CHaRMS

CENTRE FOR HEALTH RESEARCH AT THE MANAGEMENT SCHOOL

WORKING PAPER SERIES

MIXED BINARY-CONTINUOUS COPULA REGRESSION MODELS WITH APPLICATION TO ADVERSE BIRTH OUTCOMES

Nadja Klein (Humboldt University of Berlin)
Thomas Kneib (Georg-August-University Göttingen)
Giampiero Marra (University College London)
Rosalba Radice (City, University of London)
Slawa Rokicki (University College Dublin)
Mark McGovern (Queen's University Belfast)

Mixed Binary-Continuous Copula Regression Models with Application to Adverse Birth Outcomes

Nadja Klein[1], Thomas Kneib[2], Giampiero Marra[3], Rosalba Radice[4], Slawa Rokicki[5], and Mark E. McGovern[6,7]

[1]Department of Economics, Applied Statistics, Humboldt University of Berlin
[2]Chair of Statistics, Georg-August-University Göttingen
[3]Department of Statistical Science, University College London
[4]Cass Business School, City University of London
[5]Geary Institute for Public Policy, University College Dublin
[6]Centre for Health Research at the Management School, Queen's University Belfast
[7]Centre of Excellence for Public Health (Northern Ireland)

October 2018

Abstract

Bivariate copula regression allows for the flexible combination of two arbitrary, continuous marginal distributions with regression effects being placed on potentially all parameters of the resulting bivariate joint response distribution. Motivated by the risk factors for adverse birth outcomes, many of which are dichotomous, we consider mixed binary-continuous responses that extend the bivariate continuous framework to the situation where one response variable is discrete (more precisely binary) while the other response remains continuous. Utilizing the latent continuous representation of binary regression models, we implement a penalized likelihood based approach for the resulting class of copula regression models and employ it in the context of modelling gestational age and the presence/absence of low birth weight. The analysis demonstrates the advantage of the flexible specification of regression impacts including nonlinear effects of continuous covariates and spatial effects. Our results imply that racial and spatial inequalities in the risk factors for infant mortality are even greater than previously suggested.

1

1 Introduction

Birth weight and gestational age are important determinants of infant and child health. Recent evidence has also shown that these factors affect long-term health throughout adulthood (Oreopoulos et al. 2008; Hack et al. 2002). In the United States about 75% of perinatal deaths occur in infants born prematurely (less than 37 weeks gestation), and over two thirds of premature deaths are a result of the 30–40% of preterm infants who are delivered before 32 weeks' gestation (Slattery and Morrison 2002). Among surviving premature and low birth weight (LBW, defined as less than 2500g) infants, rates of long-term physical and mental disability and chronic health problems are high (Slattery and Morrison 2002). Moreover, LBW is associated with poor educational and labor force outcomes in adolescence and adulthood including lower scores on academic achievement tests, lower rates of high school completion, lower income, and higher social assistance take-up (Almond and Currie 2011; Behrman and Rosenzweig 2004; Black et al. 2007; Hack et al. 2002; McGovern 2013; Oreopoulos et al. 2008).

Although both LBW and gestational age are independent predictors of future health, modelling these outcomes jointly is essential for a number of reasons. First, birth weight and gestational age are highly correlated, and confounded by factors such as intrauterine growth restriction (Slattery and Morrison 2002). Second, risk factors for LBW such as socio-economic status, smoking, and maternal age may also be the same risk factors for preterm birth. Finally, evidence suggests that the impact of LBW on health may be modified by early gestational age, and vice-versa (Hediger et al. 2002). Thus, modelling these outcomes independently would present a confounded assessment of who is most vulnerable to poor infant health and how best to intervene. A more accurate picture is revealed by modelling these outcomes jointly.

We use data from the North Carolina State Center for Health Statistics to evaluate the probability of mortality before age 1 by birth weight and weeks of gestation. The data include all births in the state from 2007 to 2013 and provide information on maternal characteristics, delivery characteristics, and infant birth and death outcomes at the county level. Figure 1 (top) shows the probability of infant mortality for premature (but not LBW) infants, LBW (but not premature) infants, premature and LBW infants, and infants without these conditions ("normal"). The probability of death is more than two times higher for infants that are both LBW and premature, compared to those that are only LBW. In addition, the risk associated with LBW and prematurity varies depending on maternal characteristics. For instance, Figure 1 (bottom left) shows that probability of death for infants born both premature and LBW infants is significantly lower for infants born to white mothers than those born to black or other race mothers, similar to the rate of infant mortality for infants that are LBW (but not premature) even though being much lower in absolute terms.

To adequately model these data, we consider bivariate copula regression models in the same vein as Marra and Radice (2017b), Radice et al. (2016) and Klein and Kneib (2016). Generally, these papers utilize copulas to construct flexible bivariate response distributions where both margins are either binary or continuous (see e.g. de Leon and K. 2013, for a general overview of mixed data analysis). Following an approach similar to Marra

Figure 1: Probability of infant mortality by infant birth category (top), by infant birth category and mother's race (bottom left) and by weeks of gestation, stratified by low birth weight status (bottom right). Error bars show the 95% confidence interval. Source: North Carolina Vital Statistics (combined birth and death records), 2007-2013.

and Radice (2017b), in this analysis we implement a simultaneous penalized likelihood method which employs copulae to combine a binary response variable (low birth weight in our case) and a continuous outcome (gestational age) while accounting for several types of non-Gaussian dependencies. To facilitate the methodological developments for the binary part of the model, we use the latent response representation of binary regression models. To go beyond simplistic mean regression settings, the marginal and copula parameters are related to regression predictors through structured additive form.

To the best of our knowledge, other existing bivariate copula regression approaches and software implementations (see, e.g., Acar et al. 2013; Gijbels et al. 2011; Kramer et al. 2012; Kraemer and Silvestrini 2015; Sabeti et al. 2014; Yan 2007) cover only parts of the flexibility of our approach. An exception is a recent paper by Vatter and Chavez-Demoulin (2015), which is based on a two-stage technique where the parameters of the marginal distributions and of the copula function are estimated separately. As shown in the simulation by Marra and Radice (2017b), estimating all the model's parameters simultaneously offers computational and efficiency gains, hence the simultaneous estimation approach adopted here. Simultaneous copula regression in a fully parametric approach to likelihood estimation is considered in de Leon and Wu (2011) and Wu and de Leon (2014) but restricted to the case of linear dependence via a Gaussian copula.

Note that the methodology developed here is most useful when the main interest is in relating the parameters of a bivariate copula distribution to covariate effects. Otherwise, semi/non-parametric extensions where, for instance, the margins and/or copula function are estimated using kernels, wavelets or orthogonal polynomials may be considered instead (e.g., Kauermann et al. 2013; Lambert 2007; Segers et al. 2014; Shen et al. 2008). While such techniques are in principle more flexible in determining the shape of the underlying bivariate distribution, in practice they are limited with regard to the inclusion of flexible covariate effects, and may require large sample sizes to produce reliable results.

Our proposed copula modelling approach builds on previous work on adverse birth outcomes. For example, our approach allows us to model gestational age continuously, rather than using a binary cut-off (Neelon et al. 2014). The standard cut-off of 37 weeks used in previous analyses may not necessarily be appropriate. For example, in Figure 1 there is no clear change in the probability of death at 37 weeks. If anything, a more useful cut-off may be at 31 weeks, as it appears that the greatest mortality risk is concentrated among infants in this category (those born before 32 weeks gestation). By modelling gestational age continuously, we are able to investigate mortality risk across the whole distribution, which may be more useful for understanding the impact of gestational age on health. The approach we propose here will be especially relevant if the goal of the analysis is to learn about group level differences, for example racial inequalities. If the mortality risk associated with prematurity varies by maternal characteristics, as suggested in Figure 1, imposing a single cut-off could bias estimates of racial inequality in early life health. Furthermore, we allow for (i) non-Gaussian dependence structures between LBW and gestational age, (ii) the copula dependence and marginal distribution parameters to be estimated simultaneously, and (iii) each parameter to be modelled using an additive predictor incorporating several types of covariate effects (e.g., linear, non-linear, random and spatial effects).

In our analyses, birth weight is modelled as a binary outcome. While the argument we have just made for gestational age could potentially be applied to birth weight, the use of the low birth weight cut-off is often specifically required for epidemiological research. For example, the World Health Organisation produces global estimates of prevalence based on the 2,500g low birth weight threshold (Wardlaw et al. 2005). Moreover, we prefer to focus on the binary-continuous case because there are many contexts in the health domain where the focus is on clinical diagnosis thresholds which are dichotomous, and there are few existing methods which allow for flexible modelling of the dependence of these outcomes alongside other continuous measures of interest. Nevertheless, we have also conducted robustness checks using the threshold for very low birth weight ($<$1,500g), and found similar results.

In summary, our paper contributes to the literature on copula regression by

- analysing a complex, high-dimensional, and challenging data set on adverse birth outcomes where we combine a binary regression model for the presence/absence of low birth weight with a flexible, continuous specification for gestational age,

- providing a generic framework for bivariate response models with mixed binary-continuous outcomes and a copula dependence structure where all parameters are

estimated simultaneously and can potentially be related to flexible functions of explanatory variables, and

- incorporating the proposed developments into the freely distributed and easy to use R package GJRM (Marra and Radice 2017a).

The rest of the paper is organized as follows. In Section 2, we introduce bivariate copula models with mixed binary-continuous marginals and flexible covariate effects. Section 3 provides technical details on the penalized maximum likelihood inferential framework employed, while in Section 4 we discuss the results of the empirical analysis of adverse birth outcomes. Section 5 summarizes the main findings. The appendix lists further empirical results and investigates the empirical performance of our joint model against a model with independent marginal specifications.

2 Bivariate Copula Models with Mixed Binary-Continuous Marginals

2.1 Building Bivariate Distributions with Copulas

Bivariate copula regression models aim at modelling the joint distribution of a pair of response variables (Y_1, Y_2) given covariates based on a copula specification for the dependence structure between the two responses. We therefore make use of the copula-based representation of the bivariate cumulative distribution function (CDF) $F_{1,2}(y_1, y_2) = P(Y_1 \leq y_1, Y_2 \leq y_2)$ as

$$F_{1,2}(y_1, y_2) = C(F_1(y_1), F_2(y_2)) \tag{1}$$

where $C : [0,1]^2 \rightarrow [0,1]$ denotes the copula (i.e. a bivariate CDF defined on the unit square with standard uniform marginals) and $F_j(y_j) = P(Y_j \leq y_j)$, $j = 1, 2$ are the marginal CDFs of the two response components Y_1 and Y_2. If both Y_1 and Y_2 are continuous, the copula $C(\cdot, \cdot)$ in (1) is uniquely determined. In copula regression, we use the representation (1) as a construction principle for defining bivariate regression models where the distribution for the response vector is determined by choosing a specific parametric copula and two (continuous) marginals. In this way, copulas provide a flexible and versatile way of constructing bivariate distributions with various forms of dependencies induced by the copula.

2.2 Mixed Binary-Continuous Copulas

In this paper, we consider the case where one of the two responses is not continuous but rather binary, such that the immediate application of the copula regression specification is not possible (more specifically, the copula is then no longer uniquely defined). To circumvent this difficulty, we make use of the latent variable representation of binary regression models. Without loss of generality, we assume that the first response variable Y_1 is binary (i.e. $Y_1 \in \{0, 1\}$) but can be related to the (unobserved) latent variable Y_1^*

via the threshold mechanism $Y_1 = 1(Y_1^* > 0)$ where $1(\cdot)$ denotes the indicator function. Note that this implies

$$P(Y_1 = 0) = P(Y_1 \le 0) = F_1(0) = F_1^*(0) = P(Y_1^* \le 0)$$

i.e. the CDF of the observed response Y_1 ($F_1(y_1)$) and the CDF of the latent variable Y_1^* ($F_1^*(y_1^*)$) coincide at $y_1 = y_1^* = 0$.

Plugging the latent variable into the copula regression specification, we obtain

$$P(Y_1 = 0, Y_2 \le y_2) = P(Y_1^* \le 0, Y_2 \le y_2) = C(F_1^*(0), F_2(y_2))$$

and

$$P(Y_1 = 1, Y_2 \le y_2) = P(Y_1^* > 0, Y_2 \le y_2) = F_2(y_2) - C(F_1^*(0), F_2(y_2)).$$

From these expressions, we can also derive the mixed binary-continuous density

$$f_{1,2}(y_1, y_2) = \left(\frac{\partial C(F_1^*(0), F_2(y_2))}{\partial F_2(y_2)} \right)^{1-y_1} \cdot \left(1 - \frac{\partial C(F_1^*(0), F_2(y_2))}{\partial F_2(y_2)} \right)^{y_1} \cdot f_2(y_2), \quad (2)$$

where $f_2(y_2) = \frac{\partial F_2(y_2)}{\partial y_2}$ is the marginal density of Y_2. Equation (2) will provide the basis for calculating the likelihood of our copula regression specification.

2.3 Specifications for the Marginal Distributions

Various options for the specification of the marginal distribution of the latent response Y_1^* have been considered; the logistic distribution leading to a marginal logit model for the binary response Y_1, the standard normal distribution leading to a probit model, and the Gumbel distribution leading to a complementary log-log (cloglog) model. We eventually chose a probit specification (see Table 1 below) although using logit and cloglog links did not lead to different conclusions. In this case, $F_1^*(y_1^*) = \Phi(y_1^* - \eta_1)$, with the standard normal CDF $\Phi(\cdot)$ and a regression predictor η_1 specified for the success probability.

For the continuous marginal, any strictly continuous CDF $F_2(y_2)$ can be employed. In this work we have considered the normal, log-normal, Gumbel, reverse Gumbel, logistic, Weibull, inverse Gaussian, gamma, Dagum, Singh-Maddala, beta, and Fisk distributions parametrised according to Rigby and Stasinopoulos (2005). Using information criteria such as the Akaike information criterion (AIC) and Bayesian information criterion (BIC), we found the Gumbel and Dagum to be the best fitting distributions (see Table 2). For the sake of simplicity, we adopted a Gumbel specification for Y_2 with CDF

$$F_2(y_2) = \exp\left(-\exp\left(-\frac{y_2 - \mu}{\sigma} \right) \right)$$

and density

$$f_2(y_2) = \frac{1}{\sigma} \exp\left(-\frac{y_2 - \mu}{\sigma} - \exp\left(-\frac{y_2 - \mu}{\sigma} \right) \right),$$

where $\mu \in (-\infty, \infty)$ and $\sigma > 0$ denote the location and the scale parameter of the Gumbel distribution. While μ corresponds to the mode of the Gumbel distribution, its expectation and variance are given by $\mu + \gamma\sigma$ and $\sigma^2\pi^2/6$, respectively, where $\gamma \approx 0.5772$ is the Euler-Mascheroni constant.

6

2.4 Copula Specifications

Our framework allows for several copulae (see the documentation in Marra and Radice (2017a) for the choices available). The copula most supported by the AIC and BIC was the Clayton rotated by 90 degrees (see Table 3 in Section 4). The Clayton copula is defined as

$$C(u_1, u_2) = (u_1^{-\theta} + u_2^{-\theta} - 1)^{-1/\theta}$$

with dependence parameter $\theta > 0$, whereas its rotated versions can be generated as

$$C_{90}(u_1, u_2) = u_2 - C(1 - u_1, u_2),$$
$$C_{180}(u_1, u_2) = u_1 + u_2 - 1 + C(1 - u_1, 1 - u_2),$$
$$C_{270}(u_1, u_2) = u_1 - C(u_1, 1 - u_2),$$

where $C(\cdot, \cdot)$ is the standard Clayton copula. The rotation allows to the shifting of the tail dependence to either of the four corners of the unit square. This results in either upper tail (rotation by 180°) or negative tail dependence (rotation by 90° to relate large values of Y_2 with small values of Y_1 and vice versa for rotation by 270°).

Note that while the classical Gaussian copula allows for positive as well as negative dependence between the marginals, it also makes the strong assumption of symmetric dependence, i.e. the strength of dependence is the same for the lower and the upper tail. In contrast, the Clayton copula and its rotated versions allow for positive and negative tail dependence. For a comprehensive introduction to the theory of copulas and their properties see, for instance, the monographs by Nelsen (2006) and Joe (1997).

2.5 Distributional Regression Framework

For statistical inference in the mixed binary-continuous copula regression model, we embed our model structure in the distributional regression framework. We therefore assume a fully parametric specification for the distribution of the bivariate response vector, where potentially all parameters of the joint distribution can be related to regression predictors formed from covariates collected in the vector $\boldsymbol{\nu}_i$ (containing, e.g., binary, categorical, continuous, and spatial variables). More precisely, we assume that for observed response vectors $\boldsymbol{y}_i = (y_{i1}, y_{i2})'$, $i = 1, \ldots, n$ (or equivalently $\boldsymbol{y}_i = (y_{i1}^*, y_{i2})'$), the conditional density $f(\boldsymbol{y}_i | \boldsymbol{\nu}_i)$ given covariates $\boldsymbol{\nu}_i$ depends on in total $K = K_1 + K_2 + K_c$ parameters $\boldsymbol{\vartheta}_i = (\vartheta_{i1}, \ldots, \vartheta_{iK})'$ comprising

- $K_1 = 1$ parameters for the binary regression model for y_{i1} (the success probability),

- K_2 parameters for the marginal of y_{i2} (i.e. $K_2 = 2$ in case of the Gumbel distribution), and

- K_c parameters for the copula $C(\cdot, \cdot)$ (in our case, we will always have $K_c = 1$).

For each of the parameters, we assume a regression specification

$$\vartheta_{ik} = h_k(\eta_{ik}), \qquad \eta_{ik} = g_k(\vartheta_{ik})$$

7

with regression predictor η_{ik}, response functions h_k mapping the real line to the parameter space and link functions $g_k = h_k^{-1}$ mapping the parameter space to the real line. The choice of the response / link function is determined by the restrictions applying to the parameter space of the corresponding parameter such that, for example, we use the probit response function for the success probability of the binary response y_{i1} and the exponential response function for non-negative parameters.

For each of the predictors η_{ik} we assume a semiparametric, additive structure (as proposed in Fahrmeir et al. 2004)

$$\eta_{ik} = \beta_0^{\vartheta_k} + \sum_{j=1}^{J_k} s_j^{\vartheta_k}(\boldsymbol{\nu}_i) \tag{3}$$

consisting of an intercept $\beta_0^{\vartheta_k}$ and an additive combination of J_k functional effects $s_j^{\vartheta_k}(\boldsymbol{\nu}_i)$ depending on (different subsets of) the covariate vector $\boldsymbol{\nu}_i$. In particular, $\boldsymbol{\nu}_i$ can be multidimensional and univariate, see the next subsection for details and specific examples relevant in our empirical study.

2.6 Predictor Components

Dropping the parameter index ϑ_k for notational simplicity, we assume that any of the functions in (3) can be written in terms of a linear combination of D_j basis functions $B_{j,d_j}(\boldsymbol{\nu}_i)$, i.e.

$$s_j(\boldsymbol{\nu}_i) = \sum_{d_j=1}^{D_j} \beta_{j,d_j} B_{j,d_j}(\boldsymbol{\nu}_i). \tag{4}$$

Equation (4) implies that the vector of function evaluations $(s_j(\boldsymbol{\nu}_1), \ldots, s_j(\boldsymbol{\nu}_n))'$ can be written as $\boldsymbol{Z}_j\boldsymbol{\beta}_j$ with regression coefficient vector $\boldsymbol{\beta}_j = (\beta_{j1}, \ldots, \beta_{j,D_j})'$ and design matrix $\boldsymbol{Z}_j$ where $\boldsymbol{Z}_j[i, d_j] = B_{j,d_j}(\boldsymbol{\nu}_i)$. This allows us to represent the predictor vector $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_n)'$ for all n observations of any distributional parameter as

$$\boldsymbol{\eta} = \beta_0 \mathbf{1}_n + \boldsymbol{Z}_1\boldsymbol{\beta}_1 + \ldots + \boldsymbol{Z}_J\boldsymbol{\beta}_J$$

where $\mathbf{1}_n$ is a vector of ones of length n. To ensure identifiability of the model, specific constraints have to be applied to the parameter vectors $\boldsymbol{\beta}_j$ and we adopt the approach described in (Wood 2006, p. 163ff).

Since the parameter vectors $\boldsymbol{\beta}_j$ are often of considerably high dimension, quadratic penalty terms $\lambda_j\boldsymbol{\beta}_j'\boldsymbol{K}_j\boldsymbol{\beta}_j$ with positive semidefinite penalty matrix $\boldsymbol{K}_j$ are typically supplemented to the likelihood of semiparametric regression models to enforce specific properties of the j^{th} function, such as smoothness. The smoothing parameter $\lambda_j \in [0, \infty)$ then controls the trade-off between fit and smoothness, and plays a crucial role in determining the shape of $\hat{s}_j(\boldsymbol{\nu})$. For instance, let us assume that the j^{th} function models the effect of a continuous variable and s_j is represented using penalized splines. A value of $\lambda_j = 0$ (i.e., no penalization is applied to $\boldsymbol{\beta}_j$ during fitting) will result in an unpenalized regression spline estimate with the likely consequence of over-fitting, while $\lambda_j \to \infty$ (i.e., the penalty has a large influence on $\boldsymbol{\beta}_j$ during fitting) will lead to a simple polynomial fit (with the degree of the polynomial depending on the construction of the penalty matrix $\boldsymbol{K}_j$).

Different model components can be obtained by making specific choices on the basis functions in (4) and the penalty matrix $\boldsymbol{K}_j$. In the following paragraphs, we discuss the examples that are relevant to our case study.

Linear and random effects   For parametric linear effects, equation (4) becomes $\boldsymbol{z}'_{ij}\boldsymbol{\beta}_j$, and the design matrix is obtained by stacking all covariate vectors $\boldsymbol{z}_{ij}$ into $\boldsymbol{Z}_j$. No penalty is typically assigned to linear effects ($\boldsymbol{K}_j = \boldsymbol{0}$). This would be the case for binary and categorical variables. However, sometimes it is desirable to penalize parametric linear effects. For instance, the coefficients of some factor variables in the model may be weakly or not identified by the data. In this case, a ridge penalty could be employed to make the model parameters estimable (here $\boldsymbol{K}_j = \boldsymbol{I}_{D_j}$ where $\boldsymbol{I}_{D_j}$ is an identity matrix). This is equivalent to the assumption that the coefficients are independent and identically distributed normal random effects with unknown variance (e.g. Chapter 4 in Ruppert et al. 2003).

Non-linear effects   For continuous variables, the smooth functions are represented using the regression spline approach popularized by Eilers and Marx (1996) where the $B_{jd_j}(z_{ij})$ are known spline basis functions. The design matrix $\boldsymbol{Z}_j$ then comprises the basis function evaluations for each individual observation i. Note that for one-dimensional smooth functions, the choice of spline definition does not play an important role in determining the shape of $\hat{s}_j(z_j)$ (e.g. Chapters 3 and 5 in Ruppert et al. 2003). To enforce smoothness, a conventional integrated squared second derivative spline penalty is typically employed, i.e. $\boldsymbol{K}_j = \int \boldsymbol{d}_j(z_j)\boldsymbol{d}_j(z_j)'dz_j$, where the j_d^{th} element of $\boldsymbol{d}_j(z_j)$ is given by $\partial^2 B_{jd_j}(z_j)/\partial z_j^2$ and integration is over the range of z_j. The formulae used to compute the basis functions and penalties for many spline definitions are provided in Ruppert et al. (2003) and Wood (2006). For their theoretical properties see, for instance, Wojtys and Marra (2015) and Yoshida and Naito (2014). As a simple, approximate version, Eilers and Marx (1996) suggest using $\boldsymbol{K}_j = \boldsymbol{D}'_j\boldsymbol{D}_j$ where $\boldsymbol{D}_j$ is a first or second order difference matrix.

Spatial effects   When the geographic area (or country) of interest is split up into discrete contiguous geographic units (or regions) and such information is available, a Markov random field approach can be employed to exploit the information contained in neighbouring observations which are located in the same area (or country). In this case, equation (4) becomes $\boldsymbol{z}'_{ij}\boldsymbol{\beta}_j$ where $\boldsymbol{\beta}_j = (\beta_{j1}, \ldots, \beta_{jD_j})'$ represents the vector of spatial effects, D_j denotes the total number of regions and $\boldsymbol{z}_{ij}$ is made up of a set of area labels. The design matrix linking an observation i to the corresponding spatial effect is therefore defined as

$$\boldsymbol{Z}_j[i, d_j] = \begin{cases} 1 & \text{if the observation belongs to region } d_j \\ 0 & \text{otherwise} \end{cases},$$

where $d_j = 1, \ldots, D_j$. The smoothing penalty is based on the neighborhood structure of the geographic units, so that spatially adjacent regions share similar effects. That is,

$$\boldsymbol{K}_j[d_j, d_j'] = \begin{cases} -1 & \text{if } d_j \neq d_j' \wedge \ d_j \sim d_j' \\ 0 & \text{if } d_j \neq d_j' \wedge \ d_j \nsim d_j' \ , \\ N_{d_j} & \text{if } d_j = d_j' \end{cases}$$

where $d_j \sim d_j'$ indicates whether two regions d_j and d_j' are adjacent neighbors, $d_j \nsim d_j'$ indicates that d_j and d_j' are not neighbours, and N_{d_j} is the total number of neighbours for region d_j. In a stochastic interpretation, this penalty is equivalent to the assumption that $\boldsymbol{\beta}_j$ follows a Gaussian Markov random field (e.g. Chapter 2 in Rue and Held 2005).

Other effect types   Several other specifications can be employed. These include varying coefficient smooths obtained by multiplying one or more smooth components by some covariate(s), and smooth functions of two or more continuous covariates (e.g. Wood 2006; Fahrmeir et al. 2013, for several cases studies and illustrations).

Predictor specifications   When specifying the structure of the predictors in our application, we mainly follow the analysis in Neelon et al. (2014) and previous findings on relevant covariates for modelling birth weight and gestational age from the epidemiological literature (Kramer 1987). This allows us to avoid the common hurdles associated with performing model selection in a complex regression setting such as confounding of effects, collinearity and upward biases in estimated coefficients after selecting the most relevant effects. Although shrinkage or penalized regression approaches may help with these problems, we prefer to rely on theoretical arguments for covariate inclusion. Nevertheless, we provide additional insight into the role of the risk factors used in the previous literature by allowing for flexible modelling of continuous covariates through spline functions, and by allowing the dependence between birth weight and gestational age to be modified by model covariates.

For the additive predictors η of all distributional parameters (in our case π, μ, σ, θ) we use the same specification (for the reasons outlined above) and consider as the following regressors: male (yes, no), mother's race (white, Hispanic, black, other), mother's education (primary, secondary, tertiary), marital status (married, not married), whether the individual smokes (yes, no), first birth (yes, no), mother's age (in years), month of birth (from 1 to 12) and region (representing the 100 North Carolina counties). All variables except for mother's age and region enter the predictor equations parametrically. The effect of mother's age is modelled flexibly using thin plate splines with 10 bases and a second order penalty, whereas the effect of region is modelled using the Gaussian Markov field approach described above. Using different spline definitions for the smooth functions of mother's age and/or increasing the basis size did not lead to tangibly different results.

3 Penalized Maximum Likelihood Inference

In the following section, we provide some details on the penalized likelihood inferential framework employed for the proposed mixed binary-continuous copula regression models. Both variants rely on Equation (2) for constructing the likelihood of the model.

Using (2), for a random sample of n observations the log-likelihood function of the copula model can be written as

$$\ell(\boldsymbol{\beta}) = \sum_{i=1}^{n} (1 - y_{i1}) \log \left\{ F_{1|2}(0|y_{i2}) \right\} + y_{i1} \log \left\{ 1 - F_{1|2}(0|y_{i2}) \right\} + \log \left\{ f_2(y_{i2}) \right\}, \qquad (5)$$

where

$$F_{1|2}(0|y_{i2}) = \frac{\partial C\left(F_1\left(0\right), F_2(y_{i2})\right)}{\partial F_2(y_{i2})}$$

is the conditional CDF of y_1 given y_2 and the complete vector of regression coefficients given by $\boldsymbol{\beta} = (\boldsymbol{\beta}'^{\vartheta_1}, \ldots, \boldsymbol{\beta}'^{\vartheta_k}, \ldots, \boldsymbol{\beta}'^{\vartheta_K})'$ where in turn $\boldsymbol{\beta}^{\vartheta_k}$ collects all regression coefficients of one particular parameter ϑ_k. Adding the penalty terms yields the penalized log-likelihood

$$\ell_p(\boldsymbol{\beta}) = \ell(\boldsymbol{\beta}) - \frac{1}{2}\boldsymbol{\beta}'\boldsymbol{K}\boldsymbol{\beta}, \qquad (6)$$

where $\boldsymbol{K} = \mathrm{bdiag}(\lambda_1^{\vartheta_1}\boldsymbol{K}_1^{\vartheta_1}, \ldots, \lambda_{J_1}^{\vartheta_1}\boldsymbol{K}_{J_1}^{\vartheta_1}, \ldots \lambda_1^{\vartheta_K}\boldsymbol{K}_1^{\vartheta_K}, \ldots, \lambda_{J_K}^{\vartheta_K}\boldsymbol{K}_{J_K}^{\vartheta_K})$ is a block-diagonal matrix containing all penalty matrices and $\boldsymbol{\lambda}$ is the vector of all smoothing parameters $\lambda_j^{\vartheta_k}$, $k = 1, \ldots, K$, $j = 1, \ldots, J_k$. To maximize (6) we have extended the efficient and stable trust region algorithm with integrated automatic multiple smoothing parameter selection introduced by Marra and Radice (2017b) to incorporate any parametric continuous marginal response distribution, link functions for the binary equation and one-parameter copula function, and to link all parameters of the model to additive predictors. Estimation of $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}$ is carried out in a two-step fashion:

step 1 Holding the smoothing parameter vector fixed at $\boldsymbol{\lambda}^{[a]}$ and for a given parameter vector value $\boldsymbol{\beta}^{[a]}$, we seek to maximize equation (5) using a trust region algorithm. That is,

$$\breve{\ell}_p(\boldsymbol{\beta}^{[a]}) \stackrel{\mathrm{def}}{=} - \left\{ \ell_p(\boldsymbol{\beta}^{[a]}) + \boldsymbol{p}'\boldsymbol{g}_p^{[a]} + \frac{1}{2}\boldsymbol{p}'\boldsymbol{H}_p^{[a]}\boldsymbol{p} \right\},$$

$$\boldsymbol{\beta}^{[a+1]} = \arg\min_{\boldsymbol{p}} \; \breve{\ell}_p(\boldsymbol{\beta}^{[a]}) + \boldsymbol{\beta}^{[a]}, \quad \text{so that} \quad \|\boldsymbol{p}\| \leq r^{[a]}$$

where a is an iteration index, $\boldsymbol{g}_p^{[a]} = \boldsymbol{g}^{[a]} - \boldsymbol{K}\boldsymbol{\beta}^{[a]}$, $\boldsymbol{H}_p^{[a]} = \boldsymbol{H}^{[a]} - \boldsymbol{K}$, vector $\boldsymbol{g}^{[a]}$ consists of

$$\boldsymbol{g}^{\vartheta_1[a]} = \partial\ell(\boldsymbol{\beta})/\partial\boldsymbol{\beta}^{\vartheta_1}|_{\boldsymbol{\beta}^{\vartheta_1} = \boldsymbol{\beta}^{\vartheta_1[a]}}, \ldots, \partial\ell(\boldsymbol{\beta})/\partial\boldsymbol{\beta}^{\vartheta_k}|_{\boldsymbol{\beta}^{\vartheta_k} = \boldsymbol{\beta}^{\vartheta_k[a]}}, \ldots \partial\ell(\boldsymbol{\beta})/\partial\boldsymbol{\beta}^{\vartheta_K}|_{\boldsymbol{\beta}^{\vartheta_K} = \boldsymbol{\beta}^{\vartheta_K[a]}}$$

and the Hessian matrix has elements

$$\boldsymbol{H}^{l,h[a]} = \partial^2\ell(\boldsymbol{\beta})/\partial\boldsymbol{\beta}^l\partial\boldsymbol{\beta}'^h|_{\boldsymbol{\beta}^l = \boldsymbol{\beta}^{l[a]}, \boldsymbol{\beta}^h = \boldsymbol{\beta}^{h[a]}},$$

where $l, h = \vartheta_1, \ldots, \vartheta_K$, $\| \cdot \|$ denotes the Euclidean norm and $r^{[a]}$ is the radius of the trust region which is adjusted through the iterations.

step 2 For a given smoothing parameter vector value $\boldsymbol{\lambda}^{[a]}$ and holding the main parameter vector value fixed at $\boldsymbol{\beta}^{[a+1]}$, solve the problem

$$\boldsymbol{\lambda}^{[a+1]} = \arg\min_{\boldsymbol{\lambda}} \ \mathcal{V}(\boldsymbol{\lambda}) \stackrel{\text{def}}{=} \|\mathcal{Z}^{[a+1]} - \boldsymbol{A}^{[a+1]}_{\boldsymbol{\lambda}^{[a]}}\mathcal{Z}^{[a+1]}\|^2 - \check{n} + 2\text{tr}(\boldsymbol{A}^{[a+1]}_{\boldsymbol{\lambda}^{[a]}}), \quad (7)$$

where, after defining $\boldsymbol{F}^{[a+1]}$ as $-\boldsymbol{H}^{[a+1]}$, $\mathcal{Z}^{[a+1]} = \sqrt{\boldsymbol{F}^{[a+1]}}\boldsymbol{\beta}^{[a+1]} + \sqrt{\boldsymbol{F}^{[a+1]}}^{-1}\boldsymbol{g}^{[a+1]}$, $\boldsymbol{A}^{[a+1]}_{\boldsymbol{\lambda}^{[a]}} = \sqrt{\boldsymbol{F}^{[a+1]}}\left(\boldsymbol{F}^{[a+1]} + \boldsymbol{K}\right)^{-1}\sqrt{\boldsymbol{F}^{[a+1]}}$, $\text{tr}(\boldsymbol{A}^{[a+1]}_{\boldsymbol{\lambda}^{[a]}})$ represents the number of effective degrees of freedom (edf) of the penalized model and $\check{n} = Kn$. Problem (7) is solved using the automatic efficient and stable approach described in Wood (2004).

The two steps are iterated until the algorithm satisfies the criterion $\frac{\left|\ell(\boldsymbol{\beta}^{[a+1]}) - \ell(\boldsymbol{\beta}^{[a]})\right|}{0.1 + \left|\ell(\boldsymbol{\beta}^{[a+1]})\right|} <$ $1e^{-07}$. The use of a trust region algorithm in step 1 and of (7) in step 2 are justified in Marra et al. (2017). It is worth remarking that a trust region approach is generally more stable and faster than its line-search counterparts (such as Newton-Raphson), particularly for functions that are, for example, non-concave and/or exhibit regions that are close to flat (Nocedal and Wright 2006, Chapter 4). Also, since $\boldsymbol{H}$ and $\boldsymbol{g}$ are obtained as a by-product of the estimation step for $\boldsymbol{\beta}$, little computational effort is required to set up the quantities required for the smoothing step. Starting values for the parameters of the marginals are obtained by fitting separate binary and continuous models with additive predictors. An initial value for the copula parameter is obtained by using a transformation of the unadjusted Kendall's association between the responses.

The proposed algorithm uses the analytical score and Hessian of $\ell(\boldsymbol{\beta})$ which have been derived in a modular fashion. For instance, for the copula Bernoulli-Gumbel distribution, the score vector for $\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)$, represented by π, μ, σ^2, and θ respectively, is made up of

$$\frac{\partial\ell(\boldsymbol{\beta})}{\partial\boldsymbol{\beta}^\pi} = \sum_{i=1}^{n}\left(\frac{1 - y_{i1}}{F_{1|2}(0|y_{i2})} - \frac{y_{i1}}{1 - F_{1|2}(0|y_{i2})}\right)\frac{\partial F_{1|2}(0|y_{i2})}{\partial F_1(0)}\frac{\partial F_1(0)}{\partial \eta_i^\pi}\boldsymbol{Z}^\pi[i,],$$

$$\frac{\partial\ell(\boldsymbol{\beta})}{\partial\boldsymbol{\beta}^\mu} = \sum_{i=1}^{n}\left\{\left(\frac{1 - y_{i1}}{F_{1|2}(0|y_{i2})} - \frac{y_{i1}}{1 - F_{1|2}(0|y_{i2})}\right)\frac{\partial F_{1|2}(0|y_{i2})}{\partial F_2(y_{i2})}\frac{\partial F_2(y_{i2})}{\partial \mu_i}+\right.$$

$$\left.\frac{1}{f_2(y_{i2})}\frac{\partial f_2(y_{i2})}{\partial \mu_i}\right\}\frac{\partial \mu_i}{\partial \eta_i^\mu}\boldsymbol{Z}^\mu[i,], \quad (8)$$

$\partial\ell(\boldsymbol{\beta})/\partial\boldsymbol{\beta}^{\sigma^2}$ whose expression is similar to (8), and

$$\frac{\partial\ell(\boldsymbol{\beta})}{\partial\boldsymbol{\beta}^\theta} = \sum_{i=1}^{n}\left(\frac{1 - y_{i1}}{F_{1|2}(0|y_{i2})} - \frac{y_{i1}}{1 - F_{1|2}(0|y_{i2})}\right)\frac{\partial F_{1|2}(0|y_{i2})}{\partial \theta_i}\frac{\partial \theta_i}{\partial \eta_i^\theta}\boldsymbol{Z}^\theta[i,],$$

where $\boldsymbol{Z}^\pi$, $\boldsymbol{Z}^\mu$ and $\boldsymbol{Z}^\theta$ are the overall design matrices corresponding to the equations for π, μ and θ. Looking, for example, at equation (8), we see that there are two components which depend only on the chosen copula, three terms which are marginal distribution dependent and one derivative whose form will depend on the adopted link function between

μ_i and η_i^μ. However, the main structure of the equation will be unaffected by the specific choices made. It will therefore be easy to extend the algorithm to other copulas and marginal distributions not discussed in this work as long as their CDFs and probability density functions are known and their derivatives with respect to their parameters exist. If a derivative is difficult and/or computationally expensive to compute then appropriate numerical approximations can be used.

Further considerations At convergence, reliable point-wise confidence intervals for linear and non-linear functions of the model coefficients (e.g., smooth components, copula parameter, joint and conditional predicted probabilities) are obtained using the Bayesian large sample approximation $\boldsymbol{\beta} \overset{a}{\sim} \mathcal{N}(\hat{\boldsymbol{\beta}}, -\hat{\boldsymbol{H}}_p^{-1})$. The rationale for using this result is provided in Marra and Wood (2012) for GAM, whereas some examples of interval construction are given in Radice et al. (2016). For general smooth models, such as the ones considered in this paper, this result can be justified using the distribution of $\mathcal{Z}$ discussed in Marra et al. (2017), making the large sample assumption that $\boldsymbol{F}$ can be treated as fixed, and making the usual Bayesian assumption on the prior of $\boldsymbol{\beta}$ for smooth models (Silverman 1985; Wood 2006). Note that this result neglects smoothing parameter uncertainty. However, as argued by Marra and Wood (2012), this is not problematic provided that heavy over-smoothing is avoided (so that the bias is not too large a proportion of the sampling variability) and in our experience we found that this result works well in practice. To test smooth components for equality to zero, the results discussed in Wood (2013a) and Wood (2013b) can be adapted to the current context. However, for our empirical analysis we do not deem this necessary as argued in the previous section. Proving consistency of the proposed estimator is beyond the scope of this paper but the results presented for instance in Wojtys and Marra (2015) can be easily adapted to the current context.

4 Empirical analysis

Before commenting on the results of the empirical study, we describe succinctly the process used for building the bivariate copula model and the R code employed to fit the final model. For the analysis, we restrict the data to infants born in 2008 in North Carolina.

To simplify the process, we exploited the fact that in a copula context the specification of margins and copula can be viewed as separate but related issues. We first chose the marginal distributions with fixed covariate specifications based on the AIC and BIC values (see Table 2). Then we proceeded to the choice of copula. Specifically, we began with the Gaussian and then, based on the (negative or positive) sign of the dependence, we considered the alternative specifications that were consistent with this initial finding. In this case, the values for the correlation coefficients were found to be negative. Therefore, we only considered copula which were consistent with this sign of dependence (see Table 3).

	Y_1: low birth weight	
Inverse link function	AIC	BIC
Probit	52816.28	53249.79
Complementary log-log	52833.92	53273.13
Logit	52828.86	53267.76

Table 1: Comparison of AIC and BIC values for the candidate marginal distributions for low birth weight Y_1.

	Y_2: gestational age	
Distribution	AIC	BIC
Normal	466124.1	468094.7
Gumbel	440225.2	441942.9
Reverse Gumbel	518570.4	520754.7
Logistic	452021.4	453842.5
Log-normal	473938.8	475957.8
Weibull	441611.7	443345.4
Inverse Gaussian	473678.4	475678.2
Gamma	500382.1	501186.1
Dagum	439730.9	441992.1
Singh-Maddala	3526100	3526607
Fisk	457351.8	457907.3

Table 2: Comparison of AIC and BIC values for the candidate marginal distributions for gestational age Y_2.

The chosen model, written below in terms of R syntax based on the GJRM package, is

```
eq1 <- lbw     ~ male + race + educ + marital + smokes + firstbirth +
                 dobmonth + s(age) + s(county, bs = "mrf", xt = xt)
eq2 <- wksgest ~ male + race + educ + marital + smokes + firstbirth +
                 dobmonth + s(age) + s(county, bs = "mrf", xt = xt)
eq3 <-         ~ male + race + educ + marital + smokes + firstbirth +
                 dobmonth + s(age) + s(county, bs = "mrf", xt = xt)
eq4 <-         ~ male + race + educ + marital + smokes + firstbirth +
                 dobmonth + s(age) + s(county, bs = "mrf", xt = xt)

f.l <- list(eq1, eq2, eq3, eq4)
outC90 <- gjrm(f.l, margins = c("probit", "GU"), BivD = "C90",
            Model = "B", data = dat)
```

where the s(age) are the smooth effects of age represented using thin plate regression splines with 10 basis functions and penalties based on second order derivatives, the s(county, bs = "mrf", xt = xt) represent spatial effects with neighborhood structure information stored in xt, arguments margins and BivD have the obvious interpretations and Model = "B" stands for bivariate model (other options are available in the package).

Copula	AIC	BIC
Normal	477846.8	480871.1
Frank	479754.1	482632.9
Farlie-Gumbel-Morgenstern	486614.1	488980.5
Plackett	477198.4	479899.4
Ali-Mikhail-Haq	486950.1	489309.0
Clayton 90°	475933.5	478427.8
Clayton 270°	482370.4	482598.1
Joe 90°	483057.9	486311.6
Joe 270°	475940.2	478437.9
Gumbel 90°	478033.6	481054.5
Gumbel 270°	476016.7	478506.8

Table 3: Comparison of AIC and BIC values under different copula assumptions.

Using a 2.20-GHz Intel(R) Core(TM) computer running Windows 7, computing time was about 25 minutes for a sample size of $n = 109,380$ observations with 524 parameters, hence highlighting the efficiency of the numerical implementation.

The R summary output in the Appendix presents the regression coefficients from the bivariate model for low birth weight (a binary indicator for being born with a weight less than 2,500 grams), and gestational age (measured in continuous weeks). As outlined above, flexible splines are used to model the association between maternal age and each outcome, and a Markov random field smoother is applied to the county indicators for mother's residence.

4.1   Results

Effects of Covariates   As in Neelon et al. (2014), we find the expected associations between the covariates and outcomes of interest, shown in the Appendix. As noted previously (Kramer 1987), some factors have adverse associations with both intrauterine growth and gestational duration, while others contribute positively to one and negatively to the other. For example, male babies are less likely to be low birth weight than female babies, but more likely to be born earlier. First births are born later than higher-order births, but are more likely to be LBW. In contrast, the impact of maternal smoking is unambiguously harmful. It increases the risk of low birth weight, and reduces gestational age. The most substantial coefficient for gestational duration in terms of magnitude relates to race. Babies born to black mothers are more likely to be low birth weight, and more likely to be born earlier than babies born to white mothers. In contrast, babies born to Hispanic mothers are less likely to be low birth weight and more likely to be born later than babies born to white mothers. After adjusting for other covariates, education does not appear to have a consistent impact on either outcome.

Results for the flexibly modelled impact of maternal age on these outcomes are shown in Figure 2. There is clear evidence of non-linearity, supporting our implementation of spline functions, with the risk of low birth weight increasing after the mid-30s, and gestational

age reaching its maximum in the mid-20s.

These results point to both modifiable and non-modifiable risk factors for poor early life health. The main modifiable risk factor is smoking, which emphasises the importance of public health campaigns to reduce smoking prevalence for promoting infant health. Non-modifiable risk factors, such as maternal age, race, and place of residence (county) are also important from a policy perspective because they provide a basis for identifying cohorts and individuals at risk. Because of the adverse consequences associated with low birth weight and early gestational age in both the short and long run (Black et al. 2007; McGovern 2013), an understanding of who is likely to experience these outcomes can help direct public resources. Moreover, as we argue above, because of the multiplicative risk of adverse outcomes associated with jointly being preterm and low birth weight, it is important to identify babies in this category. By simultaneously modelling the outcomes under study, and allowing for flexible dependence between birth weight and gestational age, the statistical model described in this paper permits estimation of the relevant joint and marginal quantities of interest.

Predicting joint relatives risks Figure 3 shows a bar plot with relative risk ratios for being low birth weight and being born before 37 weeks and 32 weeks, stratified by race, maternal smoking status and education. It is apparent from this figure that babies born to black mothers are at greatest risk, in that the predicted probability of joint occurrence of LBW and being born before 37 weeks is roughly twice that of babies born to mothers of other races. A similar risk penalty is apparent when stratifying by maternal smoking status. Here again, the relative risk of joint occurrence of LBW and being born before 37 weeks is almost twice as great for babies born to mothers who smoke compared to mothers who do not. Because there are few observations in the primary education category, the relevant comparison in mother's education is between secondary and tertiary, where mothers with tertiary education are at slightly lower risk for joint LBW and prematurity than those with only secondary education.

In Figure 4 (top, left), the relative risk of joint occurrence of LBW and being born before 37 weeks are shown by county of residence (of which there are 100) in North Carolina. Risk of an infant being born both LBW and preterm varies across counties , from 3% to 8%. The relative risk of joint occurrence for the highest risk counties of residence is twice that of the lowest risk counties of residence. The least favourable places to be born are clustered in the northeast of the state, specifically Hertford, Northampton, Halifax, Warren, Vance, Edgecombe, Bertie, and Washington counties. These results clearly indicate substantial inequality at birth dependent on a number of background characteristics including place of residence, smoking status, and race.

As we outlined in the introduction, adverse outcomes may be associated with a particular cut-off in the gestational age distribution. Previous research has focused on the standard definition of preterm (before 37 weeks) (Neelon et al. 2014), however, certain types of risk, for example mortality, are much more concentrated lower down the gestational age distribution, in this case prior to 32 weeks of gestational age (Slattery and Morrison 2002). Therefore, it is important to be able to incorporate different cut-offs in the modelling of joint probabilities so as to be able to focus on the most at risk part of the population.

Identifying those most likely to be born low birth weight and with a gestational age less than a given cut-off (that we can vary depending on the particular application of interest) will assist with policy orientation. Therefore, motivated by the data presented in the introduction, we also compute relative risk ratios from our statistical model for being born low birth weight and with gestational age less than 32 weeks. This allows us to predict the characteristics of babies among whom risk of mortality is most concentrated. The second set of bars in Figure 3 shows the joint probability of LBW and gestational age less than 32 weeks, stratified by maternal race, smoking, and education. Compared to white mothers, babies born to black mothers are 2.75 times more likely to be low birth weight and fall below the 32 week cut-off, as opposed to 2 times more likely to be low birth weight and be born before 37 weeks. This suggests that racial inequality is even greater than would be expected from examining the standard 37 week cut-off. This same pattern for being born before 32 weeks in relation to county of residence is apparent from Figure 4 (top, right), with higher relative risk ratios than for 37 weeks. The same counties exhibit the highest risk, but in this case babies in these locations are up to 3 times more likely to fall below the relevant cut-offs. The second row shows the resulting joint probabilities using two independent models based on `gam()` from the `mgcv R` package for the binary equation and `gamlss()` from the `gamlss R` for the second equation. These results clearly deviate from the findings of our proposed joint model, providing an indication for a lack of fit when falsely assuming an independent model specification.

Conditional dependence of outcomes Figure 5 shows the estimates for Kendall's τ (Nelsen 2006) by county obtained using the 90 degrees Clayton model. After accounting for covariates, the association between the two responses is non-zero and heterogeneous across counties. This clearly supports the presence of dependence between the binary and continuous outcomes after accounting for covariates, and hence that prediction of the relative risks has to be based on the joint model rather than the two marginal models.

Our overall estimate for the association between low birth weight and gestational age is negative (the copula parameter is estimated to be -0.296 (95% CI $-0.323, -0.271$) indicating that those who are low birth weight are more likely to be born earlier. However, our model allows us to examine whether this expected negative dependence is modified by the covariates we include. For example, male babies exhibit weaker dependence than female babies (the relevant parameter in the equation for dependence is negative), as do babies born to black mothers and mothers with tertiary education. In contrast, babies born to mothers of Hispanic origin and mothers who smoke exhibit even more negative dependence, indicating that if the baby is low birth weight they are also more likely to be born early.

5   Discussion

We have developed an inferential framework for fitting flexible bivariate regression copula models with binary and continuous margins with the aim of examining the risk factors of adverse birth outcomes. Parameter estimation is carried out within a penalized maximum

likelihood estimation framework with integrated automatic multiple smoothing parameter selection, and the proposed model can be easily implemented via the `R` package `GJRM`.

The empirical results highlight the flexibility of the statistical approach, and its potential to be applied to policy-relevant questions in population health. We are able to predict the joint occurrence of low birth weight and early gestational age, and identify the covariates predicting which babies most at risk of adverse outcomes. The joint prediction is important because of the multiplicative risk associated with both outcomes, for example, as we discussed in relation to infant mortality. Our results have a number of policy implications. First, we identify substantial racial inequalities in that babies born to black mothers are more than twice as likely to be both low birth weight and be born earlier than babies born to mothers of other races. In addition, these disparities are also evident across county of residence, where the relative risk of joint occurrence for the worst counties is around twice that of the best counties. Second, we find that these disparities only increase as gestational age declines, and are present for other cutoffs besides the standard one at 37 weeks gestation.. Given the concentration in mortality risk before 32 weeks of gestation, the ability to model how family background is associated with these factors across the relevant distribution is likely to be important for policy intervention. For example, confirmation of these racial and spatial inequalities in early life health may provide a basis for action to target these disparities. Given the evidence supporting the short and long run impact of early childhood conditions, this is likely to be a policy priority (Deaton 2013; McGovern 2018). This model is widely applicable because it can be used to flexibly model both the impact of covariates of interest (including non-linearity in continuous predictors and spatial dependence), and other outcomes of interest, one of which is binary and one of which is continuous, which are expected to be correlated.

Figure 2: Estimated smooth function of mother's age and regional effects and respective 95% confidence intervals for the probit equation (row 1), location (row 2), and scale parameters (row 3) of the Gumbel distribution, and association copula coefficient.
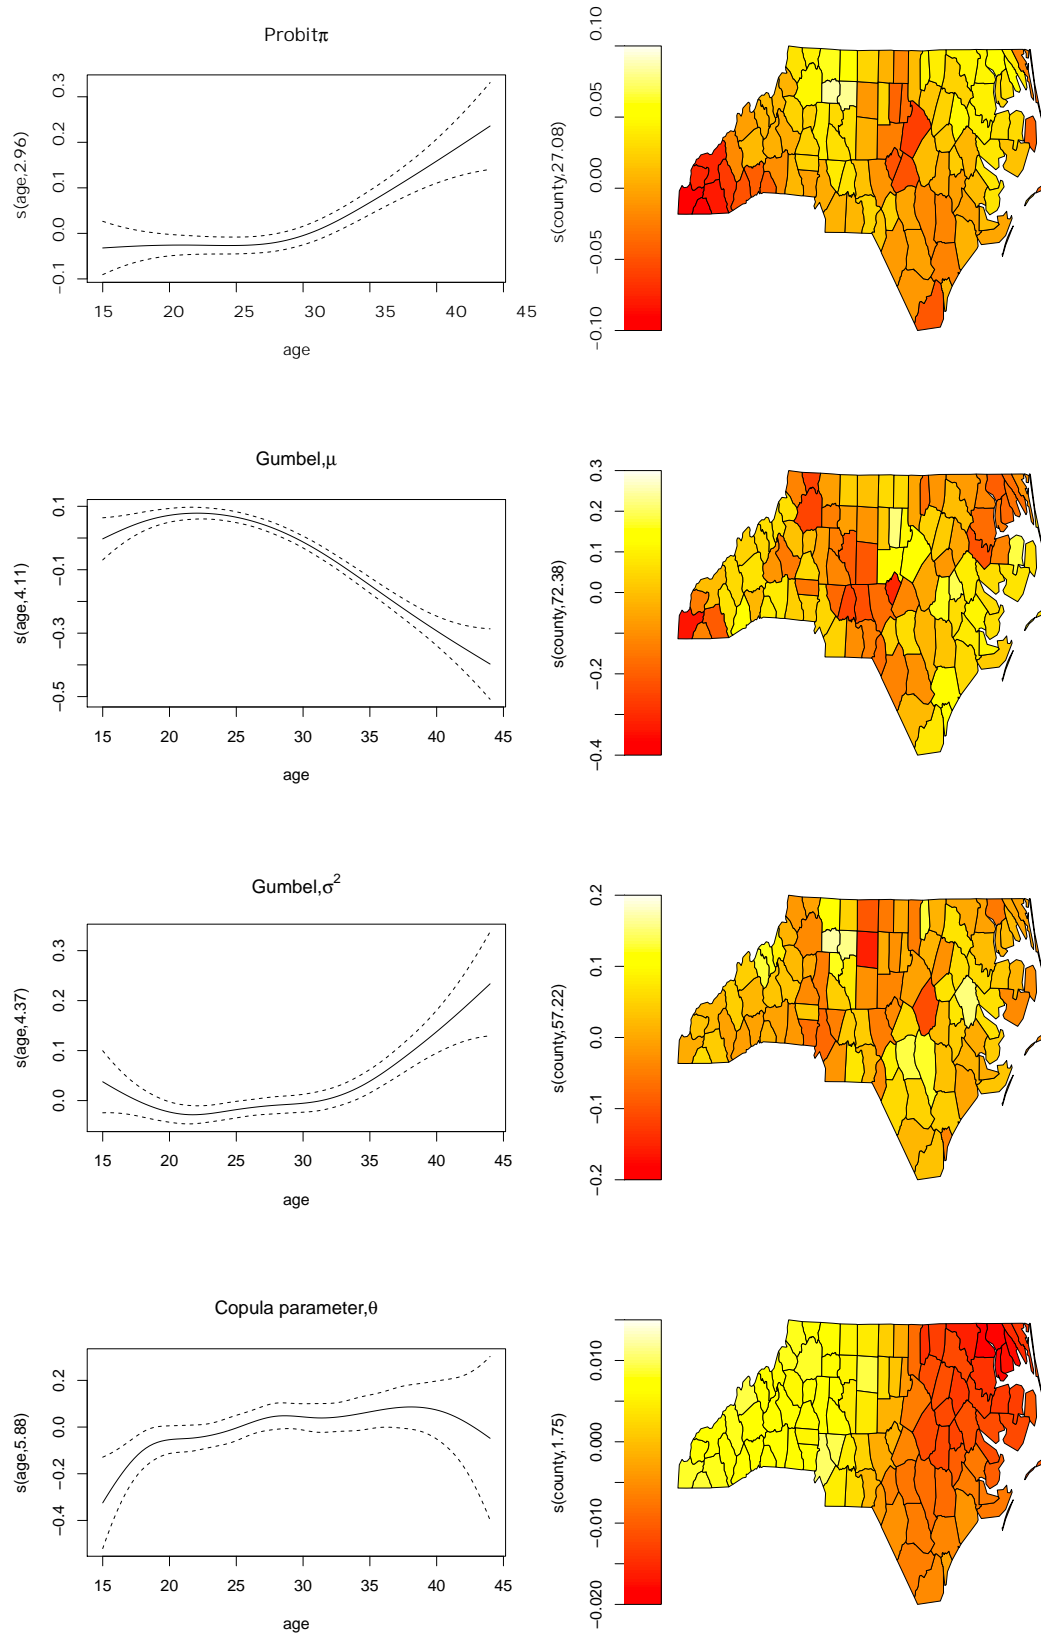
Figure 3: Joint probabilities for being low birth weight and being born before 37 weeks and 32 weeks by race, smoking behavior and education. The results are presented in terms of relative risk ratios where the baseline categories are white, non-smoker and primary, respectively. Variable age was set to its average, whereas the other factor variables were set to the respective categories with highest frequencies.
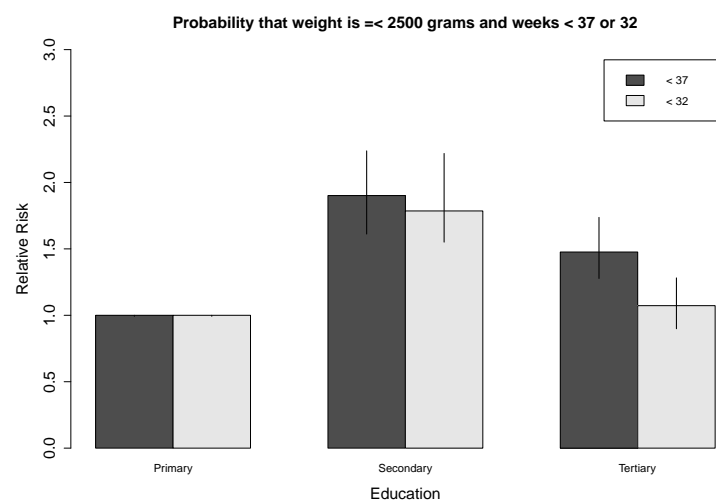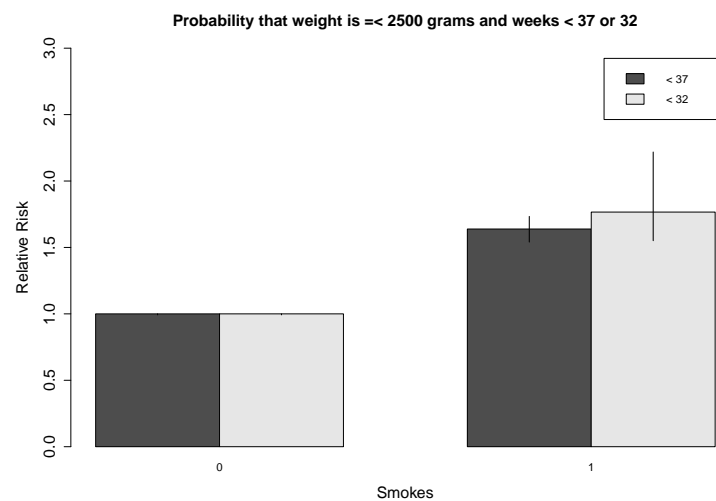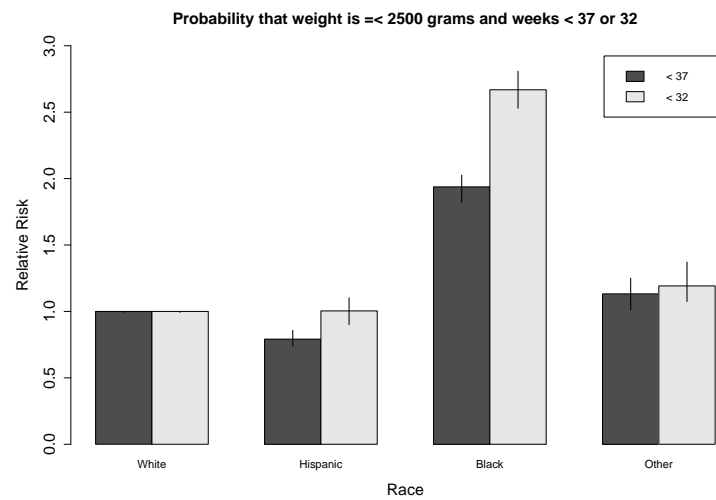
Figure 4: Joint probabilities for being low birth weight and being born before 37 weeks (left) and 32 weeks (right) by county of residence obtained by applying gjrm (top row) and gam + gamlss (bottom row). The results are displayed in terms of relative risk ratios where the baseline counties are those exhibiting the lowest probabilities.
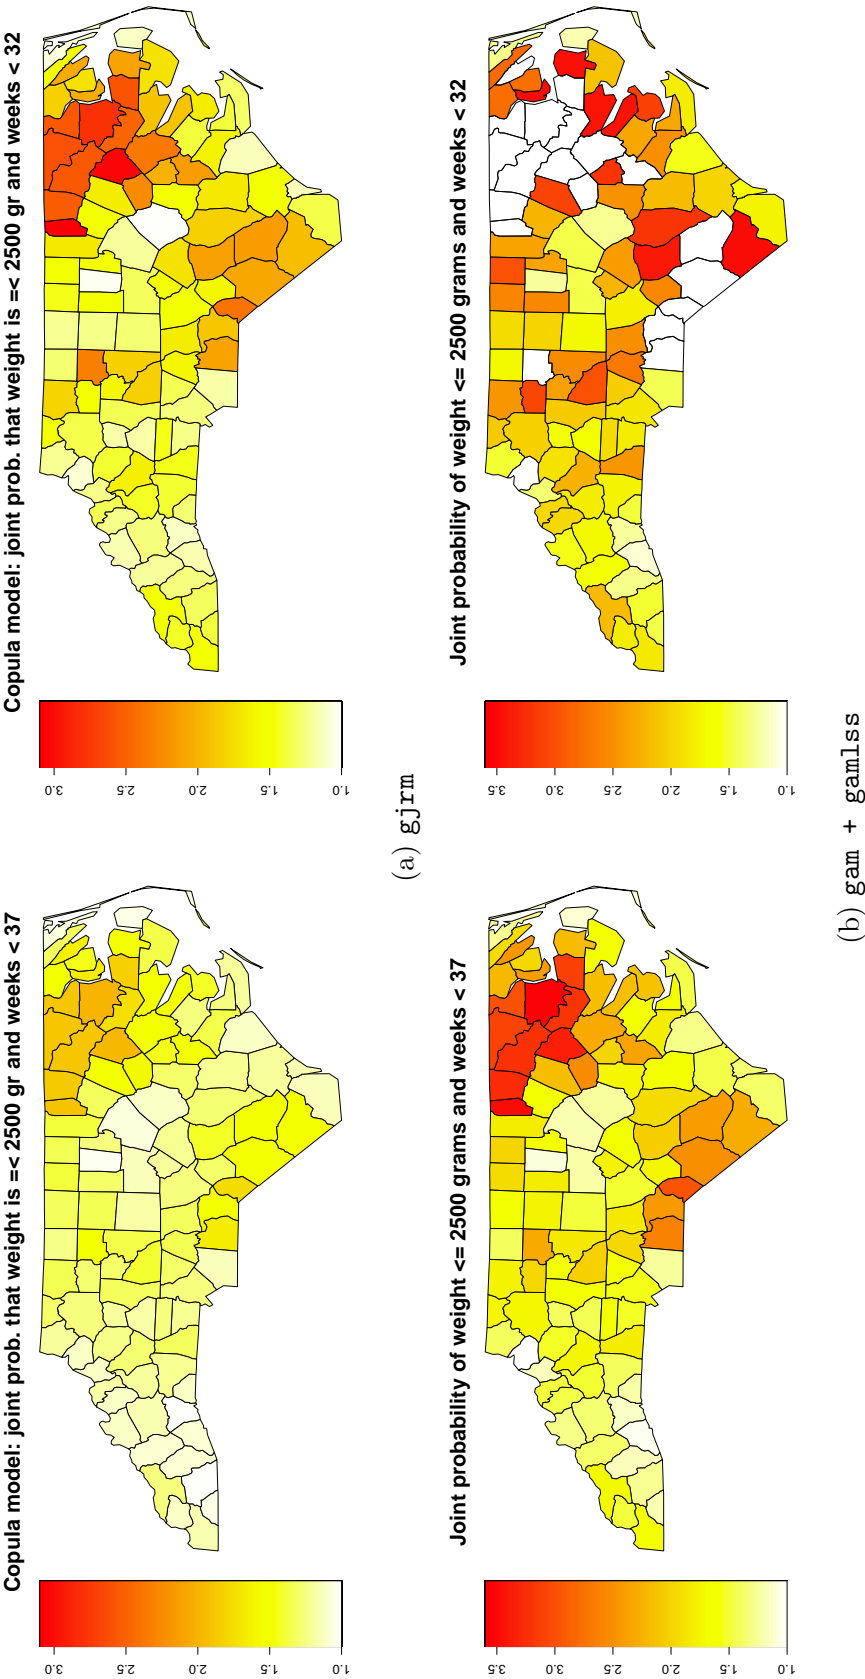
Figure 5: Estimates for Kendall's τ obtained using the 90 degrees Clayton model. These have been averaged by the covariate values, within each county in North Carolina.
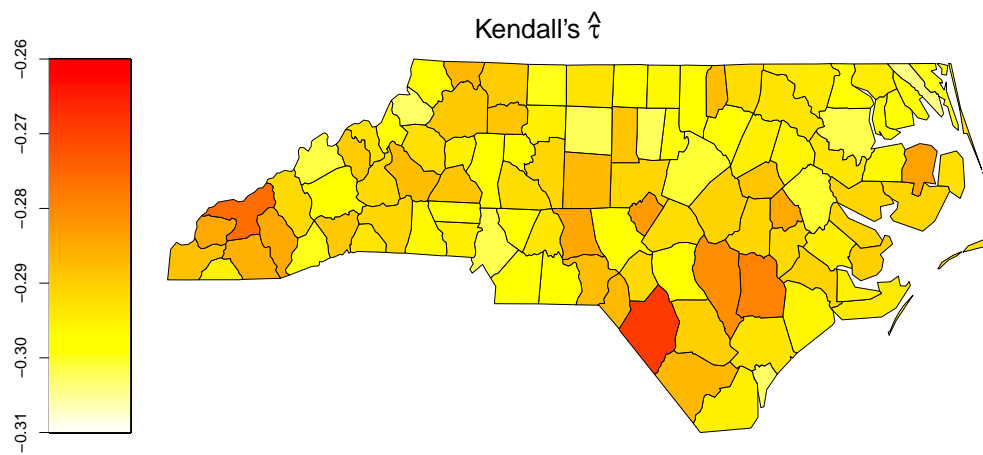
References

E. F. Acar, V. R. Craiu, and F. Yao. Statistical testing of covariate effects in conditional copula models. *Electronic Journal of Statistics*, 7:2822–2850, 2013.

D. Almond and J. Currie. Killing Me Softly: The Fetal Origins Hypothesis. *Journal of Economic Perspectives*, 25:153–172, 2011.

J. R. Behrman and M. R. Rosenzweig. Returns to birthweight. *Review of Economics and Statistics*, 86:586–601, 2004.

S. E. Black, P. J. Devereux, and K. G. Salvanes. From the Cradle to the Labor Market? The Effect of Birth Weight on Adult Outcomes. *Quartely Journal of Economics*, 122:409–439, 2007.

A. R. de Leon and C. C. K. Analysis of mixed data: Anb overview. In A. R. de Leon and C. C. K., editors, *Analysis of Mixed Data: Methods & Applications*. Chapman & Hall, New York, 2013.

A. R. de Leon and B. Wu. Copula-based regression models for a bivariate discrete and continuous outcome. *Statistics in Medicine*, 30(2):175–185, 2011.

A. Deaton. What does the empirical evidence tell us about the injustice of health inequalities? In E. Nir, S. Hurst, O. Norheim, and D. Wiclker, editors, *Inequalities in health: concepts measures, and ethics*, pages 263–281. Oxford University Press, Oxford, 2013.

P. H. C. Eilers and B. D. Marx. Flexible smoothing with B-splines and penalties. *Statistical Science*, 11:89–121, 1996.

L. Fahrmeir, T. Kneib, and S. Lang. Penalized structured additive regression for space-time data: A Bayesian perspective. *Statistica Sinica*, 14:731–761, 2004.

L. Fahrmeir, T. Kneib, S. Lang, and B. Marx. *Regression - Models, Methods and Applications*. Springer, Berlin, 2013.

J. E. Gentle. *Random number generation and Monte Carlo methods*. Springer-Verlag, London, 2003.

I. Gijbels, N. Veraverbeke, and M. Omelka. Conditional copulas, association measures and their applications. *Computational Statistics & Data Analysis*, 55:1919–1932, 2011.

M. Hack, D. J. Flannery, M. Schluchter, L. Cartar, E. Borawski, and N. Klein. Outcomes in young adulthood for very-low-birth-weight infants. *New England Journal of Medicine*, 346:149–157, 2002.

M. L. Hediger, M. D. Overpeck, W. J. Ruan, and J. F. Troendle. Birthweight and gestational age effects on motor and social development. *Paediatric and Perinatal Epidemiology*, 16:33–46, Jan. 2002.

H. Joe. *Multivariate Models and Dependence Concepts*. Chapman & Hall/CRC, London, 1997.

G. Kauermann, C. Schellhase, and D. Ruppert. Flexible copula density estimation with penalized hierarchical b-splines. *Scandinavian Journal of Statistics*, 40:685–705, 2013.

N. Klein and T. Kneib. Simultaneous estimation in structured additive conditional copula regression models: A unifying Bayesian approach. *Statistics and Computing*, 26:841–860, 2016.

N. Kraemer and D. Silvestrini. *CopulaRegression: Bivariate Copula Based Regression Models*, 2015.

M. S. Kramer. Intrauterine growth and gestational duration determinants. *Pediatrics*, 80: 502–511, 1987.

N. Kramer, E. C. Brechmann, D. Silvestrini, and C. Czado. Total loss estimation using copula-based regression models. *Insurance: Mathematics and Economics*, 53:829–839, 2012.

P. Lambert. Archimedean copula estimation using bayesian splines smoothing techniques. *Computational Statistics & Data Analysis*, 51:6307–6320, 2007.

G. Marra and R. Radice. *GJRM: Generalised Joint Regression Modelling*, 2017a.

G. Marra and R. Radice. Bivariate copula additive models for location, scale and shape. *Computational Statistics and Data Analysis*, 112:99–113, 2017b.

G. Marra and S. Wood. Coverage properties of confidence intervals for generalized additive model components. *Scandinavian Journal of Statistics*, 39:53–74, 2012.

G. Marra, R. Radice, T. Bärnighausen, S. N. Wood, and M. E. McGovern. A simultaneous equation approach to estimating HIV prevalence with non-ignorable missing responses. *Journal of the American Statistical Association*, 12:484–496, 2017.

M. E. McGovern. Still unequal at birth: birth weight, socio-economic status and outcomes at age 9. *The Economic and Social Review*, 44:53–84, 2013.

M. E. McGovern. How Much Does Birth Weight Matter for Child Health in Developing Countries? Estimates from Siblings and Twins. *Health Economics*, Forthcoming, 2018.

B. Neelon, R. Anthopolos, and M. L. Miranda. A spatial bivariate probit model for correlated binary data with application to adverse birth outcomes. *Statistical Methods in Medical Research*, 23:119–133, 2014.

R. Nelsen. *An Introduction to Copulas*. New York: Springer, 2006.

J. Nocedal and S. J. Wright. *Numerical Optimization*. New York: Springer-Verlag, 2006.

P. Oreopoulos, M. Stabile, R. Walld, and L. L. Roos. Short-, Medium-, and Long-Term Consequences of Poor Infant Health An Analysis Using Siblings and Twins. *Journal of Human Resources*, 43:88–138, 2008.

R. Radice, G. Marra, and M. Wojtys. Copula regression spline models for binary outcomes. *Statistics and Computing*, 26:981–995, 2016.

R. A. Rigby and D. M. Stasinopoulos. Generalized additive models for location, scale and shape (with discussion). *Journal of the Royal Statistical Society, Series C*, 54:507–554, 2005.

H. Rue and L. Held. *Gaussian Markov Random Fields*. Chapman & Hall/CRC, New York/Boca Raton, 2005.

D. Ruppert, M. P. Wand, and R. J. Carroll. *Semiparametric Regression*. Cambridge University Press, New York, 2003.

A. Sabeti, M. Wei, and R. V. Craiu. Additive models for conditional copulas. *Stat*, 3:300–312, 2014.

J. Segers, R. van den Akker, and B. J. M. Werker. Semiparametric Gaussian copula models: Geometry and efficient rank-based estimation. *Annals of Statistics*, 42:1911–1940, 2014.

X. Shen, Y. Zhu, and L. Song. Linear B-spline copulas with applications to nonparametric estimation of copulas. *Computational Statistics & Data Analysis*, 52:3806–3819, 2008.

B. W. Silverman. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society Series B*, 47:1–52, 1985.

M. M. Slattery and J. J. Morrison. Preterm delivery. *The Lancet*, 360:1489–1497, Nov. 2002.

M. Stasinopoulos, R. Rigby, C. Akantziliotou, V. Voudouris, G. Heller, R. Ospina, N. Motpan, F. McElduff, M. Djennad, M. Enea, A. Ghalanos, and C. Argyropoulos. *gamlss.dist: Distributions for Generalized Additive Models for Location Scale and Shape*, 2017a.

M. Stasinopoulos, R. Rigby, G. Heller, V. Voudouris, and F. D. Bastiani. *Flexible Regression and Smoothing: Using GAMLSS in R*. Chapman & Hall/CRC, London, 2017b.

T. Vatter and V. Chavez-Demoulin. Generalized additive models for conditional dependence structures. *Journal of Multivariate Analysis*, 141:147–167, 2015.

T. Wardlaw, A. Blanc, J. Zupan, and E. Ahman. United Nations Children's Fund and World Health Organization. *Low Birthweight: Country, Regional and Global Estimates*, 2005.

M. Wojtys and G. Marra. Copula based generalized additive models with non-random sample selection. *arXiv:1508.04070*, 2015.

S. N. Wood. Stable and efficient multiple smoothing parameter estimation for generalized additive models. *Journal of the American Statistical Association*, 99:673–686, 2004.

S. N. Wood. *Generalized Additive Models: An Introduction With R*. Chapman & Hall/CRC, London, 2006.

S. N. Wood. On p-values for smooth components of an extended generalized additive model. *Biometrika*, 100:221–228, 2013a.

S. N. Wood. A simple test for random effects in regression models. *Biometrika*, 100:1005–1010, 2013b.

B. Wu and A. R. de Leon. Gaussian copula mixed models for clustered mixed outcomes, with application in developmental toxicology. *Journal of Agricultural, Biological & Environmental Statistics*, 19(1):39–56, 2014.

J. Yan. Enjoy the joy of copulas: With a package copula. *Journal of Statistical Software*, 21: 1–21, 2007.

T. Yoshida and K. Naito. Asymptotics for penalized splines in generalized additive models. *Journal of Nonparametric Statistics*, 26:269–289, 2014.

6 Appendix

6.1 Further Empirical Results

Summary of parametric and non-parametric effects from the 90 degrees rotated Clayton copula model with Bernoulli (with probit link) and Gumbel margins. Sample size is 109 380, whereas the overall estimate for θ with CI is $-0.847(-0.96, -0.749)$. The corresponding Kendall's tau is $-0.296(-0.323, -0.271)$. Total number of effective degrees of freedom is 260.

```
COPULA:   90 Clayton
MARGIN 1: Bernoulli
MARGIN 2: Gumbel

EQUATION 1
Link function for mu.1: probit
Formula: lbw ~ male + race + educ + marital + smokes + firstbirth + dobmonth +
    s(mage, k = 12, bs = "ps", m = c(2, 2)) + s(county, bs = "mrf",
    xt = xt)

Parametric coefficients:
              Estimate Std. Error z value Pr(>|z|)
(Intercept)   -1.569347   0.045614 -34.405  < 2e-16 ***
male1         -0.089513   0.011734  -7.629 2.38e-14 ***
raceHispanic  -0.059786   0.021324  -2.804  0.00505 **
raceBlack      0.339162   0.015379  22.053  < 2e-16 ***
raceOther      0.169090   0.028753   5.881 4.08e-09 ***
educSecondary  0.070154   0.038956   1.801  0.07172 .
educTertiary  -0.033780   0.040390  -0.836  0.40296
marital1      -0.101551   0.014805  -6.859 6.93e-12 ***
smokes1        0.352365   0.018065  19.506  < 2e-16 ***
firstbirth1    0.106756   0.013304   8.025 1.02e-15 ***
dobmonth2     -0.039269   0.029054  -1.352  0.17650
dobmonth3     -0.004533   0.028483  -0.159  0.87354
dobmonth4     -0.015908   0.029016  -0.548  0.58352
dobmonth5     -0.009279   0.028750  -0.323  0.74690
dobmonth6      0.010851   0.028527   0.380  0.70366
dobmonth7     -0.016584   0.028284  -0.586  0.55765
dobmonth8     -0.018836   0.028315  -0.665  0.50590
dobmonth9     -0.037382   0.028612  -1.307  0.19138
dobmonth10    -0.007949   0.028268  -0.281  0.77856
dobmonth11    -0.034410   0.029093  -1.183  0.23690
dobmonth12     0.013941   0.028166   0.495  0.62063
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Smooth components' approximate significance:
            edf Ref.df Chi.sq  p-value
s(mage)   2.964  3.664  54.91 6.05e-11 ***
```

```
s(county) 27.079 99.000  61.36 2.61e-07 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1



EQUATION 2
Link function for mu.2: identity
Formula: wksgest ~ male + race + educ + marital + smokes + firstbirth +
    dobmonth + s(mage, k = 12, bs = "ps", m = c(2, 2)) + s(county,
    bs = "mrf", xt = xt)

Parametric coefficients:
              Estimate Std. Error  z value Pr(>|z|)
(Intercept)  39.613609   0.035462 1117.058  < 2e-16 ***
male1        -0.073412   0.009735   -7.541 4.67e-14 ***
raceHispanic  0.034584   0.016479    2.099 0.035839 *
raceBlack    -0.317929   0.014651  -21.700  < 2e-16 ***
raceOther    -0.082286   0.024253   -3.393 0.000692 ***
educSecondary -0.035389  0.029332   -1.207 0.227620
educTertiary  0.031329   0.030589    1.024 0.305742
marital1     -0.003031   0.012776   -0.237 0.812469
smokes1      -0.110066   0.018665   -5.897 3.70e-09 ***
firstbirth1   0.221671   0.011081   20.004  < 2e-16 ***
dobmonth2    -0.017297   0.024262   -0.713 0.475882
dobmonth3    -0.008670   0.024012   -0.361 0.718048
dobmonth4     0.023724   0.024190    0.981 0.326722
dobmonth5    -0.027441   0.023975   -1.145 0.252391
dobmonth6    -0.037190   0.023954   -1.553 0.120527
dobmonth7     0.017129   0.023580    0.726 0.467591
dobmonth8     0.008875   0.023343    0.380 0.703799
dobmonth9     0.072997   0.023279    3.136 0.001714 **
dobmonth10    0.022596   0.023598    0.958 0.338300
dobmonth11    0.030304   0.024061    1.259 0.207865
dobmonth12    0.015868   0.023954    0.662 0.507696
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1


Smooth components' approximate significance:
            edf Ref.df Chi.sq p-value
s(mage)   4.108  4.983  307.7  <2e-16 ***
s(county) 72.382 99.000  586.4  <2e-16 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1



EQUATION 3
Link function for sigma2: log
Formula: ~male + race + educ + marital + smokes + firstbirth + dobmonth +
    s(mage, k = 12, bs = "ps", m = c(2, 2)) + s(county, bs = "mrf",
```

```
    xt = xt)

Parametric coefficients:
              Estimate Std. Error z value Pr(>|z|)
(Intercept)    0.878581   0.033827  25.973  < 2e-16 ***
male1          0.045561   0.009311   4.893 9.91e-07 ***
raceHispanic   0.018671   0.015810   1.181 0.237622
raceBlack      0.260953   0.013208  19.757  < 2e-16 ***
raceOther      0.068929   0.023282   2.961 0.003070 **
educSecondary  0.010272   0.027917   0.368 0.712896
educTertiary  -0.103331   0.029156  -3.544 0.000394 ***
marital1      -0.105515   0.011945  -8.833  < 2e-16 ***
smokes1        0.147066   0.016612   8.853  < 2e-16 ***
firstbirth1   -0.013728   0.010693  -1.284 0.199198
dobmonth2      0.021599   0.022993   0.939 0.347551
dobmonth3      0.027197   0.022744   1.196 0.231768
dobmonth4      0.020124   0.022993   0.875 0.381470
dobmonth5      0.011315   0.022893   0.494 0.621148
dobmonth6      0.009636   0.022781   0.423 0.672307
dobmonth7      0.010346   0.022458   0.461 0.645017
dobmonth8     -0.027985   0.022339  -1.253 0.210293
dobmonth9     -0.061163   0.022553  -2.712 0.006689 **
dobmonth10    -0.008476   0.022602  -0.375 0.707643
dobmonth11    -0.023756   0.023012  -1.032 0.301910
dobmonth12     0.027743   0.022658   1.224 0.220802
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Smooth components' approximate significance:
            edf Ref.df Chi.sq  p-value
s(mage)    4.37  5.278  57.32 9.25e-11 ***
s(county) 57.22 99.000 351.06  < 2e-16 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1


EQUATION 4
Link function for theta: log(- )
Formula: ~male + race + educ + marital + smokes + firstbirth + dobmonth +
    s(mage, k = 12, bs = "ps", m = c(2, 2)) + s(county, bs = "mrf",
    xt = xt)

Parametric coefficients:
              Estimate Std. Error z value Pr(>|z|)
(Intercept)   -0.443353   0.106572  -4.160 3.18e-05 ***
male1          0.055999   0.026091   2.146 0.031851 *
raceHispanic  -0.182236   0.047815  -3.811 0.000138 ***
raceBlack      0.090711   0.032552   2.787 0.005325 **
raceOther     -0.164939   0.065248  -2.528 0.011476 *
```

```
educSecondary  0.142429   0.093100    1.530 0.126054
educTertiary   0.212788   0.095637    2.225 0.026085 *
marital1       0.044000   0.032003    1.375 0.169172
smokes1       -0.079724   0.041510   -1.921 0.054782 .
firstbirth1    0.031688   0.029183    1.086 0.277550
dobmonth2      0.112302   0.065288    1.720 0.085413 .
dobmonth3      0.092269   0.063786    1.447 0.148026
dobmonth4     -0.043052   0.066085   -0.651 0.514745
dobmonth5      0.038112   0.064422    0.592 0.554121
dobmonth6      0.147006   0.064017    2.296 0.021655 *
dobmonth7     -0.012702   0.064170   -0.198 0.843093
dobmonth8      0.153894   0.064453    2.388 0.016955 *
dobmonth9      0.009381   0.064231    0.146 0.883881
dobmonth10    -0.004853   0.063891   -0.076 0.939453
dobmonth11    -0.003268   0.065797   -0.050 0.960393
dobmonth12     0.093968   0.063081    1.490 0.136319
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1


Smooth components' approximate significance:
           edf Ref.df Chi.sq p-value
s(mage)   5.878  6.912 19.983  0.0071 **
s(county) 1.752 99.000  2.219  0.1817
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1



n = 109380  sigma2 = 2.43(2.26,2.61)
theta = -0.847(-0.96,-0.749)  tau = -0.296(-0.323,-0.271)
total edf = 260
```

6.2   Simulation

6.2.1   Settings and Design

The aim of this section is to assess the empirical properties of the proposed modelling approach. We consider a scenario where data are generated using probit and Gumbel margins, and the Joe copula. In this case, all the parameters of the bivariate distribution depend on additive predictors. That is,

$$\Phi^{-1}(\pi) = \alpha_0 + \alpha_1 x_1 + s_1(x_2) + \alpha_2 x_3,$$
$$\mu = \beta_0 + s_2(x_2) + \beta_1 x_3,$$
$$\log(\sigma^2) = \beta_0^{\sigma^2} + \beta_1^{\sigma^2} x_3,$$
$$\log(\theta) = \beta_0^{\theta} + \beta_1^{\theta} x_1 + s_3(x_2),$$

where $\alpha_0 = -0.8$, $\alpha_1 = -1.3$, $\alpha_2 = 1$, $\beta_0 = 0.1$, $\beta_1 = -0.9$, $\beta_0^{\sigma^2} = 0.5$, $\beta_1^{\sigma^2} = 1$, $\beta_0^{\theta} = 1.1$, $\beta_1^{\theta} = -1.4$, $s_1(x) = x + \exp(-30(x - 0.5)^2)$, $s_2(x) = \sin(2\pi x)$ and $s_3(x) = 0.6(e^x + \sin(2.9x))$. Recall that the marginal distributions are parametrised according to Stasinopoulos et al. (2017b).

The simulated data sets consists of one binary response and one continuous response related to

one binary covariate and two continuous regressors. Sample sizes were set to 1,000 and 5,000, and the number of replicates was set to 1,000. The models were fitted using `gjrm()` in GJRM and, following a reviewer's suggestion, using two independent models based on `gam()` from the `mgcv` R package for the binary response and `gamlss()` from the `gamlss` R for the Gumbel response. In both approaches, each smooth function is represented using a penalized low rank thin plate spline with second order difference penalty and 10 basis functions. For each replicate, smooth function estimates are constructed using 200 equally spaced fixed values in the $(0, 1)$ range (e.g., Radice et al. 2016). The code used to generate the data is given below.

```
library(copula)
library(gamlss.add)
library(GJRM)

cor.cov <- matrix(0.5, 3, 3); diag(cor.cov) <- 1

s1 <- function(x) x+exp(-30*(x-0.5)^2)
s2 <- function(x) sin(2*pi*x)
s3 <- function(x) 0.6*(exp(x) + sin(2.9*x))

datagen <- function(cor.cov, s1, s2, s3){

 cov <- rMVN(1, rep(0,3), cor.cov)
 cov <- pnorm(cov)

 z1 <- cov[, 1]
 z2 <- cov[, 2]
 z3 <- round(cov[, 3])

 eta_pi  <- -.8 - 1.3*z1 + s1(z2) + z3
 eta_mu  <- 0.1          + s2(z2) - 0.9*z3
 eta_si2 <- 0.5 + z3
 eta_the <- 1.1 - 1.4*z1 + s3(z2)

 Cop <- archmCopula(family = "joe", dim = 2, param = exp(eta_the) + 1 + 1e-07)

 speclist1 <- list( mu = eta_pi, sigma = 1)
 speclist2 <- list( mu = eta_mu, sigma = sqrt(exp(eta_si2)))

 spec <- mvdc(copula = Cop, c("NO", "GU"), list(speclist1, speclist2) )

 resp    <- rMvdc(1, spec)
 resp[1] <- resp[1] > 0

 c(resp, z1, z2, z3)

}
```

The package `copula` (Yan 2007) contains functions `archmCopula()`, `mvdc()` and `rMvdc()` which allow us to simulate from the desired copula. Package `gamlss.dist` (Stasinopoulos et al. 2017a)

contains all the functions required to simulate the marginals adopted here, and `rMVN()` (from `GJRM`) allows us to simulate Gaussian correlated variables. The correlation matrix used to associate the three simulated Gaussian covariates is `cor.cov`, whereas `cov <- pnorm(cov)` yields Uniform(0,1) correlated covariates (e.g., Gentle 2003). A balanced binary regressor is created using `round(cov[, 3])`. Functions `s1`, `s2`, `s3` produce curves with different degrees of complexity. The various `eta` refer to the model's additive predictors. If necessary, these are transformed in `archmCopula`, `speclist1` and `speclist2` to ensure that the restrictions on the parameters' spaces of the employed bivariate distribution are maintained.

6.2.2 Results

Figures 6 and 7 provide a pictorial representation of the accuracy and precision of the linear and non-linear effect estimates, whereas Table 4 reports the bias and root mean squared error for the two approaches.

Overall, the results show that both methods estimate all the model components fairly well, and that the estimates improve as the sample size increases. Specifically, all mean estimates are very close to the true values and, as expected, their variability decreases as the sample size increases. The components in the additive predictor of the dependence parameter are estimated less precisely than those of the others. This indicates that the effects of covariates on the association between the two equations require more information to be estimated reliably.

Our proposal generally delivered somewhat less biased and slightly more efficient estimates, although the differences are not striking. It may be argued that ignoring dependence may still deliver estimates that are similar to those obtained from a simultaneous model. To this end, we would like to stress that the main advantage of the proposed joint model is to enable the user to study and account for the association between outcomes and how such dependence is modified by covariates. As illustrated in our empirical analysis, this can reveal interesting mechanisms (which would not otherwise be accessible using a model assuming independence) and allows one to obtain joint probability statements that account for dependence between the outcomes.

We also explored the ability of information criteria such as AIC and BIC in selecting the correct model (in this case the one accounting for dependence). Unsurprisingly, the correct model was always chosen by both criteria.

(a) `gjrm` (b) `gam + gamlss`

Figure 6: Linear coefficient estimates obtained by applying `gjrm` (black lines) and `gam + gamlss` (grey lines) to data simulated as described in this section. Vertical bars in the left panel refer to $n = 1,000$ whereas the ones in the right panel belong to $n = 5,000$. Circles indicate mean estimates while bars represent the estimates' ranges resulting from 5% and 95% quantiles. True values are indicated by black solid horizontal lines.

(a) `gjrm`

(b) `gam + gamlss`

Figure 7: Smooth function estimates obtained by applying `gjrm` (top plots) and `gam +
gamlss` (bottom plots) to data simulated as described in this section. True functions are
represented by black solid lines, mean estimates by dashed lines and point-wise ranges
resulting from 5% and 95% quantiles by shaded areas. In each group of plots, the first
row refers to $n = 1000$ and the second to $n = 5000$.

		Bias		RMSE	
		`gjrm`	`gam + gamlss`	`gjrm`	`gam + gamlss`
$n = 1000$	α_1	-0.001	-0.010	0.115	0.185
	α_2	0.014	0.001	0.090	0.100
	β_1	0.018	0.004	0.123	0.123
	$\beta_1^{\sigma^2}$	-0.004	0.002	0.094	0.099
	β_1^{θ}	-0.065	-	0.429	-
	s_1	0.028	0.030	0.109	0.110
	s_2	0.029	0.031	0.118	0.117
	s_3	0.056	-	0.184	-
$n = 5000$	α_1	0.000	0.001	0.046	0.085
	α_2	-0.001	-0.001	0.038	0.046
	β_1	0.001	0.003	0.055	0.058
	$\beta_1^{\sigma^2}$	0.002	-0.001	0.044	0.044
	β_1^{θ}	-0.012	-	0.162	-
	s_1	0.012	0.014	0.041	0.053
	s_2	0.014	0.014	0.049	0.054
	s_3	0.024	-	0.081	-

Table 4: Bias and root mean squared error (RMSE) obtained by applying the `gjrm` and `gam + gamlss` parameter estimators to data simulated as described in this section. Bias and RMSE for the smooth terms are calculated, respectively, as $n_s^{-1} \sum_{i=1}^{n_s} |\bar{\hat{s}}_i - s_i|$ and $n_s^{-1} \sum_{i=1}^{n_s} \sqrt{n_{rep}^{-1} \sum_{rep=1}^{n_{rep}} (\hat{s}_{rep,i} - s_i)^2}$, where $\bar{\hat{s}}_i = n_{rep}^{-1} \sum_{rep=1}^{n_{rep}} \hat{s}_{rep,i}$, n_s is the number of equally spaced fixed values in the $(0,1)$ range, and n_{rep} is the number of simulation replicates. In this case, $n_s = 200$ and $n_{rep} = 1000$. The bias for the smooth terms is based on absolute differences in order to avoid compensating effects when taking the sum.