

Identifiability and estimation of meta-elliptical copula generators

Derumigny, A.; Fermanian, J. D.

DOI

[10.1016/j.jmva.2022.104962](https://doi.org/10.1016/j.jmva.2022.104962)

Publication date

2022

Document Version

Final published version

Published in

Journal of Multivariate Analysis

Citation (APA)

Derumigny, A., & Fermanian, J. D. (2022). Identifiability and estimation of meta-elliptical copula generators. *Journal of Multivariate Analysis*, 190, Article 104962. <https://doi.org/10.1016/j.jmva.2022.104962>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Identifiability and estimation of meta-elliptical copula generators

A. Derumigny^a, J.-D. Fermanian^{b,*}

^a Department of Applied Mathematics, Delft University of Technology, Netherlands

^b Crest-Ensaie, 5 av. Henry le Chatelier, 91764 Palaiseau cedex, France

ARTICLE INFO

Article history:

Received 31 July 2021

Received in revised form 31 January 2022

Accepted 31 January 2022

Available online 16 February 2022

AMS 2020 subject classifications:

primary 62H05

secondary 62H12

Keywords:

Identifiability

Meta-elliptical copulas

Elliptical generator

Recursive algorithm

ABSTRACT

Meta-elliptical copulas are often proposed to model dependence between the components of a random vector. They are specified by a correlation matrix and a map g , called density generator. While the latter correlation matrix can easily be estimated from pseudo-samples of observations, the density generator is harder to estimate, especially when it does not belong to a parametric family. We give sufficient conditions to non-parametrically identify this generator. Several nonparametric estimators of g are then proposed, by M-estimation, simulation-based inference, or by an iterative procedure available in the R package `ElliptCopulas`. Some simulations illustrate the relevance of the latter method.

© 2022 Elsevier Inc. All rights reserved.

1. Introduction

Elliptically contoured distributions are usual semi-parametric extensions of multivariate Gaussian or Student distributions. They correspond to continuous distributions on \mathbb{R}^d whose isodensity curves (with respect to the Lebesgue measure) are ellipsoids: see, e.g., [3,10,17,22]. To be specific, let \mathbf{X} be a random vector in \mathbb{R}^d whose elliptical distribution is parameterized by a vector $\boldsymbol{\mu} \in \mathbb{R}^d$, a positive definite matrix $\boldsymbol{\Omega} = [\Omega_{i,j}]_{1 \leq i,j \leq d}$ and a measurable function $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+ \cup \{+\infty\}$. Its density with respect to the Lebesgue measure in \mathbb{R}^d is

$$f_{\mathbf{X}}(\mathbf{x}) = |\boldsymbol{\Omega}|^{-1/2} g((\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Omega}^{-1} (\mathbf{x} - \boldsymbol{\mu})), \quad \mathbf{x} \in \mathbb{R}^d. \quad (1)$$

This distribution is denoted $\mathcal{E}_d(\boldsymbol{\mu}, \boldsymbol{\Omega}, g)$ and its cumulative distribution function by $H_{g, \boldsymbol{\Omega}, d}$. To specify the law of the random vector \mathbf{X} , we will write $\mathbf{X} \sim \mathcal{E}_d(\boldsymbol{\mu}, \boldsymbol{\Omega}, g)$. The map g is called “density generator”, or simply “generator”. By integration of (1), a density generator of an elliptical vector satisfies the constraint

$$s_d \int_0^\infty r^{d-1} g(r^2) dr = s_d \int_0^\infty t^{d/2-1} g(t) dt / 2 = 1, \quad (2)$$

where $s_d := 2\pi^{d/2} / \Gamma(d/2)$ is the surface area of the unit ball in \mathbb{R}^d , $d \geq 2$ ($s_1 = 2$). Conversely, any nonnegative function g that satisfies (2) can be used as the density generator of an elliptical distribution. For example, the density generator of a Gaussian distribution is $g_{\text{Gauss}}(u) := \exp(-u/2)/(2\pi)^{d/2}$.

* Corresponding author.

E-mail addresses: A.F.F.Derumigny@tudelft.nl (A. Derumigny), jean-david.fermanian@ensae.fr (J.-D. Fermanian).

Remark 1. It is possible to define elliptical distributions with singular matrices Ω : see [3]. In such cases, $\lambda^\top \mathbf{X} = 0$ a.s. for some vector λ in \mathbb{R}^d , and the law of \mathbf{X} is supported on an affine subspace of \mathbb{R}^d . Such “degenerate” cases will not be considered in this paper.

A d -dimensional copula C is said meta-elliptical (or simply “elliptical”) if there exists an elliptical distribution in \mathbb{R}^d whose copula is C . Due to the invariance of copulas by location-scale transforms, an elliptical copula depends on a generator and a correlation matrix only. Indeed, for any $\mathbf{X} \sim \mathcal{E}_d(\mathbf{0}, \Omega, g)$, set $\mathbf{Y} := (X_1/\sqrt{\Omega_{1,1}}, \dots, X_d/\sqrt{\Omega_{d,d}})^\top$. Then, $\mathbf{Y} \sim \mathcal{E}_d(\mathbf{0}, \Sigma, g)$ for a correlation matrix $\Sigma := [\Omega_{i,j}/(\Omega_{i,i}\Omega_{j,j})^{1/2}]$. Obviously, the meta-elliptical copulas of \mathbf{X} and \mathbf{Y} are the same. Thus, for any correlation matrix Σ , denote by $\mathcal{ME}_d(\Sigma, g)$ the (unique) meta-elliptical copula that corresponds to the elliptical distribution $\mathcal{E}_d(\mathbf{0}, \Sigma, g)$. More generally, this copula corresponds to the elliptical distributions $\mathcal{E}_d(\mu, \Omega, g)$ for any μ, Ω , such that Σ is the correlation matrix associated to Ω .

A trans-elliptical distribution [29,30] is a distribution whose copula is meta-elliptical. Trans-elliptical distributions [9] extend elliptical distributions, by allowing the associated margins to be arbitrarily specified. For any correlation matrix Σ , we will denote by $\mathcal{TE}_d(\Sigma, g, F_1, \dots, F_d)$ the trans-elliptical distribution whose copula is $\mathcal{ME}_d(\Sigma, g)$ and marginal cdfs’ are F_1, \dots, F_d .

The probabilistic properties of meta-elliptical copulas and their statistical analysis have been studied in several papers in the literature: their conditional distributions and dependence measures [9], the estimation of the correlation matrix Σ [42], stochastic ordering [1], sampling methods [41], their tail dependence function [23–25], their use semi-parametric regressions [43] or some goodness-of-fit tests [20,21], their relationships with partial and conditional correlations [26], etc., are notable contributions. Thanks to Sklar’s theorem, meta-elliptical copulas can be used as key components of many flexible multivariate models, far beyond elliptically-distributed random vectors. Moreover, the literature has considered parametric families of generators that include the popular Gaussian and/or Student copulas as particular cases, with practical applications in hydrology [13,37], risk management [6,12], econometrics [35], biology, etc.

To the best of our knowledge, all these papers assume the generator of a meta-elliptical copula is known, possibly up to finite dimensional parameter. Thus, the problem of estimating g itself is bypassed. For instance, [13] proposed a graphical tool to select a “well-suited” generator among a finite set of potential generators. Approximations of meta-elliptical copulas are proposed in [38] through projection pursuit techniques but the consistency of the proposed algorithm seems to occur only under very restrictive conditions. Actually, a general nonparametric estimation of the generator g is problematic. As noticed in Genest et al. [13]: “The estimation of g is more complex, considering that it is a functional parameter. Indeed, a rigorous approach to this problem has yet to be developed. Financial applications to date have simply treated g as fixed; however, several possible choices of g have often been considered to assess the robustness of the conclusions derived from the model”. Therefore, until now, no nonparametric consistent estimator of g seems to be available in the literature. This should not be surprising. Indeed, a preliminary point would be to state the identifiability of g from the knowledge of the underlying copula. This result is far from obvious and is one of the main contributions of our work.

Let us explain why this is the case. As recalled in the supplementary material, all margins of a distribution $\mathcal{E}_d(\mathbf{0}, \Sigma, g)$ have the same density f_g , where, for every $t \in \mathbb{R}$,

$$f_g(t) = \frac{\pi^{(d-1)/2}}{\Gamma((d-1)/2)} \int_0^{+\infty} g(t^2 + s)s^{(d-3)/2} ds. \tag{3}$$

Note that f_g is even. For notational convenience, we do not write the dependency of f_g on the dimension d . Set its marginal cdfs’ $F_g(x) = \int_{-\infty}^x f_g(t) dt$ for every real number x , and its quantile function $Q_g(u) = \inf\{x; F_g(x) \geq u\}$, $u \in [0, 1]$. By Sklar’s theorem, our meta-elliptical distribution $C \sim \mathcal{ME}_d(\Sigma, g)$ is given by

$$C(\mathbf{u}) = H_{g, \Sigma, d}(Q_g(u_1), \dots, Q_g(u_d)) = |\Sigma|^{-1/2} \int_{-\infty}^{Q_g(u_1)} \dots \int_{-\infty}^{Q_g(u_d)} g(\mathbf{x}^\top \Sigma^{-1} \mathbf{x}) d\mathbf{x}, \tag{4}$$

for every $\mathbf{u} \in (0, 1)^d$. Hence, the associated meta-elliptical copula density with respect to the Lebesgue measure exists and may be defined as

$$c(\mathbf{u}) := \frac{g(\vec{Q}_g(\mathbf{u}) \Sigma^{-1} \vec{Q}_g(\mathbf{u})^\top)}{|\Sigma|^{1/2} \prod_{k=1}^d f_g(Q_g(u_k))}, \text{ where } \vec{Q}_g(\mathbf{u}) := [Q_g(u_1), \dots, Q_g(u_d)], \tag{5}$$

for every $\mathbf{u} \in [0, 1]^d$. The latter density and cumulative distribution function will be denoted by c_g and C_g respectively, when we want to stress its dependence with respect to g .

Concerning the inference of meta-elliptical copulas, the usual estimator of the matrix Σ has been known for a long time and is based on empirical Kendall’s tau (see below). When one observes i.i.d. realizations of elliptically-distributed random vectors, several estimators of the generator g have been proposed in the literature (see our supplementary material). This paper is related to the same purpose, but for trans-elliptical distributions, or, equivalently, for meta-elliptical copulas. In such a case, the inference of g is more difficult than for elliptical laws because copula densities depend on g through a highly nonlinear and complex relationship. Moreover, it is not known whether the mapping $g \mapsto c_g$ is one-to-one. To

the best of our knowledge, this problem has never been tackled in the literature. Authors only rely on ad-hoc chosen generators, or on parametric families of generators.

In Section 2, we give sufficient conditions for the identifiability of the generator of a meta-elliptical copula. Estimation procedures of g are proposed in Section 3. They allow the nonparametric estimation of an assumed trans-elliptical distribution, because Σ and its margins are easily estimated beside. Section 3.1 proposes a penalized M-estimation of g , when a simulation-based estimator is given in Section 3.2. Section 3.3 states a numerical iterative procedure to evaluate g . It will be called MECIP, as “Meta-Elliptical Copula Iterative Procedure”. Its performances are evaluated in Section 4 and it is implemented in the R package `ElliptCopulas` [5].

2. Identifiability of meta-elliptical copulas

Consider a meta-elliptical copula $C = \mathcal{ME}_d(\Sigma, g)$ where Σ is a correlation matrix and g satisfies the usual normalization constraint (2). The question is to know whether the latter parameterization is unique. Strictly speaking, this is the same question as for trans-elliptical distributions when their margins are known.

For any meta- or trans-elliptical distribution, the correlation matrix Σ is identifiable. Indeed, it is well-known there exists a nice relationship between its components and the underlying Kendall’s tau: $\Sigma_{k,l} = \sin(\pi \tau_{k,l}/2)$, $k \neq l$, where $\tau_{k,l}$ is the Kendall’s tau between X_k and X_l . See [42] and the references therein, for instance. Since every $\tau_{k,l}$ is uniquely defined by the underlying copula of \mathbf{X} , this is still the case for Σ too.

Proposition 1. *If $\mathbf{U} \sim \mathcal{ME}_d(\Sigma, g)$ and $\mathbf{U} \sim \mathcal{ME}_d(\tilde{\Sigma}, \tilde{g})$ where Σ and $\tilde{\Sigma}$ are two correlation matrices, then $\Sigma = \tilde{\Sigma}$.*

Note that it is possible to have $g \neq \tilde{g}$ due to the non-uniqueness of meta-elliptical copula generators (see Proposition 2 below). If \mathbf{X} is a trans-elliptical random vector, the inference of the matrix Σ can be done independently of the margins. The estimated matrix $\hat{\Sigma} := [\hat{\Sigma}_{k,l}]$ is given by $\hat{\Sigma}_{k,l} := \sin(\pi \hat{\tau}_{k,l}/2)$, $k \neq l$, and $\hat{\Sigma}_{k,k} = 1$, introducing empirical Kendall’s tau

$$\hat{\tau}_{k,l} := \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} \text{sign}(X_{i,k} - X_{j,k}) \times \text{sign}(X_{i,l} - X_{j,l}), \quad k \neq l.$$

Since it is not guaranteed that $\hat{\Sigma}$ is a correlation matrix, this can be imposed by projection techniques ([34, Section 8.7.2.1], e.g.). Hereafter, we require that $\hat{\Sigma}$ is invertible.

Having tackled the identifiability and the estimation problem of Σ , the problem is reduced to the following one: let g and \tilde{g} be two density generators of meta-elliptical copulas on $[0, 1]^d$ such that $c_g = c_{\tilde{g}}$, with the previous notations. Does it imply that $g = \tilde{g}$ almost everywhere?

Remark 2. There is a one-to-one mapping between generators g of meta-elliptical copulas and the so-called univariate densities f_g , as given in (3), once the underlying copula density c is known. Indeed, since Σ^{-1} is definite positive, its diagonal elements are positive. Then, invoke (5) with $\tilde{Q}_g(\mathbf{u}) := [x, 0, \dots, 0]$ for some arbitrary $x \in \mathbb{R}$. Since f_g is even, $F_g(0) = 1/2$ and $f_g(0) \neq 0$ by (3). This yields $g(\gamma x^2) = |\Sigma|^{1/2} f_g(0)^{d-1} f_g(x) c(F_g(x), 1/2, \dots, 1/2)$ for every x and some known positive constant γ . This means the map $g \mapsto f_g$ is invertible, restricting ourselves to meta-elliptical copula generators. In other words, since copula densities are nonparametrically identifiable, the identifiability problem of g or of f_g are the same. But, since f_g can be (nonparametrically) identified only in the case of elliptical distributions, this does not prove the identifiability of g for general meta-elliptical/trans-elliptical distributions.

Recall that elliptical distributions $\mathcal{E}_d(\mu, \Omega, g)$ are not identifiable in general without any identifiability constraint (see Proposition 1 in the supplementary material). Therefore, most authors impose a condition such as $\text{Cov}(\mathbf{X}) = \Omega$ (when \mathbf{X} has finite second moments) or $\text{Tr}(\Omega) = 1$ (in the general case). To deal with meta-elliptical copulas, we are facing similar problems. Indeed, such distributions are never identifiable without identifiability constraints, as proven in the next proposition.

Proposition 2. *Let Σ be a positive definite correlation matrix and g be a density generator of a meta-elliptical copula on $[0, 1]^d$. Then, for any $a > 0$, $\mathcal{ME}_d(\Sigma, g) = \mathcal{ME}_d(\Sigma, g_a)$ by setting $g_a(t) := a^{d/2} g(a \times t)$.*

Proof of Proposition 2. By (3), we easily get $f_{g_a}(t) = \sqrt{a} f_g(\sqrt{a}t)$ for every t . Deduce $F_{g_a}(x) = F_g(\sqrt{a}x)$ and $Q_{g_a}(u) = Q_g(u)/\sqrt{a}$ for every t and u . Then, applying (5), we check that the copula densities associated to g and g_a respectively are the same. \square

Therefore, for a given meta-elliptical copula, its generator g has to satisfy at least another constraint in addition to (2), to be uniquely defined. We will prove that the generator of any elliptical copula is identifiable under some regularity conditions. This is one of the main contributions of our work. Before, we need our density generators to be sufficiently regular so that the associated univariate densities f_g are differentiable. This is guaranteed by the next assumption.

Condition 1. Set $a_g := \sup\{t \mid t > 0, g(t^2) > 0\} \in \bar{\mathbb{R}}_+$. The map $t \mapsto g(t^2)$ from \mathbb{R} to \mathbb{R}^+ is strictly positive and differentiable on $(-a_g, a_g)$. Moreover, the map $t \mapsto \int_0^{+\infty} g(t^2 + r^2) r^{d-2} dr$ is finite and differentiable on \mathbb{R} .

As a consequence, $\{t \mid g(t) > 0\}$ (resp. $\{t \mid f_g(t) > 0\}$) is equal to the interval $(-a_g^2, a_g^2)$ (resp. $(-a_g, a_g)$), possibly including the boundaries. Thus, we do not allow generators whose supports exhibit “holes”, such as sums of indicator functions that are related to disjoint subsets. The forbidden models correspond to meta-elliptical copulas whose densities are zero in some “cavities” that look like “distorted rings” (when plotted on $[0, 1]^d$). Such features can easily be identified by plotting nonparametric estimates of c_g as a preliminary stage. They are also unlikely to happen in practical applications.

Denote by \mathcal{G} the set of density generators g that satisfy **Condition 1**, in addition to (2). They will be called “regular density generators”. Moreover, denote by $c_{i,j}$ the map from $[0, 1]^2$ to \mathbb{R} that is equal to the copula density c , when all its arguments are equal to $1/2$, except the i th and the j th.

Proposition 3. Consider a meta-elliptical random vector $\mathbf{U} \sim \mathcal{ME}_d(\Sigma, g)$, where the correlation matrix Σ is not the identity matrix \mathbf{I}_d . Let i and j be two different indices in $\{1, \dots, d\}$ for which the (i, j) -component of Σ^{-1} is not zero. Assume that $c_{i,j}$ has finite first-order partial derivatives on $(0, 1)^2$. Moreover, the maps $u \mapsto \partial_k \ln c_{i,j}(u, 1/2)$, $k \in \{1, 2\}$, are locally Lipschitz on $(0, 1)$. Consider two couples (Σ_k, g_k) , $g_k \in \mathcal{G}$, $k \in \{1, 2\}$, that induce the same law of \mathbf{U} and that both satisfy

$$b := \frac{\pi^{(d-1)/2}}{\Gamma((d-1)/2)} \int_0^{+\infty} g_k(s) s^{(d-3)/2} ds, \tag{6}$$

for a given positive real number b . Then these couples are essentially unique: $\Sigma_1 = \Sigma_2$; the supports of g_1 and g_2 are the same interval (except possibly at its boundaries), and $g_1 = g_2$ on the interior of their common support.

Therefore, b can be arbitrarily chosen. For any generator g , it is always possible to find a version \tilde{g} such that $\mathcal{ME}_d(\Sigma, g) = \mathcal{ME}_d(\Sigma, \tilde{g})$ for all matrices Σ , and such that $b(\tilde{g}) = 1$. An explicit procedure to build \tilde{g} is given in Algorithm 1. Note that (6) means $f_{g_1}(0) = f_{g_2}(0) = b$. By (3), $f_{g_1} = f_{g_2}$ and $F_{g_1} = F_{g_2}$ everywhere, implying $C_{g_1} = C_{g_2}$ everywhere by (4). **Condition 1** implies $g(0) > 0$, that is equivalent to $c(1/2, 1/2, \dots, 1/2) > 0$ (see (13) in the proof). The latter condition seems to be weak, particularly under a practical perspective: copulas that have no mass at the center of their support can be considered as “pathological”.

Therefore, for a given elliptical copula and if $\Sigma \neq \mathbf{I}_d$, there “most often” exists a unique regular density generator g for which the constraints (6) and $s_d \int_0^\infty t^{d/2-1} g(t) dt / 2 = 1$ are both satisfied. In other words, two moment-type conditions are sufficient to uniquely identify the generator of an elliptical copula “most of the time”. **Proposition 3** shows that, when $d > 2$, the generator of a d -dimensional elliptical copula is uniquely defined by these two moment conditions in addition to the knowledge of a single map $(u, v) \mapsto c_{i,j}(u, v)$, i and j being two indices in $\{1, \dots, d\}$ for which the (i, j) -component of Σ^{-1} is not zero. In other words, it is not necessary to know the whole copula density c on $[0, 1]^d$ to recover the density generator of an elliptical copula. Only a single bivariate cross-section is sufficient.

Example 1. Consider the particular case of meta-elliptical copulas whose generator is given by $g(t) = P(t) \exp(-\lambda t^\beta)$ for some positive constants λ, β and some polynomial P such that $P(0) > 0$ and $P(t) \geq 0$ for every $t \geq 0$. They are linear combinations of the generators associated to the family of symmetric Kotz-type distributions (see [9], Example 2.1), including Gaussian copulas as particular cases. They satisfy **Condition 1** and then **Proposition 3** applies to them.

Example 2. When a meta-elliptical copula generator is $g(t) = K_{m,N} (1 + t/m)^{-N}$ for some positive constants $K_{m,N}, m$ and $N > 1$, this yields the copulas associated to the family of symmetric bivariate Pearson type VII distributions (see [9], Example 2.2), including Student distributions when $N = m/2 + 1$. Check that **Condition 1** is fulfilled when $N > 1 + d/2$.

Proof of Proposition 3. First consider the bivariate case $d = 2$. Denote by $\rho \neq 0$ the extra-diagonal component of Σ . With the same notations as above, the copula density c of \mathbf{U} with respect to the Lebesgue measure satisfies

$$c(F_g(x), F_g(y)) f_g(x) f_g(y) = \frac{g((x^2 + y^2 - 2\rho xy)/(1 - \rho^2))}{\sqrt{1 - \rho^2}}, \tag{7}$$

for almost every $(x, y) \in \mathbb{R}^2$, by Sklar’s theorem. Here, we clearly see that the maps $g \mapsto F_g$ and $g \mapsto f_g$ are one-to-one for a given copula density c : setting $x = \rho y$, Eq. (7) yields

$$g(y^2) = \sqrt{1 - \rho^2} c(F_g(\rho y), F_g(y)) f_g(\rho y) f_g(y), \tag{8}$$

for every $y \in \mathbb{R}$. Therefore, the knowledge of f_g (or F_g , equivalently) provides a single generator g that induces the given copula c . Now, it is sufficient to prove the identifiability of f_g . By **Condition 1**, $(-a_g, a_g)$ is the support of f_g (possibly including the boundaries) and $f_g(0) = b$ is positive. Setting $y = 0$ in (7), we get

$$\sqrt{1 - \rho^2} c(F_g(x), 1/2) f_g(x) b = g(x^2/(1 - \rho^2)),$$

for every $x \in \mathbb{R}$. Since g is non zero and continuous at zero, there exists an open neighborhood of zero V_0 for which $c(F_g(x), 1/2) > 0$ when $x \in V_0$.

We will restrict ourselves to the couples $(x, y) \in \mathbb{R}^2$ such that the non-negative number $\{x^2 + y^2 - 2\rho xy\}/(1 - \rho^2)$ belongs to $[0, a_g^2]$. Denote by $\mathcal{X}(g)$ the set of such couples. Note that $\mathcal{X}(g)$ contains an open neighborhood of $(0, 0)$ and that $f_g(x) f_g(y) > 0$ for such couples, due to (7).

Condition 1 means that f_g is differentiable on \mathbb{R} . Since it is even, $f'(0) = 0$. By differentiating (7) with respect to x , we get

$$\partial_1 c(F_g(x), F_g(y))f_g^2(x)f_g(y) + c(F_g(x), F_g(y))f_g'(x)f_g(y) = g' \left(\frac{x^2 + y^2 - 2\rho xy}{1 - \rho^2} \right) \frac{2(x - \rho y)}{(1 - \rho^2)^{3/2}}, \tag{9}$$

for any $(x, y) \in \mathcal{X}(g)$. Set $y = x/\rho$, canceling the right-hand side of (9). This yields

$$\partial_1 \ln c(F_g(x), F_g(x/\rho))f_g^2(x) + f_g'(x) = 0,$$

for every $x \in (-a_g, a_g)$. As a consequence, f_g' is continuous on the latter interval.

By independently differentiating (7) with respect to y and comparing with (9), we obtain

$$\begin{aligned} & (y - \rho x) \{ \partial_1 c(F_g(x), F_g(y))f_g^2(x)f_g(y) + c(F_g(x), F_g(y))f_g'(x)f_g(y) \} \\ &= (x - \rho y) \{ \partial_2 c(F_g(x), F_g(y))f_g(x)f_g^2(y) + c(F_g(x), F_g(y))f_g(x)f_g'(y) \}, \end{aligned} \tag{10}$$

for every $(x, y) \in \mathcal{X}(g)$. Consider the particular value $y = 0$, for which $f_g(y) = b > 0$ and $f_g'(y) = 0$. By symmetry, $F_g(x) = 1 - F_g(-x)$ for every real number x and $F_g(0) = 1/2$. Then, (10) can be rewritten as follows:

$$\rho \{ \partial_1 \ln c(F_g(x), 1/2)f_g^2(x) + f_g'(x) \} = -\partial_2 \ln c(F_g(x), 1/2)f_g(x)b, \tag{11}$$

for every x in a sufficiently small neighborhood of zero such that $c(F_g(x), 1/2) > 0$ (such as V_0 above, for instance). Thus, we have obtained an ordinary differential equation, whose solution $z := F_g$ would be a function of x when x belongs to a neighborhood of zero. The latter differential equation can be rewritten as

$$z'' + \partial_1 \ln c(z, 1/2)(z')^2 + \partial_2 \ln c(z, 1/2) \frac{z'b}{\rho} = 0. \tag{12}$$

Setting the bivariate map $\vec{z} = [z, z']$, we are facing the usual Cauchy problem: find \vec{z} , a function of x , such that $d\vec{z} = H(\vec{z}) dx$ and that satisfies $\vec{z}(0) = [1/2, b]$. Here, the latter map $H : [0, 1] \times \mathbb{R} \mapsto \mathbb{R}^2$ is

$$H(z_1, z_2) := \left[z_2; -\partial_1 \ln c(z_1, 1/2)z_2^2 - \partial_2 \ln c(z_1, 1/2) \frac{z_2 b}{\rho} \right].$$

By assumption, this map H is Lipschitz on any subset $[\alpha, \beta] \times [b - \epsilon, b + \epsilon]$, when $0 < \alpha < \beta < 1$ and $\epsilon > 0$. In particular, this is the case when $\alpha < 1/2 < \beta$. By the Cauchy–Lipschitz Theorem, we deduce there exists a unique solution \vec{z} in an open neighborhood of $x = 0$. Note that this solution satisfies (6) by construction.

Therefore, consider a (global) solution $\vec{z} = [F_g, f_g]$ of (12) on some maximum interval S_g on the real line that contains zero. We can impose the latter solution is associated to a regular generator that satisfies (6). Now, assume there are two different regular generators g and \bar{g} that induce the same copula. Set $z = [F_g, f_g]$, $\bar{z} = [F_{\bar{g}}, f_{\bar{g}}]$ and $\mathcal{A} := \{t \in S_g \cap S_{\bar{g}}; z(t) = \bar{z}(t)\}$. We have proved that \mathcal{A} contains an open ball around zero. Define $t^* := \sup\{t; t \in \mathcal{A}\}$ and assume that t^* is finite. Assume $t^* < \min(a_g, a_{\bar{g}})$. By the continuity of the considered cdfs' and densities (Condition 1), $z(t^*) = \bar{z}(t^*)$ is a known value v^* . Then, we can apply again the Cauchy–Lipschitz Theorem to the differential equation (12), with the condition $z(t^*) = v^*$. This yields a unique solution of (12) in an open neighborhood of t^* . As a consequence, $z(t^* + \epsilon) = \bar{z}(t^* + \epsilon)$ for some $\epsilon > 0$. This contradicts the definition of t^* . Assuming w.l.o.g. $a_g \leq a_{\bar{g}}$, this implies $t^* \geq a_g$. In other words, f_g (resp. F_g) and $f_{\bar{g}}$ (resp. $F_{\bar{g}}$) coincide on $(-a_g, a_g)$, implying $g = \bar{g}$ on the $(-a_g^2, a_g^2)$ (recall (8)). To satisfy (2) with \bar{g} , this requires $\bar{a}_g = \bar{a}_{\bar{g}}$, i.e., $a_g = a_{\bar{g}}$. Thus, $g = \bar{g}$ on the interior of their common support. This proves the result when $d = 2$.

Second, for an arbitrary dimension $d > 2$, there exists a non-zero extra-diagonal element in Σ^{-1} by assumption. W.l.o.g., assume it is corresponding to the couple of indices $(1, 2)$. Let us fix the other arguments of the copula density c at the value $1/2$, i.e., we focus on the points $(x, y, 0, \dots, 0)$ in \mathbb{R}^d . This implies there exist two non zero real numbers (θ, γ) such that

$$c \left(F_g(x), F_g(y), \frac{1}{2}, \dots, \frac{1}{2} \right) = \frac{g(\gamma x^2 + \gamma y^2 + 2\theta xy)}{|\Sigma|^{1/2} f_g(x) f_g(y) b^{d-2}}, \tag{13}$$

for every $(x, y) \in \mathbb{R}^2$. By differentiation with respect to x and y respectively, we get an ordinary differential equation that is strictly similar to (12), apart from different non zero constants. By exactly the same arguments as in the bivariate case, we can prove there exists a unique global solution on the real line of this differential equation. As a consequence, g is uniquely defined by c (except at the boundaries of its support), proving the result. \square

Unfortunately, the limiting case $\Sigma = \mathbf{I}_d$ cannot be managed similarly by considering differential equations and some initial conditions at the particular point $x = 0$. This is due to the nullity of both sides of (11): differentiate (7) with respect to y , set $y = 0$, and deduce that $\partial_2 \ln c(F_g(x), 1/2) = 0$ for every x in a neighborhood of zero. Nonetheless, we can provide partial answers to this problem by imposing some conditions at $+\infty$. This requires restricting ourselves to a smaller class of generators. To this end, we introduce a measurable map $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}$.

Condition 2. The density generator $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ belongs to \mathcal{G} , with $a_g = +\infty$. Moreover, for every $x \in \mathbb{R}$, the map $y \mapsto (g'/g)(x^2 + y^2)$ has a finite limit $\psi(x^2)$, when $y \rightarrow +\infty$.

Denote by $\tilde{\mathcal{G}}_\psi$ the set of density generators g that satisfy **Condition 2**, in addition to (2). Note that such g are assumed to be strictly positive on \mathbb{R}^+ . For a lot of reasonable generators, we can hope $\psi(x) = 0$. This is the case for all generators that are sums of maps of the form $P(x)\exp(-\lambda x^\beta)$ for some polynomials P and some constants $\lambda > 0$ and $\beta \in (0, 1)$. When $\beta = 1$, we get a family of ‘‘Gaussian-type’’ generators, for which $\psi(x) = -\lambda$ for every x . But **Condition 2** is not fulfilled for such generators when $\beta > 1$.

Proposition 4. Consider a meta-elliptical random vector $\mathbf{U} \sim \mathcal{ME}_d(\mathbf{I}_d, g)$. Assume that, for every $u \in (0, 1)$, the map $v \mapsto \partial_1 \ln c(u, v)$ exists and has a finite limit $\chi(u)$ when $v \rightarrow 1, v < 1$. Moreover, χ is locally Lipschitz on $(0, 1)$. For a given map ψ , consider two couples $(\Sigma_k, g_k), g_k \in \tilde{\mathcal{G}}_\psi, k \in \{1, 2\}$, that induce the same law of \mathbf{U} and that satisfy (6). Then these couples are essentially unique: $\Sigma_1 = \Sigma_2$ and $g_1 = g_2$ on their common support \mathbb{R}^+ .

Proof of Proposition 4. Let us assume first that $d = 2$. By definition, the copula density c satisfies

$$c(F_g(x), F_g(y))f_g(x)f_g(y) = g(x^2 + y^2), \tag{14}$$

for every $(x, y) \in \mathbb{R}^2$. Since the support of g is \mathbb{R}_+ , the support of f_g is the whole real line. Differentiating the latter equation with respect to x and dividing the new one by both members of (14), we deduce

$$\partial_1 \ln c(F_g(x), F_g(y))f_g(x) + \frac{f'_g}{f_g}(x) = 2x\left(\frac{g'}{g}\right)(x^2 + y^2),$$

for every real numbers x and y . Now, let us make y tend to $+\infty$. From **Condition 2**, we deduce, for every $x \in \mathbb{R}$,

$$\chi(F_g(x))f_g(x) + \frac{f'_g}{f_g}(x) = 2x\psi(x).$$

The latter equation is a second-order differential equation with respect to the unknown function $F_g =: z$, i.e.,

$$\chi(z)z' + \frac{z''}{z}(x) = 2x\psi(x), \text{ or } z'' = 2xz'\psi(x) - \chi(z)(z')^2.$$

As in the proof of **Proposition 5**, consider the initial conditions $z(0) = 1/2$ and $z'(0) = b$. By a similar reasoning (Cauchy–Lipschitz theorem), we prove the result when $d = 2$.

When $d > 2$, we consider the map $(u, v) \mapsto c(u, v, 1/2, \dots, 1/2)$ that satisfies

$$c(F_g(x), F_g(y), \frac{1}{2}, \dots, \frac{1}{2})f_g(x)f_g(y)b^{d-2} = g(x^2 + y^2),$$

for every $(x, y) \in \mathbb{R}^2$. The same reasoning as for the case $d = 2$ proves the result. \square

It can be checked that the meta-elliptical copulas of **Example 1** satisfy **Condition 2** when $\beta \leq 1$ and **Proposition 4** applies to them. This is still the case for the copulas of **Example 2**, for any value of $(m, N), N > 1 + d/2$.

Finally, as shown in [1], Proposition 1.1, the identifiability of g may be obtained in the particular case of Gaussian copulas. Let us extend the latter result in dimension $d \geq 2$.

Proposition 5. Let $\mathbf{U} \sim \mathcal{ME}_d(\Sigma, g)$ and $\mathbf{U} \sim \mathcal{ME}_d(\mathbf{I}_d, g_{\text{Gauss}})$ where Σ is a correlation matrix. Then $\Sigma = \mathbf{I}_d$ and $g = g_{\text{Gauss}}$ a.s.

Proof of Proposition 5. The first part of the proposition is the result of the identifiability of the correlation matrix Σ . Since $\mathbf{U} \sim \mathcal{ME}_d(\mathbf{I}_d, g_{\text{Gauss}})$, check that U_1, \dots, U_d are mutually independent. Thus, $Q_g(U_1), \dots, Q_g(U_d)$ are independent variables and their joint law is an elliptical distribution $\mathcal{E}_d(\mathbf{0}, \mathbf{I}_d, g)$. Lemma 5 in [22] implies that $(Q_g(U_1), \dots, Q_g(U_d)) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, or $\mathcal{E}_d(\mathbf{0}, \mathbf{I}_d, g_{\text{Gauss}})$ equivalently. Using Proposition 1 in the supplementary material, this yields $g = g_{\text{Gauss}}$. \square

3. Inference of density generators of meta-elliptical copulas

In this section, we define three inference strategies to evaluate g , since we now know that such generators are nonparametrically identifiable under some regularity conditions and two moment-type constraints.

Let \mathbf{X} be a random vector whose distribution is trans-elliptical $\mathcal{TE}_d(\Sigma, g, F_1, \dots, F_d)$ for a correlation matrix Σ . Let $(\mathbf{X}_1, \dots, \mathbf{X}_n)$ be an i.i.d. sample of realizations of \mathbf{X} . As a particular case, its law could be elliptical $\mathcal{E}_d(\mathbf{0}, \Sigma, g)$ when all its margins F_k are equal to F_g . Moreover, if its margins are uniformly distributed on $[0, 1]$, then the law of \mathbf{X} is given by a meta-elliptical copula $\mathcal{ME}_d(\Sigma, g)$. We assume there exists a single generator g such that **Condition (2)** and **Condition (6)** are fulfilled, for some given constant $b > 0$, i.e.,

$$s_d \int_0^{+\infty} t^{d/2-1} g(t) dt = 2, \quad s_{d-1} \int_0^{+\infty} t^{d/2-3/2} g(t) dt = 2b. \tag{15}$$

As proven in Section 2, this is in particular the case when the conditions of Proposition 3 (when $\Sigma \neq \mathbf{I}_d$) or Proposition 4 (when $\Sigma = \mathbf{I}_d$) are satisfied. Any candidate for the underlying density generator may be normalized to satisfy the two latter conditions, for instance through a transform $t \mapsto \alpha g(\beta t)$ with two conveniently chosen positive constants α and β . This leads to the “normalizing” Algorithm 1. In practical terms, the choice of b does not really matter. We simply advise to set $b = 1$ by default, our choice hereafter.

Algorithm 1: Normalization of a meta-elliptical copula generator

Input: An estimate \hat{g} of the generator for a meta-elliptical copula of dimension $d \geq 2$.

Compute $\mathcal{I}_1 = \int_0^{+\infty} t^{d/2-1} \hat{g}(t) dt$ and $\mathcal{I}_2 = \int_0^{+\infty} t^{d/2-3/2} \hat{g}(t) dt$;

Set $\beta = (b s_d \mathcal{I}_1 / (s_{d-1} \mathcal{I}_2))^2$ and $\alpha = 2\beta^{d/2} / (s_d \mathcal{I}_1)$;

Calculate $\tilde{g} := \{t \mapsto \alpha \times \hat{g}(\beta \times t)\}$;

Output: A modified version \tilde{g} satisfying the normalization and identification constraints (Eq. (2) and (6)).

As usual, the marginal distributions $F_j, j \in \{1, \dots, d\}$ of $\mathbf{X} \sim \mathcal{TE}_d(\Sigma, g, F_1, \dots, F_d)$ will be consistently estimated by their empirical counterparts \hat{F}_j : for every $x \in \mathbb{R}$ and $j, \hat{F}_j(x) := n^{-1} \sum_{i=1}^n \mathbf{1}(X_{i,j} \leq x)$, where $\mathbf{X}_i := (X_{i,1}, \dots, X_{i,d})$. As announced, the goal is now to propose nonparametric estimators of the generator g , assuming that g is identifiable. To the best of our knowledge, this paper is the first to propose solutions to this problem in a well-suited rigorous theoretical framework.

We will use the following notations:

- $U_k := F_k(X_k)$ and $\hat{U}_{i,k} := \hat{F}_k(X_{i,k})$ for $k \in \{1, \dots, d\}, i \in \{1, \dots, n\}$. Set $\mathbf{U} := (U_1, \dots, U_d)$ and $\hat{\mathbf{U}} := (\hat{U}_{i,1}, \dots, \hat{U}_{i,d}), i \in \{1, \dots, n\}$.
- the sample of (unobservable) realizations of \mathbf{U} is $\mathcal{U} := (\mathbf{U}_1, \dots, \mathbf{U}_n)$; the sample of pseudo-observations $\hat{\mathbf{U}} := (\hat{\mathbf{U}}_1, \dots, \hat{\mathbf{U}}_n)$.

Since $\mathbf{X} \sim \mathcal{TE}(\Sigma, g, F_1, \dots, F_d)$, note that the law of $\mathbf{U} \sim \mathcal{ME}(\Sigma, g)$ is the meta-elliptical copula C of \mathbf{X} .

3.1. Penalized M-estimation

Without any particular parametric assumption, regular generators g are living in the infinite dimensional functional space \mathcal{G} . Its subset of generators that satisfy the two identifiability constraints (2) and (6) will be denoted by \mathcal{I} . In practical terms, we could approximate \mathcal{G} by finite dimensional parametric families $\mathcal{G}_m = \{g_\theta; \theta \in \Theta_m\}$, where the dimension of Θ_m is denoted by p_m . Most of the time, $p_m \rightarrow \infty$ with m , and the family (\mathcal{G}_m) is increasing: $\mathcal{G}_m \subset \mathcal{G}_{m+1}$, even if this requirement is not mandatory. Moreover, it is usual that $m = m_n$ and m_n tends to infinity with n , as in the “method of sieves” for inference purpose (see the survey [4], e.g.). Nonetheless, we do not impose the latter constraint again. Ideally, $\cup_m \mathcal{G}_m$ is dense in \mathcal{G} for a convenient norm (typically, in a L^r space). A less demanding requirement would be to assume $g \in \cup_m \mathcal{G}_m$, a condition that is sufficient for our purpose. Therefore, a general estimator of θ would be

$$\hat{\theta}_{n,m} := \arg \min_{\theta \in \Theta_m} \mathbb{G}_n(\theta, \hat{\mathcal{U}}) + \mathbf{p}_n(\lambda_n, \theta), \tag{16}$$

for some empirical loss function \mathbb{G}_n , some penalty $\mathbf{p}_n(\cdot, \cdot)$ and some tuning parameter λ_n .

Typically, the loss function is an average of the type

$$\mathbb{G}_n(\theta, (\mathbf{u}_1, \dots, \mathbf{u}_n)) = \frac{1}{n} \sum_{i=1}^n \ell_n(\theta, \mathbf{u}_i),$$

for some map ℓ_n from $\Theta_n \times \mathbb{R}^d$ to \mathbb{R} . For instance, for the penalized canonical maximum likelihood method, set $\mathbb{G}_n(\theta, \hat{\mathcal{U}}) = -\sum_{i=1}^n \ln c_{g_\theta}(\hat{\mathbf{U}}_i) / n$. Many other examples of loss functions could be proposed, based on L^r -type distances between cdfs’, densities or even characteristic functions.

Concerning the choice of \mathcal{G}_m , an omnibus strategy could be to rely on Bernstein approximations: for any $a > 0$, define the family of polynomials

$$\mathcal{G}_{m,a} := \{g \in \mathcal{I} : g(x) = \sum_{k=0}^m b_k (a+x)^k (a-x)^{m-k} \mathbf{1}(x \in [0, a]), \quad b_k \in \mathbb{R}^+ \quad \forall k\}.$$

Any continuous map g can be uniformly approximated on $[0, a]$ by Bernstein polynomials that are members of $\mathcal{G}_{m,a}$, for m sufficiently large (see [27], e.g.). If $g(t) \rightarrow 0$ when t tends to infinity, then, for every $\epsilon > 0$, there exist $a(\epsilon) > 0$, an integer $m(\epsilon)$ and a map g_ϵ in $\mathcal{G}_{m(\epsilon), a(\epsilon)}$ such that $\|g - g_\epsilon\|_\infty < \epsilon$. If g is compactly supported, simply set $a(\epsilon)$ as the upper bound of g ’s support. If $g \in L^r, r > 0$ and is continuous, then there exists a similar approximation in L^r . Therefore, the set \mathcal{G}_m in (16) may be chosen as $\mathcal{G}_{m(\epsilon), a(\epsilon)}$ for a given $\epsilon \ll 1$, obtained with prior knowledge about the true underlying density generator.

Alternatively, if $g \in L^2(\mathbb{R})$, introduce an orthonormal basis $(h_n)_{n \geq 0}$ of the latter Hilbert space, say the Hermite functions. Then, g can be decomposed as $g = \sum_{k \geq 0} \langle g, h_k \rangle h_k$ and \mathcal{G}_m could be defined as

$$\mathcal{G}_m := \{h^+ = \max(h, 0) : h(x) = \sum_{k=0}^m b_k h_k(x), b_k \in \mathbb{R} \forall k\}.$$

Note that every latter map h^+ is not differentiable at a finite number of points, the roots of h . Thus, the members of \mathcal{G}_m as defined above do not satisfy Condition 1. This can be seen as a theoretical drawback and this can be removed by conveniently smoothing the functions of \mathcal{G}_m . For instance, replace all latter maps h^+ by a smoothed approximation $x \mapsto h^+ * \phi_N(x)$, $\phi_N(t) := N \exp(-\pi N^2 t^2)$, and for some $N \gg 1$ and any t .

Another alternative family of generators supported on \mathbb{R}^+ could be

$$\mathcal{G}_m = \left\{ g \in \mathcal{I} : g(x) = \sum_{\ell=1}^m Q_\ell^2(x) \exp(-(x - \mu_\ell)^{\alpha_\ell} / \sigma_\ell), \text{ for some polynomials } Q_\ell, \right. \\ \left. \deg(Q_\ell) \leq q_m, \text{ and constants } (\mu_\ell, \alpha_\ell, \sigma_\ell) \in \mathbb{R} \times \mathbb{R}_+^* \times \mathbb{R}_+, \ell \in \{1, \dots, m\} \right\},$$

where (q_m) denotes a sequence of integers that tends to infinity with m . The dimension p_m of \mathcal{G}_m (or \mathcal{O}_m , similarly) is then $p_m = m(q_m + 4)$. We do not know whether $\cup_m \mathcal{G}_m$ is dense in L^r for any $r > 0$ and well-chosen sequences (q_m) . Nonetheless, we conjecture that most “well-behaved” density generators can be accurately approximated in some L^r spaces by some elements of \mathcal{G}_m , at least when m and q_m are sufficiently large.

If we had observed true realizations of \mathbf{U} , i.e., if \mathcal{U} replaces $\widehat{\mathcal{U}}$ in (16), then one could apply some well established theory of penalized estimators: see Fan and Li [7], Fan and Peng [8] (asymptotic properties), Loh [31] (finite distance properties), among others. When the loss \mathbb{G}_n is the empirical likelihood and there is no penalty, $\hat{\theta}_{n,m}$ is called the Canonical Maximum Likelihood estimator of θ_0 [14,36], assuming the true density $g = g_{\theta_0}$ belongs to \mathcal{G}_m . Tsukahara [39,40] has developed the corresponding theory in the wider framework of rank-based estimators. When the parameter dimension is fixed (p_m is a constant), the limiting law of $\hat{\theta}_{n,m}$ can be deduced as a consequence of the weak convergence of an empirical copula process, here the empirical process associated to $\widehat{\mathcal{U}}$: see [2,11,15,16]. Nonetheless, to the best of our knowledge, the single existing general result that is able to simultaneously manage pseudo-observations and penalizations is [32]. The latter paper extends [31] to state finite distance boundaries for some norms of the difference between the estimated parameter and the true one. In this work, we slightly extend their results to deal with (16), i.e., with a sequence of parameter spaces (\mathcal{O}_m) . This yields the finite distance properties of $\hat{\theta}_{n,m}$: see Appendix B in the supplementary material.

3.2. Simulation-based inference

The previous inference strategy is mainly of theoretical purpose and would impose difficult numerical challenges. In particular, the log-likelihood criteria require the evaluation of copula densities through g, f_g and F_g . Unfortunately, F_g is a complex map that is not known analytically in general. Here, we propose a way of avoiding the numerical calculations of F_g, Q_g or f_g , by using the fact that it is very easy to simulate elliptical random vectors.

To be specific, consider a trans-elliptical random vector $\mathbf{X} \sim \mathcal{TE}_d(\Sigma, g, F_1, \dots, F_d)$. Its copula C is then meta-elliptical $\mathcal{ME}_d(\Sigma, g)$. By definition and with our notations, C must satisfy

$$C(F_g(x_1), \dots, F_g(x_d)) = H_{g, \Sigma, d}(x_1, \dots, x_d),$$

for every $\mathbf{x} := (x_1, \dots, x_d)^\top$. Note that all margins are the same because Σ is a correlation matrix. The goal is still to estimate g , with a sufficient amount of flexibility. To fix the idea, assume that g belongs to $\cup_m \mathcal{G}_m$, with the same notations as in Section 3.1. Therefore, as in Section 3.1, the idea would be to approximate the true generator g by some map that belongs to a parametric family \mathcal{G}_m , for some “large” m .

First, let us estimate the copula C non-parametrically, for instance by the empirical copula C_n based on the sample $(\mathbf{X}_i)_{i \in \{1, \dots, n\}}$.

Second, assume for the moment that (g, Σ) is known. For any arbitrarily large integer N , let us draw a N sample $(\mathbf{Y}_1, \dots, \mathbf{Y}_N)$ of independent realizations of $\mathbf{Y} \sim H_{g, \Sigma, d}$. This is easy thanks to the polar decomposition of elliptical vectors: $\mathbf{Y} \stackrel{\text{law}}{=} \mathbf{R} \mathbf{A}^\top \mathbf{V}$, where $\mathbf{A}^\top \mathbf{A} = \Sigma$, \mathbf{V} is uniformly distributed on the unit ball in \mathbb{R}^d , and R has a density that is a simple function of g . All margins of \mathbf{Y} have the same distributions F_g , and denote by \hat{F}_g an empirical counterpart: for every $y \in \mathbb{R}$,

$$\hat{F}_g(y) := \frac{1}{dN} \sum_{l=1}^d \sum_{k=1}^N \mathbf{1}(Y_{k,l} \leq y).$$

Moreover, denote by $\hat{H}_{g, \Sigma}$ the joint empirical cdf of \mathbf{Y} , i.e., $\hat{H}_{g, \Sigma}(\mathbf{y}) = N^{-1} \sum_{l=1}^N \mathbf{1}(\mathbf{Y}_l \leq \mathbf{y})$. Then, we expect we approximately satisfy the relationship

$$C_n \circ \vec{F}_g(\mathbf{x}) := C_n(\hat{F}_g(x_1), \dots, \hat{F}_g(x_d)) \simeq \hat{H}_{g, \Sigma}(x_1, \dots, x_d), \mathbf{x} \in \mathbb{R}^d.$$

Obviously, since we do not know g nor Σ , we cannot draw the latter sample $(\mathbf{Y}_i)_{i \in \{1, \dots, N\}}$ strictly speaking. Then, we will replace Σ by a consistent estimator $\hat{\Sigma}$. Moreover, for any current parameter value θ , we can generate the N sample $(\mathbf{Y}_{\theta,1}, \dots, \mathbf{Y}_{\theta,N}), \mathbf{Y}_{\theta,l} \sim \mathcal{E}_d(\mathbf{0}, \hat{\Sigma}, g_\theta)$. The associated empirical marginal and joint cdfs are denoted by \hat{F}_{g_θ} and $\hat{H}_{g_\theta, \hat{\Sigma}}$. Therefore, for a fixed m , an approximation of g will be given by $g_{\hat{\theta}_m}$, where

$$\hat{\theta}_m := \arg \min_{\theta \in \Theta_m} \mathcal{D}(C_n \circ \bar{F}_{g_\theta}; \hat{H}_{g_\theta, \hat{\Sigma}}),$$

for some discrepancy \mathcal{D} between cdfs' on \mathbb{R}^d and some parameter set Θ_m in \mathbb{R}^{p_m} . Note that $\hat{\theta}_m$ implicitly depends on n and N . For instance, consider

$$\mathcal{D}(C_n \circ \bar{F}_{g_\theta}; \hat{H}_{g_\theta, \hat{\Sigma}}) := \int \left\{ C_n(\hat{F}_{g_\theta}(x_1), \dots, \hat{F}_{g_\theta}(x_d)) - \hat{H}_{g_\theta, \hat{\Sigma}}(\mathbf{x}) \right\}^2 w(\mathbf{x}) d\mathbf{x},$$

for some weight function w . To avoid the calculation of d -dimensional integrals, it is possible to choose some ‘‘chi-squared’’ type discrepancies instead, as

$$\mathcal{D}_\chi(C_n \circ \bar{F}_{g_\theta}; \hat{H}_{g_\theta, \hat{\Sigma}}) := \sum_{l=1}^L \left| \int_{B_l} dC_n(\hat{F}_{g_\theta}(x_1), \dots, \hat{F}_{g_\theta}(x_d)) - \hat{H}_{g_\theta, \hat{\Sigma}}(d\mathbf{x}) \right|,$$

for some partition (B_1, \dots, B_L) of \mathbb{R}^d . Alternatively, we could replace the measure $w(\mathbf{x}) d\mathbf{x}$ with the empirical law of our observations, yielding

$$\mathcal{D}_{emp}(C_n \circ \bar{F}_{g_\theta}; \hat{H}_{g_\theta, \hat{\Sigma}}) := \frac{1}{N} \sum_{i=1}^N \left\{ C_n(\hat{F}_{g_\theta}(X_{1,i}), \dots, \hat{F}_{g_\theta}(X_{d,i})) - \hat{H}_{g_\theta, \hat{\Sigma}}(\mathbf{X}_i) \right\}^2,$$

and this modification avoids suffering from the curse of dimensionality. Unfortunately, whatever the chosen criterion, stating the limiting law of $\hat{\theta}_m$ when N and n tend to infinity (or even the limiting law of $g_{\hat{\theta}_m} - g$ when m, N and n tend to infinity) seems to be a particularly complex task that lies beyond the scope of this paper.

3.3. An iterative algorithm: MECIP, or ‘‘Meta-Elliptical Copula Iterative Procedure’’

In this section, we propose a numerical recursive procedure called MECIP that allows the estimation of g .

The supplementary material refers to several estimation procedures of the density generator of an elliptical distribution. We select one of them, that will be called A . Formally, A is the operator that maps an i.i.d. dataset $(\mathbf{Z}_1, \dots, \mathbf{Z}_n) \in \mathbb{R}^{d \times n}$ generated from an elliptical distribution $\mathcal{E}_d(\boldsymbol{\mu}, \Sigma, g)^{\otimes n}$ to a map $\hat{g} := A(\mathbf{Z}_1, \dots, \mathbf{Z}_n)$:

$$A : \begin{array}{ll} \mathbb{R}^{d \times n} & \longrightarrow \mathcal{F} \\ (\mathbf{Z}_1, \dots, \mathbf{Z}_n) & \longmapsto \hat{g}, \end{array} \tag{17}$$

where \mathcal{F} is the set of all possible density generators of elliptical distributions. Note that any g in \mathcal{F} has to satisfy (2), but not (6).

In our case, we do not have access to a dataset following an elliptical distribution. Assume that we observe a dataset $\mathbf{U}_1, \dots, \mathbf{U}_n$ following a meta-elliptical copula $\mathcal{ME}(\Sigma, g)$. If we knew the true generator g , we could compute the univariate quantile function Q_g , and, as a consequence, $\bar{Q}_g(\mathbf{U}_i) := (Q_g(U_{i,1}), \dots, Q_g(U_{i,d}))^\top \sim \mathcal{E}_d(\mathbf{0}, \Sigma, g)$ for every $i \in \{1, \dots, n\}$. Therefore, we could define an ‘‘oracle estimator’’ of g by

$$\hat{g}^{oracle} := A(\bar{Q}_g(\mathbf{U}_1), \dots, \bar{Q}_g(\mathbf{U}_n)) =: A_n(g). \tag{18}$$

In practice, two issues arise that prevent us from using this oracle estimator. First, we do not have access to the true distributions F_1, \dots, F_d , but only to empirical cdfs $\hat{F}_1, \dots, \hat{F}_d$. As usual, it is possible to replace the ‘‘unobservable’’ realizations $U_{i,k} = F_k(X_{i,k})$ by pseudo-observations $\hat{U}_{i,k} = \hat{F}_k(X_{i,k})$ in Eq. (18). Second and more importantly, we need Q_g , i.e., g itself, to compute the oracle estimator \hat{g}^{oracle} . This looks like an impossible task.

To solve the problem, we propose an iterative algorithm as follows. We fix a first estimate $\hat{g}^{(0)}$ of g , so that we can compute $\hat{\mathbf{Z}}_i^{(1)} = \bar{Q}_{\hat{g}^{(0)}}(\hat{\mathbf{U}}_i)$, for every $i \in \{1, \dots, n\}$. From this first guess, we can compute an estimator $\hat{g}^{(1)} := A(\hat{\mathbf{Z}}_1^{(1)}, \dots, \hat{\mathbf{Z}}_n^{(1)})$. Note that this estimator should be normalized in order to satisfy the necessary condition that is related to the identifiability of g . At this stage, we impose Condition (6), for a fixed constant b , and invoke Algorithm 1. Iteratively, for any $N \in \mathbb{N}$, we define $\hat{\mathbf{Z}}_i^{(N)} = \bar{Q}_{\hat{g}^{(N-1)}}(\hat{\mathbf{U}}_i)$ and $\hat{g}^{(N)} := A(\hat{\mathbf{Z}}_1^{(N)}, \dots, \hat{\mathbf{Z}}_n^{(N)})$. This procedure MECIP is detailed in Algorithm 2 below. See Fig. 1 too.

To summarize, for a fixed sample size n , the latter recursive algorithm is classical in the domain of fixed points analysis: we assume there exists a generator such that $g_n^* = A_n(g_n^*)$ and we approximate g_n^* by a recursion $\hat{g}_n^{N+1} = A_n(\hat{g}_n^N)$. When $N \rightarrow \infty$, we hope that $\hat{g}_n^N \simeq g_n^*$. This is the case if A_n is a contraction (Banach Contraction Principle), but the latter property is not guaranteed. Moreover, with a large n , we hope that g_n^* tends to the true underlying generator g , because of the consistency of the estimation procedure A_n . In other words, our algorithm is a mix between fixed point search procedures

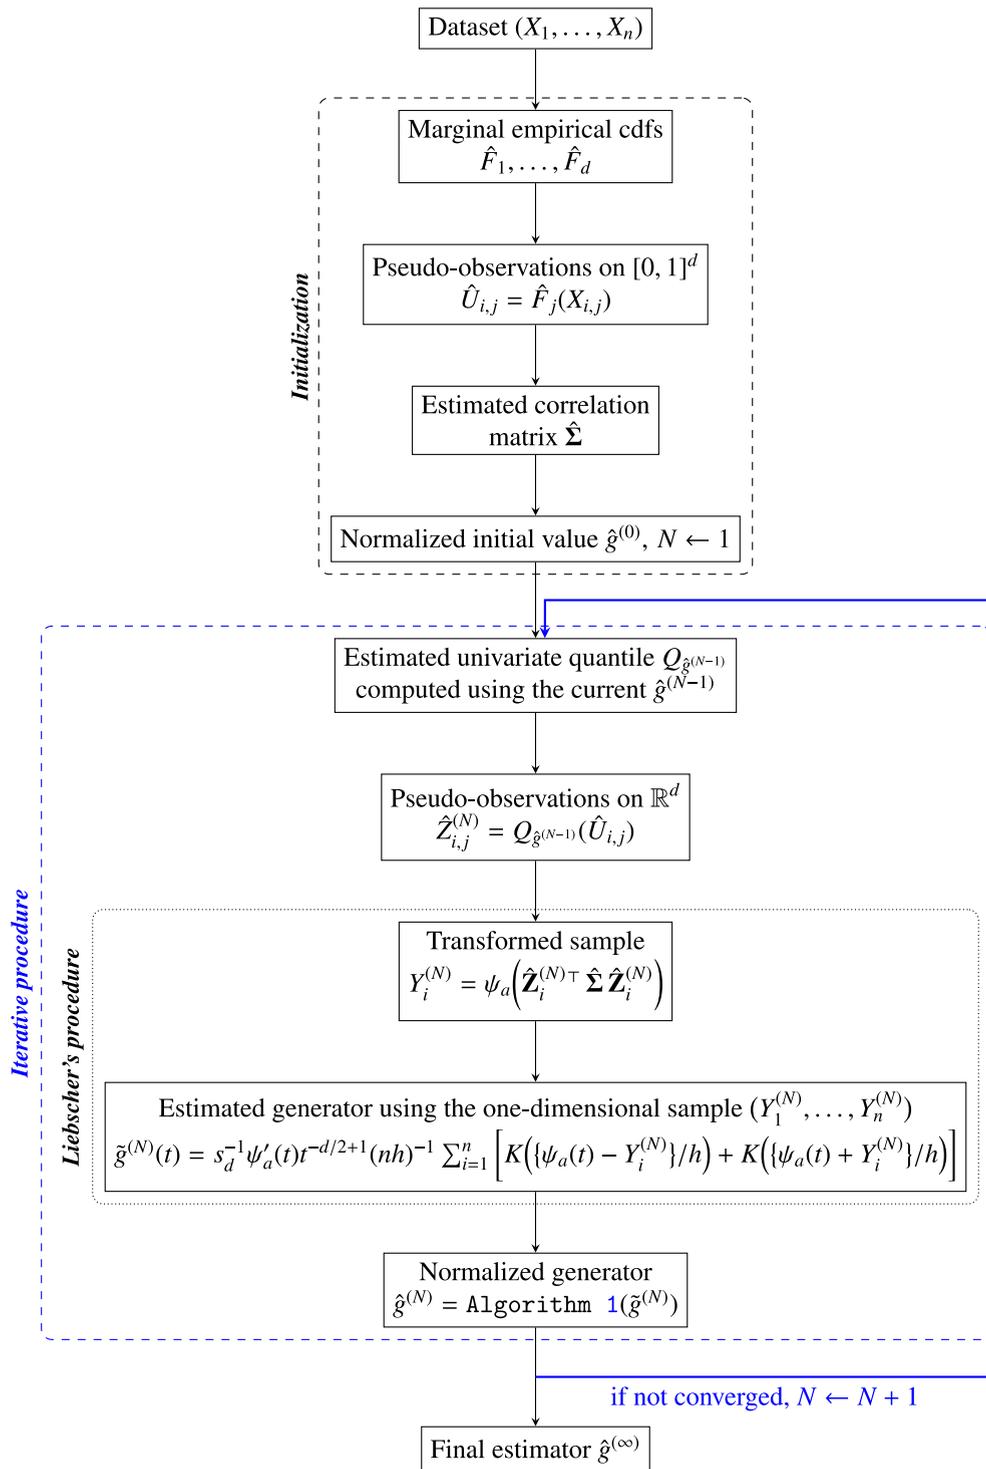


Fig. 1. Simplified flowchart of the iterative estimation procedure MECIP.

and classical nonparametric inference. The theoretical study of its convergence properties appears as particularly complex, due to its multiple nonlinear stages.

Hereafter, A will be chosen as Liebscher's estimation procedure [28] to evaluate the generator of an elliptical distribution. It requires the introduction of the instrumental map ψ_a defined by $\psi_a(x) := -a + (a^{d/2} + x^{d/2})^{2/d}$ for any

$x \geq 0$ and some constant $a > 0$. Since Liebscher’s method is non-parametric, we need a usual univariate kernel K (here, the Gaussian kernel) and a bandwidth $h = h_n, h_n \rightarrow 0$ when n tends to infinity. See the exact formula of Liebscher’s estimator below, in Algorithm 2.

We will propose three ways of initializing the algorithm:

- (i) the “Gaussian” initialization, where $\hat{g}^{(0)} := e^{-x}$, suitably normalized (with the notations of Algorithm 1, $\mathcal{I}_1 = \Gamma(d/2)$ and $\mathcal{I}_2 = \Gamma(d/2 - 1/2)$);
- (ii) the “identity” initialization, in which $\hat{g}^{(0)} := A((\hat{U}_{i,j})_{1 \leq i \leq n, 1 \leq j \leq d})$ as if, in lines 8–9 of Algorithm 2, the quantile function $Q_{\hat{g}^{(N-1)}}$ were replaced by the identity map;
- (iii) the “ $A \sim \Phi^{-1}$ ” initialization, where $\hat{g}^{(0)} := A((\Phi^{-1}(\hat{U}_{i,j}))_{1 \leq i \leq n, 1 \leq j \leq d})$ and Φ denotes the cdf of a Gaussian $\mathcal{N}(0, 1)$ distribution.

Solution (ii) may seem a brutal approximation and can be considered as an “uninformative prior”. It can actually be well-suited, compared to (i), when the true generator is far from the Gaussian generator: see Fig. 3. The main motivation of (iii) is to put U whose support is $[0, 1]$ back on the whole real line \mathbb{R} , through a usual numerical trick.

The implementation of the algorithm of MECIP is available in the R package `ELliptCopulas` [5].

Algorithm 2: Basic iteration-based procedure MECIP for the estimation of the elliptical copula density generator g with a given method A

Input: A dataset $(\mathbf{X}_1, \dots, \mathbf{X}_n), n > 0$.

Input: An estimation method A for elliptical distribution density generators.

for $j \leftarrow 1$ **to** d **do**

Compute the empirical cdf \hat{F}_j and the pseudo-observations $\hat{U}_{i,j} := \hat{F}_j(X_{i,j})$ for all $i \in \{1, \dots, n\}$;

end

Compute an estimator of the correlation matrix $\hat{\Sigma}$ of the elliptical copula using Kendall’s taus ;

Initialize $N := 1$. Initialize g to a value $\hat{g}^{(0)}$;

Normalize $\tilde{g}^{(0)} := \text{Algorithm 1}(\hat{g}^{(0)})$;

repeat

Compute the univariate quantile function $Q_{\tilde{g}^{(N-1)}}$ associated with the elliptical distribution $\mathcal{E}_d(\cdot, \mathbf{I}_d, \tilde{g}^{(N-1)})$;

For every $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, d\}$, compute $\hat{Z}_{i,j}^{(N)} := Q_{\tilde{g}^{(N-1)}}(\hat{U}_{i,j})$;

Update $\hat{g}^{(N)} := A(\hat{\mathbf{Z}}_1^{(N)}, \dots, \hat{\mathbf{Z}}_n^{(N)}; \hat{\Sigma})$, corresponding to the following step when A is Liebscher’s procedure with a given $a > 0$;

begin

/* Liebscher’s procedure */

For $i \in \{1, \dots, n\}$, let $Y_i = \psi_a(\hat{\mathbf{Z}}_i^{(N)\top} \hat{\Sigma}^{-1} \hat{\mathbf{Z}}_i^{(N)})$;

Let $\tilde{g}^{(N)}(t) = s_d^{-1} \psi'_a(t) t^{-d/2+1} (nh)^{-1} \sum_{i=1}^n \left[K\left(\frac{\psi_a(t) - Y_i}{h}\right) + K\left(\frac{\psi_a(t) + Y_i}{h}\right) \right]$;

end

Normalize $\hat{g}^{(N)} := \text{Algorithm 1}(\hat{g}^{(N)})$;

Update $N = N + 1$;

until convergence of \hat{g} ;

Output: A normalized estimator $\hat{g}^{(\infty)}$ of g .

3.4. Adjustments of the iterative algorithm MECIP in the presence of missing values

Whenever a dataset contains missing values, the previous numerical procedure can be adapted to estimate the correlation matrix Σ and the generator of the underlying elliptical copula. We consider the simplest case of missing at random observations. Note that many other missing patterns may exist, but a complete treatment of these cases is left for future work.

When missing values arise, the previous Algorithm 2 will be adjusted as follows:

1. Each empirical cdf \hat{F}_j is estimated using all non-missing observations for the j th variable.
2. Pseudo-observations $\hat{U}_{i,j} := \hat{F}_j(X_{i,j})$ are defined as “NA” whenever $X_{i,j} = \text{NA}$ (i.e., is missing).
3. Kendall’s taus are estimated using pairwise complete observations. In other words, for two variables $1 \leq j_1 \neq j_2 \leq d$, the Kendall’s tau between X_{j_1} and X_{j_2} is estimated using the set of observations $\{i \in \{1, \dots, n\} : X_{i,j_1} \neq \text{NA} \text{ and } X_{i,j_2} \neq \text{NA}\}$.

If the correlation matrix is not positive semi-definite, it is projected on the (convex) set of positive semi-definite matrices using the R function `nearPD`, which implements the method proposed in [18]. See alternative methods in [19,33], among others.

4. The pseudo-observations $\hat{Z}_{i,j}^{(N)} := Q_{\hat{g}^{(N-1)}}(\hat{U}_{i,j})$ at iteration N are defined to be NA whenever $\hat{U}_{i,j} = \text{NA}$.
5. At each step of the main loop, for every $i \in \{1, \dots, n\}$ such that some $\hat{Z}_{i,j}^{(N)}$ is NA, we complete the vector $\hat{\mathbf{Z}}_i^{(N)}$ in the following way: let $\text{miss}(i) = \{j \in \{1, \dots, d\} : \hat{Z}_{i,j}^{(N)} = \text{NA}\}$ be the set of the missing components for the i th observation. The non-missing part of $\hat{\mathbf{Z}}_i^{(N)}$, denoted $\hat{\mathbf{Z}}_{i,-\text{miss}(i)}^{(N)}$, is left unchanged. If we knew the true generator g and the correlation matrix Σ , we would use the non-missing part $\hat{\mathbf{Z}}_{i,-\text{miss}(i)}^{(N)}$ of the random vector $\hat{\mathbf{Z}}_i^{(N)}$ to complete the other entries by using their conditional law. Indeed, if a vector \mathbf{Z} follows an elliptical distribution $\mathcal{E}_d(\boldsymbol{\mu}, \Sigma, g)$, then, for any subset $I \subset \{1, \dots, d\}$, the law of \mathbf{Z}_I given $(\mathbf{Z}_{-I} = \mathbf{z}_{-I})$ is still elliptical $\mathcal{E}_{|I|}(\boldsymbol{\mu}_{\mathbf{z}_{-I}}, \Sigma_{\mathbf{z}_{-I}}, g_{\mathbf{z}_{-I}})$. Some explicit expressions for these three conditional parameters are given in [3, Corollary 5]. Therefore, an ‘‘oracle’’ way of generating $\hat{\mathbf{Z}}_{i,\text{miss}(i)}^{(N)}$ is to draw

$$\hat{\mathbf{Z}}_{i,\text{miss}(i)}^{(N)} \sim \mathcal{E}_{|\text{miss}(i)|}(\boldsymbol{\mu}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}, \Sigma_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}, g_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}),$$

neglecting the fact that $\hat{\mathbf{Z}}_i^{(N)}$ follows only approximately an elliptical distribution (unless $\hat{g}^{(N)} = g$, which is rather unlikely). However, Σ and g are unknown; then, we propose to replace Σ by its empirical counterpart and g by its most recent estimate $\tilde{g}^{(N)}$. Finally, we obtain the updated ‘‘feasible’’ generating formula

$$\hat{\mathbf{Z}}_{i,\text{miss}(i)}^{(N)} \sim \mathcal{E}_{|\text{miss}(i)|}(\hat{\boldsymbol{\mu}}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}, \hat{\Sigma}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}, \hat{g}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}^{(N)}), \tag{19}$$

for some approximate conditional mean $\hat{\boldsymbol{\mu}}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}$ based on $\hat{\Sigma}$, some approximate conditional correlation matrix $\hat{\Sigma}_{\hat{\mathbf{z}}_{i,-\text{miss}(i)}^{(N)}}$ based on $\hat{\Sigma}$ and some approximate conditional generator based on $\tilde{g}^{(N)}$ and on $\hat{\Sigma}$.

Algorithm 3: Improved version of the iteration-based procedure MECIP for the estimation of the elliptical copula density generator g with a given method A

Input: A dataset $(\mathbf{X}_1, \dots, \mathbf{X}_n)$, $n > 0$.

Input: An estimation method A for elliptical distribution density generators.

for $j \leftarrow 1$ **to** d **do**

Compute the empirical cdf \hat{F}_j using available data for the j -th variable and the pseudo-observations $\hat{U}_{i,j} := \hat{F}_j(X_{i,j})$ for all $i \in \{1, \dots, n\}$;

end

Compute an estimator of the correlation matrix $\hat{\Sigma}$ of the elliptical copula using Kendall’s taus estimated on pairwise complete observations ;

Initialize $N := 1$. Initialize g to a value $\tilde{g}^{(0)}$;

Normalize $\hat{g}^{(0)} := \text{Algorithm 1}(\tilde{g}^{(0)})$;

repeat

Compute the univariate quantile function $Q_{\hat{g}^{(N-1)}}$ associated with the elliptical distribution $\mathcal{E}_d(\cdot, \mathbf{I}_d, \hat{g}^{(N-1)})$;

For every $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, d\}$, compute $\hat{Z}_{i,j}^{(N)} := Q_{\hat{g}^{(N-1)}}(\hat{U}_{i,j})$;

For every $i \in \{1, \dots, n\}$: \mathbf{Z}_i contains missing values at entries $\text{miss}(i)$, Simulate the missing part $\mathbf{Z}_{i,\text{miss}(i)}^{(N)}$ of $\mathbf{Z}_i^{(N)}$ using Eq. (19) Update $\hat{g}^{(N)} := A(\hat{\mathbf{Z}}_1^{(N)}, \dots, \hat{\mathbf{Z}}_n^{(N)}; \hat{\Sigma})$, corresponding to the following step when A is Liebscher’s procedure with a given $a > 0$;

begin

/* Liebscher’s procedure */

For $i \in \{1, \dots, n\}$, let $Y_i = \psi_a(\hat{\mathbf{Z}}_i^{(N)\top} \hat{\Sigma}^{-1} \hat{\mathbf{Z}}_i^{(N)})$;

Let $\tilde{g}^{(N)}(t) = s_d^{-1} \psi_a'(t) t^{-d/2+1} (nh)^{-1} \sum_{i=1}^n \left[K(\{\psi_a(t) - Y_i\}/h) + K(\{\psi_a(t) + Y_i\}/h) \right]$;

end

Normalize $\hat{g}^{(N)} := \text{Algorithm 1}(\tilde{g}^{(N)})$;

Update $N = N + 1$;

until convergence of \hat{g} ;

Output: A normalized estimator $\hat{g}^{(\infty)}$ of g .

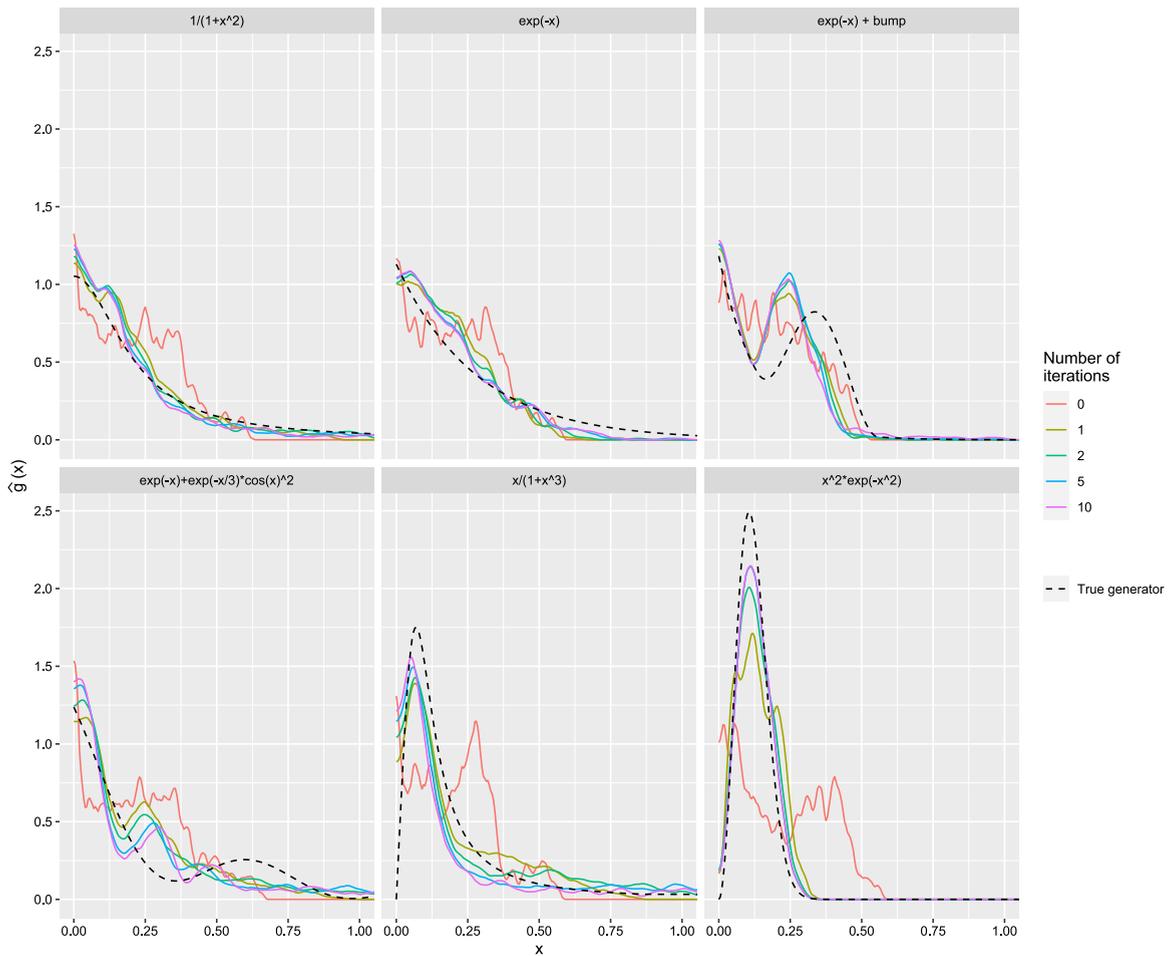


Fig. 2. Estimated generators by the iterative method MECIP, with $n = 1000$, $d = 2$, $\rho = 0.2$, $a = 1$, $h = 0.05$ and starting point = “identity”.

4. Numerical results for the iteration-based method MECIP

4.1. Simulation study in dimension 2

For this simulation study, fix the dimension $d = 2$ and the sample size $n = 1000$. The values of the estimated generators are calculated on a grid of the interval $[0, 10]$ with the step size 0.005. The correlation matrix is chosen as $\Sigma = \begin{pmatrix} 1 & 0.2 \\ 0.2 & 1 \end{pmatrix}$. We use uniform marginal distributions, but still estimate them nonparametrically using the empirical distribution function as if they were unknown.

We consider the normalized versions of six possible generators: $g(x) = 1/(1 + x^2)$, $g(x) = e^{-x}$, $g(x) = e^{-x} + \text{bump}(x)$, $g(x) = e^{-x} + e^{-x/3} \cos^2(x)$, $g(x) = x/(1+x^3)$, and $g(x) = x^2 e^{-x^2}$, where $\text{bump}(x) = \mathbf{1}\{x \in [1, 1+\pi]\}(x-1)(1+\pi-x) \sin(x-1)$ is a smooth function supported on $[1, 1 + \pi]$. The estimated generators obtained with the iterative method (Algorithm 2 using Liebscher’s procedure) are plotted in Fig. 2. In general and after less than $N = 10$ iterations, our estimated generators yield convenient approximations of the true underlying generators, even if $g(0) = 0$ (a case that was excluded by Condition 1). Nonetheless, when g is highly non monotonic, as for “double-bump” generators, the iterative algorithm is less performing and larger sample sizes are required. Note that, when $d = 2$, the parameter a has no influence on the result, since for any $x > 0$, $\psi_a(x) := -a + (a^{d/2} + x^{d/2})^{2/d} = x$.

In Fig. 3, the mean integrated squared error $\text{MISE}(\hat{g}_h) := E[\|\hat{g}_h - g\|_2^2]$ of our iterative estimator is plotted as a function of the bandwidth h , for different true generators g and initialization strategies. These MISE are computed using N equal to ten and 100 replications of each experiment. As expected, a clear-cut optimal bandwidth can be empirically identified for most generators.

Restricting ourselves to the case of the Gaussian generator, we then study the joint influence of the sample size n and of the bandwidth h on the MISE of our estimators, using the “identity” initialization: see Fig. 4. We find the same

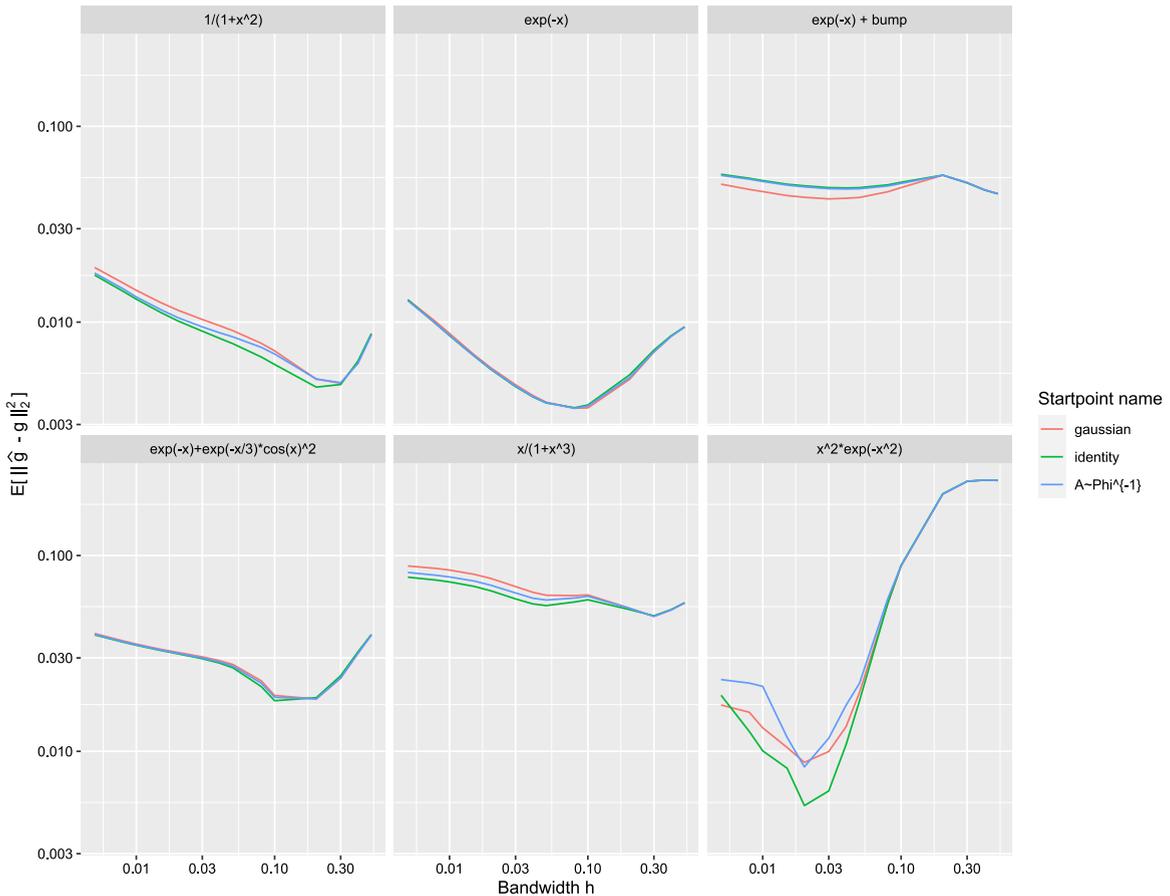


Fig. 3. MISE of the estimates given by Algorithm 2 for $n = 1000$, different choices of the generator g , the bandwidth h and the initialization method.

behaviors as for usual kernel-based estimators: empirically, the optimal bandwidths are “closely” linear functions of $\ln(n)$. The computation time for the three initialization methods are compared in Fig. 6. They are mostly similar. The “Gaussian” initialization is the fastest method as it is not data-dependent.

4.2. Simulation study for higher dimensions

Here, we consider the same sample size $n = 1000$ as before, but the dimension d varies between 3 and 11. The correlation matrix is chosen as $\Sigma_{i,j} = 0.2$ when $i \neq j$. For the Gaussian generator, we then study the performance of our algorithm as a function of the tuning parameters a and h . The results are displayed in Fig. 5. We observe that the MSE increases with the dimension, even for the best choices of the tuning parameters a and h . When $d \geq 5$, avoid choosing a less than one, while the influence of the bandwidth h seems to be less crucial.

Computation times increase only slowly with the dimension d : see Fig. 7. This is because the generator is a univariate function regardless the dimension of the random vector. Therefore, most steps in our algorithms are invariant with respect to the dimension d , except the transformation of the sample $Y_i := \psi_a(\hat{Z}_i^T \hat{\Sigma} Z_i)$. The latter step costs at most $(4d^2 + 2d + 3)n$ elementary operations, a reasonable amount when d is moderate.

Remark 3. Actually, it is possible to bypass the problem of high dimensions d . Indeed, any subvector of an elliptical distribution is itself elliptically distributed. As a consequence, if the copula of a random vector \mathbf{Y} is elliptical with generator g_d , then the copula of a subvector $\mathbf{Y}_{(m)}$ of m components of \mathbf{Y} is still an elliptical copula whose generator g_m is given by Eq. (A.2) in the supplementary material. Therefore, it is possible to estimate the generator of an elliptical copula by using only a sample of m -dimensional subvectors. By a numerical inversion of Eq. (A.2), one would get a generator that corresponds to the copula of the whole vector \mathbf{Y} .

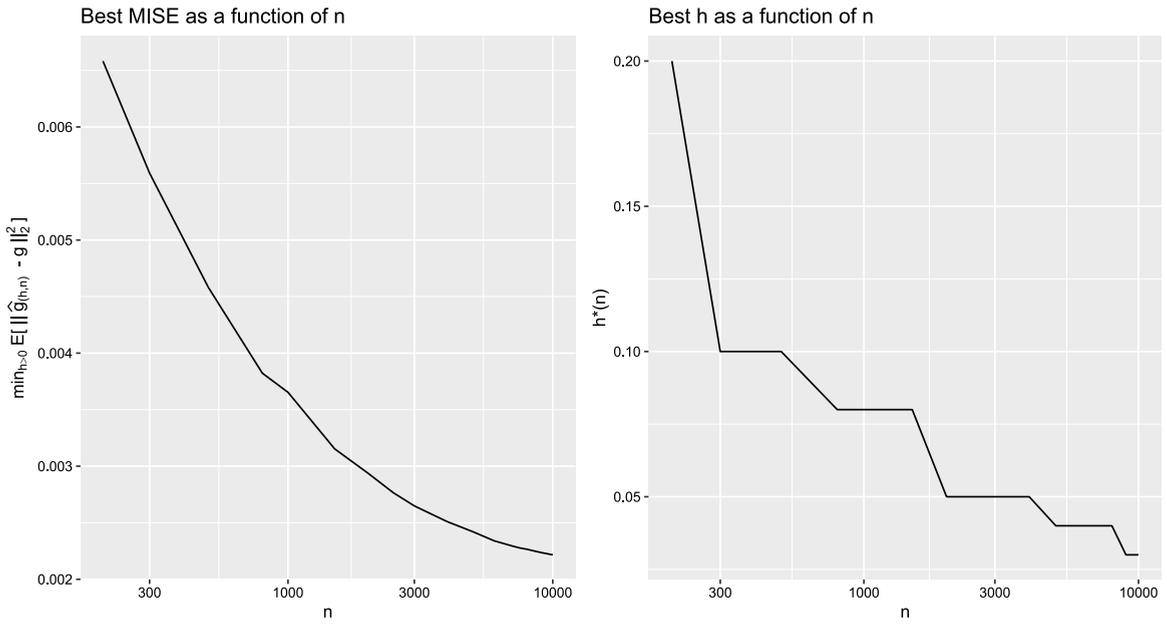


Fig. 4. MISE of the estimate given by Algorithm 2 as a function of n , for the best bandwidth $h = h^*(n)$, and $h^*(n)$ as a function of n . Both plots are in semi-log scale. The true generator is the Gaussian $g(x) = e^{-x}$ with the initialization "identity".

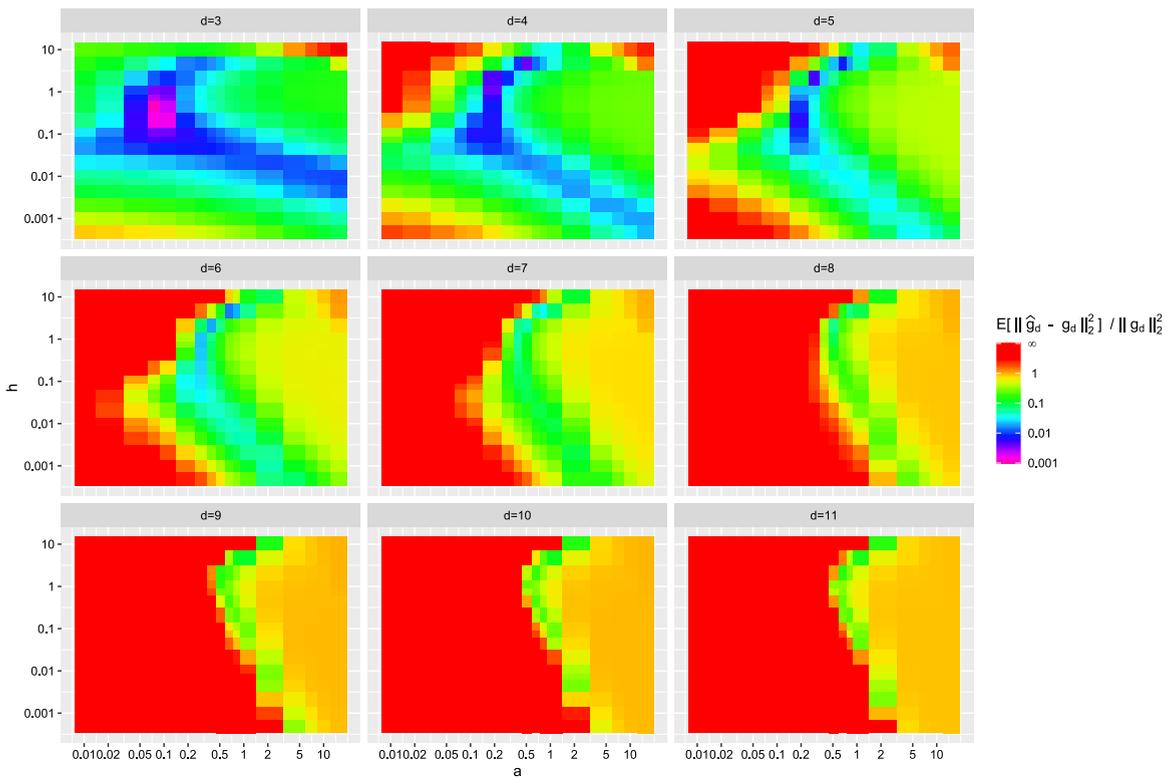


Fig. 5. MISE of the estimate given by Algorithm 2 as a function of a, h , for $n = 1000$, and different choice of the dimension d . The true generator is the Gaussian $g(x) = e^{-x}$ with the initialization "identity".

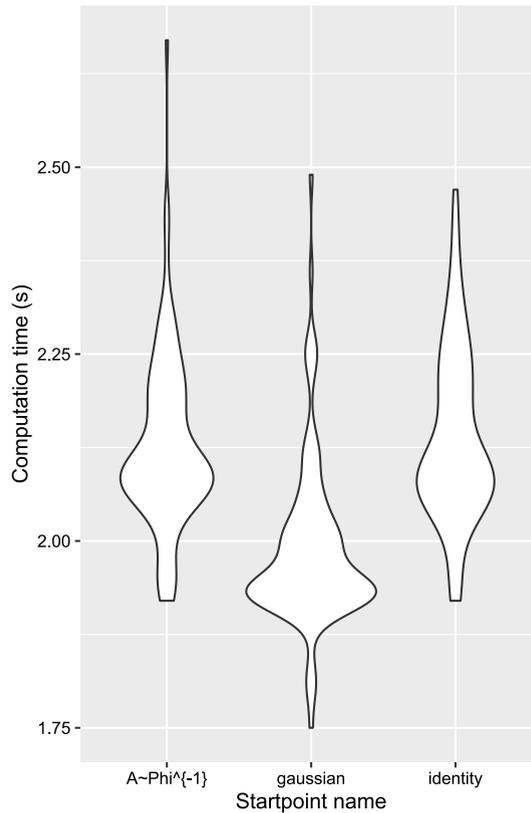


Fig. 6. Density plot of the computation time as a function of the starting point, for $d = 2$ and $n = 1000$.

4.3. When Σ is non exchangeable and almost non invertible

We consider the same setting as in the previous section, and we aim for measuring the effect of the lack of exchangeability in the matrix Σ on the estimation of g . For this, two different frameworks are considered.

- (a) First framework: the dimension is $d = 3$, $h = a = 0.2$ and the correlation matrix is $\Sigma_{(3)}(\rho_{12}) = \begin{pmatrix} 1 & \rho_{12} & 0.2 \\ \rho_{12} & 1 & 0.2 \\ 0.2 & 0.2 & 1 \end{pmatrix}$.
- (b) Second framework: the dimension is $d = 10$, $a = 1$, $h = 0.1$ and the correlation matrix is

$$\Sigma_{(10)}(\rho_{12}) := \begin{pmatrix} 1 & \rho_{12} & \rho_{12} & 0.2 & \cdots & 0.2 \\ \rho_{12} & 1 & \rho_{12} & 0.2 & \cdots & 0.2 \\ \rho_{12} & \rho_{12} & 1 & 0.2 & \cdots & 0.2 \\ 0.2 & 0.2 & 0.2 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \ddots & 0.2 \\ 0.2 & 0.2 & 0.2 & \cdots & 0.2 & 1 \end{pmatrix}$$

Note that there exists a value $\underline{\rho}_{12}$ for which $\Sigma_{(3)}(\rho_{12})$ and $\Sigma_{(10)}(\rho_{12})$ are positive semi-definite if and only if $\rho_{12} > \underline{\rho}_{12}$. For the first and the second frameworks, $\underline{\rho}_{12} \approx -0.919$ and $\underline{\rho}_{12} \approx -0.432$ respectively.

The MISE is computed for both frameworks as a function of ρ_{12} and is displayed respectively in Figs. 8 and 9. On these figures, the MISE stays stable whenever ρ_{12} is not too close to the boundary value $\underline{\rho}_{12}$. When ρ_{12} is close to $\underline{\rho}_{12}$, the estimator of the correlation matrix becomes unreliable, degrading the performance of the estimator \hat{g} . This deterioration is stronger in the second framework where $d = 10$. Note that, as in the previous section, the performance for $d = 10$ is always worse than for $d = 3$ even far away from the boundary.

4.4. Simulation study with missing values

In this case, we choose $n = 1000$, $d = 3$, and the same correlation matrix as before. Contrary to the previous simulation experiments, we introduce some missing values in the dataset (represented by NA in the R environment). The missing

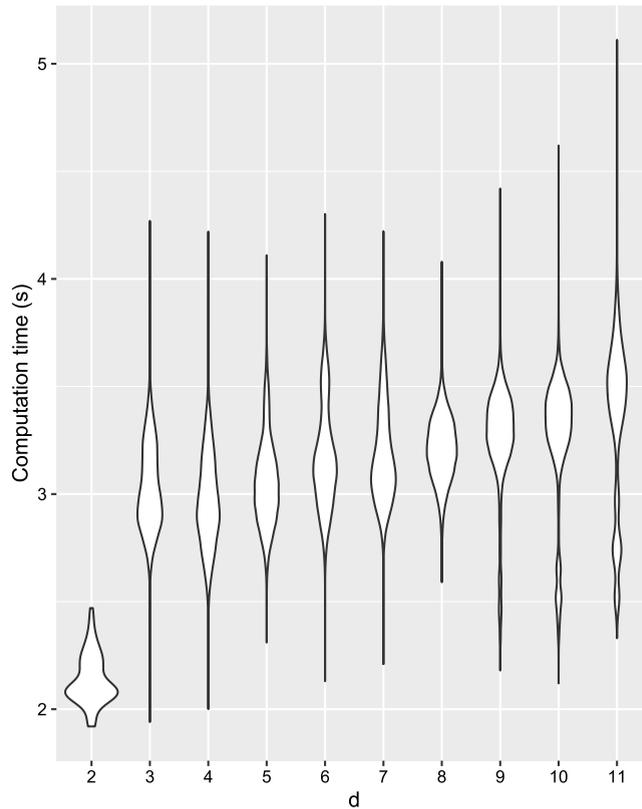


Fig. 7. Density plot of the computation time as a function of the dimension, for the initialization method “identity” and $n = 1000$.

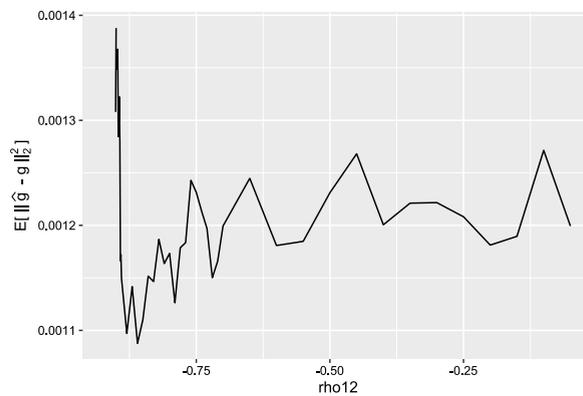


Fig. 8. MISE of the estimate given by Algorithm 2 as a function of ρ_{12} for $d = 3$ and $n = 1000$, where the correlation matrix is $\Sigma_{(3)}(\rho_{12})$.

values are generated as follows: we fix a number $N_{missing}$ of observations that are affected by missing values; we randomly draw a number $N_{missing,1}$ of observations (uniformly between 0 and $N_{missing}$) for which a single component is missing; let $N_{missing,2} := N_{missing} - N_{missing,1}$ be the number of observations for which two components are missing. Because $d = 3$, considering the case of three missing components does not make sense because this would induce an empty vector (which should rather induce a smaller sample size rather than a missing value issue if it were the case). We select $N_{missing,1}$ (respectively $N_{missing,2}$) observations at random, and replace the missing values by NA.

For the estimation procedure, we choose $a = 0.08$ and $h = 0.2$, the optimal choices according to Fig. 5. The corresponding MISE are displayed in Fig. 10. They show that the MISE deteriorates as the number of missing values increases, but this tendency is not very strong.

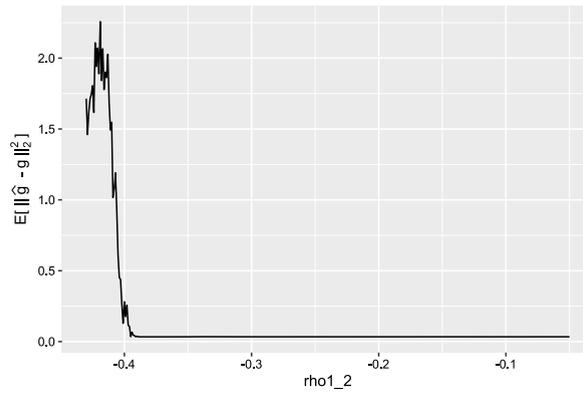


Fig. 9. MISE of the estimate given by Algorithm 2 as a function of ρ_{12} for $d = 10$ and $n = 1000$, where the correlation matrix is $\Sigma_{(10)}(\rho_{12})$.

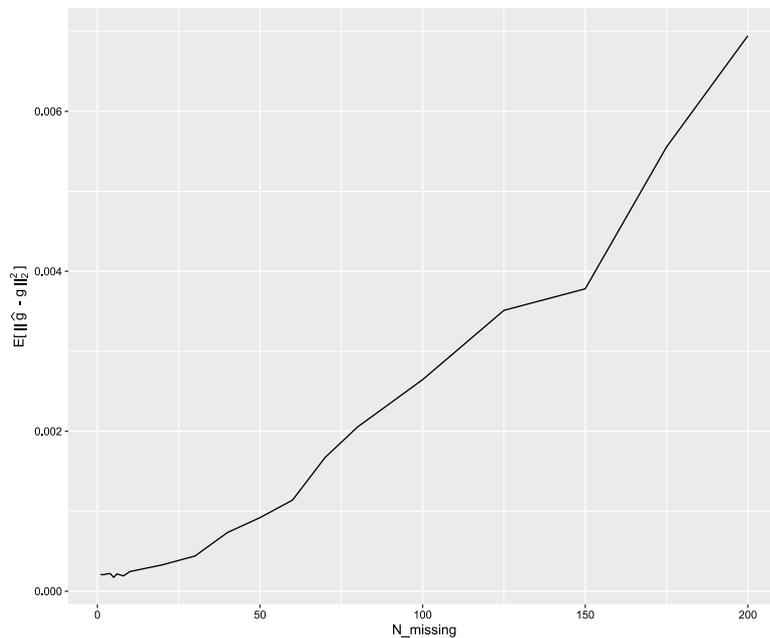


Fig. 10. MISE of the estimate given by Algorithm 2 as a function of the number of observations that contains missing values. The sample size is $n = 1000$, the dimension is $d = 3$ and the true generator is the Gaussian $g(x) = e^{-x}$ with the initialization “identity”.

5. Conclusion

We have stated some sufficient conditions to obtain the identifiability of the generator g of a meta-elliptical copula. In many standard practical situations, they are satisfied, particularly when the corresponding correlation matrix Σ is not the identity matrix. Some inference procedures have been discussed, in particular an iterative method called MECIP that yields satisfying empirical results.

Among the avenues for further studies, a theoretically sound data-driven bandwidth selector would be welcome. The theoretical properties of the iterative algorithm (consistency, rate of convergence) remain unknown, even if the procedure seems to behave conveniently in our experiments. The proof of such results would be particularly challenging, due to the highly nonlinear analytical features of the maps we use in our algorithm MECIP. Moreover, it would be nice to weaken the conditions for the identifiability of g when $\Sigma = \mathbf{I}_d$ (Proposition 4). Finally, we conjecture that our results apply even when $g(0) = 0$, a situation that was excluded for the sake of clarity in the theoretical developments, but that seems to be conveniently managed in our numerical experiments.

Acknowledgment

Jean-David Fermanian has been supported by the labex Ecodec (reference project ANR-11-LABEX-0047).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jmva.2022.104962>.

References

- [1] B. Abdous, C. Genest, B. Rémillard, Dependence properties of meta-elliptical distributions, in: *Statistical Modeling and Analysis for Complex Data Problems*, Springer, 2005, pp. 1–15.
- [2] B. Berghaus, A. Bücher, S. Volgushev, Weak convergence of the empirical copula process with respect to weighted metrics, *Bernoulli* 23 (1) (2017) 743–772.
- [3] S. Cambanis, S. Huang, G. Simons, On the theory of elliptically contoured distributions, *J. Multivariate Anal.* 11 (3) (1981) 368–385.
- [4] X. Chen, Large sample sieve estimation of semi-nonparametric models, in: *Handbook of Econometrics*, Vol.6, Elsevier, 2007, pp. 5549–5632.
- [5] A. Derumigny, J.-D. Fermanian, R package: *ElliptCopulas*, version 0.1.1, 2022, Available at <https://cran.r-project.org/package=ElliptCopulas>.
- [6] P. Embrechts, A. McNeil, D. Straumann, Correlation and dependence in risk management: properties and pitfalls, in: *Risk Management: Value At Risk and beyond*, Cambridge University Press, Cambridge, 2002, pp. 176–223.
- [7] J. Fan, R. Li, Variable selection via nonconcave penalized likelihood and its oracle properties, *J. Amer. Statist. Assoc.* 96 (456) (2001) 1348–1360.
- [8] J. Fan, H. Peng, Nonconcave penalized likelihood with a diverging number of parameters, *Ann. Statist.* 32 (3) (2004) 928–961.
- [9] H.-B. Fang, K.-T. Fang, S. Kotz, The meta-elliptical distributions with given marginals, *J. Multivariate Anal.* 82 (1) (2002) 1–16.
- [10] K.-T. Fang, S. Kotz, K.-W. Ng, *Symmetric Multivariate and Related Distributions*, Springer, New-York, 1990.
- [11] J.-D. Fermanian, D. Radulovic, M. Wegkamp, Weak convergence of empirical copula processes, *Bernoulli* 10 (5) (2004) 847–860.
- [12] G. Frahm, M. Junker, A. Szimayer, Elliptical copulas: applicability and limitations, *Statist. Probab. Lett.* 63 (3) (2003) 275–286.
- [13] C. Genest, A.-C. Favre, J. Béliveau, C. Jacques, Metaelliptical copulas and their use in frequency analysis of multivariate hydrological data, *Water Resour. Res.* 43 (9) (2007).
- [14] C. Genest, K. Ghoudi, L.-P. Rivest, A semiparametric estimation procedure of dependence parameters in multivariate families of distributions, *Biometrika* 82 (3) (1995) 543–552.
- [15] K. Ghoudi, B. Rémillard, Empirical processes based on pseudo-observations, in: *Asymptotic Methods in Probability and Statistics*, Elsevier, 1998, pp. 171–197.
- [16] K. Ghoudi, B. Rémillard, Empirical processes based on pseudo-observations II: the multivariate case, *Fields Inst. Commun.* 44 (2004) 381–406.
- [17] E. Gómez, M.A. Gómez-Villegas, J.M. Marín, A survey on continuous elliptical vector distributions, *Revista Matemática Complutense* 16 (1) (2003) 345–361.
- [18] N.J. Higham, Computing the nearest correlation matrix—A problem from finance, *IMA J. Numer. Anal.* 22 (3) (2002) 329–343.
- [19] N.J. Higham, N. Strabic, V. Segó, Restoring definiteness via shrinking, with an application to correlation matrices with a fixed block, *SIAM Rev.* 58 (2) (2016) 245–263.
- [20] M. Jaser, S. Haug, A. Min, A simple non-parametric goodness-of-fit test for elliptical copulas, *Depend. Model.* 5 (1) (2017) 330–353.
- [21] M. Jaser, A. Min, On tests for symmetry and radial symmetry of bivariate copulas towards testing for ellipticity, *Comput. Statist.* (2020) 1–26.
- [22] D. Kelker, Distribution theory of spherical distributions and a location-scale parameter generalization, *Sankhyā: Indian J. Stat. Ser. A* (1970) 419–430.
- [23] C. Klüppelberg, G. Kuhn, L. Peng, Estimating the tail dependence function of an elliptical distribution, *Bernoulli* 13 (1) (2007) 229–251.
- [24] C. Klüppelberg, G. Kuhn, L. Peng, Semi-parametric models for the multivariate tail dependence function—the asymptotically dependent case, *Scand. J. Stat.* 35 (4) (2008) 701–718.
- [25] K. Kostadinov, Non-parametric estimation of elliptical copulae with application to credit risk, 2005, Preprint.
- [26] D. Kurowicka, J. Misiewicz, R. Cooke, Elliptical copulae, in: *Proc of the International Conference on Monte Carlo Simulation—Monte Carlo*, 2000, pp. 209–214.
- [27] K.M. Levasseur, A probabilistic proof of the weierstrass approximation theorem, *Amer. Math. Monthly* 91 (4) (1984) 249–250.
- [28] E. Liebscher, A semiparametric density estimator based on elliptical distributions, *J. Multivariate Anal.* 92 (1) (2005) 205–225.
- [29] H. Liu, F. Han, C.-h. Zhang, Transelliptical graphical models, in: *Advances in Neural Information Processing Systems*, 2012, pp. 809–817.
- [30] H. Liu, J. Mulvey, T. Zhao, A semiparametric graphical modelling approach for large-scale equity selection, *Quant. Finance* 16 (7) (2016) 1053–1067.
- [31] P.-L. Loh, Statistical consistency and asymptotic normality for high-dimensional robust M -estimators, *Ann. Statist.* 45 (2) (2017) 866–896.
- [32] B. Poignard, J.-D. Fermanian, The finite sample properties of sparse M -estimators with pseudo-observations, *Ann. Inst. Statist. Math.* (2021) 1–31.
- [33] H. Qi, D. Sun, A quadratically convergent Newton method for computing the nearest correlation matrix, *SIAM J. Matrix Anal. Appl.* 28 (2) (2006) 360–385.
- [34] B. Remillard, *Statistical Methods for Financial Engineering*, CRC Press, 2013.
- [35] A. Sancetta, A. Nikandrova, Forecasting and prequential validation for time varying meta-elliptical distributions, *J. Time Ser. Econom.* 1 (2) (2009) 1–41.
- [36] J.H. Shih, T.A. Louis, Inferences on the association parameter in copula models for bivariate survival data, *Biometrics* (1995) 1384–1399.
- [37] S. Song, V.P. Singh, Meta-elliptical copulas for drought frequency analysis of periodic hydrologic data, *Stoch. Environ. Res. Risk Assess.* 24 (3) (2010) 425–444.
- [38] J. Touboul, Goodness-of-fit tests for elliptical and independent copulas through projection pursuit, *Algorithms* 4 (2) (2011) 87–114.
- [39] H. Tsukahara, Semiparametric estimation in copula models, *Canad. J. Statist.* 33 (3) (2005) 357–375.
- [40] H. Tsukahara, Erratum to "semiparametric estimation in copula models", *Canad. J. Statist.* 39 (4) (2011) 734–735.
- [41] X. Wang, J. Yan, Practical notes on multivariate modeling based on elliptical copulas, *Journal de la Société Française de Statistique* 154 (1) (2013) 102–115.
- [42] M. Wegkamp, Y. Zhao, Adaptive estimation of the copula correlation matrix for semiparametric elliptical copulas, *Bernoulli* 22 (2) (2016) 1184–1226.
- [43] Y. Zhao, C. Genest, Inference for elliptical copula multivariate response regression models, *Electron. J. Stat.* 13 (1) (2019) 911–984.