Flexible Bayesian estimation of elliptical copulas

Panfeng Liang
Department of Mathematical Sciences, University of Texas at El Paso
and
Ori Rosen
Department of Mathematical Sciences, University of Texas at El Paso

October 6, 2025

Abstract

Elliptical copulas provide flexibility in modeling the dependence structure of a random vector. They are often parameterized with a correlation matrix and a scalar function, called generator. The estimation of the generator can be challenging, because it is a functional parameter. In this paper, we provide a rigorous approach to estimating the generator in a Bayesian framework, which is simpler, more robust, and outperforms existing estimation methods in the literature. A major contribution of this paper is a robust method of evaluating the elliptical copula likelihood by using mixtures of B-spline densities. The Matlab code used for the simulation study is available in the supplementary material.

Keywords: Elliptical copulas, B-spline densities, MCMC.

1 Introduction

Elliptical copulas are derived from the family of elliptical distributions, and are often parameterized with a correlation matrix and a scalar function known as a generator. In this paper, we propose a method for estimating generators of elliptical copulas in a Bayesian framework using mixtures of B-spline densities.

According to Sklar (1959), a multivariate distribution with continuous marginals can be uniquely represented by a copula and a set of marginal distributions. Specifically, the joint cumulative distribution function (CDF) of $d \geq 2$ continuous random variables $(Z_1, \dots, Z_d)$, can be represented by

$$F(z_1, \dots, z_d) = C\{F_1(z_1), \dots, F_d(z_d)\},$$

where C is a copula and $F_1, \dots, F_d$ are marginal CDFs. If these marginals are continuous, then the copula C is unique. The copula C describes the dependence structure of the multivariate distribution, and in itself is a multivariate distribution with uniform margins. The copula itself can be expressed by inversion as follows.

$$C(u_1, \dots, u_d) = F\left\{F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\right\},\ u_1, \dots, u_d \in (0,1), \tag{1}$$

where $F_1^{-1}, \dots, F_d^{-1}$ are the inverse CDFs, see Nelsen (2006). If the copula C is absolutely continuous with respect to Lebesgue's measure, the copula probability density function (PDF) is obtained by differentiating Equation (1).

$$c(u_1, \dots, u_d) = \frac{f\left\{F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\right\}}{\prod_{k=1}^{d} f_k\left\{F_k^{-1}(u_k)\right\}},\ u_1, \dots, u_d \in (0,1),$$

where f is the PDF corresponding to F and $f_1, \dots, f_d$ are the marginal PDFs. For a detailed introduction to copula theory, see Nelsen (2006).

Elliptical copulas, also known as meta-elliptical copulas and originally introduced by Fang et al. (2002), are extensions of Gaussian copulas. Elliptical copulas induce the dependence in elliptical distributions by analogy to the way Gaussian copulas characterize the dependence in multivariate Gaussian distributions. A d-dimensional continuous random vector $\boldsymbol{X} = (X_1, ..., X_d)$ has an elliptical distribution $\mathcal{E}_d(\boldsymbol{\mu}, \Sigma, g)$ if its PDF is of the form

$$f(\boldsymbol{x} \mid \boldsymbol{\mu}, \Sigma, g) = \det(\Sigma)^{-1/2} g\left\{(\boldsymbol{x}-\boldsymbol{\mu})^\top \Sigma^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right\}, \tag{2}$$

where $\boldsymbol{\mu}$ is a mean vector, Σ is a covariance matrix, and g is a generator, which is a non-negative function defined on $[0, \infty)$. A typical example of an elliptical distribution is the multivariate Gaussian distribution, whose generator satisfies $g(t) \propto \exp(-t/2)$. Generators of other common elliptical distributions, such as the multivariate Student-t and the multivariate logistic distributions, can be found, for example, in Genest et al. (2007) and Lemonte and Patriota (2011).

Since our article is focused on elliptical copulas, without loss of generality, we only consider elliptical distributions $\mathcal{E}(\mathbf{0}, \Omega, g)$, where Ω is a correlation matrix. In this case, the elliptical distribution has identical margins. If a d-variate random vector $\boldsymbol{X}$ has elliptical distribution $\mathcal{E}(\mathbf{0}, \Omega, g)$, it can be expressed in the following stochastic representation (see Fang et al. (2002) and Fang et al. (2005))

$$\boldsymbol{X} = RA\boldsymbol{U}, \tag{3}$$

where A is a Cholesky factor such that $AA^\top = \Omega$, $\boldsymbol{U}$ is a random vector uniformly distributed on the unit sphere in $\mathbb{R}^d$, and R is a continuous nonnegative random variable, independent of $\boldsymbol{U}$, with PDF

$$h(r) = \frac{2\pi^{d/2}}{\Gamma(d/2)} r^{d-1} g(r^2). \tag{4}$$

Elliptical distributions generalize *spherical* distributions. In particular, if $\boldsymbol{X}$ has an elliptical distribution, $\boldsymbol{Y} = A^{-1}\boldsymbol{X}$ has a spherical distribution whose stochastic form is $\boldsymbol{Y} = R\boldsymbol{U}$. It is straightforward to show that R in Equation (3) is equal to $(\boldsymbol{Y}^\top\boldsymbol{Y})^{1/2}$, which is the Euclidean distance from $\boldsymbol{Y}$ to the origin. The isodensity lines of a two-dimensional spherical distribution are circles, where R can intuitively be viewed as their "radius". For a detailed introduction to spherical distributions, see Kelker (1970), Cambanis et al. (1981) and Steerneman and van Perlo-ten Kleij (2005). The geometric interpretation of R gives insight into the role of the generator g in Equation (2), since g and the distribution of R are closely related. This is discussed in Section 2.

The PDF of an *elliptical copula* can be derived from the density of an elliptical distribution $\mathcal{E}(\mathbf{0}, \Omega, g)$ via the inversion method, resulting in

$$c(\boldsymbol{u} \mid \Omega, g) = \frac{\det(\Omega)^{-1/2} g\left\{F_g^{-1}(\boldsymbol{u})^\top \Omega^{-1} F_g^{-1}(\boldsymbol{u})\right\}}{\prod_{j=1}^d f_g\left\{F_g^{-1}(u_j)\right\}}, \quad \boldsymbol{u} \in [0,1]^d, \tag{5}$$

where the numerator is the PDF of $\mathcal{E}(\mathbf{0}, \Omega, g)$, and f_g and F_g are its marginal PDF and CDF, respectively. The PDF f_g and the CDF F_g are given by (See Abdous et al. (2005))

$$f_g(x) = \frac{\pi^{(d-1)/2}}{\Gamma\left(\frac{d-1}{2}\right)} \int_{x^2}^{\infty} (y - x^2)^{\frac{d-1}{2}-1} g(y) dy, \tag{6}$$

and

$$F_g(x) = \frac{1}{2} + \frac{\pi^{(d-1)/2}}{\Gamma\left(\frac{d-1}{2}\right)} \int_0^x \int_{u^2}^{\infty} (y - u^2)^{\frac{d-1}{2}-1} g(y) dy du. \tag{7}$$

We denote the elliptical copula in (5) by $\mathcal{C}(\Omega, g)$, and refer to $\mathcal{E}(\mathbf{0}, \Omega, g)$ as the "associated" elliptical distribution, which means that it is the distribution from which the copula is derived.

Though elliptical copulas have been around for a long time, research on the estimation of elliptical copulas is fairly scarce. Related articles are mostly focused on the application of specific types of elliptical copulas, such as Gaussian or Student-t copulas, where a known generator g is assumed. Genest et al. (2007) assume that the generator g is unknown but is one option from a fixed list. Genest et al. (2007) provide goodness-of-fit tests to select the best generator from a list for a given data set. The main challenge in estimating an elliptical copula without assuming a pre-determined generator is the evaluation of the likelihood function. Assuming a single observation $\boldsymbol{u} = (u_1, \ldots, u_d)$ drawn from a d-dimensional elliptical copula with PDF of the form (5) and a known correlation matrix Ω, evaluating the likelihood

$$\ell_1(\hat{g} \mid \boldsymbol{u}) = \frac{\det(\Omega)^{-1/2}\hat{g}\left\{F_{\hat{g}}^{-1}(\boldsymbol{u})^\top \Omega^{-1} F_{\hat{g}}^{-1}(\boldsymbol{u})\right\}}{\prod_{j=1}^{d} f_{\hat{g}}\left\{F_{\hat{g}}^{-1}(u_j)\right\}} \tag{8}$$

for a given estimate, $\hat{g}$, of the generator, requires the evaluation of $f_{\hat{g}}$ and $F_{\hat{g}}$ according to equations (6) and (7) and inverting $F_{\hat{g}}$, which may be infeasible. Recently, an iterative procedure was proposed by Derumigny and Fermanian (2022) which evaluates the generator of an elliptical copula on a grid. The algorithm iteratively converts the copula data into pseudo-data having an elliptical distribution and then estimates the generator of the elliptical distribution using the procedure in Liebscher (2005). Doing so, Derumigny and Fermanian (2022) avoid the evaluation of the likelihood in Equation (8). Rather, they estimate an elliptical distribution from the transformed data by evaluating the likelihood

$$\ell_2(\hat{g} \mid \boldsymbol{u}) = \det(\Omega)^{-1/2}\hat{g}\left(\boldsymbol{x}^{*\top}\Omega^{-1}\boldsymbol{x}^*\right), \tag{9}$$

where $\boldsymbol{x}^* = F_{\hat{g}}^{-1}(\boldsymbol{u})$ are the pseudo-data. The drawbacks to the procedure of Derumigny and Fermanian (2022) are two-fold. First, they compute the likelihood in Equation (9) rather than that of Equation (8), hindering the understanding of theoretical properties of their algorithm. Second, their simulation studies do not show good results in most of the simulation settings. In this paper, we model elliptical copula generators by using B-spline PDFs, which eases the computation of f_g and F_g^{-1}, facilitates direct evaluation of the likelihood (8) and Bayesian computation via Markov chain Monte Carlo methods. Our method returns a *function* as the estimate of g unlike the estimated values of g on a grid which result from the algorithm of Derumigny and Fermanian (2022).

Our paper is organized as follows. Section 2 introduces properties of generators of elliptical copulas as well as the difference between generators of elliptical *copulas* and those of elliptical *distributions*. Section 3 presents the proposed methodology, while Section 4 introduces meta-elliptical distributions (whose dependence structure is induced by elliptical copulas). Sections 5 and 6 present the results of the simulation study for the bivariate case and for higher-dimensional settings, respectively. Section 7 illustrates the method with an application to wine data. Section 8 concludes with a discussion.

2 The Generator

Equation (4) expresses the PDF of R in the stochastic representation (3) as a function of the generator g. The generator g in turn is given by

$$g(t) = \frac{\Gamma(d/2)}{2\pi^{d/2}} t^{(1-d)/2} h\left(t^{1/2}\right),\ t > 0, \tag{10}$$

see Genest et al. (2007) or Liebscher (2005). Since $h(r)$ is a PDF, it follows that the generator g must satisfy

$$\int_0^\infty r^{d-1} g\left(r^2\right) dr = \frac{\Gamma(d/2)}{2\pi^{d/2}}. \tag{11}$$

Other constraints on g are detailed in Derumigny and Fermanian (2022) but do not play an important role in what follows. If a function does not satisfy Equation (11), it can be transformed to do so. We refer to this process as “normalization”, because it amounts to normalizing the corresponding PDF of R.

The elliptical copula $\mathcal{C}(\Omega, g)$ is parameterized in the same way as its associated elliptical distribution $\mathcal{E}(\mathbf{0}, \Omega, g)$. However, generators for elliptical *copulas* are not unique, see Derumigny and Fermanian (2022). In particular, the marginal PDF and CDF of $\mathcal{E}(\mathbf{0}, \Omega, g)$ for a given generator g are given in equations (6) and (7) from which it is seen that an elliptical distribution's generator contains not only information about the dependence structure, but also information about the marginal distributions. This can result in non-identifiability when using g to parameterize an elliptical copula. Specifically, assume that two elliptical distributions, $\mathcal{E}(\mathbf{0}, \Omega, g_a)$ and $\mathcal{E}(\mathbf{0}, \Omega, g_b)$, share the same dependence structure but differ in their marginal distributions. In this case, these distributions have different generators, g_a and g_b, but their copulas are identical, which may be denoted by either $\mathcal{C}(\Omega, g_a)$ or $\mathcal{C}(\Omega, g_b)$. The correlation matrix Ω is unique and identifiable for elliptical copulas, because it contains only information about the linear dependence. Derumigny and Fermanian (2022) propose a transformation for modifying the generator of an elliptical copula that preserves its dependence structure. In particular, let Ω be a positive-definite correlation matrix and g a generator. Then, for any positive value α, $\mathcal{C}(\Omega, g) = \mathcal{C}(\Omega, g_\alpha)$, where

$$g_\alpha(t) = \alpha^{d/2} g(\alpha t). \tag{12}$$

That is, the dependence structure induced by g is invariant under the transformation (12), although F_g and f_g have become F_{g_α} and f_{g_α} due to the dependence of the marginal PDF and CDF on the generator in equations (6) and (7). Transformation (12) is able to make the generator unique and constrain the marginal distribution to a have a specific form. According to Derumigny and Fermanian (2022), for an elliptical copula, there "most often" exists a unique normalized generator g satisfying

$$\frac{\pi^{\frac{d-1}{2}}}{\Gamma\left(\frac{d-1}{2}\right)} \int_0^\infty s^{\frac{d-3}{2}} g(s) ds = b, \tag{13}$$

where b is a positive number that is referred to as a standardization constant. In other words, each elliptical copula has only *one* generator satisfying the constraint in Equation (13). The quantity b is in fact $f_g(0)$. Fixing $f_g(0) = b$ specifies f_g, it is not necessary to fix the values of f_g everywhere. For a detailed proof, see Proposition 3 and Appendix A in Derumigny and Fermanian (2022). We provide below Algorithm 1 for standardizing the generator of an elliptical copula. It has been modified from the algorithm in Derumigny and Fermanian (2022).

Algorithm 1 Standardizing the Generator

Input: A generator g that is already normalized.

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

2: Set $\beta = \{b/(s_{d-1}\,\mathcal{I}_2)\}^2$, where b is any positive constant, and $s_{d-1} = \frac{\pi^{(d-1)/2}}{\Gamma\{(d-1)/2\}}$;

3: Calculate $g_s = \beta^{d/2}\, g(\beta\, t)$.

Output: A modified version g_s satisfying the normalization and identification constraints.

3 The model, prior and sampling scheme

3.1 The model

Liebscher (2005) proposes a method for constructing models for the generator of an elliptical *distribution*, and our model is built in this framework. In particular, let Ψ be a transformation function such that $Y = \Psi(R)$ where

$$\Psi(R) = -a + \left(a^{d/2} + R^d\right)^{2/d}. \tag{14}$$

In Equation (14), a is a positive constant, such as 1, and R is the "radius" in the stochastic form (3). In Section 6, which discusses the case $d > 2$, we take advantage of the special property of ellipical copulas, i.e., identical margins, and select any two variates to fit the copula. This renders the choice of a specific value of a unnecessary, since in this case a cancels out from Equation (14). We model the PDF of Y, $f_Y(y)$, via a mixture of $(K + 4)$ B-spline densities (Ghidey et al. (2004)) with a knot sequence $0 = T_1,\ T_2,\ \ldots,\ T_K,\ T_{K+1},\ T_{K+2},\ T_{K+3} = y_{\max}$, given by

$$f_Y(y) = \sum_{j=-1}^{j=K+2} w_j \tilde{B}_j(y), \tag{15}$$

where the weights $w_{-1}, \ldots, w_{K+2}$ are of the form

$$w_j = \frac{\exp(\nu_j)}{\sum_{\ell=-1}^{K+2} \exp(\nu_\ell)}. \tag{16}$$

In Equation (16), the ν_j are unknown unconstrained parameters, and for identifiability, ν_{K+2} is set to zero. The $\tilde{B}_j(x)$ are B-spline PDFs, which are normalized B-splines that integrate to 1. These B-splines can be constructed by first using the Cox–de Boor recursion formula (de Boor (1978)) and then normalizing the resulting B-splines such that they integrate to 1. For details, see the Appendix, where we use the formulation of B-spline PDFs from Staudenmayer et al. (2008). In our experiments, using up to 30 B-spline PDFs is usually adequate for estimating the PDF f_Y. The support of f_Y is taken to be $[0, y_{\max}]$, which simplifies the computation of the likelihood. Recall that the generator g is a function of $h(r)$, as is evident from Equation (10). The PDF $h(r)$ in turn can be expressed as a function of f_Y via a change of variables. Specifically, the transformation Jacobian $Y \to R$ is given by $J_{Y\to R} = 2r^{d-1}(a^{d/2} + r^d)^{2/d-1}$ from which it follows that

$$h(r) = 2r^{d-1}(a^{d/2} + r^d)^{2/d-1} f_Y\{-a + (a^{d/2} + r^d)^{2/d}\}.$$

Plugging the expression for $h(r)$ into Equation (10) and expressing $f_Y(y)$ as in (15) results in

$$g(t) = \frac{\Gamma(d/2)}{\pi^{d/2}} (a^{d/2} + t^{d/2})^{2/d-1} \sum_{i=-1}^{i=K+2} w_i \tilde{B}_i\{-a + (a^{d/2} + t^{d/2})^{2/d}\}. \tag{17}$$

In model (17), the generator g is defined on the interval $\left[0, t_{\max} = \{(y_{\max} + a)^{d/2} - a^{d/2}\}^{2/d}\right]$ instead of $[0, \infty)$, which simplifies the evaluation of the likelihood for given weights $(w_{-1}, \dots, w_{K+2})$ in Equation (15).

3.2 Priors

Prior on $\boldsymbol{\nu} = \{\nu_{-1}, \dots, \nu_{K+1}\}$:

We follow Rosen and Thompson (2015) in placing a smoothing prior on ν. This prior is in fact a Bayesian version of the smoothing penalty advocated by Eilers and Marx (1996, 2021) which also gave rise to Bayesian P-splines (Lang and Brezger (2004)). In particular, defining $\Delta\nu_j = \nu_j - \nu_{j-1}$ and $\Delta^2\nu_j = \Delta\nu_j - \Delta\nu_{j-1}$, Lang and Brezger (2004) let

$$\Delta^2\nu_j = \delta_j,\ j = 1, \dots, K+2,$$

where $\delta_j \sim \mathcal{N}(0, \tau^2)$ and the variance parameter τ^2 controls the amount of smoothness. To avoid impropriety of this prior, Chib and Jeliazkov (2006) place a joint normal prior on (ν_{-1}, ν_0), i.e., $(\nu_{-1}, \nu_0) \sim \mathcal{N}_2(\mathbf{0}, c\tau^2 I_2)$, where c is a fixed constant. The resulting prior on $\boldsymbol{v}$ is

$$p(\boldsymbol{\nu} \mid \tau^2) \propto \left(\tau^2\right)^{-\frac{1}{2}(K+3)} \exp\left(-\frac{1}{2\tau^2} \boldsymbol{\nu}^\top P^* \boldsymbol{\nu}\right),$$

where $P = D^\top D$, D is the following $(K+1) \times (K+3)$ second-order difference matrix

$$D = \begin{bmatrix} 1 & \text{-2} & 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & \text{-2} & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & & & & \vdots \\ 0 & 0 & \dots & & 1 & \text{-2} & 1 \end{bmatrix}$$

and $P^*_{\ell,\ell} = P_{\ell,\ell} + c^{-1}$ for $\ell = 1, 2$.

Prior on τ:

The parameter τ is assumed to follow a Half-t distribution (Gelman (2006)), with PDF

$$p(\tau \mid n_\tau, A) \propto \left(1 + \frac{\tau^2}{A^2 n_\tau}\right)^{-\frac{n_\tau+1}{2}},$$

where the hyperparameters n_τ and A are assumed known.

3.3 The likelihood

Based on N independent observations $\boldsymbol{U}_1, \dots, \boldsymbol{U}_N$ from an elliptical copula $\mathcal{C}(\Omega, g)$ (Equation (5)), where the correlation matrix Ω is assumed known, the likelihood function is given by

$$\mathcal{L}(w_{-1}, \dots, w_{K+2} \mid \boldsymbol{U}_1, \dots, \boldsymbol{U}_N) = \prod_{i=1}^{N} \left[\frac{\det(\Omega)^{-1/2} g_s\{F_{g_s}^{-1}(\boldsymbol{U}_i)^\top \Omega^{-1} F_{g_s}^{-1}(\boldsymbol{U}_i)\}}{\prod_{j=1}^{d} f_{g_s}\{F_{g_s}^{-1}(U_{ij})\}} \right], \tag{18}$$

where g_s is the standardized version of g via Algorithm 1. Recall that f_g and F_g are the marginal PDF and CDF of the associated elliptical distribution $\mathcal{E}(\mathbf{0}, \Omega, g)$, and F_g^{-1} is the inverse CDF, all of which can be derived from g using equations (6) and (7). The functions f_{g_s}, F_{g_s} and $F_{g_s}^{-1}$ in Equation (18) can be derived from f_g, F_g and F_g^{-1}. Computing F_g and f_g for a given g involves complicated integrals which makes the evaluation of the likelihood challenging. Using numerical integration to evaluate F_g and f_g and root finding methods to obtain the inverse function F_g^{-1} can be very time consuming and often inaccurate. Instead, we use B-spline smoothing to approximate f_g, F_g and F_g^{-1}.

Obtaining the marginal density f_g: Instead of directly deriving f_g from the generator g via Equation (6), this task can be accomplished in two steps. The marginal distributions of an elliptical distribution are also elliptical, having their own generators. Recall that the distribution $\mathcal{E}(\mathbf{0}, \Omega, g)$ has identical marginal distributions. The generator g_1 of this marginal distribution can be derived by

$$g_1(u) = \frac{\pi^{(d-1)/2}}{\Gamma\left\{(d-1)/2\right\}} \int_0^\infty g(u+s)s^{(d-3)/2}ds. \tag{19}$$

The marginal PDF f_g can in turn be directly expressed as

$$f_g(x) = g_1(x^2). \tag{20}$$

For more details, see Appendix A of Derumigny and Fermanian (2022) and Theorem 6 of Gómez et al. (2003). To approximate f_g, we approximate g_1 first. Since g is derived from a mixture of B-spline densities, its support is the interval $[0, t_{\max} = \{(y_{\max} + a)^{d/2} - a^{d/2}\}^{2/d}]$. Thus, Equation (19) reduces to

$$g_1(u) = \frac{\pi^{(d-1)/2}}{\Gamma\left\{(d-1)/2\right\}} \int_0^{t_{\max}-u} g(u+s)s^{(d-3)/2}\mathbb{I}_{[0,t_{\max}]}(u)ds,$$

where $\mathbb{I}_{[a,b]}(\cdot)$ is an indicator function. Note that both the integrand and the integral limits change with u, which makes the computation of the integral difficult. Suppose we wish to compute $g_1(u_i)$, where u_i is a real number in the interval $[0, t_{\max}]$. We need to evaluate the integral

$$\int_0^{t_{\max}-u_i} g(u_i+s)s^{(d-3)/2}ds.$$

First, we divide the interval $[0, t_{\max} - u_i]$ into n subintervals, each of length $\Delta_x = (t_{\max} - u_i)/n$. This yields a grid of equally spaced points:

$$X = (0,\ \Delta_x,\ 2\Delta_x,\ \ldots,\ t_{\max} - u_i).$$

Next, we evaluate the integrand $g(u_i + s)s^{(d-3)/2}$ on X, resulting in

$$Y = g(u_i + X)X^{(d-3)/2}.$$

Given the sequence X and Y, the integral in $g_1(u_i)$ can be approximated using the Trapezoidal or Simpson's rules. Next, we approximate the values of g_1 at an equally-spaced grid consisting of n points, $u_0 = 0, \ \ldots, \ u_n = t_{\max}$. (Note that the letter n is used to denote the number of points in a grid. A grid of length n is constructed multiple times for function approximation through spline smoothing and interpolation. The sample size is denoted by N.) The approximation process of $g_1(u_0), \ldots, g_1(u_n)$ is described in Table 1. This approach is also taken by Derumigny and Fermanian (2022).

Table 1: Approximating g_1 via the Trapezoid rule.

u_i	X	$Y = g(u_i + X)X^{(d-3)/2}$	$g_1(u_i) \propto \text{Trapz}(X, Y)$
$u_0 = 0$	$0, \ \ldots, \ t_{\max}$	$\ldots$	$\ldots$
$\ldots$	$\ldots$	$\ldots$	$\ldots$
u_j	$0, \ \ldots, \ t_{\max} - u_j$	$\ldots$	$\ldots$
$\ldots$	$\ldots$	$\ldots$	$\ldots$
$u_n = t_{\max}$	0	0	0

Given the sample points $\{(u_0, g_1(u_0)), \ldots, (u_n, g_1(u_n))\}$, Equation (20) is used to evaluate f_g at an equally spaced grid: $\{-{u_n}^{1/2}, \ \ldots, \ -{u_1}^{1/2}, \ 0, \ {u_1}^{1/2}, \ \ldots, \ {u_n}^{1/2}\}$, resulting in the sample points

$$f_g(-{u_n}^{1/2}), \ \ldots, f_g(-{u_1}^{1/2}), \ f_g(0), \ f_g({u_1}^{1/2}), \ \ldots, \ f_g({u_n}^{1/2}). \tag{21}$$

The sequence in (21) is used to fit a curve for approximating f_g. The curve fitting is performed using the MATLAB functions `fit` and `cfit`, with the fit type option set to `smoothingspline`. A detailed introduction to spline fitting can be found in Eilers and Marx (2021). Note that f_g is a probability density function (PDF), and so is its approximation. The fitted curve must be normalized if it does not integrate to one over its support, $[-t_{\max}^{1/2}, t_{\max}^{1/2}]$, which is determined by the domain of g_1 and Equation (20).

Obtaining the marginal cumulative distribution function F_g: The function F_g is the integral of f_g which is already approximated by a spline function. Its integral

$$F_g(x) = \int_{-t_{\max}^{1/2}}^{x} f_g(x)\mathbb{I}_{[-t_{\max}^{1/2}, t_{\max}^{1/2}]}(x)dx$$

is obtained via the function `integrate` in Matlab.

Obtaining the marginal quantile function F_g^{-1}: First, the function F_g is evaluated at an equally spaced grid

$$\{x_1 = -t_{\max}^{1/2},\ \ldots,\ x_n = t_{\max}^{1/2}\}$$

resulting in the sequence

$$F_g(x_1), \ldots, F_g(x_n). \tag{22}$$

Next, the x- and y-axes of the sample in (22) are swapped, and interpolation is applied to obtain a functional approximation of F_g^{-1}. For simplicity, we use linear interpolation, following the approach of Derumigny and Fermanian (2022). The interpolation is performed using the `interp1` function in MATLAB.

Standardization of related functions: Given a standardizing constant b and a generator as in Equation (17), the standardized generator resulting from Algorithm 1 can be expressed as

$$g_s = \beta^{d/2}\, g(\beta\, t),\ t \in [0, t_{\max}/\beta],$$

where β is computed from the constant b, as shown in Algorithm 1. It follows that

$$\begin{aligned}
g_{1s}(u) &= \frac{\pi^{(d-1)/2}}{\Gamma\{(d-1)/2\}} \int_0^\infty g_s(u+s)s^{(d-3)/2}ds \\
&= \frac{\pi^{(d-1)/2}}{\Gamma\{(d-1)/2\}} \int_0^\infty \beta^{d/2} g(\beta\, u + \beta\, s)s^{(d-3)/2}ds \\
&= \frac{\pi^{(d-1)/2}}{\Gamma\{(d-1)/2\}} \int_0^\infty \beta^{1/2} g(\beta\, u + S)S^{(d-3)/2}dS \\
&= \beta^{1/2} g_1(\beta u),
\end{aligned}$$

where we have used the substitution $s = S/\beta$. Next,

$$\begin{aligned} f_{g_s}(x) &= g_{1_s}(x^2) \\ &= \beta^{1/2} g_1(\beta x^2) \\ &= \beta^{1/2} f_g\left(\beta^{1/2} x\right),\ x \in [-(t_{\max}/\beta)^{1/2}, (t_{\max}/\beta)^{1/2}], \end{aligned}$$

$$\begin{aligned} F_{g_s}(x) &= \int_{-\infty}^{x} f_{g_s}(u)\, du \\ &= \int_{-\infty}^{x} \beta^{1/2} f_g\left(\beta^{1/2} u\right) du \\ &= F_g(\beta^{1/2} x),\ x \in [-(t_{\max}/\beta)^{1/2}, (t_{\max}/\beta)^{1/2}], \end{aligned}$$

and

$$F_{g_s}^{-1}(x) = F_g^{-1}(x)/\beta^{1/2},$$

where $F_{g_s}^{-1}$ has the range $[-(t_{\max}/\beta)^{1/2}, (t_{\max}/\beta)^{1/2}]$. Model (17) leads to a closed domain or range for the related functions g_s, f_{g_s} and $F_{g_s}^{-1}$, making the evaluation of the copula likelihood (18) feasible in terms of both speed and accuracy.

3.4 Sampling Scheme

The posterior distribution does not admit simple posterior conditional distributions, and for this reason we use Metropolis steps to sample all the parameters.

The proposal distributions for the ν_j in Equation (16) are all uniform. In particular, a new value $\nu_j^{(p)}$ is proposed by $\nu_j^{(p)} \sim \mathcal{U}(\nu_j^{(c)} - \delta_\nu, \nu_j^{(c)} + \delta_\nu)$, where δ_ν is a fixed tuning parameter and $\nu_j^{(c)}$ is the current value of ν_j. The entries of $\boldsymbol{\nu} = (\nu_{-1}, \ldots, \nu_{K+1})$ are sampled jointly. The proposed vector $\boldsymbol{\nu}^{(p)} = (\nu_{-1}^{(p)}, \ldots, \nu_{K+1}^{(p)})$ is accepted with probability $\min\{1, p_{\boldsymbol{\nu}}(\boldsymbol{\nu}^{(p)} \mid \cdots)/p_{\boldsymbol{\nu}}(\boldsymbol{\nu}^{(c)} \mid \cdots)\}$, where $p_{\boldsymbol{\nu}}(\boldsymbol{\nu} \mid \cdots)$ is the conditional posterior PDF of $\boldsymbol{\nu} = (\nu_{-1}, \ldots, \nu_{K+1})$ conditional on the current values of the other parameters.

Since τ is a positive number, we set $\tau_* = \ln(\tau)$. A new value $\tau_*^{(p)}$ is proposed by $\tau_*^{(p)} \sim \mathcal{U}(\tau_*^{(c)} - \delta_\tau, \tau_*^{(c)} + \delta_\tau)$, where δ_τ is a fixed tuning parameter and $\tau_*^{(c)}$ is the current value of τ_*. The proposed value τ_* is accepted with probability $\min\{1, p_{\tau_*}(\tau_*^{(p)} \mid \cdots)/p_{\tau_*}(\tau_*^{(c)} \mid \cdots)\}$, where $p_{\tau_*}(\tau_* \mid \cdots)$ is the conditional posterior PDF of τ_*, conditional on the current values of $\boldsymbol{\nu}$. Before starting the MCMC loop, we tune the values of δ_ν and δ_τ by running 500 iterations and observing the acceptance rate. If the acceptance rate turns out to fall between 0.3 and 0.6, then the values of δ_v and δ_τ are retained for running a subsequent MCMC chain.

4 Meta-elliptical Distribution

A typical application of elliptical copulas is to construct meta-elliptical distributions (Fang et al. (2002)). The marginal distributions of elliptical distributions are also elliptical. However, in practice, for modeling multi-dimensional data with arbitrary continuous margins, a good choice is to use meta-elliptical distributions. Such a distribution consists of an elliptical copula and a collection of arbitrarily chosen continuous margins, i.e., the dependence structure is induced by an elliptical copula while the marginal distributions can be chosen by the data analyst.

A d-dimensional meta-elliptical distribution is denoted by $\mathcal{ME}(\Omega, g; F_1, \ \ldots, \ F_d)$, where Ω is a correlation matrix, g is a copula generator and $F_1, \ldots, F_d$ are marginal CDFs. The joint CDF of a meta-elliptical distribution is

$$F_{\mathcal{ME}}(x_1, \ldots, x_d) = C\left\{F_1(x_1), \ldots, F_d(x_d)\right\},$$

where C is the CDF of the elliptical copula $\mathcal{C}(\Omega, g)$. Given that C is absolutely continuous, the PDF of $\mathcal{ME}$ is given by

$$f_{\mathcal{ME}}(x_1, \ldots, x_d) = c\left\{F_1(x_1), \ldots, F_d(x_d)\right\} \prod_{\ell=1}^{d} f_\ell(x_\ell),$$

where c is the PDF of $\mathcal{C}(\Omega, g)$ and the f_ℓ are the marginal PDFs. The estimation of a meta-elliptical distribution consists of estimating the marginal distributions and the underlying copula separately. Suppose we wish to analyze a data set with N observations from a meta-elliptical distribution $\mathcal{ME}(\Omega, g; F_1, \ \ldots, \ F_d)$, where the data are denoted as follows,

$$\begin{bmatrix} \boldsymbol{x}_1^\top \\ \vdots \\ \boldsymbol{x}_N^\top \end{bmatrix} = \begin{bmatrix} x_{11} & \ldots & x_{1d} \\ \vdots & \ddots & \vdots \\ x_{N1} & \ldots & x_{Nd} \end{bmatrix}. \tag{23}$$

The marginal distributions are estimated first. There are many ways to estimate $F_1, \ldots, F_d$, such as normal mixtures, Gamma mixtures, kernel densities, empirical cumulative distribution functions, etc. In the application of Section 7, we use normal mixtures and kernel densities to estimate the marginal distributions. Given the estimated marginal CDFs, $\hat{F}_1, \ldots, \hat{F}_d$, we apply the probability integral transform to the data (23),

$$u_{ij} = \hat{F}_j(x_{ij}), \quad \{j = 1, \ldots, d; \ i = 1, \ldots, N\}$$

to obtain the pseudo data

$$\begin{bmatrix} \boldsymbol{u}_1^\top \\ \vdots \\ \boldsymbol{u}_N^\top \end{bmatrix} = \begin{bmatrix} u_{11} & \ldots & u_{1d} \\ \vdots & \ddots & \vdots \\ u_{N1} & \ldots & u_{Nd} \end{bmatrix}. \tag{24}$$

The copula data (24) can be viewed as coming from an elliptical copula $\mathcal{C}(\Omega, g)$, which is the copula of $\mathcal{ME}(\Omega, g; F_1, \ \ldots, \ F_d)$. The rest of the estimation is performed on the copula data. Our simulation studies in Section 5 are performed directly on copula data.

5 Simulation Study for the Bivariate Case

For the bivariate case, we use the Student-t copula as well as five other elliptical copulas from Derumigny and Fermanian (2022) having the following generators (before normalization and standardization):

1. $g(t) = 1/(1+t^2)$,

2. $g(t) = \exp(-t)$,

3. $g(t) = \exp(-t) + \text{bump}(t)$, where $\text{bump}(x) = (x-1)(1+\pi-x)\sin(x-1)\mathbb{I}_{[1,1+\pi]}(x)$ is a smooth function supported on $[1, 1+\pi]$,

4. $g(t) = \exp(-t) + \exp(-t/3)\cos^2(t)$,

5. $g(t) = (1+t/\nu)^{-(\nu+d)/2}$, where $\nu = 5$, which is a Student-t copula with 5 degrees of freedom,

6. $g(t) = t^2 \exp(-t^2)$.

The correlation matrix is fixed throughout the simulations and is given by

$$\Omega = \begin{bmatrix} 1 & 0.2 \\ 0.2 & 1 \end{bmatrix}$$

for all six cases.

To sample from an elliptical copula, we sample from its associated elliptical distribution first, and then apply the probability integral transform using the marginal distribution function. Each simulation study consists of drawing 100 independent data sets of size $N = 1000$ each from a given copula. This yields 100 independent estimates of the generator for each of the six aforementioned elliptical copulas. For each replicate, the MCMC scheme is run for $10,000$ iterations, the first 5000 of which are used as burn-in. The posterior mean, $\hat{g}(t)$, is obtained by averaging the iterates after burn-in. The posterior mean as well as the true underlying generator are evaluated on the interval $[0, 1]$. The posterior means from the 100 replicates are plotted along with the true generators in Figure 1. For all replicates, the differences between $\hat{g}(t)$ and $g(t)$ are calculated on a grid $(t_1,\ t_2,\ t_3,\ t_4,\ t_5)$. Figure 2 presents boxplots of the 100 absolute differences at each of these grid values.

Figure 1 demonstrates good perfomance of the proposed methodology, showing close proximity between the true generators and their Bayesian estimates. It is seen from the

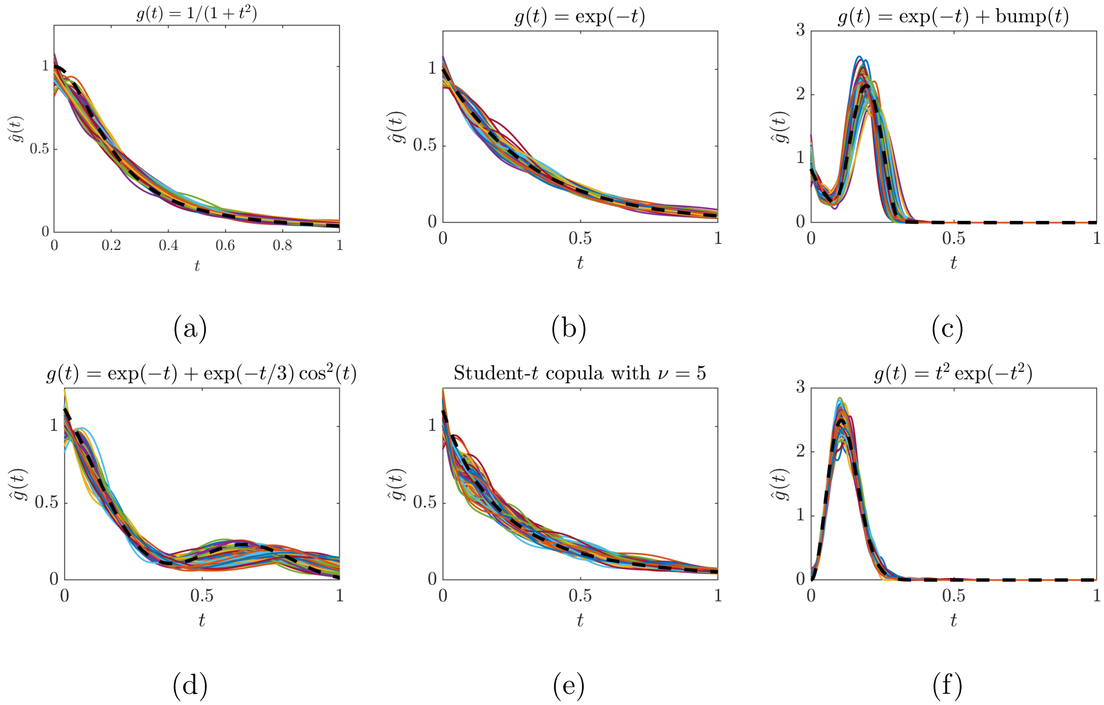

Figure 1: The true generators, normalized and standardized (dashed lines), along with 100 posterior means, based on 100 random samples.

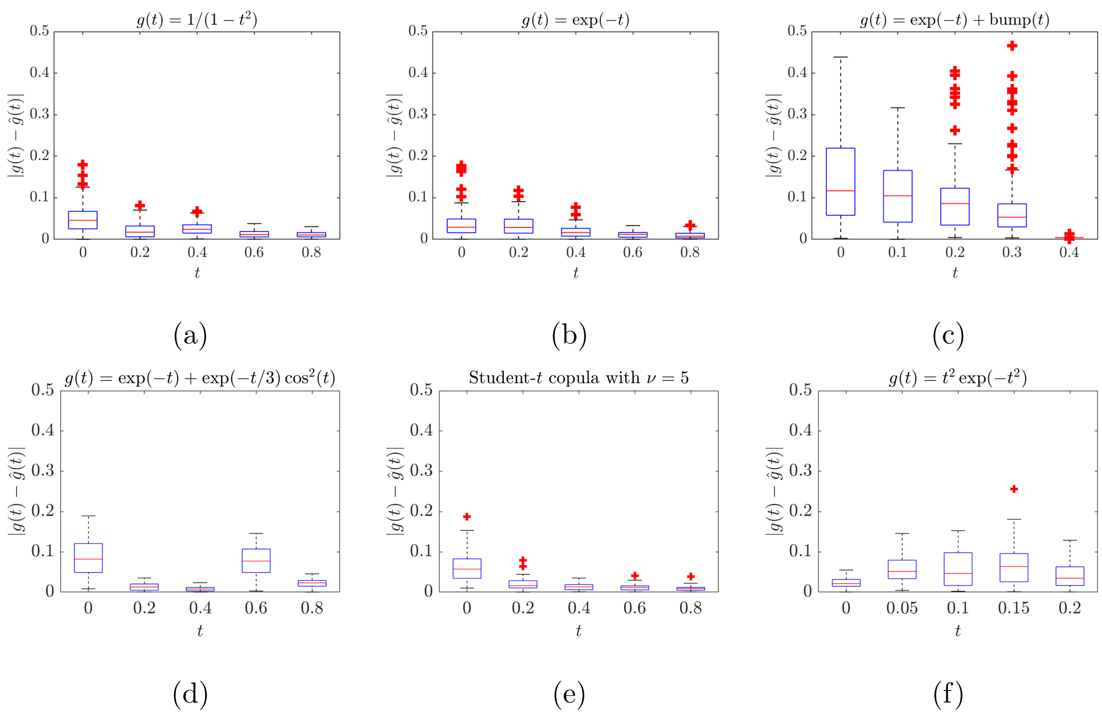

Figure 2: Boxplots of $|g(t) - \hat{g}(t)|$ at five grid values, based on 100 random samples.

boxplots of Figure 2 that the values of $|g(t) - \hat{g}(t)|$ are larger when $g(t)$ has sharp turns, such as in (c) and (f). Overall, the values of $|g(t) - \hat{g}(t)|$ are small, indicating good fits.

Comparison with Derumigny's method:

Using the same datasets, we also evaluated the approach implemented in the R-package `EllipCopulas` by Derumigny and Fermanian (2022). Figures 3 and 4 present the estimated generator $\hat{g}(t)$ from both methods side by side.

It can be observed that the estimates obtained from `EllipCopulas` exhibit greater bias, particularly in the cases of $g(t) = \exp(-t)+\text{bump}(t)$ and $g(t) = \exp(-t)+\exp(-t/3)\cos^2(t)$. In the case of $g(t) = t^2 \exp(-t^2)$, the performance of `EllipCopulas` is very unsatisfactory which is also shown in the paper by Derumigny and Fermanian (2022). Table 2 reports the mean integrated squared error (MISE) of the estimated generator functions for both methods. The column MISE1 corresponds to our proposed approach, while MISE2 represents the estimates obtained using `EllipCopulas`.

The results clearly demonstrate that our approach outperforms the alternative, particularly in scenarios where the generator exhibits complex patterns.

Table 2: Comparison of the Mean Integrated Squared Error (MISE) for different $g(t)$s using our approach (MISE1) and the method by Derumigny and Fermanian (2022) (MISE2).

$g(t)$	MISE1	MISE2
$1/(1+t^2)$	0.0016	0.0052
$\exp(-t)$	0.0007	0.0041
$\exp(-t) + \text{bump}(t)$	0.0139	0.0851
$\exp(-t) + \exp(-t/3)\cos^2(t)$	0.0032	0.0172
Student-t copula with $\nu = 5$	0.0021	0.0076
$t^2 \exp(-t^2)$	0.0035	0.0362

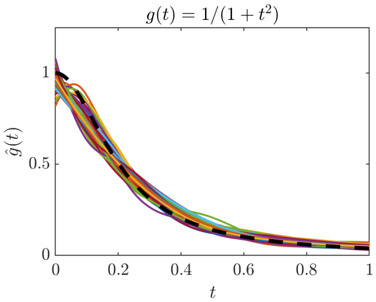

(a)

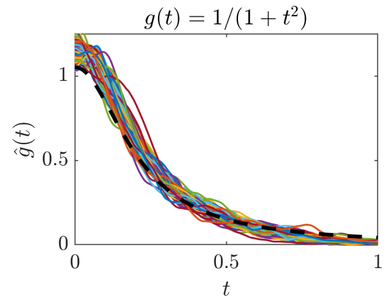

(b)

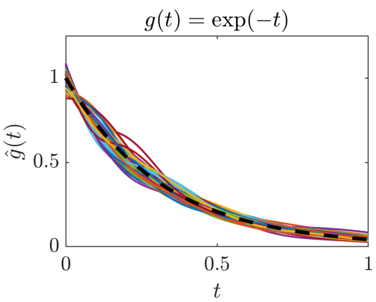

(c)

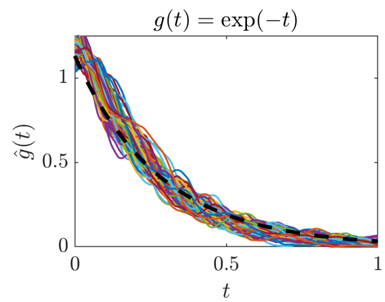

(d)

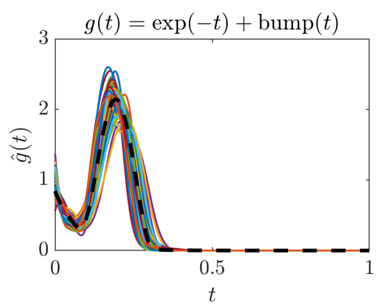

(e)

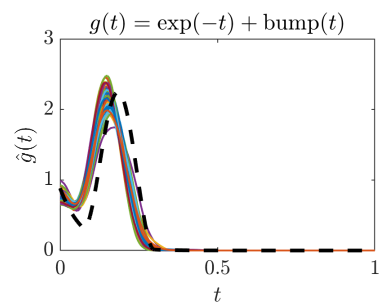

(f)

Figure 3: Comparison of the estimated function $\hat{g}$ using our proposed approach (left) and the method by Derumigny and Fermanian (2022) (right). Each row corresponds to a different functional form of $g(t)$.

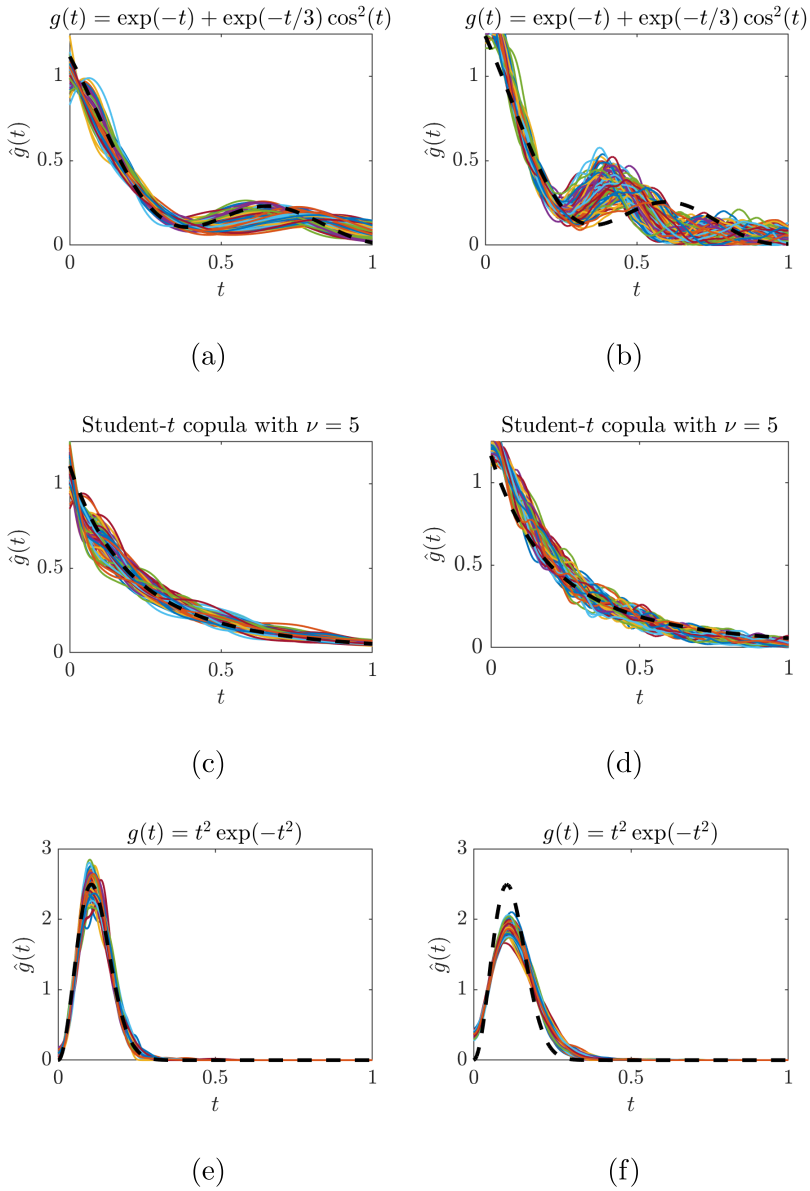

Figure 4: Comparison of the estimated function $\hat{g}$ using our proposed approach (left) and the method by Derumigny and Fermanian (2022) (right). Each row corresponds to a different functional form of $g(t)$.

Comparison with Genest's method:

In the approach proposed by Genest et al. (2007), the data analyst creates a list of candidate generators and tests them on a given data set to determine which one performs best in a goodness-of-fit test. The main drawback of this method is that the true generator, g, may not be included in the list of common generators, leading to a potentially incorrect selection. To illustrate this issue, we present a small example. We apply Genest's method to two copula data sets used previously, i.e., the one with $g(t) = 1/(1+t^2)$ and the one with $g(t) = \exp(-t) + \exp(-t/3)\cos^2(t)$. We assume the copula is either Gaussian or Student-t with some degrees of freedom. For estimation, we use the Matlab function `copulafit`, which operates similarly to Genest's method. The generator $g(t) = 1/(1+t^2)$ is similar to the generators of the Gaussian copula and the Student-t copula. In this case, we correctly identify the generator, as shown by Figure 5 (a), where the estimated generators and the true one almost overlap. However, when the true generator is $g(t) = \exp(-t) + \exp(-t/3)\cos^2(t)$, which is not included in our candidate list, the resulting estimate is very inaccurate, as shown in Figure 5 (b). By contrast, our method yields a much better estimate, as shown in Figure 1 (d).

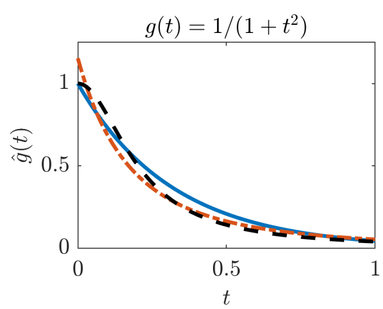

(a)

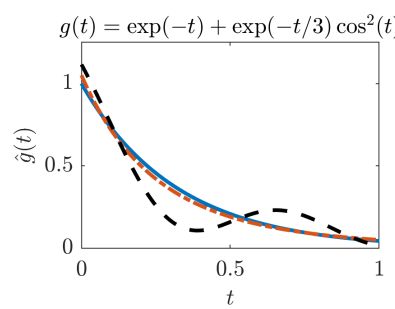

(b)

Figure 5: The true generators, normalized and standardized (dashed lines), along with estimates with a Gaussian copula (solid line) and a Student-t copula (dash-dotted line),

6 Higher Dimensions

When applied directly, our approach does not produce satisfactory results in simulation studies for cases where $d > 2$. A possible explanation is that certain regions of f_Y exhibit very large first derivatives (in absolute value) as the dimension of the elliptical copula increases, as illustrated in Figure 6. This issue cannot be mitigated by adjusting the value of a in Equation (14). However, this may not be a significant concern, as noted in Remark 3 of Derumigny and Fermanian (2022). If $\mathbf{U}_d \sim C(\Omega, g)$, then any two-dimensional subvector of $\mathbf{U}_d$ follows an elliptical copula with the same generator g. Therefore, given data from an elliptical copula of a higher dimension, we can estimate its generator using only a 2-dimensional subset of the data. Next, we apply this method to estimate the generator of a 5-dimensional Gaussian copula.

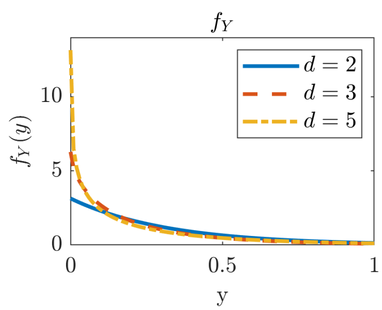

Figure 6: The PDF $f_Y(t)$ for Gaussian copulas of different dimensions.

We generate $N = 1000$ observations from a 5-dimensional Gaussian copula with correlation matrix

$$\Omega = \begin{bmatrix} 1 & 0.2 & 0 & 0 & 0 \\ 0.2 & 1 & 0.2 & 0 & 0 \\ 0 & 0.2 & 1 & 0.2 & 0 \\ 0 & 0 & 0.2 & 1 & 0.2 \\ 0 & 0 & 0 & 0.2 & 1 \end{bmatrix}$$

assumed known. We then obtain a 2-dimensional subset of the copula data to estimate the generator g. The MCMC settings are as in Section 5 except that we have only used 50 independent data sets this time. From the "spaghetti" plot and boxplot in Figure 7, we can see that the fit is very good.

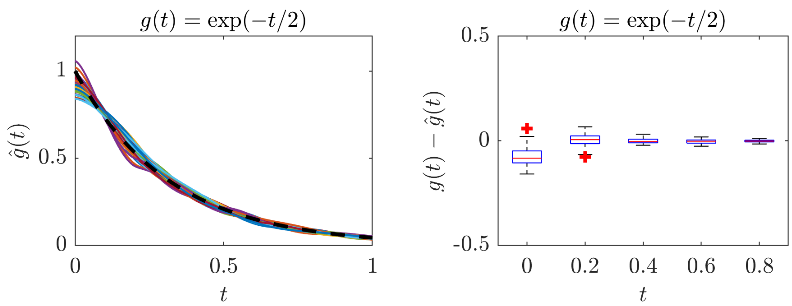

(a) (b)

Figure 7: The true generator of a 5-dimensional Gaussian copula, normalized and standardized (dashed black lines), along with 50 posterior means, and boxplots of $|g(t) - \hat{g}(t)|$ at five grid values, based on 50 random samples.

7 Application

7.1 Wine Data

The wine quality data (Cortez and Reis, 2009) contain two datasets, having to do with red and white variants of the Portuguese "Vinho Verde" wine. The data are taken from the UC irvine Machine Learning Repository. They contain various variables such as citric acid level, pH level, density, and sulphates, measured from distinct 1599 red and 4898 white wine examples, from May 2004 to February 2007. The data are also available in the supplementary material. In our analysis we focus on modeling the dependence between variables of the red wine data. In particular, we use an elliptical copula to model the dependence between citric acid (`citricAcid`) and pH level (`pH`) where the joint distribution of `citricAcid` and `pH` is assumed to be meta-elliptical. The histograms in Figure 8a and 8b show that the marginal distribution of citric acid is multi-modal while that of pH is bell-shaped.

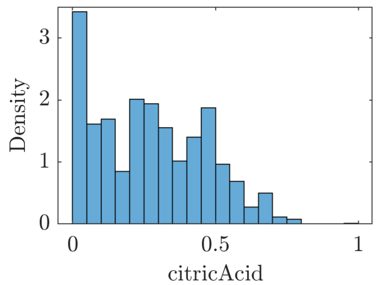

(a) Citric acid.

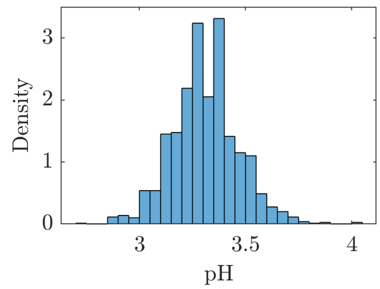

(b) pH.

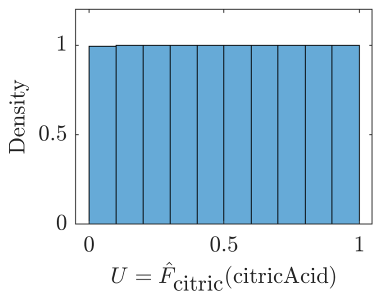

(c) citric acid copula data.

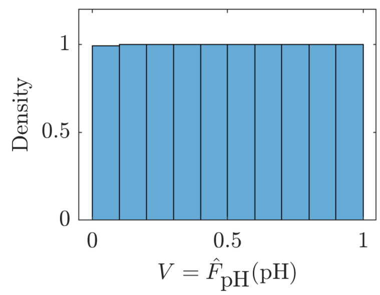

(d) pH copula data.

Figure 8: Histograms of citric acid and pH along with their copula data.

7.2 Estimation and Results

We estimate the marginal distributions of `citricAcid` and `pH` first, using empirical cdfs, and then use the probability integral transform to obtain copula data. In particular,

$$U = \hat{F}_{\text{citric}}(\text{citricAcid}),\ V = \hat{F}_{\text{pH}}(\text{pH}),$$

where $\hat{F}_{\text{citric}}$ and $\hat{F}_{\text{pH}}$ are empirical CDFs. The marginal histograms of the copula data appear in Figure 8c and 8d. These histograms look close to uniform, as expected. A bivariate histogram is presented in Figure 9a.

To estimate the correlation matrix Ω, we first compute the sample Kendall's tau and then convert it into the Pearson correlation using the transformation $\rho = \sin(\pi * \tau/2)$. This

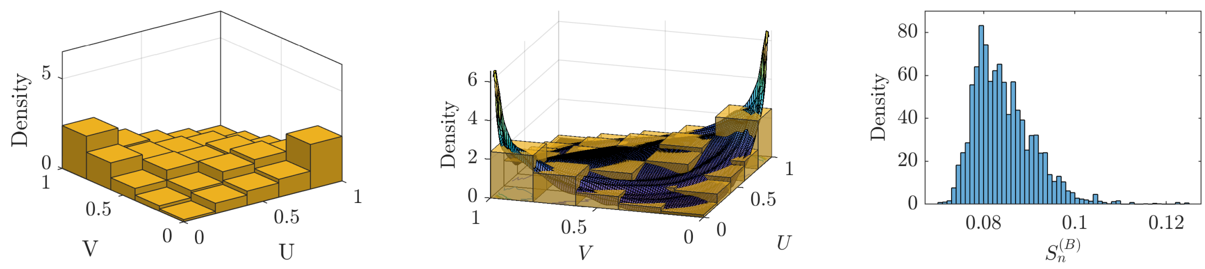

(a) Histogram of the copula data.

(b) Fitted copula density overlaid on the histogram of the copula data.

(c) Density histogram of $S_n^{(B)}$.

Figure 9: Density histogram of the copula data, fitted copula density and a density histogram of the goodness of fit measure.

results in

$$\hat{\Omega} = \begin{bmatrix} 1 & -0.57 \\ -0.57 & 1 \end{bmatrix}.$$

Thus, in the ensuing MCMC scheme, we only estimate the copula generator. The MCMC scheme is run for $10,000$ iterations where the first 5000 are used as burn-in. The results are presented in Figure 9b. The estimated copula PDF and generator are the posterior means computed based on the iterations after burn-in. It is evident that the method provides a good fit to the data.

We also evaluate goodness of fit using the $S_n^{(B)}$ statistic from Genest et al. (2009), which is computed via Rosenblatt's probability integral transform (Rosenblatt (1952)). Figure 9c shows a density histogram of the $S_n^{(B)}$ values computed at each MCMC iteration after burn-in. The 99%th empirical quantile of these values is equal to 0.1049, which indicates a good fit of the estimated elliptical copula to the wine data.

8 Conclusion and discussion

We have proposed a new inferential procedure to estimate generators of elliptical copulas in a Bayesian framework, which exhibits good results in simulation studies. A more theoretical study of the formulation of elliptical copulas may be a topic of further research. Elliptical copula generators are not unique and identifiable because they contain abundant information on the marginal distributions. Elliptical copulas may have a better parameterization which only includes information on dependence structures. Future work may examine more elegant formulations for elliptical copulas instead of simply utilizing their associated elliptical distributions and the inversion method.

Disclosure

The authors report there are no competing interests to declare.

SUPPLEMENTARY MATERIAL

Appendix: A pdf file, describing the construction of B-spline densities and the algorithm to sample an elliptical copula.

EllipCopEst2D: A zipped file, containing Matlab code and a "ReadMe" file to illustrate the methods described in the article.

Wine quality data set: Data set used in the application. (.csv file)

References

Abdous, B., Genest, C., and Rémillard, B. (2005), "Dependence properties of meta-elliptical distributions," in *Statistical modeling and analysis for complex data prob-*

lems, volume 1 of *GERAD 25th Anniv. Ser.*, Springer, New York, 1–15, URL https://doi.org/10.1007/0-387-24555-3_1.

Cambanis, S., Huang, S., and Simons, G. (1981), "On the theory of elliptically contoured distributions," *J. Multivariate Anal.*, 11, 368–385, URL https://doi.org/10.1016/0047-259X(81)90082-8.

Chib, S. and Jeliazkov, I. (2006), "Inference in semiparametric dynamic models for binary longitudinal data," *J. Amer. Statist. Assoc.*, 101, 685–700, URL https://doi.org/10.1198/016214505000000871.

Cortez, Paulo, C. A. A. F. M. T. and Reis, J. (2009), "Wine Quality," UCI Machine Learning Repository. DOI: https://doi.org/10.24432/C56S3T.

de Boor, C. (1978), *A practical guide to splines*, volume 27 of *Applied Mathematical Sciences*, Springer-Verlag, New York-Berlin.

Derumigny, A. and Fermanian, J.-D. (2022), "Identifiability and estimation of meta-elliptical copula generators," *J. Multivariate Anal.*, 190, Paper No. 104962, 19, URL https://doi.org/10.1016/j.jmva.2022.104962.

Eilers, P. and Marx, B. (2021), *Practical smoothing: the joys of p-splines*, Cambridge University Press, URL https://books.google.com/books?id=ez0QEAAAQBAJ.

Eilers, P. H. C. and Marx, B. D. (1996), "Flexible smoothing with B-splines and penalties," *Statist. Sci.*, 11, 89–121, URL https://doi.org/10.1214/ss/1038425655. With comments and a rejoinder by the authors.

Fang, H.-B., Fang, K.-T., and Kotz, S. (2002), "The meta-elliptical distributions with given marginals," *J. Multivariate Anal.*, 82, 1–16, URL https://doi.org/10.1006/jmva.2001.2017.

— (2005), “Corrigendum to: “The meta-elliptical distributions with given marginals” [J. Multivariate Anal. **82** (2002), no. 1, 1–16; MR1918612],” *J. Multivariate Anal.*, 94, 222–223, URL `https://doi.org/10.1016/j.jmva.2004.10.001`.

Gelman, A. (2006), “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper),” *Bayesian Anal.*, 1, 515–533, URL `https://doi.org/10.1214/06-BA117A`.

Genest, C., Favre, A.-C., Béliveau, J., and Jacques, C. (2007), “Metaelliptical copulas and their use in frequency analysis of multivariate hydrological data,” *Water Resources Research*, 43.

Genest, C., Rémillard, B., and Beaudoin, D. (2009), “Goodness-of-fit tests for copulas: a review and a power study,” *Insurance Math. Econom.*, 44, 199–213, URL `https://doi.org/10.1016/j.insmatheco.2007.10.005`.

Ghidey, W., Lesaffre, E., and Eilers, P. (2004), “Smooth random effects distribution in a linear mixed model,” *Biometrics*, 60, 945–953, URL `https://doi.org/10.1111/j.0006-341X.2004.00250.x`.

Gómez, E., Gómez-Villegas, M. A., and Marín, J. M. (2003), “A survey on continuous elliptical vector distributions,” *Rev. Mat. Complut.*, 16, 345–361, URL `https://doi.org/10.5209/rev_REMA.2003.v16.n1.16889`.

Kelker, D. (1970), “Distribution theory of spherical distributions and a location-scale parameter generalization,” *Sankhyā Ser. A*, 32, 419–438.

Lang, S. and Brezger, A. (2004), “Bayesian P-splines,” *J. Comput. Graph. Statist.*, 13, 183–212, URL `https://doi.org/10.1198/1061860043010`.

Lemonte, A. J. and Patriota, A. G. (2011), “Multivariate elliptical models with general parameterization," *Stat. Methodol.*, 8, 389–400, URL https://doi.org/10.1016/j.stamet.2011.03.001.

Liebscher, E. (2005), "A semiparametric density estimator based on elliptical distributions," *J. Multivariate Anal.*, 92, 205–225, URL https://doi.org/10.1016/j.jmva.2003.09.007.

Nelsen, R. B. (2006), *An introduction to copulas*, Springer Series in Statistics, Springer, New York, second edition, URL https://doi.org/10.1007/s11229-005-3715-x.

Rosen, O. and Thompson, W. K. (2015), "Bayesian semiparametric copula estimation with application to psychiatric genetics," *Biom. J.*, 57, 468–484, URL https://doi.org/10.1002/bimj.201300130.

Rosenblatt, M. (1952), "Remarks on a multivariate transformation," *Ann. Math. Statistics*, 23, 470–472, URL https://doi.org/10.1214/aoms/1177729394.

Sklar, A. (1959), "Fonctions de répartition à n dimensions et leurs marges," *Publ. Inst. Statist. Univ. Paris*, 8, 229–231.

Staudenmayer, J., Ruppert, D., and Buonaccorsi, J. P. (2008), "Density estimation in the presence of heteroscedastic measurement error," *J. Amer. Statist. Assoc.*, 103, 726–736, URL https://doi.org/10.1198/016214508000000328.

Steerneman, A. G. M. and van Perlo-ten Kleij, F. (2005), "Spherical distributions: Schoenberg (1938) revisited," *Expo. Math.*, 23, 281–287, URL https://doi.org/10.1016/j.exmath.2005.01.002.