k = K while proposing a new copula $C_{\theta^{(K)}}$ associated to a R-vine structure V_K mostly composed of independent pair-copulas.

At each iteration k, we denote by Ω_k the *selected pairs* which are considered non-trivial (non-independent) due to their influence on the quantile minimization. Let $\Omega_{-k} = \Omega \backslash \Omega_k$ be the *candidate pairs*, which were not the remaining pairs, which influence on the minimization is still to be tested and are still considered independent. We also consider \mathbf{B} as a set of candidate copula families. The pseudo-code of Algorithm 2 shows in detail how this iterative exploration and building is conducted. More algorithms in Appendix B.2 described how to construct a vine structure with a given list of indexed pairs of variable.

3.3.3 Example

Consider the four-dimensional (d = 4) situation such as $\mathbf{X} = (X_1, \dots, X_4)$ where, for to the sake of simplicity, all marginal distributions of \mathbf{X} are assumed to be uniform on [0, 1]. We consider a simple additive model described by

$$\eta(\mathbf{X}) = 30X_1 + 10X_3 + 100X_4. \tag{21}$$

For an additive model and uniform margins, the output quantile is monotonic with the dependence parameters (see Subsection 2 in Appendix C) which locates the minimum quantile at the edge of Θ . Thus, Step 1.b. of Algorithm 2 is simplified by considering only Fréchet-Hoeffding copulas in the exploration.

In this illustration we consider $\alpha=0.1$ and we select n=300,000 large enough in order to have a great quantile estimation and the algorithm stops at K=3. Figure 3 shows, for each iteration k, the p-k vine structures that have been created by the algorithm. The red nodes and edges are the candidate pairs $(i,j)\in\Omega_{-k}$ and the blue nodes and edges are the selected pairs Ω_k . At iteration k=0, the selected pair is (1,4) with an estimated minimum quantile of -56.03. At iteration k=1, the second selected pair is (2,4) with an estimated minimum quantile of -56.03. At iteration k=2, the third selected pair is (2,4) with an estimated minimum quantile of -56.23.

Algorithm 2: Minimization of the output quantile and estimation of $\theta^{(K)}$ over an increasing family of R-vine structures.

```
1 Initialization:
  2 Iteration: k = 0;
  3 Selected pairs: \Omega_0 = \emptyset;
 4 Selected families: \mathbf{B}_0 = \emptyset;
  5 while k \le K do
             Copula parameter space of the selected pairs: \Theta_k = \prod_{(i,j) \in \Omega_k} \Theta_{i,j};
             1. Explore the set of candidate pairs \Omega_{-k};
  7
             for (i, j) \in \Omega_{-k} do
 8
                    a. Create a vine structure V_{(i,i)} using the procedure of Section B.2 applied to the list \Omega_k \cup (i,j);
                    b. Explore the set of candidate families B;
10
                    for \mathfrak{B} \in \mathbf{B} do
 11
                           Apply Algorithm 1 with the pair-copula families \mathcal{B} \cup \mathbf{B}_k;
12
                             (i) Define a (k+1)-dimensional grid \Delta_{i,j} of \Theta_k \times \Theta_{i,j} with
                                   cardinality N_k;
                            (ii) Select the minimum over the grid \Delta_{i,j}:
                                                                                                          \hat{\boldsymbol{\theta}}_{\mathcal{B}} = \underset{\boldsymbol{\theta}_{\mathcal{B}} \in \boldsymbol{\Delta}_{i,i}}{\operatorname{argmin}} \left\{ \widehat{G}_{\boldsymbol{\theta}_{\mathcal{B}}}^{-1}(\alpha) \right\}.
                    c. Select the minimum among B
13
                                                                                   \mathcal{B}_{i,j} = \underset{\mathcal{B} \in \mathbf{B}}{\operatorname{argmin}} \left\{ \widehat{G}_{\hat{\boldsymbol{\theta}}_{\mathcal{B}}}^{-1}(\alpha) \right\}\hat{\boldsymbol{\theta}}_{i,j} = \hat{\boldsymbol{\theta}}_{\mathcal{B}_{i,j}}
             2. Select the minimum among \Omega_{-k}
14
                                                                            (i,j)^{(k)} = \underset{(i,j)\in\Omega_{-k}}{\operatorname{argmin}} \left\{ \widehat{G}_{\widehat{\boldsymbol{\theta}}_{i,j}}^{-1}(\alpha) \right\},
15
             3. Check the stopping condition;
             if \widehat{G}_{\widehat{\pmb{\theta}}^{(k)}}^{-1}(\alpha) \ge \widehat{G}_{\widehat{\pmb{\theta}}^{(k-1)}}^{-1}(\alpha) then
16
                  K = k - 1:
17
18
                    Extend the list of selected pairs: \Omega_k = \Omega_k \cup (i,j)^{(k)} and families: \mathbf{B}_k = \mathbf{B}_k \cup \mathcal{B}^{(k)};
19
                    if k < K and computational budget not reached then
20
                           New iteration: k = k + 1;
21
```

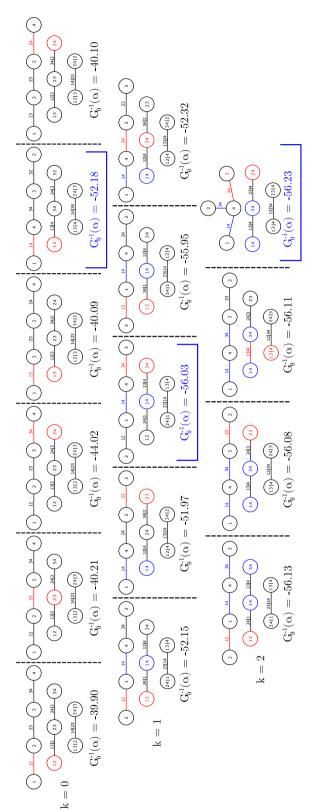


Figure 3: Illustration of the vine structures created during the 3 iterations of the algorithm for the example of Section 3.3.3. The candidate and selected pairs are respectively represented in red and blue. The quantile associated to the selected pair of each iteration is written in blue.

We observe that X_4 appears in all the selected pairs. This is not surprising since X_4 is the most influential variable with the largest coefficient in (21). The algorithm considers D-vines by default, but this is important for the first iterations since most of the pairs are independent. When it is possible, the algorithm creates a vine such as the selected pairs and the candidate pair are in the first trees. For example, the fourth vine at iteration k = 2 with the candidate pair (2, 4) shows a R-vine structure that respects the ranking of the listed pairs. However, the third vine at iteration k = 2 for the candidate pair (1, 3) along with the selected pairs $\{(1, 4), (3, 4) \text{ could respect the ranking and set all the pairs in the first tree altogether. Thus, using Algorithm 3 in Appendix, a valid vine structure is determined by placing the candidate pair (1, 3) in the next tree.$

3.3.4 Computational cost

The number of model evaluations is influenced by several characteristics from the probabilistic model and from the algorithm. Let $|\mathbf{B}|$ be the number of family candidates. The total number of runs is

$$N = |\mathbf{B}| \frac{n}{2} \sum_{k=0}^{K-1} N_k \times (d(d-1) - 2k).$$
 (22)

The sum corresponds to the necessary iterations to determine the influential pairs. The maximum possible cost is if all the pairs are equivalently influential (i.e., K = p = d(d-1)/2), which would be extremely high. The term N_k is the cost from the grid-search quantile minimization at step **2.** of the algorithm. The greater N_k is and the better the exploration of $\Theta_k \cup \Theta_{i,j}$. Because the dimension of Θ_k increases at each iteration k, it is normal that N_k should also increases with k (e.g. $N_k = \gamma^{\beta k}$, where γ and β are constants). Also, the greater n is and the better the quantile estimations. The second term is the cost from the input dimension d which influences the number of candidate pairs Ω_{-k} at each iteration k. Table 1 illustrates the number of quantile estimations in function of the dimension d and the number of influential pairs K. We observe that the computational cost increases significantly faster when K = 1 than when K = 3.

Table 1: Example of the computational cost with d (columns) and K (rows) in function of n for $|\mathbf{B}| = 1$ and $N_k = 10 \times (k+1)^2$.

K/d	3	4	5	6	7	8	9	10
1	30n	60n	100n	150n	210n	280n	360n	450n
2	80n	200n	360n	560n	800n	1080n	1400n	1760n
3	90n	360n	720n	1170n	1710n	2340n	3060n	3870n

Extensions can be implemented to reduce the computational cost such as removing from Ω , the pairs that are not sufficiently improving the minimization.

4 Applications

The previously proposed methodology is applied to a toy example and a real industrial case-study. It is worth to mention that these experiments (and future ones) can be conducted again using the Python library dep-impact [8], in which are encoded all the procedures of estimation and optimization presented here.

4.1 Numerical example

We consider and extend the portfolio example provided and illustrated in Subsection 2 of Appendix C. The numerical model η is now defined by the weighted sum

$$Y = \eta(\mathbf{X}) = -\boldsymbol{\beta}\mathbf{X}^T = -\sum_{j=1}^d \beta_j X_j,$$

where the $\beta = (\beta_1, \dots, \beta_d)$ is a vector of constant weights. The margins of the random vector **X** follow the same generalized Pareto distribution with scale σ and shape parameter ξ . Note that the bivariate example in Subsection 2 of Appendix C considers $\beta = 1$ and the distribution parameters as $\sigma = 1$ and $\xi = 1$. In the following examples, we aim at minimizing the median ($\alpha = 0.5$) of the output distribution. We chose to fix the marginal distribution's parameters at $\sigma = 10$ and $\xi = 0.75$, and we set the constant vector $\boldsymbol{\beta}$ to a base-10 increasing sequence such that $\beta = (10^{1/d}, 10^{2/d}, \dots, 10)$. This choice of weights aims to give more influence to the latest components of **X** on Y. Thus, some pairs of variables should be more important in the minimization of the output quantile, as required by the sparsity constraint. We also took n large enough to estimate the output quantile with correct precision on this particular case-study (i.e. n = 300,000).

For all these experiments the results from the different methods can be compared.

- Method 1: the grid-search approach with an optimized LHS sampling [46] inside $\boldsymbol{\Theta}$ and a random vine structure,
- Method 2: the iterative algorithm with an increasing grid-size of $N_k = 25(k+1)^2$.

Method 1 is established with the same computational budget as Method 2.

4.1.1 Dimension 3

In a three dimensional problem, only three pairs of variables (p = 3) are involved in the dependence structure. The sampling size of $\boldsymbol{\theta}$ in Method 1 is set to 400, which is great enough to explore a three dimensional space. The results are displayed on Figure 4: the estimated quantiles from Method 1 (blue dots) with a convex hull (blue dot line) and the quantile at independence (dark point) are provided. It also highlights the minimum estimated quantiles from Methods 1 and 2 which are respectively represented in blue and red points. We also show in green point, the minimum quantile by considering only the Fréchet-Hoeffding bounds. For each minimum, the 95 % bootstrap confidence intervals is displayed in dot lines.

This low dimensional problem confirms the non-monotonic form of the quantile with the dependence parameter, in particular for the variation of the quantile in function of $\tau_{2,3}$. As expected, the pair $X_2 - X_3$ is more influential on the output quantile due to the large weights on X_2 and X_3 . The minimum values obtained by each method are still lower that the results given by an independent configuration. The minimum using Fréchet-Hoeffding bounds is also provided to show that the minimum is not at the boundary of Θ . Method 1 and 2 have very similar minimum results.

4.1.2 Dimension 10

To illustrate the advantages of the iterative procedure, we now consider d = 10. In this example, we chose to only consider a Gaussian family for the set of pair-copula family candidates. The sampling size for the exploration of $\boldsymbol{\Theta}$ in Method 1 is set to 6, 000. Experimental results are summarized over Figure 5, by displaying the minimum quantiles in function of the iteration *k* of Method 2. The quantile at independence is shown in dark line, the minimum estimated quantile from Method 1 is shown in blue line and the other lines are the minimum quantiles at each iteration of the algorithm, all with their 95% bootstrap confidence interval. We display at each iteration the minimum quantiles of each candidate pair in small dots.

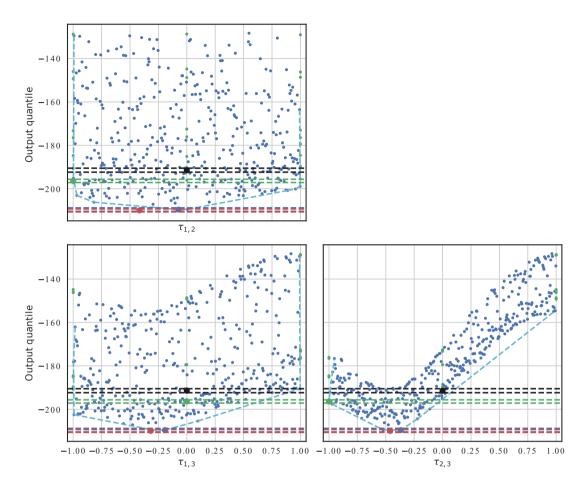


Figure 4: Matrix plot of the output median in function of the Kendall coefficient of each pair. The blue dots represents the estimated quantiles of Method 1. The black point is the quantile at independence and the minimum of Method 1 and 2 are the red and blue points, which are equivalent here. The green point is the the minimum with only Fréchet-Hoeffding bounds. The 95 % bootstrap confidence interval are displayed in dot lines.

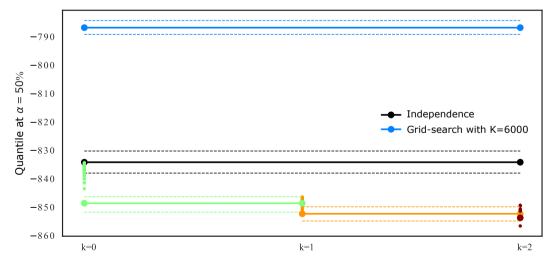


Figure 5: Minimization of the output quantile in function of the iteration *k*. The green and the orange lines respectively show the minimum quantiles obtained from the first and second iterations from Method 2. The quantile at independence is shown in black and the quantile obtained using Method 1 is shown in blue. The 95 % bootstrap confidence intervals are also displayed for the independence and each minimums.

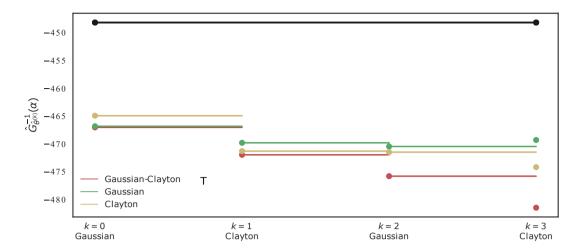


Figure 6: Quantile minimization for different set of family candidates **B**. The black line shows the quantile at independence. The minimum at each iteration for the family candidates sets \mathbf{B}^1 , \mathbf{B}^2 and \mathbf{B}^3 respectively in red, green and yellow.

The minimum result from Method 1 is even higher than the quantile at independence. This is due to the very large number of pairs (p=45) that makes the exploration of $\boldsymbol{\Theta}$ extremely difficult. On the other hand, Method 2 (iterative algorithm) is definitely better and significantly decreases the quantile value even at the first iteration (for only one dependent pair). The results are slightly improved with the iterations. We observe at the last iteration that the results from the candidate pairs are slightly higher than the minimum from the previous iteration. It is due to the choice of N_k which does not increases enough with the iterations to correctly explore $\boldsymbol{\Theta}_k$, which also increases with the iterations.

4.1.3 Using multiple pair-copula family candidates

To show the importance of testing multiple copula families, we consider d=6 and three tests of Method 2 (iterative procedure). The figure 6 shows the minimum from the iterative results using three sets of family candidates: a set of Gaussian and Clayton in red ($\mathbf{B}^1 = \{G, C\}$), Gaussian only in green ($\mathbf{B}^2 = \{G\}$), and Clayton only in yellow ($\mathbf{B}^3 = \{C\}$). We also display below the iteration number, the selected family for \mathbf{B}^1 .

At iteration k = 0, the algorithm with the set \mathbf{B}^1 has selected the Gaussian copula as the selected pair and the result is as expected equivalent as for the set \mathbf{B}^2 . The next iteration, a Clayton copula has been selected for algorithm with the set \mathbf{B}^1 , which slightly improves the minimization compared to the others. The improvement start at iteration k = 2 where the Algorithm with the set \mathbf{B}^1 minimizes more the output quantile than the other sets with only one copula family. At the last iteration, the algorithm with set \mathbf{B}^1 selected a mix between Gaussian an Clayton families. This diversity seems to lead to better results than using only one family for every pairs. Testing multiple families is an interesting feature of the algorithm and is something that cannot be feasible for the grid-search approach. However, the cost for \mathbf{B}^1 is twice larger than for the other methods.

4.2 Industrial Application

4.2.1 Context

We consider an industrial component belonging to a production unit. This component must maintain its integrity even in case of an accidental situation. For this reason, it is subject to a justification procedure by regulation authorities, in order to demonstrate its ability to withstand severe operating conditions. This undesirable event consists in the concomitance of three different factors:

- the potential existence of small and undetectable manufacturing defects;
- the exposition of the structure to an ageing phenomenon harming the material which progressively diminishes its mechanical resistance throughout its lifespan;
- the occurrence of an accidental event generating severe constraints on the structure.

If combined, these three factors might lead to the initiation of a crack within the structure. Since no failure was observed until now, a structural reliability study should be conducted to check the safety of the structure. To do so a thermal-mechanical code $\eta:\mathbb{R}^d\to\mathbb{R}^+$ was used, which calculates the ratio between the resistance and the stress acting on the component during a simulated accident. The numerical model depends on parameters affected by uncertainties quantified throughout numerous mechanical tests. Nevertheless, these experiments are mostly established individually and only few experiments involve simultaneously two parameters.

4.2.2 Probabilistic model

For this problem, we introduce d=6 random variables with predefined marginal distributions $(P_j)_{j=1...d}$. The dependence structure is however unknown. From the 15 pairs of variables, only the dependencies of two pairs are known: one is independent and the other follows a Gumbel copula with parameter 2.27. Therefore, we consider p=13 pairs of variables with unknown dependencies.

Given expert feedbacks, we restricted the exploration space Θ by defining bounds for each pair of variables $(i, j) \in \Omega$ such that

$$T_{c_{i,j}}(\tau_{i,j}^-) \leq \theta_{i,j} \leq T_{c_{i,j}}(\tau_{i,j}^+),$$

where $\tau_{i,j}^-$ and $\tau_{i,j}^+$ are respectively the upper and lower Kendall's correlation coefficient bounds for the dependence of the pair (i,j) and $T_{c_{i,j}}$ is the transformation from Kendall's tau value to the copula parameter for the associated copula $c_{i,j}$. This choice enables to explore only realistic dependence structures. For these experiments we only considered Gaussian copulas.

Remark 1. Note that such bounds can arise from expert knowledge based on a phenomenologic law linking two variables: based on a physical property as the ideal or real gas laws, or the weakest link fracture theory linking material toughness to temperature, a variable is known to be either negatively or positively correlated to another one. Methodologies have been developed in various fields to refine such assessments [14, 67], but were not considered in this research work.

4.2.3 Results

We consider the quantile at $\alpha=0.01$ as a quantity of interest. A first experiment is established with the incomplete probability structure: only the two pairs with a known dependence structure and all others at independence. Two other experiments are established: an exhaustive grid-search approach with a given vine structure and an iterative procedure with a maximum budget equivalent to the grid-search. A grid-size of 1000 is chosen with n=20,000.

The results are displayed in Figure 7. and has the same description as Figure 5. The quantile for the incomplete probability structure is approximately at 1.8. The grid-search and the iterative approaches found dependence structures leading to output quantile values close to 1.2 and 1.1 respectively. The minimum quantile from the iterative procedure is slightly lower than the grid-search approach. The problem dimension is not big enough to create make a significant difference between the methods. However, the resulting dependence structure from the iterative method is greatly simplified with only four pairs of variables, in addition to the already known pair.

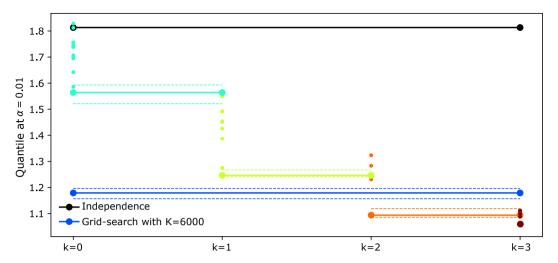


Figure 7: Minimization of the output quantile using a grid-search and the iterative procedure for $\alpha = 1\%$. The description is the same as in Figure 5.

This result highlights the risk of having an incomplete dependence structure in a reliability problem. In this application, the critical limit (*safety margin*) of the considered industrial component is 1. With the incomplete distribution of **X**, the output quantile is very high compared to the critical limit and states a high reliability of the component. Unfortunately, if we consider worst-case dependence structures, the output quantile is significantly minimized and becomes closer to the critical limit. Thus, if the true dependence structure is close to the obtained worst case dependence structure, the risk of over estimating the output quantile can be important.

5 Conclusion and discussion

This article considers the situation where one wishes to detect and determine a penalizing correlation structure between the input random variables of a deterministic computer model, when no information is available, through data, on this correlation structure. The meaning given to the notion of penalization is as follows: a low-order quantile of the output distribution of the computer model must be as small as possible, under the constraint that the marginal distributions of the inputs are known. This work takes a first step towards this goal, by proposing a greedy, heuristic algorithm that explores a set of possible dependencies, taking advantage of the pair-copula construction (PCC) of multivariate probability distributions. Results of experiments conducted on toy and a real models illustrate the good behavior of the procedure: in situations where the monotonicity of the output quantile with respect to the inputs is postulated, a minimum quantile value is obtained using Fréchet-Hoeffding bounds. Nonetheless, it is possible to exhibit situations where the algorithm detect more penalizing dependence structures.

The tools proposed in this paper may be particularly useful in some structural risk and reliability studies, where the inputs typically correspond to independently measured physical parameters, the computer model implements complex structural equations, and the resulting quantile is compared to a safety limit (or a historical quantile value). The correlation structure obtained, associated with the marginal distributions, then allows the generation of scenarios exploring the behavioural limits of the global model. We believe that the methodology proposed here can also be applied to other supervised types of decision-helping studies conducted in the more general context of uncertainty quantification. For instance, many problems of natural resource management and sustainable development (where the correlations between influential parameters are little or poorly known) use such quantiles as key indicators and dimensional quantities of resource extraction [10, 32].

However, it is essential that the correlation structure obtained remains phenomenologically relevant and not biased by the interactions between input variables within the model. In the case of structural reliability problems, the input parameters can be roughly separated between *loads* (forcing conditions, e.g. an injection of cooling water) and *capacities* (structural resistances, e.g. material toughness) [25], which are independent. This independence is known a priori and this constraint must be integrated into the exploration of correlation structures. The use of similar, trivial bounds on the values of correlation coefficients (§ 4.2.2) makes it possible to respect such a constraint.

The methodology presented in this article requires a number of hypotheses and approximations that pave the way for future research. Besides, some perspectives arise from additional technical results.

It would be interesting to improve the statistical estimation of the minimum quantile, given a dependence structure, by checking the hypotheses underlying the convergence results of Theorem 1. Checking and relaxing these hypotheses should be conducted in relation with expert knowledge on the computer model η and, possibly, a numerical exploration of its regularity properties. The grid search estimation strategy promoted in Section 2.3 arises from the lack of information about the convexity and the gradient of $\theta \to G_{\theta}^{-1}(\alpha)$. However, the method remains basic and stochastic recursive algorithms, such as the Robbins-Munro algorithm [54], can be proposed and tested as possibly more powerful (faster) alternatives. Moreover, the empirical quantile was considered for theoretical reasons but can be costly for large α . Another estimator could be consider to reduce the estimation cost. Methods like importance sampling (see for instance [55]) are commonly used to estimate small failure probabilities but its use can be more challenging in the case of quantile estimation.

A significant issue is the computational cost of the exploration of possible dependence structures. Reducing this cost while increasing the completeness of this exploration should be a main concern of future works. Guiding the exploration in the space of conditional bivariate copulas using enriching criteria and possible expert knowledge can facilitate the minimization. The Algorithm 2 can also be improved using nonparametric bootstrap. This would quantify the estimation quality of the selected minimum quantile of each iteration. Note however that a seducing feature of an iterative procedure is the a priori possibility of its adaptation to situations where the computational model η is time-consuming. In such cases, it is likely that Bayesian global optimization methods based on replacing the computer model by a surrogate model (e.g., a kriging-based meta-model) [52] should be explored, keeping in mind that nontrivial conservative correlations – losses of quantile monotonicity– can be due to edge effects (e.g., discontinuities) characterizing the computational model itself. In this sense, mixing dynamic exploration designs with surrogate approaches, in the continuity of the active learning approach introduced in [22, 23], seems a relevant research avenue.

We noticed in our experiments on real case-studies that expert knowledge remains difficult to incorporate otherwise that using association and concordance measures. Precise bounds on association measures between two variables are sometimes complex to obtain from the experts due to the lack of bivariate observations. However it is often possible to derive crude upper and lower bounds for association measures, for instance positive/negative correlation or a weak/strong association, and reducing greatly the exploration of the space of correlations. The guidance provided by such expert knowledge can certainly be improved, depending on the situation considered, by adapting dedicated methodologies evoked in Remark 1. Especially, some of these methodologies are based on representation tools of the features of sampled multivariate distributions, that provide intelligible diagnostics, and participate in expert training, a key methodological requirement for Bayesian analysis [5]. Doing so, these new bounds should be considered as probabilistic rather than deterministic, which would lead to adapt the algorithmic core proposed in the present article.

Finally, another approach to consider could be to address the optimization problem (2) within the more general framework of optimal transport theory, and to take advantage of the many ongoing works in this area. Indeed, the problem (2) can be seen as a multi-marginal optimal transport problem (see [53] for an overview). When d = 2, it corresponds respectively to the classical optimal transport problems of Monge and Kantorovich [66]. However, the multimarginal theory is not as well understood as for the bimarginal case, and developing efficient algorithms for solving this problem remains also a challenging issue [53].

A Proof of the consistency result

The consistency of the estimator $\hat{\boldsymbol{\theta}}$ requires some regularity of the function $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}$. This regularity can be also expressed in term of modulus of increase of the function $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}^{-1}(\alpha)$, on which some useful definitions and connections with the modulus of continuity are reminded.

A.1 Modulus of increase of a cumulative distribution function

Let us recall that a modulus of continuity is any real-extended valued function $\omega : [0, \infty) \mapsto [0, \infty)$ such that $\lim_{\delta \to 0} \omega(x) = \omega(0) = 0$. The function $f : \mathbb{R} \to \mathbb{R}$ admits ω as modulus of continuity if for any $(x, x') \in \mathbb{R}^2$,

$$|f(x)-f(x')|\leq \omega(|x-x'|).$$

Similarly, for some $x \in \mathbb{R}$, the function f admits ω as a *local* modulus of continuity if for any $x' \in \mathbb{R}^2$,

$$|f(x)-f(x')|\leq \omega(|x-x'|).$$

To control the deviation of the empirical quantile in the proof of Proposition 1 further, we consider the modulus of continuity of the quantile functions $G^{-1}:[0,1]\to\mathbb{R}$ where G is a distribution function on \mathbb{R} . The quantile function being an increasing function, the *exact local modulus of continuity* of the quantile function G^{-1} at $\alpha \in (0,1)$ can be defined as

$$\omega_{G^{-1}}(\epsilon,\alpha) := \max \left(G^{-1}(\alpha+\epsilon) - G^{-1}(\alpha), G^{-1}(\alpha) - G^{-1}(\alpha-\epsilon) \right).$$

In the proof of Proposition 1, we note that the continuity of a quantile function G^{-1} can be connected to the increase of the distribution function G (see also for instance Section A in [9]). Using the fact that the distribution function is increasing, we introduce the *local modulus of increase* of the distribution function ϵ_G at $v = G^{-1}(\alpha) \in \mathbb{R}$ as:

$$\epsilon_G(\delta, \gamma) := \min (G(\gamma + \delta) - G(\gamma), G(\gamma) - G(\gamma - \delta)).$$

A.2 Proofs

The estimator $\hat{\boldsymbol{\theta}}$ defined in (12) is an extremum-estimator (see for instance Section 2.1 of [51]). The main ingredient to prove the consistency of this estimator is the uniform convergence in probability of the families of the empirical quantiles $(\hat{G}_{\boldsymbol{\theta}}^{-1}(\alpha))_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}}$ over the family of grids $\boldsymbol{\Theta}_{K_n}$.

Proposition 1. Let Θ_{K_n} be defined as in Theorem 1. Let assume that **B** and **C** are both satisfied. Then, for all $\epsilon > 0$,

$$\mathbb{P}[\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_n}}|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|\geq\epsilon]\xrightarrow{n\to\infty}0.$$

Proof 1 (Proof of Proposition 1). We first make the connection between the local continuity of the quantile function G_{θ}^{-1} and the local increase of the distribution function G_{θ} . According to Assumption **C**, we have that for any $\epsilon \in (0, \max((1 - \alpha), \alpha))$, for any $\delta > 0$ and for any $\theta \in \Theta$,

$$(\star): \left\{ \begin{array}{l} G_{\boldsymbol{\theta}}^{-1}(\alpha+\epsilon) - G_{\boldsymbol{\theta}}^{-1}(\alpha) < \delta \\ G_{\boldsymbol{\theta}}^{-1}(\alpha) - G_{\boldsymbol{\theta}}^{-1}(\alpha-\epsilon) < \delta \end{array} \right. \Longrightarrow \left\{ \begin{array}{l} G_{\boldsymbol{\theta}}\left(G_{\boldsymbol{\theta}}^{-1}(\alpha+\epsilon)\right) < G_{\boldsymbol{\theta}}\left(G_{\boldsymbol{\theta}}^{-1}(\alpha) + \delta\right) \\ G_{\boldsymbol{\theta}}\left(G_{\boldsymbol{\theta}}^{-1}(\alpha) - \delta\right) < G_{\boldsymbol{\theta}}\left(G_{\boldsymbol{\theta}}^{-1}(\alpha-\epsilon)\right) \end{array} \right. .$$

Next, using basic properties of quantile functions (see for instance point ii of Lemma 21.1 in [65]) together with Assumption B, we find that

$$G_{\theta}\left(G_{\theta}^{-1}(\alpha+\epsilon)\right) = \alpha+\epsilon = G_{\theta}\left(G_{\theta}^{-1}(\alpha)\right)+\epsilon$$

and

$$G_{\theta}\left(G_{\theta}^{-1}(\alpha-\epsilon)\right)=\alpha-\epsilon=G_{\theta}\left(G_{\theta}^{-1}(\alpha)\right)-\epsilon.$$

Thus,

$$(\star) \Longrightarrow \begin{cases} G_{\theta} \left(G_{\theta}^{-1}(\alpha) + \delta \right) - G_{\theta} \left(G_{\theta}^{-1}(\alpha) \right) > \epsilon \\ G_{\theta} \left(G_{\theta}^{-1}(\alpha) \right) - G_{\theta} \left(G_{\theta}^{-1}(\alpha) - \delta \right) > \epsilon \end{cases}.$$

We have shown that any $\epsilon \in (0, \max((1 - \alpha), \alpha))$, for any $\delta > 0$ and for any $\theta \in \Theta$,

$$\omega_{G_{\theta}^{-1}}(\epsilon, \alpha) > \delta \Longrightarrow \epsilon_{G_{\theta}}(\delta, G_{\theta}^{-1}(\alpha)) < \epsilon.$$
 (23)

We now prove the proposition. For any $n \ge 1$ and any $\epsilon > 0$, we have

$$\mathbb{P}\left(\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_{n}}}|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|>\epsilon\right) = \mathbb{P}\left(\bigcup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_{n}}}\{|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|>\epsilon\}\right) \\
\leq \sum_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_{n}}}P_{\boldsymbol{\theta}}\left(|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|>\epsilon\right). \tag{24}$$

Let ξ_1, \ldots, ξ_n be n i.i.d. uniform random variables. The uniform empirical distribution function is defined by

$$\mathbb{U}(t) = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\xi_i \leq t} \quad \text{for } 0 \leq t \leq 1.$$

The inverse uniform empirical distribution function is the function

$$\bigcup_{n=0}^{\infty} (u) = \inf\{t \mid \mathbb{G}_n(t) > u\} \quad \text{for } 0 \le u \le 1.$$

The empirical distribution function \hat{G}_{θ} can be rewritten as (see for instance [65]):

$$\widehat{G}_{\boldsymbol{\theta}}(y) \stackrel{\mathcal{L}}{=} \mathbb{U}_n(G_{\boldsymbol{\theta}}(y))$$

and as well for the quantile function,

$$\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha) \stackrel{\mathcal{L}}{=} G_{\boldsymbol{\theta}}^{-1}(\mathbb{U}_n^{-1}(\alpha)).$$

From Inequality (24), we obtain

$$\sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} P_{\boldsymbol{\theta}} \left(|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha) - G_{\boldsymbol{\theta}}^{-1}(\alpha)| > \epsilon \right) = \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} P_{\boldsymbol{\theta}} \left(|G_{\boldsymbol{\theta}}^{-1}(\mathbb{U}_n^{-1}(\alpha)) - G_{\boldsymbol{\theta}}^{-1}(\alpha)| > \epsilon \right)$$
(25)

By definition of the local modulus of continuity $\omega_{G_{\theta}^{-1}}$ of the quantile function G_{θ}^{-1} at α , we have

$$|G_{\theta}^{-1}(\mathbb{U}_{n}^{-1}(\alpha)) - G_{\theta}^{-1}(\alpha)| \le \omega_{G_{\theta}^{-1}}(|\mathbb{U}_{n}^{-1}(\alpha) - \alpha|, \alpha).$$
(26)

Therefore, by replacing (26) in (25) and using (23), we obtain

$$\begin{split} \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} P_{\boldsymbol{\theta}} \left(\left| G_{\boldsymbol{\theta}}^{-1}(\mathbb{U}_n^{-1}(\boldsymbol{\alpha})) - G_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\alpha}) \right| > \epsilon \right) & \leq \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} P_{\boldsymbol{\theta}} \left(\omega_{G_{\boldsymbol{\theta}}^{-1}}(|\mathbb{U}_n^{-1}(\boldsymbol{\alpha}) - \boldsymbol{\alpha}|, \boldsymbol{\alpha}) > \epsilon \right) \\ & \leq \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} P_{\boldsymbol{\theta}} \left(\epsilon_{G_{\boldsymbol{\theta}}}(\boldsymbol{\epsilon}, G_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\alpha})) < |\mathbb{U}_n^{-1}(\boldsymbol{\alpha}) - \boldsymbol{\alpha}| \right). \end{split}$$

Assumption C then yields

$$\mathbb{P}\left(\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_n}}|\hat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|>\epsilon\right)\leq K_n\mathbb{P}\left(|\mathbb{U}_n^{-1}(\alpha)-\alpha|>\underline{\epsilon}_{\boldsymbol{\Theta}}(\epsilon)\right). \tag{27}$$

The DKW inequality [21] gives an upper bound of the probability of an uniform empirical process $\{|U_n(\alpha) - \alpha|\}$. As well for an uniform empirical quantile process $\{|U_n^{-1}(\alpha) - \alpha|\}$ (see for example Section 1.4.1 of [16]), such as $\forall \lambda > 0$:

$$\mathbb{P}(\sup_{\alpha\in[0,1]}|\mathbb{U}_n^{-1}(\alpha)-\alpha|)\geq\lambda)\leq C\exp(-2n\lambda^2).$$

Moreover, [45] proved that one can take C = 2. Therefore, Equation (27) can be bounded using the DKW and

$$\mathbb{P}\left(\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_n}}|\widehat{G}_{\boldsymbol{\theta}}^{-1}(\alpha)-G_{\boldsymbol{\theta}}^{-1}(\alpha)|>\epsilon\right)\leq 2K_n\exp\left[-2n\underline{\epsilon}_{\boldsymbol{\theta}}^2(\epsilon)\right]\xrightarrow{n\to\infty}0$$

since $K_N \lesssim n^{\beta}$.

A second requirement to get the consistency of the extremum estimator is the regularity of $\theta \mapsto G_{\theta}^{-1}(\alpha)$. This is shown in the next proposition.

Proposition 2. Under Assumptions A, B and C, the function

$$\Theta \longrightarrow \overline{Im(\eta)}$$

$$\boldsymbol{\theta} \longrightarrow G_{\boldsymbol{\theta}}^{-1}(\alpha)$$

is continuous in θ over Θ .

Proof 2 (Proof of Proposition 2). According to Assumption **A**, for any $\theta \in \Theta$, the distribution P_{θ} admits a density function f_{θ} with respect to the Lebesgue measure on \mathbb{R}^d such that

$$f_{\boldsymbol{\theta}}(x_1,\ldots,x_d) = c_{\boldsymbol{\theta}}(F_1(x_1,\ldots,x_d))f_1(x_1)\ldots f_d(x_d),$$

where f_j is the marginal density function of X_j , for j = 1, ..., d and the Lebesgue measure on \mathbb{R} . Moreover, for any $\mathbf{x} \in \mathbb{R}^d$, the function $\mathbf{\theta} \to f_{\mathbf{\theta}}(x)$ is continuous in $\mathbf{\theta}$ over $\mathbf{\Theta}$.

The domain $\Theta \times [0, 1]^p$ is a compact set and according to Assumption A, there exists a constant \bar{c} such that $\forall (\theta, \mathbf{u}) \in \Theta \times [0, 1]^d$, $c_{\theta}(\mathbf{u}) \leq \bar{c}$. Consequently, we have

$$|f_{\theta}(x_1, \dots, x_d)| \le \bar{c} \prod_{i=1}^d f_i(x_i).$$
 (28)

For $\theta \in \Theta$ and for any h > 0, we denote $y_h = G_{\theta+h}^{-1}(\alpha)$. According to Assumption **B** we have $\alpha = G_{\theta+h}(y_h)$ and thus,

$$G_{\theta}^{-1}(\alpha) - G_{\theta+h}^{-1}(\alpha) = G_{\theta}^{-1}(\alpha) - y_h$$

$$= G_{\theta}^{-1}(G_{\theta+h}(y_h)) - G_{\theta}^{-1}(G_{\theta}(y_h))$$
(29)

Now, using Assumption C, we have that G_{θ} is strictly increasing in the neighborhood of $G_{\theta}^{-1}(\alpha)$ and thus G_{θ}^{-1} is continuous in the neighborhood of α . Note that

$$|G_{\boldsymbol{\theta}+h}(y_h) - G_{\boldsymbol{\theta}}(y_h)| = \left| \int_{\mathbb{R}^d} \mathbb{1}_{\eta(\boldsymbol{x}) \leq y_h} dF_{\boldsymbol{\theta}+h}(\boldsymbol{x}) - \int_{\mathbb{R}^d} \mathbb{1}_{\eta(\boldsymbol{x}) \leq y_h} dF_{\boldsymbol{\theta}}(\boldsymbol{x}) \right|$$

$$= \left| \int_{\mathbb{R}^d} \left[f_{\boldsymbol{\theta}+h}(\boldsymbol{x}) - f_{\boldsymbol{\theta}}(\boldsymbol{x}) \right] \mathbb{1}_{\eta(\boldsymbol{x}) \leq y_h} d\lambda(\boldsymbol{x}) \right|$$

$$\leq \int_{\mathbb{R}^d} |f_{\boldsymbol{\theta}+h}(\boldsymbol{x}) - f_{\boldsymbol{\theta}}(\boldsymbol{x})| d\lambda(\boldsymbol{x})$$

We then apply a standard dominated convergence theorem using (28) to get that

$$G_{\boldsymbol{\theta}+h}(y_h) - G_{\boldsymbol{\theta}}(y_h) \xrightarrow{h \to 0} 0.$$

This, with (29) and with the continuity of $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}$, shows that

$$G_{\boldsymbol{\theta}}^{-1}(\alpha) - G_{\boldsymbol{\theta}+h}^{-1}(\alpha) \xrightarrow{h \to 0} 0.$$

We are now in position to prove Theorem 1.

Proof 3 (Proof of Theorem 1). *Under Assumptions* **B** *and* **C**, *Proposition 1 directly gives that for any* $\varepsilon > 0$,

$$P\left(\left|\inf_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_n}}\widehat{G}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\alpha})-\inf_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_{K_n}}G_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\alpha})\right|>\varepsilon\right)\xrightarrow{n\to\infty}0$$

which means that

$$P\left(\left|\widehat{G}_{\widehat{\boldsymbol{\theta}}}^{-1}(\alpha) - \inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} G_{\boldsymbol{\theta}}^{-1}(\alpha)\right| > \varepsilon\right) \xrightarrow{n \to \infty} 0. \tag{30}$$

If Assumption **A** is also satisfied, Proposition 2 together with (11) give that $\inf_{\theta \in \Theta_{K_n}} G_{\theta}^{-1}(\alpha)$ tends to $\inf_{\theta \in \Theta} G_{\theta}^{-1}(\alpha)$ as n tends to infinity. Thus

$$\inf_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_{K_n}} G_{\boldsymbol{\theta}}^{-1}(\alpha) \xrightarrow{n \to \infty} G_{\boldsymbol{C}}^{-1}(\alpha) = G_{\boldsymbol{\theta}_{\boldsymbol{C}}^*}^{-1}(\alpha)$$
 (31)

We then derive (13) from (30) and (31).

We now assume that Assumption **D** is also satisfied. Let $\boldsymbol{\theta}^*$ be the unique minimizer of $\boldsymbol{\theta} \mapsto G_{\boldsymbol{\theta}}^{-1}(\alpha)$. Let h > 0 such that $B(\boldsymbol{\theta}^*, h)^c := \{\boldsymbol{\theta} \in \boldsymbol{\Theta} : \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \ge \varepsilon\}$ is not empty. According to Proposition 2 and using the fact that $\boldsymbol{\Theta}$ is compact, we have

$$\sup_{\boldsymbol{\theta} \in B(\boldsymbol{\theta}^*, h)^c} |G_{\boldsymbol{\theta}}^{-1}(\alpha) - G_{\boldsymbol{\theta}^*}^{-1}(\alpha)| > 0.$$
(32)

Consequently, for any $\forall h > 0$ small enough, there exists $\epsilon > 0$ such that

$$|G_{\boldsymbol{\theta}}^{-1}(\alpha) - G_{\boldsymbol{\theta}^*}^{-1}(\alpha)| \le \epsilon \Longrightarrow |\boldsymbol{\theta} - \boldsymbol{\theta}^*| < h \tag{33}$$

Let h > 0 and take ϵ such that (33) is satisfied for h. According to Proposition 1, $\widehat{G}_{\hat{\theta}}^{-1}(\alpha) - G_{\hat{\theta}}^{-1}(\alpha)$ tends to zero in probability as n tends to infinity. This, with (13), shows that

$$P\left(\left|G_{\hat{\theta}}^{-1}(\alpha)-G_{\theta^*}^{-1}(\alpha)\right|>\varepsilon\right)\xrightarrow{n\to\infty}0.$$

We conclude using (33).

B Vine copulas

B.1 Definition

A vine model describes a d-dimensional pair-copula construction (PCC) and is a sequence of linked trees where the nodes and edges correspond to the d(d-1)/2 pair-copulas. According to Definition 1 from [6], a vine structure is composed of d-1 trees T_1, \ldots, T_{d-1} with several conditions.

Definition 1 (R-vine). The sequence $\mathcal{V} = (T_1, \dots, T_{d-1})$ is an R-vine on n elements if

1. T_1 is a tree with nodes $N_1 = \{1, ..., d\}$ and a set of edges denoted E_1 .

- 2. For i = 2, ..., d 1, T_i is a tree with nodes $N_i = E_{i-1}$ and edges set E_i .
- 3. For $i = \ldots, d-1$ and $\{a, b\} \in E_i$ with $a = \{a_1, a_2\}$ and $b = \{b_1, b_2\}$ it must hold that $\#(a \cap b) = 1$ (proximity condition).

Each tree T_i is composed of d-i+1 nodes which are linked by d-i edges for $i=1,\ldots,d-1$. A node in a tree T_i must be an edge in the tree T_{i-1} , for $i=2,\ldots,d-1$. Two nodes in a tree T_i can be joined if their respective edges in tree T_{i-1} share a common node, for $i=2,\ldots,d-1$. The proximity condition, suggests that two nodes connected by an edge should share one variable from the conditioned set. The *conditioning* set and *conditioned set* are defined in Definition 2 along with the *complete union*. The complete union of an edge e is a set of all unique variables contained in e.

Definition 2 (Complete union, conditioning and conditioned sets of an edge). *Let* A_e *be the complete union of an edge* $e = \{a, b\} \in E_k$ *in a tree* T_k *of a regular vine* V,

$$A_e = \{v \in N_1 | \exists e_i \in E_i, i = 1, \dots, k-1, such that v \in e_i \in \dots \in e_{k-1} \in e\}.$$

The conditioning set associated with edge $e = \{a, b\}$ is $D(e) := A_a \cap A_b$ and the conditioned sets associated with edge e are $i(e) := A_a \setminus D(e)$ and $j(e) := A_b \setminus D(e)$. Here, $A \setminus B := A \cap B^c$ and B^c is the complement of B.

The conditioned and conditioning sets of an edge $e = \{a, b\}$ are respectively the symmetric difference and the intersection of the complete unions of a and b. The conditioned and conditioning sets of all edges of \mathcal{V} are collected in a set called *constraint set*. Each element of this set is composed of a pair of indices corresponding to the conditioned set and a set containing indices corresponding to the conditioning set, as shown in Definition 3.

Definition 3 (Constraint set). *The constrain set for* V *is a set:*

$$\mathcal{CV} = \{(\{i(e), j(e)\}, D_e) | e \in E_i, e = \{a, b\}, i = 1, \dots, d-1\}$$

The pair-copula in the first tree characterize pairwise unconditional dependencies, while the pair-copula in higher order trees model the conditional dependency between two variables given a set of variables. The number of conditioning variables grows with the tree order. Note that a PCC where all trees have a path-like structure define the D-vine subclass while the star-like structures correspond to C-vine subclass. All other vine structures are called regular vines (R-vines) [6].

We illustrate the concept of a vine model with a d = 5 dimensional example. For clarity reasons, we use the same simplifications as in Section 3.2 which consider for instance $f_1 = f_1(x_1)$, $f_2 = f_2(x_2)$ and so on for higher order and conditioning. One possible PCC can be written for this 5-dimensional configuration:

$$\begin{split} f(x_1, x_2, x_3, x_4, x_5) &= f_1 \cdot f_2 \cdot f_3 \cdot f_4 \cdot f_5 \text{ (margins)} \\ &\text{(unconditional pairs)} \times c_{12} \cdot c_{35} \cdot c_{34} \cdot c_{24} \\ &\text{(1st conditional pair)} \times c_{14|3} \cdot c_{23|4} \cdot c_{45|3} \\ &\text{(2nd conditional pair)} \times c_{15|34} \cdot c_{25|34} \\ &\text{(3rd conditional pair)} \times c_{12|345}. \end{split}$$

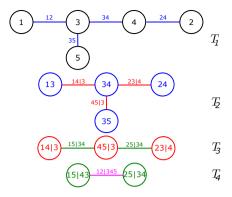


Figure 8: R-vine structure for d = 5.

The vine structure associated to (34) is illustrated in Figure 8. This graphical model considerably simplify the understanding and we observe that this model is a R-vine because there is no specific constraints on the trees.

A re-labeling of the variables can lead to a large number of different PCC. [49] calculated the number of possible vine structures with the dimension d and shows that it becomes extremely large for high dimension problems. We illustrate below, using the same d = 5 dimensional example, two other PCC densities:

$$f_{D} = f_{1} \cdot f_{2} \cdot f_{3} \cdot f_{4} \cdot f_{5}$$

$$\times c_{12} \cdot c_{23} \cdot c_{34} \cdot c_{45}$$

$$\times c_{13|2} \cdot c_{24|3} \cdot c_{35|4}$$

$$\times c_{14|23} \cdot c_{25|34}$$

$$\times c_{15|234}$$

$$f_{C} = f_{1} \cdot f_{2} \cdot f_{3} \cdot f_{4} \cdot f_{5}$$

$$\times c_{12} \cdot c_{13} \cdot c_{14} \cdot c_{15}$$

$$\times c_{23|1} \cdot c_{24|1} \cdot c_{25|1}$$

$$\times c_{34|12} \cdot c_{35|12}$$

$$\times c_{45|123}$$
(36)

where (35) and (36) respectively correspond to D-vine and C-vine structures and are represented in Figures 9a and 9b. As we can see in these examples, the D-vine have a constraint on each tree that gives a path-like arrangement of the nodes. The C-vine on the other hand only has one node connected to all others for each tree.

An efficient way to store the information of a vine structure is proposed in [48] and is called a R-vine array. The approach uses the specification of a lower triangular matrix where the entries belong to $1, \ldots, d$. Such matrix representation allows to directly derive the tree structure (or equivalently the associated PCC distribution). For more details, see [48].

B.2 Generating R-vine from an indexed list of pairs

The iterative procedure proposed in Section 3.3, described by Algorithm 2, minimizes the output quantile by iteratively determining the pairs of variables that influences the most the quantile minimization. At each iteration of the algorithm (step 1.a), a new vine structure is created by considering the list of influential pairs. The specificity of this vine creation is to consider the ranking of the list by placing the most influential pairs in the first trees of the R-vine. Thus, we describe in this section how to generate vine structure with the constraint of a given list of indexed pairs to fill in the structure.

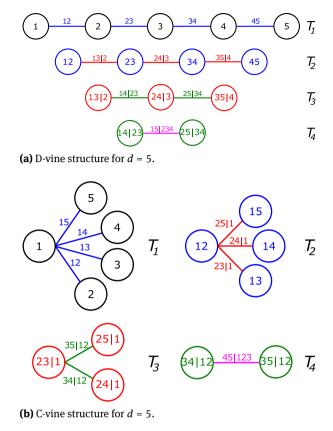


Figure 9: D-vine and C-vine structures for d = 5.

B.2.1 The algorithm

We consider the same notation as in Algorithm 2. Creating a vine structure from a given indexed list of pairs Ω_k is not straightforward. The difficulties come from respecting the ranking of Ω_k and the respect of the R-vine conditions. Indeed, the pairs cannot be append in the structure easily. The vine structure must respect these conditions, which can be sometime very restrictive. The procedure we proposed is detailed by the pseudocode of Algorithm 3 and can be greatly simplified in these few key steps:

- 1. fill \mathcal{V} with the list Ω_k ,
- 2. fill \mathcal{V} with a permutation of Ω_{-k} ,
- 3. if V is not a R-vine, then permute Ω_k and restart at step 1.

In step 1 and 2, the *filling* procedure, detailed in Algorithm 4, successively *adds* the pairs of a list in the trees of a vine structure. Adding a pair (i, j) in a tree T_l associates (i, j) with the conditioned set and determine a possible conditioning set D from the previous tree such as a possible edge is i, j|D.

In step 2, because the ordering of Ω_{-k} is not important in the filling of \mathcal{V} , the permutation of Ω_{-k} aims at finding a ranking such as \mathcal{V} leads to a R-vine.

In step 3, when the previous step did not succeeded and the resulting \mathcal{V} is not a R-vine structure, then the ranking of Ω_k is not possible and must be changed. The permutation of some elements of Ω_k must be done such as the ranking of the most influential pairs remains as close as possible to the initial one.

Algorithm 3: Generating a vine structure from a given list of indexed pairs Ω_k

```
Data: \Omega_k, d
    Result: A vine structure \mathcal{V}.
 1 \Omega_k^{init} = \Omega_k;
 2 k = 1;
 3 do
         /* initialize \mathcal V with a first empty tree
                                                                                                                                                   */
         N_1 = (1, \ldots, d);
         E_1 = ();
 5
         \mathcal{V} = ((N_1, E_1));
         /* filling {\mathcal V} with the list of selected pairs \Omega_k
         V = \text{Fill}(V, \Omega_k, d);
 7
                                                                                                                           // See Algorithm 4
         /* determining a permutation of \Omega_{-k} that fills {\mathcal V}
         for \Omega_{-k}^{\pi} \in \pi(\Omega_{-k}) do
 8
               /* filling {\mathcal V} with the candidate pairs \Omega^\pi_{-k}
               V_{\pi} = \text{Fill}(V, \Omega_{-k}^{\pi}, d);
 9
                                                                                                                           // See Algorithm 4
               if V_{\pi} is a R-vine then
10
                    /* a permutation worked \rightarrow we quit the loop
                    break
11
         \mathcal{V} = \mathcal{V}_{\pi};
12
         if \mathcal{V} is not a R-vine then
13
               /* filling did not work 
ightarrow permute initial list \Omega_k^{init}
               Get \Omega_k by inverting pairs of (\Omega_k^{init};
14
              k = k + 1:
15
16 while \mathcal{V} is not a R-vine:
```

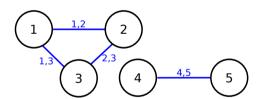


Figure 10: Example: first tree of a non valid vine structure for d = 5 that does lead to a single connected tree.

B.2.2 Example

For illustration, let's create a d=5 dimensional vine structure with the given list of pairs $\Omega_k=((1,2),(1,3),(2,3),(4,5),(2,4),(1,5))$ using Algorithm 3. Using the original list Ω_k , the Fill function may fail at line 7 of Algorithm 3, and more precisely, at line 15 of Algorithm 4. Indeed, the first tree of $\mathcal V$ does not validate the R-vine conditions. The tree is illustrated in Figure 10 and as we can see, the nodes are not all connected into one single tree. Therefore, we permuted the list Ω_k by exchanging the pairs (2,4) and (4,5), as shown in Figure 11. This permutation now leads to a vine structure that respects the new ranked list $\Omega_k = ((1,2),(1,3),(2,3),(2,4),(4,5),(1,5))$.

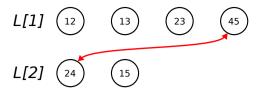


Figure 11: Example: exchange of elements of Ω_k in order to lead to a valid vine structure.

Algorithm 4: Filling a vine structure with a given list

```
1 function Fill (\mathcal{V}, \Omega_k, d):
        /* \mathcal{V}: an incomplete vine structure ,
        /* \Omega_k:, a list of indexed pairs
        /* d: the input dimension.
        l = |\mathcal{V}|;
 2
                                                                                            // number of existing trees
        (T_1,\ldots,T_l)=\mathcal{V};
 3
        k = |T_l|;
                                                                            // number of existing nodes in last tree
 4
        /* loop over the list of pairs
        for (i, j) \in \Omega_k do
 5
            D = \emptyset;
 6
            if l \ge 2 then
 7
                 /st conditioning set is only computed from T_2
                 D = \text{FindConditioningSet}((i, j), N_{l-1});
                                                                                                       // See Algorithm 5
 8
                 if D = \emptyset then
                      /* no conditioning set found \rightarrow not possible
                                                                                                                            */
                      return False
10
             E_1 = E_1 \cup i, j|D;
                                                                                                    // add new edge in E_l
11
             T_1 = (N_1, E_1);
                                                                                                   // update current tree
12
            \mathcal{V} = (T_1, \ldots, T_l);
13
            if k \ge d - l then
14
                 /* if tree T_l is complete
                 if \mathcal{V} does not fulfill the R-vine conditions then
15
                      /* the vine structure {\mathcal V} is not valid
                      return False
16
                 k = 1;
17
                 l = l + 1;
18
                 N_{l} = E_{l-1};
                                                           // nodes of next tree are the edges of previous tree
19
             else
20
                 k = k + 1;
21
        return \mathcal{V}
22
```

Algorithm 5: Gets the conditioning set of a given conditioned set

```
1 function FindConditioningSet((i, j), N_-):
       /*(i,j): the conditioned set,
       /* \emph{N}_{-}: list of nodes from the previous tree.
       D = \emptyset;
2
       for a, b \in N_-, with a \neq b do
3
           if i \in a and j \in b then
4
                if j \notin A_a and i \notin A_b then
5
                     /* See Definition 2
                    D = A_a \cap A_h;
6
                    break;
7
       return D
8
```

C Supplementary material on copulas

This supplementary material is dedicated to a preliminary exploration of the influence of copula structure on the behavior of the worst quantile, illustrated with toy examples. Especially, while it could be expected that $G_{\theta}^{-1}(\alpha)$ is a monotonic function with θ , and that the minimum can be reached for a trivial copula (i.e., reaching the Fréchet-Hoeffding bounds). Our experiments show that this behavior is not systematic.

C.1 About the copula choice

One of the most common approaches to model the dependence between random variables is to assume linear correlations feeding a Gaussian copula. In this case, the problem is reduced by determining the correlation matrix of **X** that minimizes $G_{\theta}^{-1}(\alpha)$. However, the positive semi-definite constraint on the correlation matrix makes the exploration difficult and the minimization harder when the problem dimension increases. Moreover, such a Gaussian assumption is very restrictive and is inappropriate for simulating heavy tail dependencies [44]. Still in this elliptical configuration, the t-copulas [19] can be used to counterpart these problems. Nevertheless, tail dependencies are symmetric and with equal strengths for each pair of variables. Another alternative is to consider multivariate Archimedean copulas [47] which are great tools to describe asymmetric tail dependencies. However, only one parameter governs the strength of the dependence among all the pairs, which is very restrictive and not flexible in high dimension. For a same correlation measure between two random variables, multiple copulas can be fitted and lead to a different distribution of Y.

It is clear that the copula choice of \mathbf{X} has a strong impact on the distribution of Y (see for instance [60]). Therefore, various copula types should be tested to determine the most conservative configuration. In the following, we may consider a flexible approach setting by modeling the input multivariate distribution using regular vine copulas (R-vines). The necessary basics of R-vines are introduced in Section 3 of the article and detailed in its Appendix B.

C.2 About the monotony of the quantile

For many simple case studies case studies, the worst quantile is reached for perfect dependencies (Fréchet-Hoeffding bounds). More generally, when the function has a monotonic behavior with respect to many variables, it is likely that the minimum output quantile is reached at the boundary of Θ . This phenomenon is observed for various physical systems.

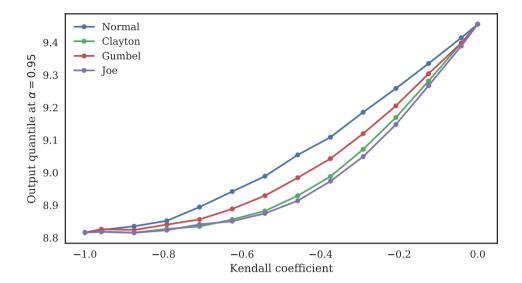


Figure 12: Variation of the quantile of the overflow distribution with the Kendall coefficient τ for $\alpha = 95\%$ and different copula families (Gaussian, Clayton, Gumbel and Joe).

To illustrate this phenomenon, we consider a simplified academic model that simulates the overflow of a river over a dike that protects industrial facilities. The river overflow *S* is described by

$$S = H_d + C_b - Z_v - H \qquad \text{with} \qquad H = \left(\frac{Q}{BK_s\sqrt{\frac{Z_m - Z_v}{L}}}\right)^{0.6}, \tag{37}$$

such as, when S < 0, a flooding occurs. The involved parameters of (37) are physical characteristics of the river and the dike (e.g., flow rate, height of the dike) which are described by random variables with known marginal distributions. See [36] for more information. For a given risk α , we aim at quantifying the associated overflow's height describe by the α -quantile of S. We extend this model by supposing that the friction (Strickler-Manning) coefficient K_S and the maximal annual flow rate Q are dependent with an unknown dependence structure. To show the influence of a possible correlation between K_S and Q on the quantile of S, we describe their dependence structure with multiple copula families.

Figure 12 shows the variation of the estimated quantile of S (with a large sample size) in function of the Kendall coefficient τ between K_S and Q for different copula families. We observe different slopes of variation for the different copula families, with lower quantile values for the copulas with heavy tail dependencies (i.e., Clayton, Joe). At independence ($\tau = 0$) and for the counter-monotonic configuration ($\tau = -1$),the quantile values of these families are obviously equivalent. This variation is slight and the quantile is still above zero, but this shows how the dependencies can influence the results of a reliability problem. This illustration shows that the minimum is reached at the boundary of the exploration space, where the two variables are perfectly correlated.

We can take advantage of this observation to speed up the algorithms presented in the next sections by exploring only the boundaries of Θ . However, assuming that the minimum is reached on the boundary of Θ is a strong assumption that can be unsatisfied in some applications. See Fallacy 3 of [24] for a highlight of this pitfall.

To illustrate this statement, we now give a counter example in the bidimensional setting. We assume uniform marginal distributions for the input such that $X_1 \sim \mathcal{U}(-3,1)$ and $X_2 \sim \mathcal{U}(-1,3)$, and we consider the model function

$$\eta(x_1, x_2) = 0.58x_1^2x_2^2 - x_1x_2 - x_1^2 - x_2^2.$$
(38)

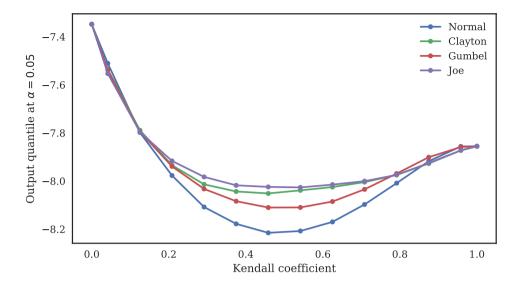


Figure 13: Variation of the output quantile with the Kendall coefficient τ for $\alpha = 5\%$ and different copula families (Gaussian, Clayton, Gumbel and Joe).

The same experience as for Figure 12 is established and the results are shown in Figure 13. The slopes of the quantile estimations with the Kendall coefficient, for each copula families, are quite different than the results of Figure 12. We observe that the quantile is not monotonic with the Kendall coefficient and its minimum is not reached at the boundary, but for $\tau \approx 0.5$. Moreover, the Gaussian copula is the family that minimizes the most the quantile. It shows that copulas with tail dependencies are not always the most penalizing.

A second example, inspired from Example 6 of [24], also shows that the worst case dependence structure in an additive problem is not necessary for perfectly correlated variables. We consider a simple portfolio optimization problem with two random variables X_1 and X_2 with generalized Pareto distributions such as $F_1(x) = F_2(x) = \frac{x}{1+x}$. We aim at maximizing the profit of the portfolio, which is equivalent as minimizing the following additive model function

$$\eta(X_1, X_2) = -(X_1 + 10X_2). \tag{39}$$

We consider the median ($\alpha=0.5$) of the output as an efficiency measure. Figure 14 shows the output median in function of the Kendall coefficient τ between X_1 and X_2 . Just like the previous example, we observe a non-monotonic slope of the median in function of τ . The variation can be significant and the minimum is obtained at $\tau\approx0.53$ for the heavy tail copula families (i.e., Clayton and Joe). The phenomenon can be explained by the marginal distributions of the random variables, which are close Pareto distributions. A large correlation seems to diminish the influence of the tails, which gives a higher quantile value. This explains why the minimum is obtained for a dependence structure other that independence or the perfect dependence.

Therefore, these examples show that the worst quantile can be reached for other configurations than the perfect dependencies.

Acknowledgements: The authors express their thanks to Gérard Biau (Sorbonne Université), Bertrand Iooss and Roman Sueur (EDF R&D) for their useful comments and advices during the writing of this article. They would also like to thank an anonymous reviewer, whose comments contributed significantly to improving the clarity of this article. This work was funded by the French National Agency for Research (ANR) and EDF Research & Development.

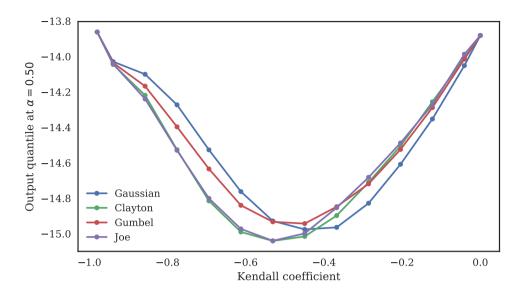


Figure 14: Variation of the portfolio median with the Kendall coefficient τ for different copula families.

References

- [1] Aas, K., C. Czado, A. Frigessi and H. Bakken (2009). Pair-copula constructions of multiple dependence. *Insurance Math. Econom.* 44(2), 182–198.
- [2] Agrawal, S., Y. Ding, A. Saberi and Y. Ye (2012). Price of correlations in stochastic optimization. Oper. Res. 60(1), 150–162.
- [3] Bayarri, M., J. O. Berger, R. Paulo, J. Sacks, J. Cafeo, J. Cavendish, C. Lin and J. Tu (2007). A framework for validation of computer models. *Technometrics* 49(2), 138–154.
- [4] Beaudoin, D. and L. Lakhal-Chaieb (2008). Archimedean copula model selection under dependent truncation. *Stat. Med.* 27(22), 4440–4454.
- [5] Bedford, T., J. Quigley and L. Walls (2006). Expert elicitation for reliable system design. Statist. Sci. 21(4), 428–450.
- [6] Bedford, T. and R. M. Cooke (2001). Probability density decomposition for conditionally dependent random variables modeled by vines. *Ann. Math. Artif. Intell.* 32(1), 245–268.
- [7] Bedford, T. and R. M. Cooke (2001). Vines: a new graphical model for dependent random variables. *Ann. Statist. 30*(4), 1031–1068.
- [8] Benoumechiara N. (2017). dep-impact: A python package for conservative estimation of an output quantity of interest in reliabilty problems. Available at https://github.com/nazben/dep-impact.
- [9] Bobkov,S. and M. Ledoux (2019). *One-dimensional Empirical Measures, Order Statistics, and Kantorovich Transport Distances*. American Mathematical Society, Providence.
- [10] Brabbia C., M.E. Conti and E. Tiezzi (2006). *Management of Natural Ressources, Sustainable Development and Ecological Hazards*. WIT Press, Southampton.
- [11] Brechmann E. C., K. Aas and C. Czado (2012). Truncated regular vines in high dimensions with application to financial data. *Canad. J. Statist.* 40(1), 68–85.
- [12] Chazal, F., P. Massart and B. Michel (2016). Rates of convergence for robust geometric inference. *Electron. J. Stat.* 10(2), 2243–2286.
- [13] Cherubini, U., E. Luciano, W. Vecchiato (2004). Copula Methods in Finance. Wiley, Chichester.
- [14] Clemen, R. T. and T. Reilly (1999). Correlations and copulas for decision and risk analysis. Manag. Sci. 45(2), 208-224.
- [15] National Research Council (2012). Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation and Uncertainty Quantification. National Academies Press, Washington.
- [16] Csorgo M. (1983). Quantile Processes with Statistical Applications. SIAM, Philadelphia PA.
- [17] Czado C. (2010). Pair-copula constructions of multivariate copulas. In P. Jaworski, F. Durante, W. K. Härdle, T. Rychlik (Eds.), *Copula Theory and its Applications*, pp. 93–109. Springer, New York.
- [18] de Rocquigny, E., N. Devictor and S. Tarantola (2008). *Uncertainty in Industrial Practice: a Guide to Quantitative Uncertainty Management.* John Wiley & Sons, Chichester.
- [19] Demarta, S. and A. J. McNeil (2005). The t copula and related copulas. Int. Stat. Rev. 73(1), 111–129.
- [20] Dissmann, J., E. C. Brechmann, C. Czado and D. Kurowicka (2013). Selecting and estimating regular vine copulae and application to financial returns. *Comput. Statist. Data Anal.* 59, 52–69.

- [21] Dvoretzky, A., J. Kiefer and J. Wolfowitz (1956). Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Ann. Math. Statist.* 27(3), 642–669.
- [22] Echard, B., N. Gayton and M. Lemaire (2011). AK-MCS: an active learning reliability method combining kriging and Monte Carlo simulation. *Struct. Saf. 33*(2), 145–154.
- [23] Echard, B., N. Gayton, M. Lemaire, N. Relun (2013). A combined importance sampling and kriging reliability method for small failure probabilities with time-demanding numerical models. *Reliab. Eng. Syst. Saf. 111*, 232–240.
- [24] Embrechts, P., A. J. McNeil and D. Straumann (2002). Correlation and dependence in risk management: properties and pit-falls. In M. A. H. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*, pp. 176–223. Cambridge University Press.
- [25] Estes, A. C. and D. M. Frangopol (2005). Load rating versus reliability analysis. J. Struct. Eng. 131(5), 843-847.
- [26] Evans, L. C. and R. F. Gariepyl (2015). Measure Theory and Fine Properties of Function. CRC Press, Boca Raton FL.
- [27] Fréchet M. (1951). Sur les tableaux de correlation dont les marges sont données. Ann. Univ. Lyon 9, 53-77.
- [28] Frees, E. W. and E. A. Valdez (1998). Understanding relationships using copulas. N. Am. Actuar. J. 2(1), 1-25.
- [29] Goda K. (2010). Statistical modeling of joint probability distribution using copula: application to peak and permanent displacement seismic demands. *Struct. Saf.* 32(2), 112–123.
- [30] Grigoriu, M. and C. Turkstra (1979). Safety of structural systems with correlated resistances. *Appl. Math. Model. 3*(2), 130–136.
- [31] Gruber, L. and C. Czado (2015). Sequential bayesian model selection of regular vine copulas. Bayesian Anal. 10(4), 937-963.
- [32] Guntenspergen G. R. (2014). Application of Threshold Concepts in Natural Resource Decision Making. Springer, New York.
- [33] Haff. I. H. (2016). How to select a good vine. International FocuStat Workshop on Focused Information Criteria and Related Themes. Available at https://www.mn.uio.no/math/english/research/projects/focustat/workshops%20and% 20conference/workshop-2016/vineselection.pdf
- [34] Helton J. C. (2011). Quantification of margins and uncertainties: conceptual and computational basis. *Reliab. Eng. Syst. Saf.* 96(9), 976–1013.
- [35] Hoeffding W.(1940). Scale-invariant correlation theory. In N.I. Fisher and P. K. Sen (Eds.), *The Collected Works of Wassily Hoeffding*, pp. 57–107. Springer, New York.
- [36] Iooss, B. and P. Lemaître (2015). A review on global sensitivity analysis methods. In G. Dellino and C. Meloni (Eds.), *Uncertainty Management in Simulation-Optimization of Complex Systems*, pp. 101–122, Springer, Boston.
- [37] Jiang, C., W. Zhang, X. Han, B. Y. Ni and L. J. Song (2015). A vine-copula-based reliability analysis method for structures with multidimensional correlation. *J. Mech. Des.* 137(6), 13 pages.
- [38] Joe H. (1994). Multivariate extreme-value distributions with applications to environmental data. *Canad. J. Statist.* 22(1), 47–64.
- [39] Joe H. (1996). Families of m-variate distributions with given margins and m (m-1)/2 bivariate dependence parameters. In L. Rüschendorf, B. Schweizer and M. D.Taylor (Eds.), *Distributions with Fixed Marginals and Related Topics*, pp. 120–141. IMS, Hayward
- [40] Karlin, S., K. Arrow and H. Scarf (1958). A min-max solution of an inventory problem. In J. K. Arrow, S. Karlin and H. Scarf (Eds.), Studies in the International Theory of Inventory and Productions, pp. 201–209. Standford University Press, Redwood.
- [41] Kendall M. G. (1938). A new measure of rank correlation. Biometrika. 30(1/2), 81–93.
- [42] Kurowicka D. (2011). Optimal truncation of vines. In D. Kurowicka and H. Joe (Eds.), *Dependence Modeling: Vine Copula Handbook*, pp. 233–247. World Scientific Publishing Co., Singapore.
- [43] Lemaire, M., A. Chateauneuf and J.-C. Mitteau (2010). Structural Reliability. Wiley, Hoboken NJ.
- [44] Malevergne, Y., D. Sornette (2003). Testing the Gaussian copula hypothesis for financial assets dependences. *Quant. Finance* 3(4), 231–250.
- [45] Massart P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. Ann. Probab. 18(3), 1269-1283.
- [46] McKay, M. D., R. J. Beckman and W. J. Conover (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 42(1), 55–61.
- [47] McNeil, A. J. and J. Nešlehoá (2009). Multivariate Archimedean copulas, d-monotone functions and l-norm symmetric distributions. *Ann. Statist.* 37(5B), 3059–3097.
- [48] Nápoles, O. M. (2010). Bayesian Belief Nets and Vines in Aviation Safety and other Applications. Ph.D. thesis, Delft University of Technology, Netherlands.
- [49] Nápoles O. M. Counting vines. In D. Kurowicka and H. Joe (Eds.), *Dependence Modeling: Vine Copula Handbook*, pp. 189–218. World Scientific Publishing Co., Singapore.
- [50] Nelsen, R. B. (2006). An Introduction to Copulas. Second edition. Springer, New York.
- [51] Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. *Handb. Econom.* 4, 2111–2245.
- [52] Osborne, M. A., R. Garnett, S. J. Roberts (2009). Gaussian processes for global optimization. *3rd International Conference on Learning and Intelligent Optimization (LION 3)*.
- [53] Pass B. (2015). Multi-marginal optimal transport: theory and applications. ESAIM Math. Model. Numer. Anal. 49(6), 1771–1790.
- [54] Robbins, H. and S. Monro (1951). A stochastic approximation method. Ann. Math. Stat. 22(3), 400-407.
- [55] Rubino, G. and B. Tuffin (2009). Rare Event Simulation Using Monte Carlo Methods. Wiley, Chichester.
- [56] Saltelli, A., K. Chan and E. M. Scott (2000). Sensitivity Analysis. Wiley, New York.

- [57] Schoelzel, C. and P. Friederichs (2008). Multivariate non-normally distributed random variables in climate research Introduction to the copula approach. Nonlin. Processes Geophys. 15(5), 761-772.
- [58] Sklar, M. (1959). Fonctions de répartition à n dimensions et leurs marges. Publ. Inst. Statist. Univ. Paris 8, 229-231.
- [59] Nilsen, T. and T. Aven (2003). Models and model uncertainty in the context of risk analysis. Reliab. Eng. Syst. Saf. 79(3), 309-317.
- [60] Tang, X. S., D. Q. Li, G. Rong, K. K. Phoon and C. B. Zhou (2013). Impact of copula selection on geotechnical reliability under incomplete probability information. Comput. Geotech. 49, 264-278.
- [61] Tang, X. S., D. Q. Li, C. B. Zhou and K. K. Phoon (2015). Copula-based approaches for evaluating slope reliability under incomplete probability information. Struct. Saf. 52, 90-99.
- [62] Thoft-Christensen, P. and J. D. Sørensen (1982). Reliability of structural systems with correlated elements. Appl. Math. Model. 6(3), 171-178.
- [63] Torre, E., S. Marelli, P. Embrechts and B. Sudret (2019). Data-driven polynomial chaos expansion for machine learning regression. J. Comput. Phys. 388, 601-623.
- [64] Torre, E., S. Marelli, P. Embrechts and B. Sudret (2019). A general framework for data-driven uncertainty quantification under complex input dependencies using vine copulas. Prob. Eng. Mech. 55, 1-16.
- [65] Van der Vaart, A. W. (2000). Asymptotic Statistics. Cambridge University Press.
- [66] Villani, C. (2008). Optimal Transport: Old and New. Springer, Berlin.
- [67] Zondervan-Zwijnenburg, M., W. van de Schoot-Hubeek, K. Lek, H. Hoijtink and R. van de Schoot (2017). Application and Evaluation of an Expert Judgment Elicitation Procedure for Correlations. Front. Psychol. 8, 1-15.