2023-01-30

Bayesian Nonparametric Modelling of Conditional Multidimensional Dependence Structures

Barone, R

http://hdl.handle.net/10026.1/20188

Bayesian Nonparametric Modelling of Conditional Multidimensional Dependence Structures

Rosario Barone [1], Luciana Dalla Valle [2]

[1] Sapienza University of Rome, Italy   [2] University of Plymouth, UK

Abstract

In recent years, conditional copulas, that allow dependence between variables to vary according to the values of one or more covariates, have attracted increasing attention. However, the literature mainly focused on the bivariate case, since the constraints on the multivariate copulas correlation matrices would make the specifications of covariates arduous. In high dimension, vine copulas offer greater flexibility compared to multivariate copulas, since they are constructed using bivariate copulas as building blocks. We present a novel inferential approach for multivariate distributions, which combines the flexibility of vine constructions with the advantages of Bayesian nonparametrics, not requiring the specification of parametric families for each pair copula. Expressing multivariate copulas using vines allows us to easily account for covariate specifications driving the dependence between response variables. We specify the vine copula density as an infinite mixture of Gaussian copulas, defining a Dirichlet process prior on the mixing measure, and performing posterior inference via Markov chain Monte Carlo sampling. Our approach is successful as for clustering as well as for density estimation. We carry out simulation studies and apply the proposed approach to analyse a veterinary dataset and to investigate the impact of natural disasters on financial development. Supplementary materials are available online.

Keywords: Conditional Copulas, Dirichlet Process Prior, Heterogeneity, MCMC, Mixtures, Vine Copulas.

1 Introduction

In many real data applications we are often required to model jointly $d \geq 3$ continuous random variables, denoted as $Y_1, \ldots, Y_d$. The multivariate distribution, which allows us to describe the joint behaviour of those variables, can be denoted as $F(Y_1, \ldots, Y_d) = P(Y_1 \leq y_1, \ldots, Y_d, \leq y_d)$. However, complex relations between data, particularly asymmetric and tail dependent associations, are often difficult to be modelled. The copula approach allows us to express the multivariate distribution of a set of variables by separating the marginals from the dependence structure. Copulas were introduced by Sklar [1959] and, since then, they have been applied in a wide variety of fields (see, for example, Kolev et al. [2006] for a review, and Genest et al. [2009] and Fan and Patton [2014] for applications in finance and economics, respectively).

In recent years, the idea of modelling the effect of covariates on the dependence structure described by copulas has attracted increasing attention. However, literature contributions concentrated on the bivariate case, since the specification of covariates for traditional multivariate copulas is challenging, due the constraints on the correlation matrix. For example, Patton [2006], Jondeau and Rockinger [2006] and Bartram et al. [2007] considered time-varying dependence copula parameters in time series analysis. Acar et al. [2011] estimated the functional relationship between copula parameters and covariates adopting a non-parametric approach. Craiu and Sabeti [2012] introduced a bivariate conditional copula model for continuous or mixed outcomes. Abegaz et al. [2012] and Gijbels et al. [2011], respectively, suggested

semiparametric and non-parametric approaches for the estimation of conditional copulas, proving the consistency and the asymptotic normality of the estimators.

For inference, several contributions in the literature follow the Bayesian nonparametric approach. Wu et al. [2014] presented a Bayesian nonparametric method for estimating multivariate copulas using a DP mixture of multivariate skew-Normal copulas, and Wu et al. [2015] proposed a DP mixture of bivariate Gaussian copulas. In both cases the authors performed posterior inference via slice sampling [Walker, 2007, Kalli et al., 2011]. Dalla Valle et al. [2018] extended the approach of Wu et al. [2015] to bivariate conditional copulas, introducing dependence from covariates and implementing Bayesian nonparametric inference via an infinite mixture model. Ning and Shephard [2018] considered a different choice of prior for unconditional bivariate copulas, adopting a multi-partition Dirichlet-based Pólya tree prior. In the conditional case, Levi and Craiu [2018] proposed to jointly estimate the marginal distributions and the copula using Gaussian process (GP) models, where the calibration function follows a priori a single-index model based on GPs, to handle high-dimensional covariates.

A different approach is followed by Grazian and Liseo [2017], who described an approximate Bayesian inference method for semiparametric bivariate copulas, based on the empirical likelihood. This approach is extended by Grazian et al. [2022], who compared several Bayesian methods to approximate the posterior distribution of functionals of the dependence including

3

covariates, using nonparametric models which avoid the selection of the copula function. However, these contributions are mostly limited to the bivariate case, due to the challenges related to extending the conditional model in higher dimension. Lu [2021], instead, focused on the nonparametric estimation of conditional copulas employing the empirical checkerboard Bernstein copula (ECBC) estimator, based on a hierarchical empirical Bayes model that enables the estimation of a smooth copula function. Lu [2021]'s family of copulas allows for both clustering and density estimation. The authors propose nonparametric multivariate copulas for an arbitrary number of dimensions, with applications, in the unconditional case, to portfolios of 3 and 10 assets. In the conditional case they only focus on the bivariate conditional copula estimation with a single covariate, underlying that the potential extensions to higher dimension and more covariates are possible. However, by construction Lu [2021]'s approach (i) requires the covariate to be continuous and (ii) the extensions to higher dimension and/or the accounting for more covariates turns out to be complex. Moreover, although the functional complexity helps for very accurate density estimations, on the other hand there is a lack of interpretability of the estimates, particularly as the dimension and the number of considered covariates increases. In this paper, we propose a method that exploits the advantages of vine copulas and Bayesian nonparametrics in terms of flexibility and, at the same time, returns easy-to-interpret results, by providing a procedure for estimating the effect of one or more - discrete or continuous - covariates on a d-dimensional dependence structure.

4

Vines are multivariate copulas constructed by using only bivariate building blocks which can be selected independently [Czado, 2019]. This class of flexible copula models has become very popular in the last years for many applications in diverse fields such as finance, insurance, hydrology, marketing, engineering, chemistry, aviation, climatology and health. The popularity of vine copulas is due to the fact that they allow, in addition to the separation of margins and dependence by the copula approach, tail asymmetries and separate multivariate component modeling [Aas et al., 2009].

Sahin and Czado [2022] formulate a vine copula mixture model for continuous data allowing all types of vine tree structures, parametric pair copulas, and margins. However, their approach treat the number of mixture components as known and does not account for covariates, mainly focusing on the model selection problems. We propose a Bayesian nonparametric inferential approach for multivariate conditional copulas, which we express using vines. Conditional vine copulas were introduced by the frequentist work of Vatter and Nagler [2018]. However, the Bayesian nonparametric approach proposed in this paper circumvents the drawbacks of the frequentist approach, not requiring the specification of copula families for each pair-copula in the vine. We consider a DP mixture of vine copulas, assuming a DP prior distribution on the mixing measure of an infinite mixture of vine copulas. Bayesian nonparametric inference for mixture copula models was adopted by Zhuang et al. [2021] to group similar dependence structures. However, the authors restricted their attention to the single-parameter unconditional copula func-

tions and the bivariate scenario. Our approach has proven to be particularly successful in the multivariate conditional case, where the effects of covariates on the dependence structures are considered. Here, the constraints on the correlation matrix to which traditional multivariate copulas are subjected to would make the specifications of covariates arduous. Our approach overcomes the limitations of multivariate copula modelling, combining the flexibility of the vine construction with the advantages of the Bayesian nonparametric approach and allowing for both clustering and density estimation.

Our method allows us to model the unobserved heterogeneity, which is often present in real datasets, in a more natural and flexible way, compared to the current state-of-the-art. Bayesian mixture models that account for heterogeneity were implemented for example by Buddhavarapu et al. [2016], who incorporated heterogeneity in the model using a finite multivariate normal mixture prior on the random parameters. Zhang and Wang [2018] adopted a similar approach to capture heterogeneity in learning styles in a dataset collected from a computer-based learning system. Following the Bayesian nonparametric strand in the literature, Green and Richardson [2001] modelled heterogeneity with DP based models illustrated in the context of univariate mixtures. DP mixtures were also employed by Turek et al. [2021] to detect heterogeneity in ecological datasets. However, to the best our knowledge, this paper is the first to apply the Bayesian nonparametric approach to conditional vine copulas mixtures for analysing heterogeneous data. The main strength of our approach compared to traditional approaches is that it allows

us to capture the intrinsic heterogeneity in a more natural and flexible way. In particular, temporal heterogeneity can be easily captured by a Bayesian nonparametric vine model where the sequence of observations at each time point is represented by a set of nodes in the vine linked by pair-copulas. The vine structure is suitable, by its nature, to model the heterogeneity generated by the temporal sequence. In addition, the presence of different bivariate copulas for each pair of variables makes the model much more flexible than traditional approaches to capture heterogeneity in the data.

Presenting two applications, we first show how our model ensures enough flexibility by providing accurate predictive samples. We also demonstrate that our method allows us to capture the unobserved heterogeneity in a real dataset that relates financial development in different countries to the occurrence of natural disasters. Our approach identifies two distinct country clusters. In the first one, the financial development temporal dependence is negatively affected by natural calamities, while in the second one we observe a positive effect. These results reflect government preparedness to face natural hazards.

The remainder of the paper is organised as follows. Section 2 introduces conditional vine copulas; Section 3 illustrates the DP mixture approach for conditional vines; Section 4 focuses on the implementation of the model using the MCMC sampler; Section 5 applies the proposed approach to simulated datasets; while analyses of real datasets and discussions of results are reported in Section 6.

2 Conditional Copulas and Vines

Let us consider $Y_1, \ldots, Y_d$, which are continuous random variables of interest and let $\mathbf{X} = (X_1, \ldots, X_p)$ be a vector of covariates that may affect the dependence between $Y_1, \ldots, Y_d$. Then, the conditional joint distribution function of $(Y_1, \ldots, Y_d)$ given $\mathbf{X} = \mathbf{x}$ is

$$F_x(y_1, \ldots, y_d) = P(Y_1 \leq y_1, \ldots, Y_d \leq y_d | \mathbf{X} = \mathbf{x}),$$

under the assumption that such conditional distribution exists (see Gijbels et al. [2012], Abegaz et al. [2012] and Acar et al. [2011]).

We denote the conditional marginals of F_x as

$$
\begin{aligned}
F_{1,x}(y_1) &= P(Y_1 \leq y_1 | \mathbf{X} = \mathbf{x}), \\
&\ldots \\
F_{d,x}(y_d) &= P(Y_d \leq y_d | \mathbf{X} = \mathbf{x}).
\end{aligned}
$$

If the marginals are continuous, then Sklar's theorem [Sklar, 1959] allows us to write

$$C_x(u_1, \ldots, u_d) = F_x \left(F_{1,x}^{-1}(u_1), \ldots, F_{d,x}^{-1}(u_d) \right)$$

where $F_{j,x}^{-1}(u_j) = \inf \{y_j : F_{j,x} \geq u_j\}$, for $j = 1, \ldots, d$, are the conditional quantile functions and $u_j = F_{j,x}(y_j)$ are called pseudo-observations or u-data. The conditional copula C_x fully describes the conditional dependence structure of $(Y_1, \ldots, Y_d)$ given $\mathbf{X} = \mathbf{x}$. Therefore, the conditional joint distribution of $(Y_1, \ldots, Y_d)$ given $\mathbf{X} = \mathbf{x}$ can be written as

$$F_x(Y_1, \ldots, Y_d) = C_x \left(F_{1,x}(y_1), \ldots, F_{d,x}(y_d) \right).$$

Let us denote the copula density corresponding to the distribution $C_x\left(F_{1x}(y_1), \ldots, F_{dx}(y_d)\right)$ as

$$c_x\left(u_1, \ldots, u_d\right) = c_{\boldsymbol{\theta}}(u_1, \ldots, u_d | \mathbf{x}) = c_{\boldsymbol{\theta}(\mathbf{x})}(u_1, \ldots, u_d),$$

where $\boldsymbol{\theta}$ is the parameter vector of the d-variate copula density. We assume that the function $\boldsymbol{\theta}(\mathbf{x})$ depends on a vector of parameters $\boldsymbol{\beta}$ such that

$$c_{\boldsymbol{\theta}(\mathbf{x})}(u_1, \ldots, u_d) = c_{\boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})}(u_1, \ldots, u_d) = c_{1:d}(u_1, \ldots, u_d \,|\, \boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})). \qquad (1)$$

The (1) can be written in terms of vines [Czado, 2019], where each pair-copula depends on the vector of covariates $\mathbf{X}$.


Figure 1: Trivariate vine representation.

For example, considering a tri-variate vine copula, depicted in Figure 1, we obtain

$$\begin{aligned}
c_{1:3}(u_1, u_2, u_3 \,|\, \boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})) = {} & c_{1,2}(F_{1,x}(y_1), F_{2,x}(y_2) \,|\, \boldsymbol{\theta}_{12}(\mathbf{x}|\boldsymbol{\beta})) \\
& \times c_{2,3}(F_{2,x}(y_2), F_{3,x}(y_3) \,|\, \boldsymbol{\theta}_{23}(\mathbf{x}|\boldsymbol{\beta})) \\
& \times c_{1,3;2}(F_{1|2,x}(y_1|y_2), F_{3|2,x}(y_3|y_2) \,|\, \boldsymbol{\theta}_{13;2}(\mathbf{x}|\boldsymbol{\beta})),
\end{aligned}$$

where $\boldsymbol{\theta}_{13;2}(\mathbf{x}|\boldsymbol{\beta})$ denotes the parameter function of the pair copula $c_{1,3;2}(\cdot, \cdot)$, $F_{1|2,x}(\cdot|\cdot)$ is the conditional distribution of Y_1 given Y_2 and $\mathbf{X} = \mathbf{x}$, $F_{3|2,x}(\cdot|\cdot)$

is the conditional distribution of Y_3 given Y_2 and $\mathbf{X} = \mathbf{x}$, $\boldsymbol{\theta}_{12}(\mathbf{x}|\boldsymbol{\beta})$ denotes the parameter function of the pair copula $c_{1,2}(\cdot, \cdot)$ and $\boldsymbol{\theta}_{23}(\mathbf{x}|\boldsymbol{\beta})$ denotes the parameter function of the pair copula $c_{2,3}(\cdot, \cdot)$.

In higher dimension, the vine representation can be generalized to special vine distribution classes, the most popular of which are D-vines (see Bedford and Cooke [2001], Aas et al. [2009] and Czado [2019]). The conditional D-vine decomposition takes the form

$$
\begin{aligned}
c_{1:d}(u_1, \ldots, u_d \,|\, \boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})) = \\
\prod_{\ell=1}^{d-1} \prod_{k=1}^{d-\ell} c_{k,\ell+k;k+1,\ldots,k+\ell-1} \big\{ F_{k|k+1,\ldots,k+\ell-1,x}(y_k|y_{k+1,\ldots,k+\ell-1}), \\
F_{\ell+k|k+1,\ldots,k+\ell-1,x}(y_{\ell+k}|y_{k+1,\ldots,k+\ell-1}) \,|\, \boldsymbol{\theta}_{k,\ell+k;k+1,\ldots,k+\ell-1}(\mathbf{x}|\boldsymbol{\beta}) \big\} \, .
\end{aligned}
$$

3 Dirichlet Process Mixture of Conditional Vine Copulas

Vine copulas require a model selection step where a copula family is selected for each bivariate pair-copula forming the vine. However, the complexity of the problem increases with the vine dimension. As a solution, we propose a hierarchical approach based on the vine construction. More specifically, we adopt a Bayesian nonparametric approach, which overcomes the need of specifying the families of each pair-copula. Moreover, since often it is hard to consider the effect of the covariates on the dependence structures, par-

ticularly when the copula matrices are unstructured, we provide a general methodology which also allows to account for covariates in multivariate dependence structures, exploiting the flexibility of the pair-copula construction. More precisely, we consider the effect of covariates on the dependence structures between the paired variable, leveraging the information deriving from the covariates for a more accurate clustering.

Given a set of N observations, let us consider the random vectors of interest $\mathbf{Y}_1, \ldots, \mathbf{Y}_d$ and the $N \times p$ matrix of observed covariates $\mathbb{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_p)$. As in Müller and Rosner [1997] and Müller and Quintana [2010], we define the covariates as random variables, such that $\mathbb{X}$ corresponds to N realizations of the independent random variables with densities $f_h(x_h)$, with $h = 1, \ldots, p$.

Starting from a simple example, let us consider the three-dimensional case with $d = 3$ where $\mathbf{Y}_1, \mathbf{Y}_2, \mathbf{Y}_3$ are random vectors of interest and $(\mathbf{X}_1, \ldots, \mathbf{X}_p)$ are covariate vectors influencing the dependence between the variables of interest. Let $F_{1,x}(y_{1i}), F_{2,x}(y_{2i}), F_{3,x}(y_{3i})$, with $i = 1, \ldots, N$, be the conditional cdfs of the variables of interest, and let $\mathbb{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_p)$ be the $N \times p$ matrix of observed covariates, such that $x_{i,h}$ corresponds to the i-th realization of the h-th covariate. Adopting the 3-dimensional D-vine specification illustrated in Figure 1, we model the dependence between the variables of interest as the product of $d = 3$ pair-copulas indexed by the vector of parameters $\boldsymbol{\beta} = \left(\beta_{0_{12}}, \beta_{1_{12}}, \ldots, \beta_{p_{12}}, \beta_{0_{23}}, \beta_{1_{23}}, \ldots, \beta_{p_{23}}, \beta_{0_{13;2}}, \beta_{1_{13;2}}, \ldots, \beta_{p_{13;2}} \right)$, where the subscripts $\{12\}, \{23\}, \{13; 2\}$ denote the pair-copulas in the vine. We assume that the covariate distributions are governed by the matrix of

parameters $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_p)$ where the dimension of $\boldsymbol{\phi}$ depends on the co-variate kernels. For example, if we assume each covariate X_h to be Normally distributed with parameters (μ_h, σ_h), for $h = 1, \ldots, p$, $\boldsymbol{\phi}$ will be a $(p \times 2)$-dimensional matrix. Let us consider a random probability measure G and a vector of parameters $\boldsymbol{\xi} = (\boldsymbol{\beta}, \boldsymbol{\phi})$ defined on the parameter space Ξ. Let us assume that $\mathbf{f}(\mathbf{x}) = \prod_{h=1}^{p} f_h(x_h)$ denotes the product of the densities of the covariates, and $c_{\boldsymbol{\xi}}(\cdot, \cdot, \cdot | \mathbf{x})$ denotes the 3-variate conditional copula density arising from the vine. We rewrite the density $\mathbf{f}_G(\mathbf{x}) \cdot c_G(\cdot, \cdot, \cdot | \mathbf{x})$ as an infinite mixture of vine copulas with kernel $\mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}) \cdot c_{\boldsymbol{\xi}}(\cdot, \cdot, \cdot | \mathbf{x})$ with respect to the mixing measure G, such that

$$\mathbf{f}_G(\mathbf{x}) c_G(F_{1,x}(y_1), F_{2,x}(y_2) | \mathbf{x}) c_G(F_{2,x}(y_2), F_{3,x}(y_3) | \mathbf{x})$$

$$\times c_G(F_{1|2,x}(y_1|y_2), F_{3|2,x}(y_3|y_2) | \mathbf{x}) =$$

$$\int \mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}) c_{\boldsymbol{\xi}}(F_{1,x}(y_1), F_{2,x}(y_2) | \mathbf{x}) c_{\boldsymbol{\xi}}(F_{2,x}(y_2), F_{3,x}(y_3) | \mathbf{x})$$

$$\times c_{\boldsymbol{\xi}}(F_{1|2,x}(y_1|y_2), F_{3|2,x}(y_3|y_2) | \mathbf{x}) dG(\boldsymbol{\xi}).$$

Generalizing, let us consider d-variables of interest $Y_1, \ldots, Y_d$. Given a set of N observations, let us assume that $F_{1,x}(y_{1i}), \ldots, F_{d,x}(y_{di})$, with $i = 1, \ldots, N$, are the conditional cdfs of the d variables of interest. The multivariate dependence structure is specified by a vine defined as the product of $\nu = d(d-1)/2$ pair copulas, indexed by the $\nu \times (q+1)$-dimensional vector of parameters $\boldsymbol{\beta}$, while $f_h(x_h)$, $h = 1, \ldots, p$, are independent random variables with parameters $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_p)$. Note that $q \geq p$ and its value depends on the chosen link function; for example if the link is linear $q = p$. Let the

vector of parameters $\boldsymbol{\xi} = (\boldsymbol{\beta}, \boldsymbol{\phi})$ be defined on the parameter space Ξ. As before, we rewrite the density $\mathbf{f}_G(\mathbf{x}) \cdot c_G(\cdot, \ldots, \cdot | \mathbf{x})$ as an infinite mixture of conditional vine copulas with kernel $\mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}) \cdot c_{\boldsymbol{\xi}}(\cdot, \ldots, \cdot | \mathbf{x})$ with respect to the mixing measure G, that in the D-vine case is

$$
\mathbf{f}_G(\mathbf{x}) \prod_{\ell=1}^{d-1} \prod_{k=1}^{d-\ell} c_G \left(F_{k|k+1,\ldots,k+\ell-1,x}(y_k | y_{k+1,\ldots,k+\ell-1}), \right.
$$

$$
\left. F_{\ell+k|k+1,\ldots,k+\ell-1,x}(y_{\ell+k} | y_{k+1,\ldots,k+\ell-1}) | \mathbf{x} \right) =
$$

$$
\int \mathbf{f}_{\boldsymbol{\xi}}(\mathbf{x}) \prod_{\ell=1}^{d-1} \prod_{k=1}^{d-\ell} c_{\boldsymbol{\xi}} \left(F_{k|k+1,\ldots,k+\ell-1,x}(y_k | y_{k+1,\ldots,k+\ell-1}), \right.
$$

$$
\left. F_{\ell+k|k+1,\ldots,k+\ell-1,x}(y_{\ell+k} | y_{k+1,\ldots,k+\ell-1}) | \mathbf{x} \right) dG(\boldsymbol{\xi}) =
$$

$$
\int \mathbf{f}_{\boldsymbol{\phi}}(\mathbf{x}) \prod_{\ell=1}^{d-1} \prod_{k=1}^{d-\ell} c_{\boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})} \left(F_{k|k+1,\ldots,k+\ell-1,x}(y_k | y_{k+1,\ldots,k+\ell-1}), \right.
$$

$$
\left. F_{\ell+k|k+1,\ldots,k+\ell-1,x}(y_{\ell+k} | y_{k+1,\ldots,k+\ell-1}) | \mathbf{x} \right) dG(\boldsymbol{\phi}, \boldsymbol{\beta}).
$$

With a Dirichlet Process (DP) prior on the mixing measure G, we get a Dirichlet Process Mixture (DPM) of conditional vine copulas, which may be represented as

$$
\mathbf{f}_{\boldsymbol{\phi}}(\mathbf{x}) c_{\boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta})}(u_1, \ldots, u_d \,|\, \mathbf{x}) = \sum_{j=1}^{\infty} \omega_j \, \mathbf{f}_{\boldsymbol{\phi}_j}(\mathbf{x}) c_{1:d}(u_1, \ldots, u_d \,|\, \boldsymbol{\theta}(\mathbf{x}|\boldsymbol{\beta}_j)),
$$

where the weights ω_j sum to 1.

The posterior distribution $\Pi(G|\mathbf{Y}, \mathbf{X})$ is a mixture of DP models, mixing with respect to the latent variables $\boldsymbol{\xi}_i$ specific to each observation i for $i = 1, \ldots, N$:

$$
G|\mathbf{Y}, \mathbf{X} \sim \int DP \left(M G_0 + \sum_{i=1}^{N} \delta_{\boldsymbol{\phi}_i \boldsymbol{\beta}_i} \right) d\Pi(\boldsymbol{\phi}, \boldsymbol{\beta} | \mathbf{y}, \mathbf{x}),
$$

where δ_t denotes the Dirac measure at t.

In general, the choice of the kernel for DPM models should consider two aspects. From a computational point of view, the conjugacy between the centering measure G_0 and the kernel $f_{\boldsymbol{\phi}} c_{\boldsymbol{\beta}}$ is particularly convenient. A more important feature is the flexibility of the chosen density. In our case, we need to specify a product of bivariate copula densities and a product of independent densities for the covariates. The choice of the covariate densities will depend on the nature of X_h, $h = 1, \ldots, p$. On the other hand, in order to model the dependence structure between the variables, we need a copula which is able to capture various kinds of dependence and may approximate different copula families. Wu et al. [2015] showed that bivariate density functions on the real plain can be arbitrarily well approximated by a mixture of a countably infinite number of bivariate normal distributions. Dalla Valle et al. [2018] proposed a Bayesian nonparametric estimation of bivariate conditional copulas, with a Gaussian copula as kernel of a DPM. Following the previous approaches, we propose as kernel the product of the density of a Gaussian vine-copula and the densities of covariates which depend on the nature of $\mathbb{X}$.

4 MCMC Sampling for DPM of Conditional Vine Copulas

We propose an MCMC sampler by using a Pólya-urn scheme for integrating out of the model the random distribution function from the Dirichlet process.

In DPM models there are two levels of conjugacy: between the DP random measure $\Pi(G)$ and its posterior $\Pi(G|\mathbf{y})$, and between the kernel element $\mathbf{f}_{\boldsymbol{\xi}}c_{\boldsymbol{\xi}}$ and the centering measure G_0, which plays the role of prior on $\boldsymbol{\xi}$. When considering non-conjugate DP mixtures, we refer to the case where $c_{\boldsymbol{\xi}}$ and G are not conjugate; a sampler for this case is the *no gaps* sampler, proposed by MacEachern and Müller [1998], which still relies on conjugacy of the DP posterior. In our case, we choose as kernel the product of the density of a Gaussian vine-copula and the densities of the covariates which depend on the nature of $\mathbb{X}$. However, we lose the conjugacy between the kernel and the centering measure G_0 because of the presence of the covariates.

Let $(u_{1i}, \ldots, u_{di}) = (F_{1,x}(y_{1i}), \ldots, F_{d,x}(y_{di}))$, for $i = 1, \ldots, N$, be pseudo-observations defined in the hypercube $\mathbf{I}^d$. Let $\mathbb{X}$ be the $N \times p$ matrix of covariates. The Gaussian bivariate copula is governed by the correlation parameter $\rho \in (-1, 1)$. Following the conditional copula literature, the correlation parameter is associated to the covariates through a link function g, such that

$$\rho(\mathbf{x}|\boldsymbol{\beta}) = g^{-1}(\eta(\mathbf{x}|\boldsymbol{\beta}))$$

where g^{-1} is the inverse link function (that we assume to be the Fisher's transform) and $\eta(\cdot)$ is a calibration function, which, following Dalla Valle et al. [2018], will need to be specified. An alternative approach is proposed by Wehrhahn et al. [2020], who consider a predictor-dependent stick-breaking prior distribution for the collection of predictor-dependent random measure, avoiding the specification of the calibration function. From now on we will

index the conditional copula functions with $\boldsymbol{\beta}$, i.e. $c_{\boldsymbol{\beta}}(F(y_i), F(y_j)|\mathbf{x})$.

In our setting, we need to define a centering measure for the parameters of the $\nu = d(d-1)/2$ pair-copulas forming the vine and the parameters of the covariate densities. The centering measure dimension will depend on the number of covariates and on the probability models assumed for the covariates. Hence, we need to define a $(\nu \times q + r)$-dimensional centering measure, where q represents the number of unknown parameters of the calibration function for each pair copula and r denotes the total number of unknown parameters of the covariate densities. Considering the general D-vine construction, we assume, as kernel density of the mixture, the product of the bivariate densities

$$
\mathbf{f}_{\boldsymbol{\phi}}(\mathbf{x}) \prod_{\ell=1}^{d-1} \prod_{k=1}^{d-\ell} c_{\boldsymbol{\beta}} \left(F_{k|k+1,\ldots,k+\ell-1,x}(y_k|y_{k+1,\ldots,k+\ell-1}), \right.
$$

$$
\left. F_{\ell+k|k+1,\ldots,k+\ell-1,x}(y_{\ell+k}|y_{k+1,\ldots,k+\ell-1})|\mathbf{x} \right).
$$

In this case, the model parameters are included in the $(\nu \times q + r)$-dimensional random vector $\boldsymbol{\xi} = (\boldsymbol{\beta}, \boldsymbol{\phi})$. Since the discreteness of the DP implies a positive probability for ties among the latent elements of the vector $\boldsymbol{\xi} = (\boldsymbol{\beta}, \boldsymbol{\phi})$, the DPM induces a probability model on clusters. The detailed illustration of the MCMC sampler for DPM of conditional vine copulas is included in the Supplementary Material.

In the next Sections we apply the proposed model to both simulated and real data.

16

5 Simulation Study

In order to assess the validity of the proposed model from different perspectives, we ran an extensive simulation study. We considered various scenarios, generating data from several multivariate frameworks, and we showed that the model is reliable for both clustering and density estimation. With simulations in high dimensions, we compared the performances of our sampler, by proving the stability of the sampler and checking for the computational times. The procedure and the results are described in detail in Section S2 of the Supplementary Material. Finally, we compared the proposed method with an alternative Bayesian nonparametric model. In particular, Lu [2021] proposed the empirical checkerboard Bernstein copula (ECBC), a flexible method for the estimation of copula functions in a multivariate framework. However, conditional copulas are only presented in the bivariate case. Therefore, we compared the performances of the ECBC with the DPM of *unconditional* vine copulas with a further simulation study, by generating data from Hierarchical Archimedean Copulae (HAC) (for more details see Okhrin and Ristig [2014]). Following Lu [2021], in the ECBC we have assumed hierarchical shifted Poisson as emprical prior distributions on the polynomials degrees, while in the DPM of vine copulas, we have defined total mass parameter $M = 1$ and flat Normal distribution centred on zero on the Fisher transform of each pair-copula parameter. We considered four complex scenarios: (i) a three-dimensional structure generated from a Hierarchical Frank Copula with parameters $\zeta = (1.5, 0.75)$, (ii) a three-dimensional Hierarchical

17

Gumbel Copula with parameters $\zeta = (1.6, 2.1)$ (iii) a three-dimensional t distribution and (iv) a three-dimensional Hierarchical Ali-Mikhail-Haq Copula with parameters $\zeta = (0.6, 0.9)$. In (i) and (ii) we generated $n = 500$ observations. The posterior predictive density estimates are presented in Figure S7 of the Supplementary materials. In (iii) and (iv) we considered two samples of dimension $n = 50$, in order to compare the performances of the models with few observations. The results are presented in Figure 2.

6   Applications and Discussion

In this section we present two applications to real datasets. The first one, a veterinary application, demonstrates the potential of the proposed approach in presence of missing data for the covariates, and its accuracy for density estimation compared to another approaches. The second application, on the other hand, focuses on economic data and mainly demonstrates the clustering potential of the proposed approach. Finally, we present a short discussion.

6.1   Veterinary Dataset

We present a first application to a Veterinary Medicine (VM) dataset, used previously by Vallée et al. [2018]. Since most New Zealand sheep flocks are not vaccinated, we investigated the impact of Leptospira serovars Hardjo and/or Pomona on the weight of the sheep. The dataset [1] comprises the weight at three different time points of $N = 983$ sheep, together with three binary

---

[1] Available at `https://data.mendeley.com`.

Figure 2: Simulated samples with Normal margins (left panels) vs posterior predictive densities of DPM of vines (middle panels, continuous contour line) and ECBC (right panels, dashed contour line). In the top panels (rows 1 and 2) the sample of $n = 50$ observations is generated from a multivariate t distribution. In the bottom panels (rows 3 and 4) the sample of $n = 50$ observations is generated from a Hierarchical Ali-Mikhail-Haq Copula.

covariates, as illustrated in Table S3 of the Supplementary Material. We defined a conditional D-Vine structure to account for the effect of the covariates on the estimation of the dependencies between the weight values over time. In particular, after transforming the variables into pseudo-observations, we defined the link function $\eta(x|\boldsymbol{\beta}) = \beta_0 + \beta_V x_V + \beta_H x_H + \beta_P x_P$, where the subscript V denotes vaccine, H Hardjo and P Pomona. Despite the presence of several missing values for the covariates, our model structure allows for easily overcoming this issue: since the model requires distributional assumptions for the covariates, missing data imputation is immediately dealt with by sampling from the posterior predictive distribution of the covariates $f(\tilde{x}|x, \Psi) = \int f(\tilde{x}|x, \phi, \Psi)\pi(\phi|x, \Psi)d\phi$. We set the total mass parameter as $M = 1$ and define the centering measure as $G_0 \equiv \text{Beta}_3(a_\phi, b_\phi) \times \mathbb{N}_{(3\times3)}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with $a = b = 1$. We drew 10000 observations of the posterior distributions of the model parameters with a burnin of 2000 iterations, estimating $\Psi = 3$ clusters. Results are reported in Table S4 of the Supplementary Material.

For comparison purposes, we also fitted the VM dataset with the Generalized additive model for pair-copula constructions (GAM-PCC) by Vatter and Nagler [2018], assuming that covariates values are missing-at-random. In Figure S8 of the Supplementary Material we compare the predictive distributions of the DPM of conditional vines and the GAM-PCC, fitted via the `gamCopula` package available in `R`, assuming linear covariates. Although the GAM-PCC is fast to implement from a computational perspective, there are two critical points to consider: (i) it does not allow to treat observations

with missing values for the regressors, causing a reduction of the sample to $N = 548$, and (ii), as it can be noticed from the predictive samples, it does not ensure the flexibility of a mixture. Indeed, the approach of Vatter and Nagler [2018] requires the specification of a probability density for each pair copula via Akaike information criterion (AIC); specifically, the model selected the the families t, Gumbel and Gaussian, respectively, for the pair copulas c_{12}, c_{13} and $c_{13|2}$.

6.2 Financial Development and Natural Disasters Data

The second application is to a heterogeneous dataset to study the impacts of worldwide natural disasters on international financial development. In recent years several authors in the literature have focused on the relationship between economic growth and natural disasters. Toya and Skidmore [2007] show that countries with higher income, higher educational attainment, greater openness, more advanced financial systems and smaller governments tend to suffer fewer losses in presence of natural disasters. Felbermayr and Gröschl [2014] built a comprehensive database of disaster events and their intensities from primary geophysical and meteorological information, revealing a substantial negative and robust average impact effect of disasters on growth. So far, researchers have mostly dealt with financial development studies using several proxies as measures of financial depth, such as the ratio of private credit to GDP or stock market capitalization to GDP. For example, Keerthiratne and Tol [2017] proposed an empirical analysis of the effect of the oc-

currence of natural disasters on financial development proxies. However, this approach does not take into account the complex nature of financial developments. The International Monetary Fund (IFM) provides, for each country, the Financial Development (FD) index, which summarizes how developed financial institutions and financial markets are in that specific country. The index considers several factors, such as size and liquidity of the markets, ability of individuals and companies to access financial services, ability of institutions to provide financial services at low cost and with sustainable revenues and level of activity of the capital markets. We constructed a dataset merging information from the IFM [2] and the Emergency Events Database [3] and we analyzed the impact of natural disasters on the dependence between the FD index values in several years. In particular, we investigated how the occurrence of a natural disaster affects the time dependence between the FD index values in the four years following the disaster.

The original data included observations of the FD index (expressed in percentages) in 181 countries from 1980 to 2019 with annual frequency. We are interested in the effect of a single natural disaster on the FD index in the 4 years following the event. Note that, if more that one natural disaster occurred in the considered 4 years interval, the observation was discarded. Based on the original data, we constructed a study dataset where each observation corresponds to a set of 4 consecutive year periods affected by a natural disaster in the first year. Hence, each country may be represented by more

than one observation. Our final study data comprises $N = 525$ observations (for $n = 86$ different countries) and four variables containing the FD indexes for each one of the 4 considered years. The dataset is structured as reported in Table S5 of the Supplementary Material, where t_i denotes the number of time periods for each country, with $j = 1, \ldots, t_i$ and i denotes the country, with $i = 1, \ldots, n = 86$. We applied the proposed Bayesian nonparametric unconditional vine mixture approach to the study dataset, constructing a 4-dimensional vine copula, with marginals denoting the FD index in 4 consecutive years. The adoption of an infinite mixture approach allows us to control for both individual and temporal heterogeneity in the dataset. Since the marginal variables $(y_{i,t_i}, y_{i,t_i+1}, y_{i,t_i+2}, y_{i,t_i+3})$, with $i = 1, \ldots, n$, are expressed in percentages, they are defined on the support $[0, 1]$. Hence, we assume the marginal distributions to be independent $\text{Beta}(a_j, b_j)$ with a_j and b_j defined a priori $\text{Gamma}(1, 1)$ for $j = 1, \ldots, 4$. Inference for the margins is separately performed via Metropolis-Hastings. Results are reported in Table S6 of the Supplementary Material. We consider the occurrence of a natural disaster to be a covariate in our conditional vine copula model. Since natural disasters are defined basing on their intensity, which is measured using the total damage as proxy variable, we constructed the binary covariate taking value $X = 1$ if the total damage is over 100 million dollars and $X = 0$ otherwise. Since the model requires a distributional assumption on the covariate, we defined $X \sim \text{Bernoulli}(\phi)$ assuming a priori $\phi \sim \text{Beta}(a_\phi, b_\phi)$. Moreover, we chose a linear calibration function for each pair copula $\eta_s(x|\boldsymbol{\beta}) = \beta_{0_s} + \beta_{1_s}x$,

where $s = 1, \ldots, 6$ denotes the pair-copulas in the vine. For the definition of kernel of the mixture we adopt a D-vine approach, since this specific type of vine better describes temporal sequences and easily takes the time ordering of consecutive events into account [Barthel et al., 2019]. We therefore consider a 4-dimensional D-vine, where the marginals are the FD indexes of the 4 consecutive years and with covariate given by the natural disaster intensity. We set the total mass parameter as $M = 1$ and define the centering measure as $G_0 \equiv \text{Beta}\,(a_\phi, b_\phi) \times \mathbb{N}_{(6 \times 2)}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and we run the MCMC sampler for 10000 iterations with a burnin of 2000.

The results of the application of our model to the data are graphically represented in Figures 3 and 4. Our model estimates two clusters, as shown in Figure 3. In Table S7 of the Supplementary Material we report the resulting mean, standard deviation and credibility intervals for the posterior densities of the calibration functions parameters. The top part of the table shows the results for the first cluster $\psi = 1$, while the bottom part shows the results of the second cluster $\psi = 2$. Figure 4 reveals that the two mixture components present substantial differences in terms of how they are impacted by natural disasters. For the first cluster ($\psi = 1$) the model estimates a general negative effect which tends to remain constant until the fourth year; instead, for the second cluster ($\psi = 2$) the model estimates a positive effect of the natural disaster on the time dependence between yearly FD indexes.

Let w_1 and w_2 be the mixture weights for the two estimated components, i.e. the variables indicating the proportion of observations belonging to each
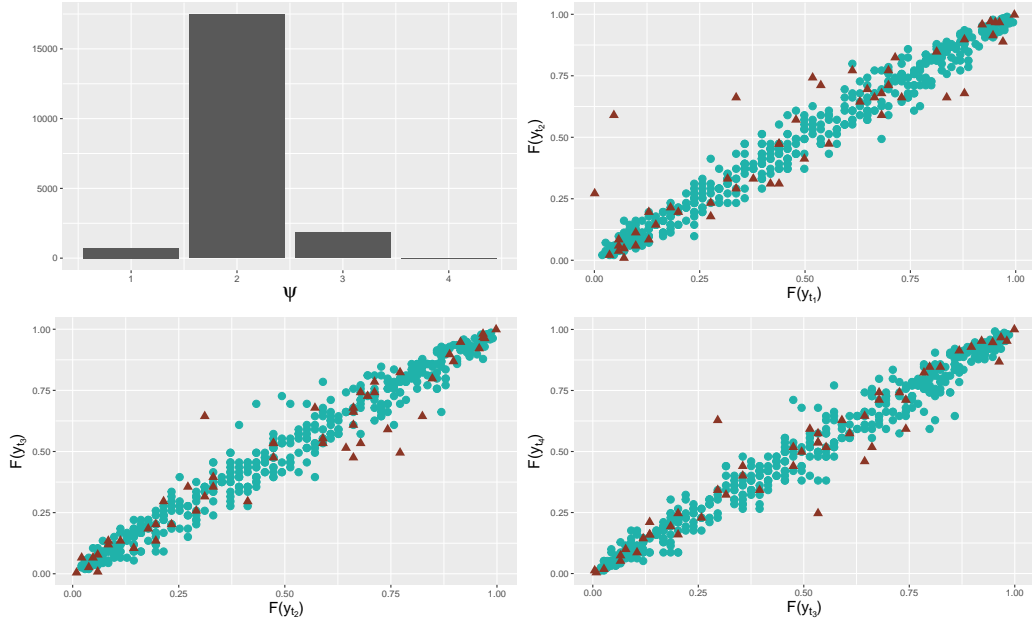
Figure 3: Results from the financial development and natural disasters analysis. The top left panel shows the barplot of the mode of the number of observed mixture components. The top right and the bottom panels show the scatterplots of the observed u-data for the first vine tree; the dots denote the observations belonging to the first cluster $\psi = 1$, while the triangles denote the observations belonging to the second cluster $\psi = 2$.
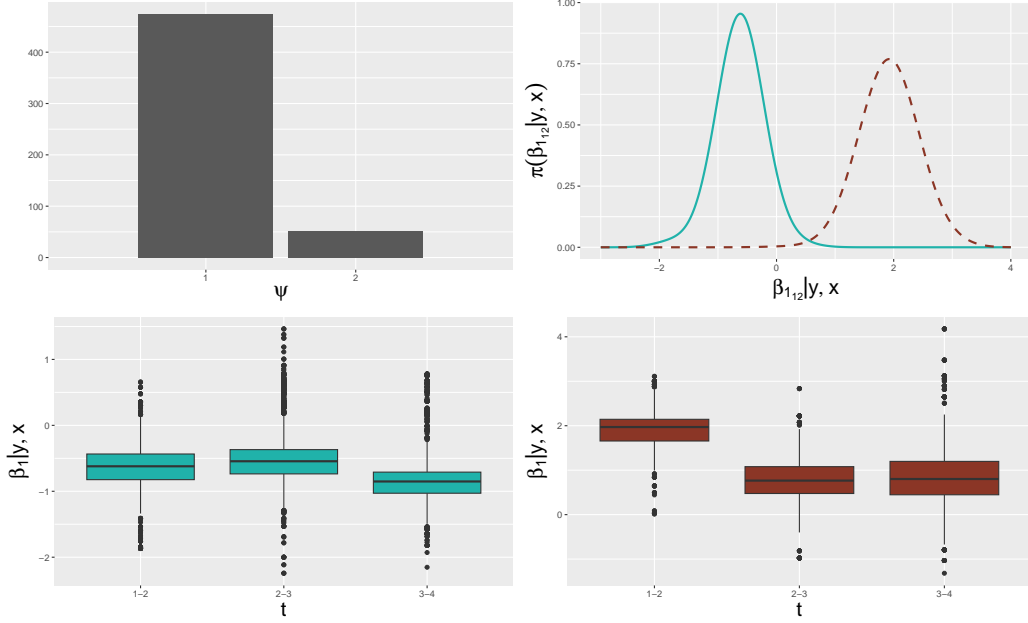
Figure 4: Results from the financial development and natural disasters analysis. The top left panel shows the barplot of the number of observations allocated to the two estimated mixture components. The top right panel compares the posterior densities of the calibration function parameter $\beta_{1_{12}}$ (which is related to the first time interval from t_i to $t_i + 1$, with $i = 1, \ldots, n$) for the first (solid line) and the second (dashed line) mixture components. The left and right bottom panels show, for the first and second mixture components, the boxplots of the calibration function parameters $\beta_{1_{12}}$ (left; first time interval from t_i to $t_i + 1$), $\beta_{1_{23}}$ (middle; second time interval from $t_i + 1$ to $t_i + 2$) and $\beta_{1_{34}}$ (right; third time interval from $t_i + 2$ to $t_i + 3$) for the first vine tree.

group. Table S8 shows, on the left, the mean, standard deviation and credibility intervals for the posterior densities of the covariate parameter for the two estimated mixture components ($\psi = 1$ in the left column, $\psi = 2$ in the right column). On the right, Table S8 shows the same statistics for the mixture weights. As reported in Table S8, and in the barplot of Figure 4, the 85% of the observations belong to the first cluster, meaning that in most of the cases natural disasters are expected to have a negative effect on the financial development. However, from Table S8 we can get another interesting point: the expected value of the posterior probability of natural disasters $E(\phi|x)$ is very low in the first and very high in the second cluster. This result may suggest a further insight: the less a natural disaster was expected, the less the government was prepared for that event. On the contrary, if the occurrence of a natural disaster was expected, governments and financial institutions were prepared for that event and were able to face the possible effects on the financial structure of the country. In Table S9 of the Supplementary Material we report the list of countries, years and type of natural disasters which had an estimated positive effect on the FD index.

6.3 Discussion

The results of the applications show that the proposed approach allows us to model individual as well as temporal heterogeneity in a natural way. Our method is based on the use of a single DP for all the copulas in the vines. An alternative approach would be to model each pair copula as a DPM,

27

by assuming a vector of independent random probability measures on the vine copulas mixing measures. This approach is introduced and discussed by Griffin and Leisen [2017] and Griffin and Leisen [2018], who use normalised versions as priors in Bayesian non-parametric mixture models. Also, Camerlenghi et al. [2019] propose a class of latent nested processes, which preserve heterogeneity a posteriori, even when distinct values are shared by different samples. This approach would ensure more flexibility, with benefits to density estimation and heterogeneity control. Vectors of independent random probability measures would allow us to extend our methodology and will be object of future work. However, there are two issues to consider: first, the computational cost would further increase, making it extremely challenging to implement this approach in high dimension or with large sample sizes; secondly, if applied to clustering problems, the interpretation of the results would be harder.

Supplementary Material

Appendix: the supplemental files include an overview of the computational methods for DPM, a detailed illustration of MCMC sampling for DPM of conditional vine copulas, extensive simulation studies and the Tables and Figures of the applications and discussion Section.

R code: the supplemental files include a folder with the R scripts implementing the methodology described in the paper and a readme file.

Acknowledgments

Declarations of interest

The authors report there are no competing interests to declare.

References

Kjersti Aas, Claudia Czado, Arnoldo Frigessi, and Henrik Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and economics*, 44(2):182–198, 2009.

Fentaw Abegaz, Irène Gijbels, and Noël Veraverbeke. Semiparametric estimation of conditional copulas. *Journal of Multivariate Analysis*, 110:43–73, 2012.

Elif F Acar, Radu V Craiu, and Fang Yao. Dependence calibration in con-

ditional copulas: A nonparametric approach. *Biometrics*, 67(2):445–453, 2011.

Nicole Barthel, Candida Geerdens, Claudia Czado, and Paul Janssen. Dependence modeling for recurrent event times subject to right-censoring with d-vine copulas. *Biometrics*, 75(2):439–451, 2019.

Söhnke M Bartram, Stephen J Taylor, and Yaw-Huei Wang. The euro and european financial market dependence. *Journal of Banking & Finance*, 31 (5):1461–1481, 2007.

Tim Bedford and Roger M Cooke. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial intelligence*, 32(1-4):245–268, 2001.

Prasad Buddhavarapu, James G Scott, and Jorge A Prozzi. Modeling unobserved heterogeneity using finite mixture random parameters for spatially correlated discrete count data. *Transportation Research Part B: Methodological*, 91:492–510, 2016.

Federico Camerlenghi, David B Dunson, Antonio Lijoi, Igor Prünster, and Abel Rodríguez. Latent nested nonparametric priors (with discussion). *Bayesian Analysis*, 14(4):1303–1356, 2019.

V Radu Craiu and Avideh Sabeti. In mixed company: Bayesian inference for bivariate conditional copula models with discrete and continuous outcomes. *Journal of Multivariate Analysis*, 110:106–120, 2012.

Claudia Czado. *Analyzing Dependent Data with Vine Copulas*. Springer, 2019.

Luciana Dalla Valle, Fabrizio Leisen, and Luca Rossini. Bayesian nonparametric conditional copula estimation of twin data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 67(3):523–548, 2018.

Yanqin Fan and Andrew J Patton. Copulas in econometrics. *Annu. Rev. Econ.*, 6(1):179–200, 2014.

Gabriel Felbermayr and Jasmin Gröschl. Naturally negative: The growth effects of natural disasters. *Journal of development economics*, 111:92–106, 2014.

Christian Genest, Michel Gendron, and Michaël Bourdeau-Brien. The advent of copulas in finance. *The European Journal of Finance*, 15(7-8):609–618, 2009.

Irène Gijbels, Noël Veraverbeke, and Marel Omelka. Conditional copulas, association measures and their applications. *Computational Statistics & Data Analysis*, 55(5):1919–1932, 2011.

Irene Gijbels, Marek Omelka, Noël Veraverbeke, et al. Multivariate and functional covariates and conditional copulas. *Electronic Journal of Statistics*, 6:1273–1306, 2012.

Clara Grazian and Brunero Liseo. Approximate bayesian inference in semiparametric copula models. *Bayesian Analysis*, 12(4):991–1016, 2017.

Clara Grazian, Luciana Dalla Valle, and Brunero Liseo. Approximate Bayesian conditional copulas. *Computational Statistics and Data Analysis*, 169:107417, 2022.

Peter J Green and Sylvia Richardson. Modelling heterogeneity with and without the dirichlet process. *Scandinavian journal of statistics*, 28(2): 355–375, 2001.

Jim Griffin and Fabrizio Leisen. Modelling and computation using ncorm mixtures for density regression. *Bayesian Analysis*, 13(3):897–916, 2018.

Jim E Griffin and Fabrizio Leisen. Compound random measures and their use in bayesian non-parametrics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(2):525–545, 2017.

Eric Jondeau and Michael Rockinger. The copula-garch model of conditional dependencies: An international stock market application. *Journal of international money and finance*, 25(5):827–853, 2006.

Maria Kalli, Jim E Griffin, and Stephen G Walker. Slice sampling mixture models. *Statistics and computing*, 21(1):93–105, 2011.

Subhani Keerthiratne and Richard SJ Tol. Impact of natural disasters on financial development. *Economics of Disasters and Climate Change*, 1(1): 33–54, 2017.

Nikolai Kolev, Ulisses dos Anjos, and Beatriz Vaz de M Mendes. Copulas: a review and recent developments. *Stochastic models*, 22(4):617–660, 2006.

Evgeny Levi and Radu V Craiu. Bayesian inference for conditional copulas using gaussian process single index models. *Computational Statistics & Data Analysis*, 122:115–134, 2018.

Lu Lu. *Multivariate Dependence Using Nonparametric Copula Models*. North Carolina State University, 2021.

Steven N MacEachern and Peter Müller. Estimating mixture of dirichlet process models. *Journal of Computational and Graphical Statistics*, 7(2): 223–238, 1998.

Peter Müller and Fernando Quintana. Random partition models with regression on covariates. *Journal of statistical planning and inference*, 140(10): 2801–2808, 2010.

Peter Müller and Gary L Rosner. A bayesian population model with hierarchical mixture priors applied to blood count data. *Journal of the American Statistical Association*, 92(440):1279–1292, 1997.

Shaoyang Ning and Neil Shephard. A nonparametric bayesian approach to copula estimation. *Journal of Statistical Computation and Simulation*, 88 (6):1081–1105, 2018.

Ostap Okhrin and Alexander Ristig. Hierarchical archimedean copulae: the hac package. *Journal of Statistical Software*, 58:1–20, 2014.

Andrew J Patton. Modelling asymmetric exchange rate dependence. *International economic review*, 47(2):527–556, 2006.

Özge Sahin and Claudia Czado. Vine copula mixture models and clustering for non-gaussian data. *Econometrics and Statistics*, 22:136–158, 2022.

A Sklar. Fonctions dé repartition à n dimension et leurs marges. *Université Paris*, 8(3.2):1–3, 1959.

Hideki Toya and Mark Skidmore. Economic development and the impacts of natural disasters. *Economics letters*, 94(1):20–25, 2007.

Daniel Turek, Claudia Wehrhahn, and Olivier Gimenez. Bayesian non-parametric detection heterogeneity in ecological models. *Environmental and Ecological Statistics*, 28(2):355–381, 2021.

Emilie Vallée, Cord Heuer, Julie M. Collins-Emerson, Jackie Benschop, Anne L. Ridler, and Peter R. Wilson. Effects of natural infection by l. borgpetersenii serovar hardjo type hardjo-bovis and l. interrogans serovar pomona, and leptospiral vaccination, on sheep growth. *Preventive Veterinary Medicine*, 159:196–202, 2018. ISSN 0167-5877. doi: https://doi.org/10.1016/j.prevetmed.2018.09.017. URL `https://www.sciencedirect.com/science/article/pii/S0167587718302113`.

Thibault Vatter and Thomas Nagler. Generalized additive models for pair-copula constructions. *Journal of Computational and Graphical Statistics*, 27(4):715–727, 2018.

Stephen G Walker. Sampling the dirichlet mixture model with slices. *Com-*

munications in Statistics—Simulation and Computation®, 36(1):45–54, 2007.

Claudia Wehrhahn, Ruth Fuentes-García, Ramsés H Mena, Fabrizio Leisen, Maria Elena González-Villalpando, and Clicerio González-Villalpando. A copula-based fully bayesian nonparametric evaluation of cardiovascular risk markers in the mexico city diabetes study. *arXiv preprint arXiv:2007.11700*, 2020.

Juan Wu, Xue Wang, and Stephen G Walker. Bayesian nonparametric inference for a multivariate copula function. *Methodology and Computing in Applied Probability*, 16(3):747–763, 2014.

Juan Wu, Xue Wang, and Stephen G Walker. Bayesian nonparametric estimation of a copula. *Journal of Statistical Computation and Simulation*, 85(1):103–116, 2015.

Susu Zhang and Shiyu Wang. Modeling learner heterogeneity: A mixture learning model with responses and response times. *Frontiers in psychology*, 9:2339, 2018.

Haoxin Zhuang, Liqun Diao, and Y Yi Grace. A bayesian nonparametric mixture model for grouping dependence structures and selecting copula functions. *Econometrics and Statistics*, 2021.