



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2023

A flexible copula regression model with Bernoulli and Tweedie margins for estimating the effect of spending on mental health

Marra, Giampiero ; Fasiolo, Matteo ; Radice, Rosalba ; Winkelmann, Rainer

DOI: <https://doi.org/10.1002/hec.4668>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-234434>

Journal Article

Published Version



The following work is licensed under a Creative Commons: Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) License.

Originally published at:

Marra, Giampiero; Fasiolo, Matteo; Radice, Rosalba; Winkelmann, Rainer (2023). A flexible copula regression model with Bernoulli and Tweedie margins for estimating the effect of spending on mental health. *Health Economics*, 32(6):1305-1322.

DOI: <https://doi.org/10.1002/hec.4668>

A flexible copula regression model with Bernoulli and Tweedie margins for estimating the effect of spending on mental health

Giampiero Marra¹ | Matteo Fasiolo² | Rosalba Radice³ | Rainer Winkelmann⁴ 

¹Department of Statistical Science,
University College London, London, UK

²School of Mathematics, University of
Bristol, Bristol, UK

³Bayes Business School City, University of
London, London, UK

⁴Department of Economics, University of
Zurich, Zurich, Switzerland

Correspondence

Rainer Winkelmann.

Email: rainer.winkelmann@econ.uzh.ch

Abstract

We develop a flexible two-equation copula model to address endogeneity of medical expenditures in a distribution regression for health. The expenditure margin uses the compound gamma distribution, a special case of the Tweedie family of distributions, to account for a spike at zero and a highly skewed continuous part. An efficient estimation algorithm offers flexible choices of copulae and link functions, including logit, probit and cloglog for the health margin. Our empirical application revisits data from the Rand Health Insurance Experiment. In the joint model, using random insurance plan assignment as instrument for spending, a \$1000 increase is estimated to reduce the probability of a low post-program mental health index by 1.9 percentage points. The effect is not statistically significant. Ignoring endogeneity leads to a spurious positive effect estimate.

KEYWORDS

binary response, copula, health expenditures, penalized regression spline, Rand experiment, simultaneous estimation, Tweedie distribution

1 | INTRODUCTION

It is quite common in empirical health economics to specify models where two or more outcome variables are potentially associated conditional on a set of covariates. Modeling such association matters for efficiency, but also for addressing endogeneity in a recursive system of equations. The choice of model is dictated by the nature of the dependent variables. For example, for two jointly determined binary responses, the bivariate probit is frequently used (see Waters (1999), Farbmacher et al. (2017), Humphreys et al. (2014), among others). Often, however, there will be a mix of outcomes. As typical example, consider a joint model for insurance status and health care utilization, as in Deb and Trivedi (2006) and Marra et al. (2020). In such cases, one can use copulae to introduce dependence between outcomes with known arbitrary margins. In fact, the copula approach allows the marginal distributions to be chosen to best fit the data.

In this paper, we are interested in expenditures for medical services (per person and year) and their effect on health, based on a joint model with associated outcomes. Modeling medical expenditure data is a challenge because they display a substantial fraction of zeros, often more than 50%, in combination with a continuous distribution of positive amounts that is highly skewed. In single equation models, which neglect the joint determination of outcomes, estimation of regression parameters does not rely on a correctly specified distribution, and determinants of expenditure have often been estimated using the Gamma pseudo maximum likelihood estimator (Manning et al., 1987), linear regression for logarithmic non-zero expenditures (Manning & Mullahy, 2001), or a variety of two-part models (Mullahy, 1998). For the copula approach, however, it is important to find a

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. Health Economics published by John Wiley & Sons Ltd.

suitable marginal distribution for expenditures. Kurz (2017) recently argued that a compound Gamma distribution can provide a good characterization of health care expenditures. Specifically, assume that expenditures Y is a sum of N independent and identically distributed Gamma random variables, where N is distributed as Poisson(ω). Then the probability of a zero expenditure is given by $\exp(-\omega)$ and the distribution of positive amounts is skewed to the right. The compound gamma distribution is a special case of the more general Tweedie family of distributions. A plot based on quantile residuals confirms the suitability of the Tweedie for modeling the marginal distribution of health care expenditures in our application.

Our main technical contribution is then to build a flexible copula regression model for the joint determination of a Tweedie variable and a Bernoulli response (in our application an indicator whether a mental health index falls below a given threshold) and describe the steps for an efficient estimation algorithm. The proposed model is flexible in the sense that it is possible to: choose several link functions for the Bernoulli margin, incorporate (linear and non-linear) covariate effects into any distributional parameter, and allow for the exploration of a wide set of copulae. Moreover, it can be made fully interdependent, recursive, or “seemingly unrelated”, by imposing zero-constraints on corresponding coefficients. A similar model has been proposed by Marra et al. (2020) who focused on binary and count margins (see, also, the references therein for alternative simpler versions, with different types of margins). Our new model has been implemented in the GJRM package (Marra & Radice, 2022), written for the programming language R (R Core Team, 2022). To the best of our knowledge, this is the first freely available implementation of a flexible copula regression model involving the Tweedie distribution.

In terms of substantive application, we use the proposed framework for estimating the causal effect of expenditure on health. We revisit data from the Rand Health Insurance Experiment (RHIE, e.g., Manning et al., 1987; Aron-Dine et al., 2013). The experiment generated exogenous variation in expenditure, by providing more or less insurance coverage to participating households. The experiment lasted for a period of up to 5 years, and health evaluations were conducted both before and after. Unfortunately, the public use files only contain information on mental rather than physical health, and that is one reason why our analysis is focused on the effect of mental health expenditures on mental health.

The other reason is that empirical studies on mental health expenditures are relatively scarce, despite their policy relevance. The OECD assesses that more spending on mental health (often outpatient psychotherapy) could massively increase productivity and well-being (OECD, 2014). The United Nations include improved mental care provision among the sustainable development goals (United Nations, 2015). Nevertheless, among the stream of papers that came out of the RHIE, only one dealt explicitly with mental health (Manning et al., 1986). A possible explanation is that mental spending in the RHIE accounted only for around 4% of overall spending. Regarding overall expenditures, earlier analyzes of the data suggested a strong “first-stage” effect (i.e., more generous insurance coverage increased expenditures) but no systematic “reduced-form” effect (of more generous insurance cover on actual health, or mental health, as in Manning et al. (1989)). A formal analysis of the direct effect of increased mental spending on mental health has, to the best of our knowledge, not yet been undertaken for these data.

The bivariate model we implement is recursive: it postulates a direct effect of the Tweedie distributed health expenditure variable on the Bernoulli outcome “mental health below x ”, but not vice-versa. Joint estimation using the copula approach allows for common, or correlated, unobserved components. Ignoring such confounders would bias the effect size. Direct reverse causation is unlikely in our application, as spending is measured in the years prior to the health outcome. To better achieve empirical identification, we instrument spending using RHIE’s random assignment of individuals to insurance plans with different levels of cost-sharing. Our approach is an alternative to ad hoc endogeneity corrections, such as plug-in or control function approaches, that are not designed for non-normal and non-linear reduced forms as evidently required for the mixed discrete-continuous variable “mental spending” (e.g., Rivers & Vuong, 1988; Wooldridge, 2010). Using the copula approach with appropriate margins for the binary mental health indicator and the semi-continuous mental spending variable, we address the endogeneity of spending in a model-consistent way, based on a specification of the joint probability function that could have generated the observed data.

Our main findings are as follows. The Tweedie margin fits the expenditure data well, as evidenced by approximately normally distributed quantile residuals. Among the eleven copulae and three binary link functions considered, the Gaussian copula with probit link exhibits the smallest Akaike information criterion (AIC) value. The instrument for mental spending (made up of several categories) is highly significant (p -value of 0.00027 using the Wald test). Substantively, we find that the probability of low mental health is strongly predicted by the initial mental health score. The probability is higher for women and for the less educated, although the latter effect is not statistically significant (p -value of 0.14). The estimated Gaussian copula parameter indicates positive dependence. Ignoring such dependence leads to a spurious estimate in the univariate model, where an increase in spending by \$1000 (or about 3500 inflation adjusted dollars in 2022) predicts a 8.3% points increase of the probability of having a mental health index below 50 (on a 0 – 100 scale). Once the endogeneity of spending is accounted for, the average partial effect (APE) switches sign. The point estimate of -1.9% points is of modest size and not statistically significant.

The paper is organized as follows. Section 2 introduces the Tweedie-Bernoulli copula model by discussing its components. Section 3 discusses the results of a simulation study. Section 4 presents the application whereas Section 5 concludes the paper. Appendix A provides a set of technical results that are required to implement the Tweedie distribution in a copula context, whereas Appendix B discusses parameter estimation and the related implementation in R.

2 | COPULA MODEL

The copula regression model introduced in this paper aims at modeling in a flexible way the joint distribution of a Bernoulli outcome variable and a semi-continuous (endogenous) variable. In the application, $Y_1 \in \{0, 1\}$ is the binary outcome “low mental health” and $Y_2 \in \mathbb{R}_0^+$ is mental spending. Assume that

$$F_{12}(y_1, y_2 | \boldsymbol{\theta}) = C(F_1(y_1 | \pi), F_2(y_2 | \mu, \sigma, \nu); \theta), \quad (1)$$

where $\boldsymbol{\theta} = (\pi, \mu, \sigma, \nu, \theta)'$, $F_1(y_1 | \pi)$ and $F_2(y_2 | \mu, \sigma, \nu)$ represent the marginal cumulative distribution functions (cdf) of Y_1 and Y_2 taking values in $(0, 1)$, $C : (0, 1)^2 \rightarrow (0, 1)$ is a two-place copula function whose specification does not depend on the marginals, and θ is a copula dependence parameter that quantifies the association between the two random variables (see, e.g., Nelsen, 2006, for further details). Variable Y_1 is modeled via a Bernoulli distribution with parameter $\pi \in [0, 1]$ (representing the probability that the outcome is equal to 1), and Y_2 using a Tweedie distribution with parameters μ, σ and ν (see Section 2.1). Note that π, μ, σ, ν and θ can be specified as functions of covariate effects as detailed in Section 2.2. Hence, the marginal model for Y_1 can be regarded as a generalized additive model and that for Y_2 as a generalized additive model for location scale and shape (e.g., Rigby & Stasinopoulos, 2005; Wood, 2017). Function C can be specified as in tab. 1 of Marra et al. (2020) which reports several copulae.

As opposed to classical copula regression settings, variable Y_2 appears as an explanatory variable in the equation for π , hence giving the model a recursive structure. This in turn implies that Y_2 is endogenous with respect to Y_1 if the dependence between the two marginals (captured by θ) is statistically significant; see Han and Vytlačil (2017), Marra et al. (2020) and references therein for some works which have adopted the same logic for copula regression models. Note that, for the model adopted in this paper, Sklar (1973)'s result can only guarantee that the copula is unique over the range of the outcomes. However, for applied purposes, in a regression context, this is not problematic as noted by several authors including Joe (2014), Nikoloulopoulos and Karlis (2010) and Trivedi and Zimmer (2017).

The joint density $f_{12}(y_1, y_2)$, required for calculating the model's log-likelihood, is built by considering the four possible combinations of values that $(Y_1, Y_2)'$ can take. That is,

$$f_{12}(y_1, y_2) = \begin{cases} C(F_1(0 | \pi), F_2(0 | \mu, \sigma, \nu); \theta) & \text{if } y_1 = 0 \text{ and } y_2 = 0 \\ F_2(0 | \mu, \sigma, \nu) - C(F_1(0 | \pi), F_2(0 | \mu, \sigma, \nu); \theta) & \text{if } y_1 = 1 \text{ and } y_2 = 0 \\ f_2(y_2 | \mu, \sigma, \nu) \frac{\partial C(F_1(0 | \pi), F_2(y_2 | \mu, \sigma, \nu); \theta)}{\partial F_2(y_2 | \mu, \sigma, \nu)} & \text{if } y_1 = 0 \text{ and } y_2 > 0 \\ f_2(y_2 | \mu, \sigma, \nu) \left\{ 1 - \frac{\partial C(F_1(0 | \pi), F_2(y_2 | \mu, \sigma, \nu); \theta)}{\partial F_2(y_2 | \mu, \sigma, \nu)} \right\} & \text{if } y_1 = 1 \text{ and } y_2 > 0 \end{cases} \quad (2)$$

where $f_2(y_2 | \mu, \sigma, \nu) = \partial F_2(y_2 | \mu, \sigma, \nu) / \partial y_2$ denotes the density function of the Tweedie distribution. The first two lines of Equation (2) show the probabilities associated with their respective events, whereas the last two lines show the densities obtained by differencing the first two lines of Equation (2) with respect to y_2 when $y_2 > 0$. For notational convenience, we dropped observation index i from the formulae above. However, it should be clear from the context of the paper that a set of n observations is assumed to be available for practical modeling.

2.1 | Tweedie distribution

The Tweedie distribution is a linear exponential dispersion model (Jørgensen, 1987) with a power mean-variance relationship, that is

$$\text{var}(Y) = \sigma \mu^\nu,$$

where $\mu = \mathbb{E}(Y)$, $\sigma > 0$ is a scale parameter and $\nu \in \mathbb{R}$ controls the shape of the relationship. Widely used distributions such as the Gaussian, Poisson and Gamma are nested by the Tweedie family and can be recovered by setting ν to the relevant value (e.g., $\nu = 0$ for the Gaussian). In this work, we are interested in the interval $\nu \in (1, 2)$, for which a Tweedie-distributed random variable can be represented as the sum of N independent Gamma-distributed random variables, where N is Poisson distributed. The resulting density is supported on the non-negative real line and has a positive mass at $y = 0$. Prior applications of the Tweedie compound Poisson-Gamma distribution are mainly known from the actuarial sciences, where the distribution is used to model insurance claim payments data (see, e.g., Smyth & Jørgensen, 2002). Kurz (2017) provides an application to total medical expenditures.

Fitting reliably the proposed copula regression model involving a Tweedie margin, using the method mentioned in Appendix B, requires the ability to compute the Tweedie probability density function and cdf as well as their first and second order derivatives with respect to μ , σ and ν . Such quantities are not trivial to compute.

The density of the Tweedie is

$$f(y|\mu, \sigma, \nu) = a(y, \sigma, \nu) \exp\left[\frac{1}{\sigma}\{y\xi - \kappa(\xi)\}\right],$$

where

$$\xi = \frac{\mu^{1-\nu}}{1-\nu} \text{ for } \nu \neq 1 \text{ and } \xi = \log \mu \text{ for } \nu = 1,$$

and

$$\kappa(\xi) = \frac{\mu^{2-\nu}}{2-\nu} \text{ for } \nu \neq 2 \text{ and } \kappa(\xi) = \log \mu \text{ for } \nu = 2.$$

As explained in Dunn and Smyth (2005), evaluating the Tweedie density requires approximating the factor $a(y, \sigma, \nu)$, which does not have a closed-form expression, using specifically designed numerical methods. Wood et al. (2017) provide methods for computing the first and second order derivatives of the log-density, while avoiding numerical problems. In Appendix A, we explain how to compute the Tweedie cdf and its derivatives.

2.2 | Additive predictor

Each of the parameters in $\boldsymbol{\theta}$ is linked to covariate effects via an additive predictor $\eta \in \mathbb{R}$ and a known monotonic one-to-one transformation function g that maps the parameter space to the real line. For the proposed model, we have $g_\pi(\pi) = \eta_\pi$, $g_\mu(\mu) = \eta_\mu$, $g_\sigma(\sigma) = \eta_\sigma$, $g_\nu(\nu) = \eta_\nu$ and $g_\theta(\theta) = \eta_\theta$, where $g_\pi(\cdot)$ can be specified using a logit, probit or cloglog link function, $g_\mu(\mu) = \log(\mu)$, $g_\sigma(\sigma) = \log(\sigma)$, $g_\nu(\nu) = \log((\nu - 1.001)/(1.999 - \nu))$, and g_θ depends on the chosen copula (e.g., for Gumbel $g_\theta(\theta) = \log(\theta - 1)$). Dropping for simplicity the subscript denoting the parameter the additive predictor belongs to, η_i can be written as

$$\eta_i = \beta_0 + \sum_{k=1}^K s_k(\mathbf{z}_{ki}), \quad i = 1, \dots, n, \quad (3)$$

where $\beta_0 \in \mathbb{R}$ is an overall intercept, \mathbf{z}_{ki} denotes the k th sub-vector of the complete covariate vector \mathbf{z}_i (potentially containing various types of covariates) and the K functions $s_k(\mathbf{z}_{ki})$ represent generic effects which are chosen according to the type of covariate(s) considered. Each $s_k(\mathbf{z}_{ki})$ can be approximated as a linear combination of J_k basis functions $b_{kj_k}(\mathbf{z}_{ki})$ and regression coefficients $\beta_{kj_k} \in \mathbb{R}$, that is, (e.g., Wood, 2017)

$$s_k(\mathbf{z}_{ki}) = \sum_{j_k=1}^{J_k} \beta_{kj_k} b_{kj_k}(\mathbf{z}_{ki}). \quad (4)$$

This formulation implies that the vector of evaluations $\{s_k(\mathbf{z}_{k1}), \dots, s_k(\mathbf{z}_{kn})\}'$ can be written as $\mathbf{Z}_k \boldsymbol{\beta}_k$ with $\boldsymbol{\beta}_k = (\beta_{k1}, \dots, \beta_{kJ_k})'$ and design matrix $Z_k[i, j_k] = b_{kj_k}(\mathbf{z}_{ki})$. This allows Equation (3) to be written as

$$\boldsymbol{\eta} = \beta_0 \mathbf{1}_n + \mathbf{Z}_1 \boldsymbol{\beta}_1 + \dots + \mathbf{Z}_K \boldsymbol{\beta}_K, \quad (5)$$

where $\mathbf{1}_n$ is an n -dimensional vector made up of ones. Equation (5) can also be written as $\boldsymbol{\eta} = \mathbf{Z}\boldsymbol{\beta}$, where $\mathbf{Z} = (\mathbf{1}_n, \mathbf{Z}_1, \dots, \mathbf{Z}_K)$ and $\boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}'_1, \dots, \boldsymbol{\beta}'_K)'$.

Each $\boldsymbol{\beta}_k$ has an associated quadratic penalty $\lambda_k \boldsymbol{\beta}'_k \mathbf{D}_k \boldsymbol{\beta}_k$ whose role is to enforce specific properties on the k th function, such as smoothness. Matrix \mathbf{D}_k only depends on the choice of basis functions. Smoothing parameter $\lambda_k \in [0, \infty)$ controls the trade-off between fit and smoothness, and plays a crucial role in determining the shape of $\hat{s}_k(\mathbf{z}_{ki})$. The overall penalty can be defined as $\boldsymbol{\beta}' \mathbf{D}_\lambda \boldsymbol{\beta}$, where $\mathbf{D}_\lambda = \text{diag}(0, \lambda_1 \mathbf{D}_1, \dots, \lambda_K \mathbf{D}_K)$. Note that the first term in \mathbf{D}_λ is set to 0 since β_0 is not penalized in estimation. However, it could be replaced with $\mathbf{0}$ should some parameter vector(s) in $\boldsymbol{\beta}$ also be unpenalized. Finally, the smooth functions are subject to centering (identifiability) constraints, which impose that $\sum_{i=1}^n s_k(\mathbf{z}_{ki}) = 0$ for every k (see Wood (2017) for more details). The above smooth function representation allows one to specify a rich variety of covariate effects (such as linear, nonlinear and spatial Markov random field effects) and we refer the reader to Wood (2017) for full details. The overall penalty is used during model fitting as explained in Appendix B.

Consider the compact form for the random vectors $\mathbf{Y}_1 = (Y_{11}, \dots, Y_{1n})'$ and $\mathbf{Y}_2 = (Y_{21}, \dots, Y_{2n})'$. Then, by some slight abuse of notation, $\mathbf{Y}_1 \sim \mathcal{D}_1(\boldsymbol{\pi})$ and $\mathbf{Y}_2 \sim \mathcal{D}_2(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\nu})$, where \mathcal{D}_1 and \mathcal{D}_2 denote the Bernoulli and Tweedie distributions, respectively, and $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)'$, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)'$, $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)'$ and $\boldsymbol{\nu} = (\nu_1, \dots, \nu_n)'$ are modeled through

$$\begin{aligned} \boldsymbol{\eta}_\pi &= g_\pi(\boldsymbol{\pi}) = \beta_{10} \mathbf{1}_n + \boldsymbol{\beta}_{\text{end}} \mathbf{Y}_2 + \mathbf{Z}_{11} \boldsymbol{\beta}_{11} + \dots + \mathbf{Z}_{1K} \boldsymbol{\beta}_{1K}, \\ \boldsymbol{\eta}_\mu &= g_\mu(\boldsymbol{\mu}) = \beta_{20} \mathbf{1}_n + \mathbf{Z}_{21} \boldsymbol{\beta}_{21} + \dots + \mathbf{Z}_{2K} \boldsymbol{\beta}_{2K}, \\ \boldsymbol{\eta}_\sigma &= g_\sigma(\boldsymbol{\sigma}) = \beta_{30} \mathbf{1}_n + \mathbf{Z}_{31} \boldsymbol{\beta}_{31} + \dots + \mathbf{Z}_{3K} \boldsymbol{\beta}_{3K}, \\ \boldsymbol{\eta}_\nu &= g_\nu(\boldsymbol{\nu}) = \beta_{40} \mathbf{1}_n + \mathbf{Z}_{41} \boldsymbol{\beta}_{41} + \dots + \mathbf{Z}_{4K} \boldsymbol{\beta}_{4K}, \end{aligned}$$

where the functions g are applied element-wise. The same specification can also be adopted for the dependence parameter vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)'$ in which case $\boldsymbol{\eta}_\theta = g_\theta(\boldsymbol{\theta}) = \beta_{50} \mathbf{1}_n + \mathbf{Z}_{51} \boldsymbol{\beta}_{51} + \dots + \mathbf{Z}_{5K} \boldsymbol{\beta}_{5K}$. Term $\boldsymbol{\beta}_{\text{end}} \mathbf{Y}_2$, the endogenous effect of the semi-continuous variable on the binary response, modeled using representation Equation (4) as for any other covariate in the model, enters equation $\boldsymbol{\eta}_\pi$, hence giving the model setup a recursive structure as explained in Section 2. In terms of specification, covariates might be common across the additive predictors, except for the equation(s) related to the endogenous variable which should include at least one variable (an instrument) that is not included in the other predictors (e.g., Han & Vytlačil, 2017). Finally, note that not all parameters have to be specified as functions of predictors.

3 | SIMULATION STUDY

This section provides simulation evidence on the practical performance of the proposed approach. The data generating process (DGP) has been designed to mimic some of the features of the data as well as results from the case study that follows in the next section. For instance, the distributions of the binary and semi-continuous variables have been simulated to look similar to their observed versions, and the values for $\boldsymbol{\sigma}$, $\boldsymbol{\nu}$ and $\boldsymbol{\theta}$ chosen to be close to the estimates obtained from the empirical analysis. The DGP has the form

$$\begin{aligned} Y_2 &\sim \text{TW with } \mu = g_\mu^{-1}(1 + s_2(z_1) + s_3(z_2) + z_3), \sigma = g_\sigma^{-1}(4.24), \nu = g_\nu^{-1}(0.22), \\ Y_1 &\sim \text{Bernoulli with } \pi = \Phi\{-3 + \beta_{\text{end}} y_2 + s_1(z_1) + z_2\}, \end{aligned}$$

where $s_1(z) = z + \exp\{-30(z - 0.5)^2\}$, $s_2(z) = 0.6\{\exp(z) + \sin(2.9z)\}$ and $s_3(z) = 0.6 \sin(2\pi z)$, and the inverses of the g functions can be worked out from the expressions reported in Section 2.2 of the paper. Parameters β_{12} and β_{21} are set to 1 and hence not displayed in the equations above. As an empirically relevant benchmark, β_{end} is set to 0. Variables z_1 , z_2 and z_3 (the instrument) are generated using a multivariate standard Gaussian with correlation parameters set at 0.5, and then transformed using the distribution function of a standard Gaussian. Regressor z_3 is dichotomised by rounding it. Associated responses are generated via function `BiCopSim()`, from the package `VineCopula` (Schepsmeier et al., 2019), using the Gaussian copula with dependence parameter specified as $\theta = \tanh(0.41)$. We considered sample sizes of 1500 and 3000, while the number of replicates was set to 500.

Using `gjrm()` we fitted models with `probit` and TW marginals, and a handful of copulae available in the package, specifically N, C0, J180, G0, G180, FGM, F and PL. We also fitted a univariate model with TW distribution (using `gamlss()` within `GJRM`), the Gaussian copula model with `probit` and TW marginals that does not account for covariate

z_2 , the Gaussian copula model with `probit` and Gaussian marginals, and that with `probit` and Gamma marginals. We also considered the two-stage residual inclusion (2SRI) approach where the first stage involved a Tweedie regression for the endogenous variable and the second step a probit regression for the binary outcome (e.g., Terza et al., 2008). The correct model is the one based on the N copula with `probit` and TW marginals and all covariates in the model, whereas all the others served to assess the impacts that various misspecifications have on the endogenous effect of interest. The smooth components in the models were represented using penalized low rank thin plate splines with second order penalty and 10 bases (Wood, 2017). For each replicate, curve estimates were constructed using 200 equally spaced fixed values in the (0, 1) range.

Figure 1 shows the estimates for the endogenous parameter of interest β_{end} under several fitted models. As expected, the simple univariate model, which does not correct for endogeneity, is clearly not able to capture the true parameter value; the estimates are upwardly biased, reflecting the positive copula dependence. When the copula is misspecified the results show some non-vanishing bias, hence suggesting that mis-modeled dependence structure plays a role in the estimation of the effect of interest. Misspecifying the regression structure does not seem to affect the effect on interest, while using the incorrect marginal distribution for the semi-continuous outcome has a detrimental impact on the estimation of β_{end} in terms of bias and variance. The 2SRI approach yields rather unreliable estimates; this was somewhat expected since, with such high share of zeros in the semi-continuous outcome, the residuals obtained from the first stage may be a too crude approximation of unobserved factors. The results obtained under the correct model (N) exhibit the best bias-variance tradeoff. We also report, for the correct model, the results for the other model's terms, namely β_{12} , β_{21} , σ , ν , θ , $s_1(\cdot)$, $s_2(\cdot)$ and $s_3(\cdot)$; these are displayed in Figures 2 and 3. Although these are not of primary interest, the findings further support the empirical effectiveness of the proposed approach since the true values are overall well recovered.

4 | APPLICATION TO MENTAL CARE SPENDING AND HEALTH

The empirical literature on the effect of health insurance on utilization and health is decades old, substantial, and still growing. The most recent example for the U.S. are numerous studies evaluating the Affordable Care Act of 2010 (see, e.g., Mazurenko et al. (2018) for a systematic review). In that reform, Medicaid coverage was extended to nonelderly adults with incomes up to 138% of the federal poverty level. After a subsequent Supreme Court ruling, states were allowed to opt out, and by 2017, eighteen states had decided to do so. Evaluations of that policy therefore mostly use a difference-in-differences design, comparing trends for treated and non-treated adults depending on the state they live in.

Earlier studies of the topic often employed data from two genuine large scale experiments, the 2008 Oregon Health Insurance Experiment and the 1974 RHIE. The RHIE data contain detailed information on medical spending, and they are freely available as public use files from <https://www.icpsr.umich.edu/icpsrweb/NACDA/studies/06439>; see Deb and Trivedi (2002) or Aron-Dine et al. (2013) for a more detailed description of data and variables.

Kurz (2017) reports for the RHIE data that the Tweedie model provides a better fit for overall spending than other single-equation models, including the Tobit. Here, we focus on the bivariate relation between mental care spending and mental health. The RHIE provides a mental health score both at baseline and on exit from the program. The exit score will be our outcome variable. In addition, mental health expenditures are listed as a separate cost category in the data. Average spending on mental health is low, mostly due to the high share of non-users (90% in our sample). The average annual amount spent among users is \$236. While the high share of zeros can be a problem for traditional modeling approaches, it is compatible with the compound Poisson-Gamma framework of the Tweedie distribution.

Random plan-assignment is used as an instrument for spending on mental care. There were a total of 14 different insurance plans that can be classified into four types of coverage: full insurance; individual deductible plans; coinsurance plans; and catastrophic coverage only. We find evidence for instrument relevance: spending is higher for plans with more generous coverage. This comes with a caveat; Aron-Dine et al. (2013) provide some evidence that attrition from the experiment is more prevalent in the catastrophic coverage group, and there is possibly some selection based on better health, indicating a violation of the usual exclusion restriction. In our set-up, we are able to condition on mental health at baseline, and additionally allow for associated confounders in the bivariate model, thereby alleviating concerns that non-random attrition may invalidate our results.

Thus, we estimate a recursive copula two-equation model for spending (a semi-continuous variable with a mass point at zero) and a binary indicator of later mental health. The dependence between spending and the distribution of mental health stemming from associated unobservables is modeled using copulae. The benchmark specification combines Tweedie marginal models with the probit Bernoulli margin, using a Gaussian copula. Other variants are explored. The effects of continuous covariates are modeled via spline functions. Our application also provides a template for related analyses, as the same

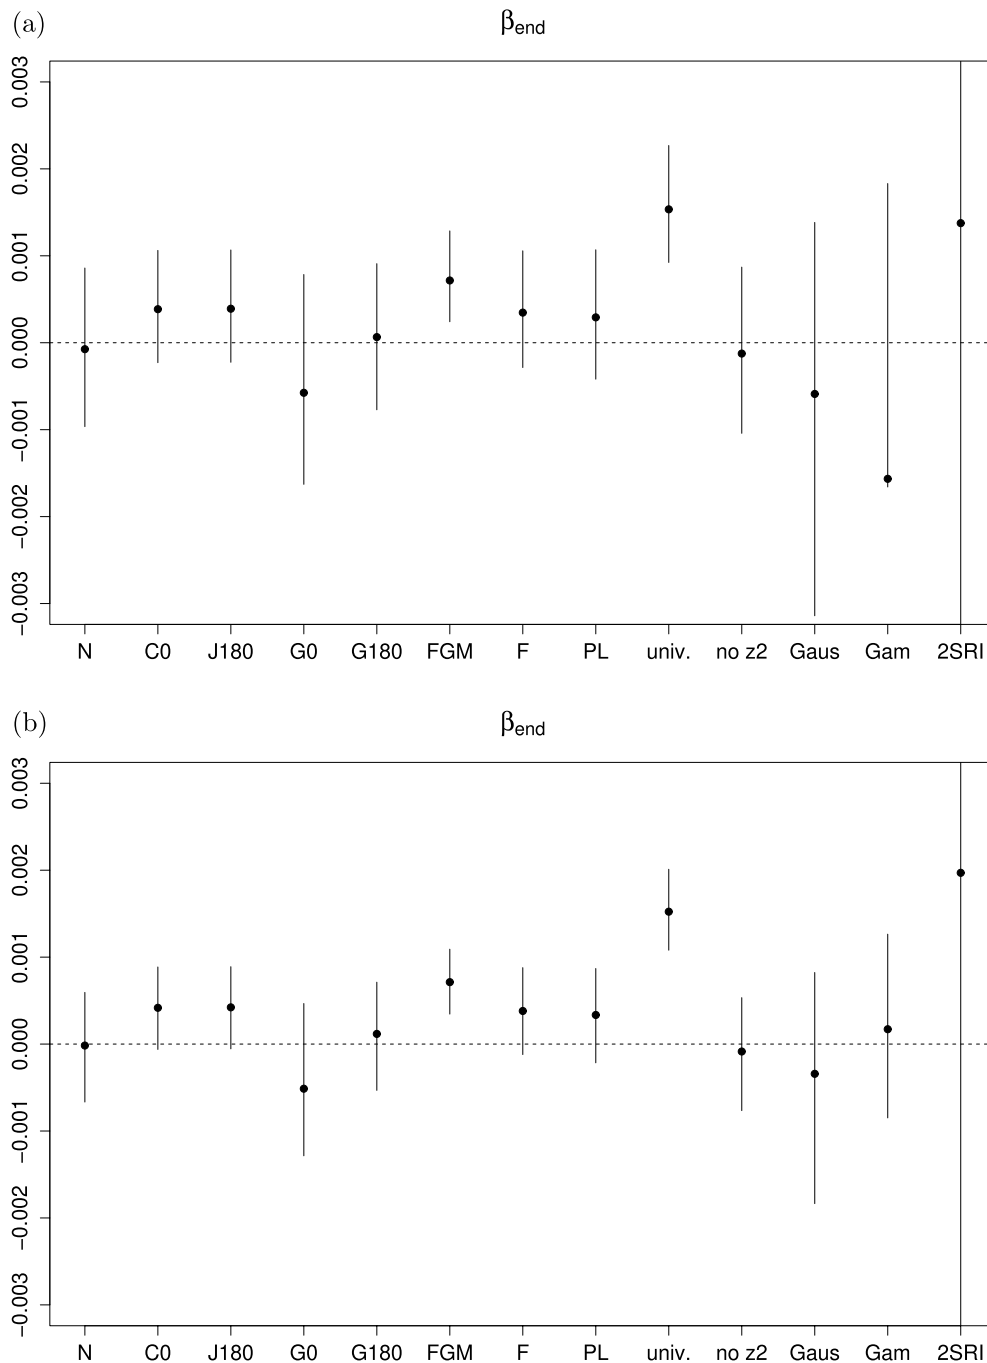


FIGURE 1 Estimates for β_{end} (the endogenous parameter of interest) obtained via GJRM by applying copula and univariate models to simulated data based on the Gaussian copula with probit and Tweedie margins. The correct model is denoted by N, while the other models are misspecified in various ways. Specifically, we have additionally considered (i) Seven more copula models (C0, J180, G0, G180, FGM, F, PL) with `probit` and `TW` margins, (ii) A univariate model with `TW` distribution, (iii) The Gaussian copula model with `probit` and `TW` marginals but without covariate z_2 , (iii) The Gaussian copula model with `probit` and Gaussian marginals, (iv) The Gaussian copula model with `probit` and Gamma marginals and (v) The two-stage residual inclusion approach. Circles indicate mean estimates while bars represent the estimates' ranges resulting from 5% to 95% quantiles. The true value is 0 and is denoted by the dashed horizontal line. (a) $n = 1500$. (b) $n = 3000$.

methodology can be used to estimate the effect of any type of health care spending on any type of binary health outcome. In addition, the order of the equations can be reversed to accommodate a situation of a Tweedie model with binary endogenous regressor.

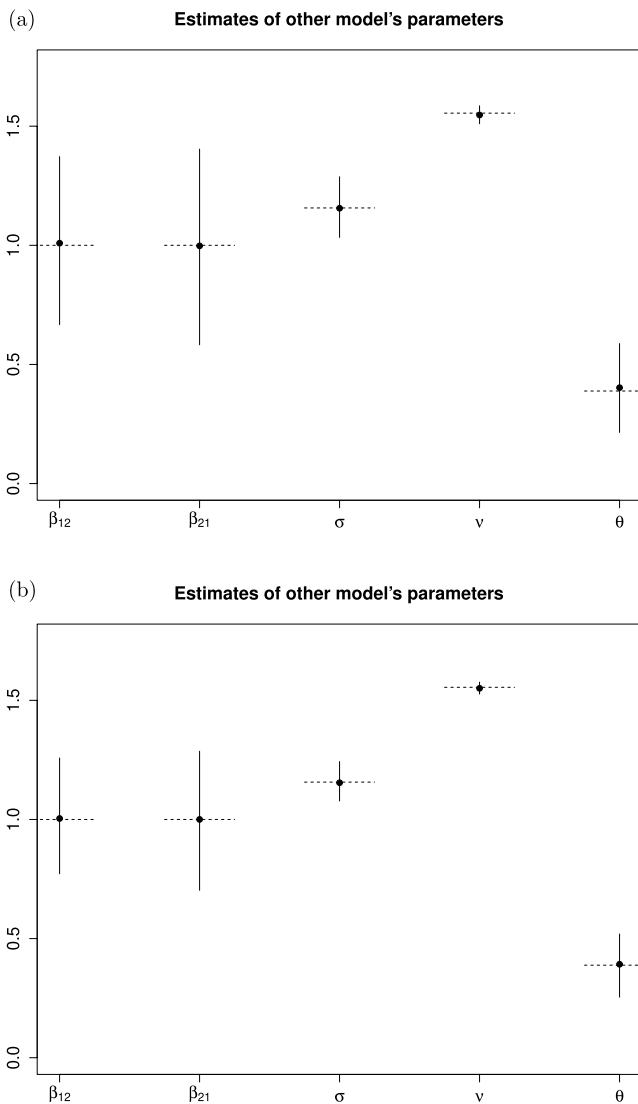


FIGURE 2 Estimates of other model's parameters obtained by applying `gjrmm()` to simulated data based on a Gaussian copula with probit and Tweedie margins when fitting correctly specified models. Circles indicate mean estimates while bars represent the estimates' ranges resulting from 5% to 95% quantiles. True values are indicated by dashed horizontal lines. (a) $n = 1500$. (b) $n = 3000$.

4.1 | Description of data and model

To obtain our final sample, we excluded all individuals under the age of 18, as well as those with missing information on any of the variables employed in the analysis. This leaves us with 2777 observations. Of those, 69% were enrolled in the program for 3 years and the remaining 31% for 5 years. The spending variable is obtained as the average yearly spending for mental care during the enrollment period (`mentalexp`).

Mental health, `mhix`, was evaluated upon exit from the program, on a 0–100 scale with mean 76; by dichotomizing `mhix` at a given point x , we obtain the Bernoulli outcome “`mhix < x`”. Thus, we effectively estimate a model for the distribution function $F(x) = P(\text{mhix} < x)$ where x can be varied. In our application, we let $x \in \{30, 40, 50, 60, 70, 80, 90\}$, and thus trace out the distribution function of `mhix` at these points. This is in the spirit of Chernozhukov et al. (2013) who study the use of distribution regressions as a flexible approach for estimating treatment effects. In order to save space, we focus our reporting on the cut-off of $x = 50$ and label all individuals falling below that threshold as having “low mental health” (`lowmhix`). This applies to 5% of the sample. We summarize the effect estimates at other thresholds in a single graph.

For the Tweedie model, only the location parameter μ is specified as a function of covariates; modeling σ and ν using additive predictors does not alter the substantive conclusions, hence we focus on the location parameter only. Therefore, the model equations for the Bernoulli and Tweedie margins are

$$\pi_{\text{lowmhix}} = g_{\pi}^{-1}(\beta_{10} + s_{11}(\text{mentalexp}) + \beta_{12}\text{female} + \beta_{13}\text{white} + \beta_{14}\text{poor} + s_{12}(\text{mhi}) + s_{13}(\text{age}) + s_{14}(\text{educ}))$$

and

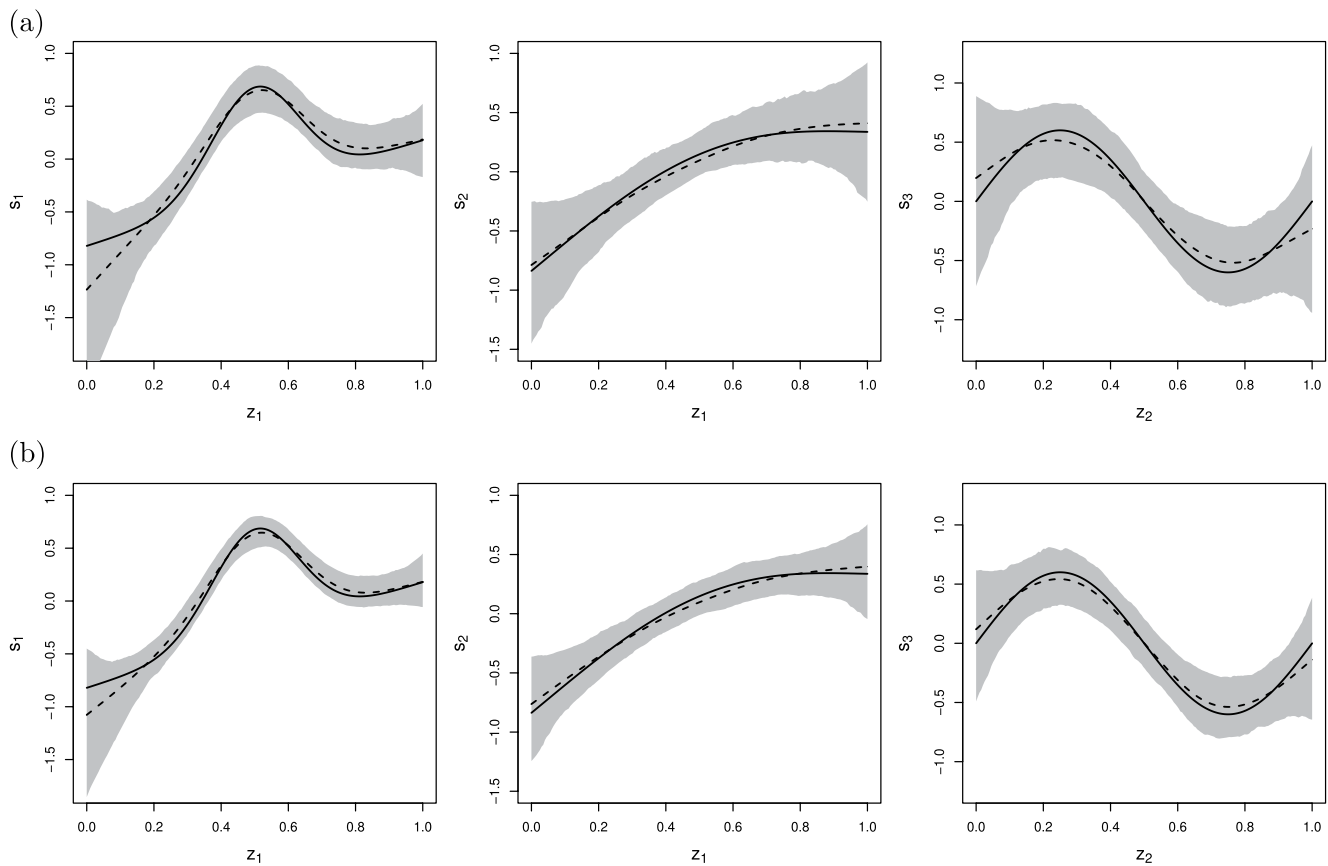


FIGURE 3 Estimates of smooth effects obtained by applying $g_{jrm}()$ to simulated data based on a Gaussian copula with probit and Tweedie margins when fitting correctly specified models. True functions are represented by black solid lines, mean estimates by dashed lines and point-wise ranges resulting from 5% to 95% quantiles by shaded areas. (a) $n = 1500$. (b) $n = 3000$.

$$\begin{aligned} \mu_{\text{mentalexp}} = & \exp(\beta_{20} + \beta_{21}\text{female} + \beta_{22}\text{white} + \beta_{23}\text{poor} + \beta_{24}\mathbb{1}(\text{plantype} = 2) + \beta_{25}\mathbb{1}(\text{plantype} = 3) \\ & + \beta_{26}\mathbb{1}(\text{plantype} = 4) + s_{21}(\text{mhi}) + s_{22}(\text{age}) + s_{23}(\text{educ})) \end{aligned}$$

were π is the binary response probability, g_{π}^{-1} is derived according to the link function chosen for π , and μ is the expected value of mental expenditure. Thus, both equations have six variables in common, namely the mental health score at enrollment (`mhi`), age, gender (`female = 1`), race (`white = 1`), the education of the household head in years (`educ`), as well as the indicator variable `poor`. Participants are classified as being poor if the per-capita household income falls into the lower tercile of the distribution.

The factor variable `plantype` is an instrument, that is, it is included in the second equation but excluded from the first equation. In our data, 33% of observations are assigned to the first, full coverage plan; 22% to the deductible plan; 27% to the coinsurance plan, and the remaining 18% to the catastrophic coverage plan. F -tests for balancing of the pre-treatment variables by plan type do not reject the null hypothesis of no difference (p -values between 0.16 and 0.64) except for the variable `poor`, which is underrepresented in the catastrophic plan, a finding echoing Aron-Dine et al. (2013).

4.2 | Results for the univariate models

Table 1 shows estimation results obtained using the univariate approach where the Bernoulli and Tweedie marginals have been estimated separately, as single equation models. For the Bernoulli margin, using information criteria, the probit link was selected.

The Tweedie coefficients determine the relative change in expected spending associated with an increase in the covariate value. For example, women are predicted to spend $[\exp(0.5692) - 1] \times 100 = 76.7\%$ more than men, keeping everything else

TABLE 1 Results from single equation models ($n = 2777$). The first two columns refer to the Tweedie margin and the last two to the Bernoulli (with probit link) margin. As for the smooth functions, we report effective degrees of freedoms (representing the degrees of complexity of the estimated curves) and p -values. 95% intervals for σ and ν are obtained by posterior simulation.

Dependent variable	mentalexp (Tweedie)		lowmhix (Bernoulli with probit link)	
	Estimate	Std. Error	Estimate	Std. Error
(Intercept)	0.2965	0.4213	-2.1539	0.1528
mentalexp $\times 10^3$			0.6444	0.2333
female	0.5692	0.1881	0.3157	0.0991
white	1.9382	0.4009	0.0717	0.1330
poor	0.4539	0.2114	-0.0025	0.1023
plantype2	0.1550	0.2279		
plantype3	-0.6482	0.2403		
plantype4	-0.8792	0.2873		
s (mhi)	1.611	$p\text{-val} < 2e-16$	2.051	$p\text{-val} < 2e-16$
s (age)	6.389	$p\text{-val} = 0.028$	1.000	$p\text{-val} = 0.551$
s (educ)	6.059	$p\text{-val} < 2e-16$	1.002	$p\text{-val} = 0.022$
σ	69.9	(63.8, 76.9)		
ν	1.55	(1.52, 1.58)		

constant. The omitted `plantype1` is “full insurance”, so `plantype4` (catastrophic insurance) cuts spending on mental care relative to full insurance by approximately 60% ($\exp(-0.8792) - 1 = -0.5849$). The `plantype` variable as a whole has a significant impact on the response (p -value = 0.00027 obtained using the Wald test), as are all the covariates as well as the three spline functions for `mhi`, `age` and `educ`.

Significant predictors of `lowmhix` are `female`, initial `mhi` as well as `educ`. Age, race and poverty status seem not to matter. Note that we initially considered a smooth function of `mentalexp`, that is $s_{11}(\text{mentalexp})$. Since this resulted in a straight line estimate, we eventually let the variable enter the equation parametrically. Spending on mental care has a positive coefficient and a z -value of 2.762. Thus, we find indeed that, in this univariate model, higher spending is associated with an increased probability of low mental health. Specifically, the approximate probit APE predicts a 8.3% points increase in low mental health per \$1000 increase in spending. A 95% interval for the APE obtained by Bayesian posterior simulation (based on the result mentioned in Appendix B) is given by [0.017, 0.189], confirming the statistical significance of the effect. This counterintuitive finding likely reflects associated “shocks”. During the enrollment period of three or 5 years, individual unobserved traits related to mental problems can drive up both `mentalexp` as well as the probability of low mental health upon exit, hence generating a spurious positive association in the univariate model. This is exactly the type of problem that estimation of the joint model can address.

We tested for interactions between `poor` and `plantype`, but found no significant evidence for heterogeneous effects. For the sake of completeness, we report in Figure 4 the histogram as well as the QQ-plot for the Tweedie quantile residuals, the latter together with 95% confidence bands. Almost all points are within the intervals, hence supporting the fact that the Tweedie distribution provides an appropriate characterization for mental spending.

4.3 | Results for the bivariate model

The great advantage of the copula approach in conjunction with flexible model margins is the modularity that gives the practitioner the opportunity to explore many unique combinations of elements in order to determine, on one hand, the best fitting model, and, on the other, the sensitivity of key results of interest to modeling assumptions. In our case, we take the Tweedie margin as well as the set of covariates as given. In contrast, we consider three link functions for the Bernoulli `lowmhix` equation (`probit`, `logit` and `cloglog`) and eleven different copulae to account for the dependence between the two equations. The copula models were based on the equations shown in the previous section and were fitted using the `GJRM` R package discussed in Appendix B. Specifying the dependence parameter as function of covariates did not, in this case, return any interesting results as all covariate effects were insignificant. Note that we used only the rotated 180° versions of the Clayton, Gumbel and Joe copulas because the data support the presence of positive dependence between the two equations, hence it would not make sense to consider rotations allowing for negative dependence. We employed the AIC and Bayesian information criterion (BIC) in order to find the best fitting model given the covariates and the Tweedie margin for spending. Results are provided in Table 2.

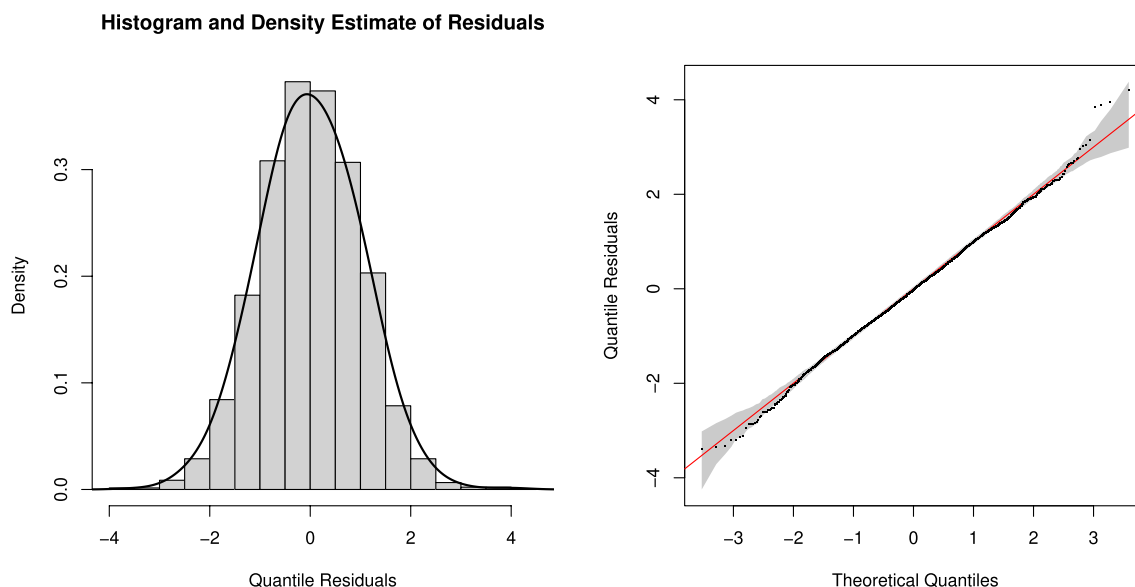


FIGURE 4 Histogram and normal Q–Q plot (with 95% confidence bands) of normalized quantile residuals for the mental spending variable produced after fitting a Tweedie univariate regression model.

The information criteria select the same copula function, namely the Gaussian, but not the same link function for the binary response (`probit` using the AIC and `cloglog` using the BIC), although the substantive conclusions are robust with respect to the selected copula and link function (see last column of Table 2). The full results of the joint Gaussian copula model with Bernoulli (with `probit` link) and Tweedie margins are given in Table 3, whereas the three estimated smooth functions (for `mhi`, `age` and `educ`) are shown in Figure 5, panel for the mental health equation).

We find that better initial mental health is associated with lower spending as well as a reduced probability of low mental health on exit from the program. The estimated smooth functions are nearly linear in both cases. The positive and significant association indicates that controlling for initial mental health takes out at least some of the individual heterogeneity that drives both spending and subsequent mental health. The effect of `age` is again similar in both equations, displaying a pronounced M-shape with estimated peaks in the late twenties and mid-fifties. The effect of education is less uniform and only imprecisely determined, especially for the spending equation.

The main result of interest is the effect of spending on mental health. The point estimate in the probit equation of the bivariate recursive copula model is given by -0.2647 , with a standard error of 0.2821 . Allowing for endogeneity of spending overturns the earlier finding of a positive and statistically significant association from the single-equation approach. The point estimate as such corresponds to an APE of -0.019 (with a Bayesian 95% confidence set equal to $[-0.043, 0.032]$), slightly negative but statistically insignificant.

Figure 6 shows the APE, together with 95% confidence interval, obtained for different values of the cut-off point used for the mental health variable, namely $x \in \{30, 40, 50, 60, 70, 80, 90\}$. We see that all effects are negative, except for $x = 90$. An increase in expenditures shifts the distribution downward over much of the range, except for the right tail, where we estimate a decrease in the complementary probability. However, none of these changes is statistically significant.

Using the 2SRI approach, we obtained an APE of 0.383 with a very wide 95% confidence set equal to $[-0.046, 0.937]$, a positive value but statistically insignificant. However, as found in the simulation study, this method may not produce reliable estimates of the effect of interest.

Not finding an effect of spending on mental health may be surprising at first. However, it is compatible with findings in the related literature studying the effect of insurance coverage on mental health. For instance, Brook et al. (1983) demonstrate for the RHIE data that there is no significant effect of insurance status on mental health. Using a regression discontinuity design that exploits the onset of Medicare when turning 65, Rhodes (2018) finds no mental health effect of better insurance coverage either.

In summary, there is strong evidence that the single-equation negative association between mental spending on mental health is driven by the endogeneity of spending, for example, due to health “shocks” occurring after the initial mental health score is determined. A copula model can address this endogeneity problem in the context of joint estimation of two non-linear outcome equations. Doing so leads to a null-finding that mirrors the consensus of the previous RHIE literature: while insurance

TABLE 2 Akaike information criterion (AIC) and Bayesian information criterion (BIC) values by different copulae and link functions for the Bernoulli margin. The probit average partial effect (APE), with the corresponding 95% interval obtained by Bayesian posterior simulation, is reported for each model.

Link function	Copula	AIC	BIC	APE
probit	<i>N</i>	5891.504	6107.676	-0.019 (-0.043, 0.032)
logit	<i>N</i>	5895.133	6100.153	-0.006 (-0.017, 0.021)
cloglog	<i>N</i>	5896.859	6095.220	-0.004 (-0.013, 0.022)
probit	<i>C0</i>	5897.271	6114.967	-0.001 (-0.029, 0.052)
logit	<i>C0</i>	5900.432	6106.005	-0.001 (-0.012, 0.026)
cloglog	<i>C0</i>	5901.629	6103.694	-0.001 (-0.010, 0.019)
probit	<i>C180</i>	5895.136	6110.731	-0.018 (-0.045, 0.051)
logit	<i>C180</i>	5898.891	6102.091	-0.003 (-0.016, 0.040)
cloglog	<i>C180</i>	5900.447	6098.598	-0.001 (-0.011, 0.033)
probit	<i>J0</i>	5899.430	6114.035	-0.014 (-0.050, 0.102)
logit	<i>J0</i>	5903.010	6105.085	0.002 (-0.014, 0.065)
cloglog	<i>J0</i>	5904.366	6101.348	0.002 (-0.011, 0.039)
probit	<i>J180</i>	5897.499	6115.236	-0.001 (-0.029, 0.054)
logit	<i>J180</i>	5900.621	6107.059	-0.001 (-0.012, 0.027)
cloglog	<i>J180</i>	5901.842	6104.033	-0.001 (-0.010, 0.019)
probit	<i>G0</i>	5897.583	6112.597	-0.017 (-0.049, 0.083)
logit	<i>G0</i>	5901.302	6104.262	-0.001 (-0.016, 0.050)
cloglog	<i>G0</i>	5902.787	6100.008	0.001 (-0.011, 0.036)
probit	<i>G180</i>	5891.693	6108.226	-0.016 (-0.040, 0.035)
logit	<i>G180</i>	5895.074	6100.828	-0.006 (-0.016, 0.020)
cloglog	<i>G180</i>	5896.893	6095.932	-0.004 (-0.012, 0.018)
probit	<i>AMH</i>	5896.485	6106.838	0.012 (-0.018, 0.064)
logit	<i>AMH</i>	5899.617	6098.679	0.005 (-0.009, 0.033)
cloglog	<i>AMH</i>	5901.260	6096.626	0.003 (-0.007, 0.023)
probit	<i>FGM</i>	5895.972	6105.752	0.013 (-0.018, 0.064)
logit	<i>FGM</i>	5899.083	6097.999	0.005 (-0.009, 0.034)
cloglog	<i>FGM</i>	5900.974	6095.574	0.004 (-0.007, 0.024)
probit	<i>T</i>	5897.080	6113.935	-0.027 (-0.051, 0.034)
logit	<i>T</i>	5901.212	6108.210	-0.004 (-0.018, 0.069)
cloglog	<i>T</i>	5903.132	6099.329	0.001 (-0.013, 0.078)
probit	<i>F</i>	5893.310	6110.670	-0.007 (-0.034, 0.043)
logit	<i>F</i>	5896.291	6103.940	-0.003 (-0.014, 0.023)
cloglog	<i>F</i>	5898.574	6098.664	-0.002 (-0.011, 0.018)

generosity had a sizable effect on spending (of any type, not only mental care), its effects on health outcomes measured upon conclusion of the experimental period were small if any.

5 | CONCLUSIONS

This paper presents an estimation approach, based on copulae, that consistently estimates the effect of an endogenous semi-continuous regressor on a binary outcome. The approach offers a convenient test for the presence of endogeneity, via the copula dependence parameter.

In our application, we reconsider data from the RHIE and estimate the effect of medical expenditures on the binary variable “low mental health”, using random insurance plan assignment as an instrument for spending. The expenditure variable has around 90% of zeros and we find that the Tweedie distribution provides a good fit to the data. The Gaussian copula dependence parameter is positive and statistically significant, indicating the presence of factors unobserved to the analyst that increase both

TABLE 3 Results from the chosen copula model ($n = 2777$). 95% intervals for σ , ν , and θ are obtained by posterior simulation.

Dependent variable	mentalexp (Tweedie)		lowmhi x (Bernoulli with probit link)	
	Estimate	Std. Error	Estimate	Std. Error
(Intercept)	0.3344	0.4168	-2.1177	0.1520
mentalexp $\times 10^3$			-0.2647	0.2821
female	0.5793	0.1868	0.3298	0.0979
white	1.8936	0.3960	0.0492	0.1308
poor	0.4685	0.2096	0.0334	0.1004
plantype2	0.1614	0.2246		
plantype3	-0.6830	0.2365		
plantype4	-0.8356	0.2825		
s (mhi)	1.461	p-val < 2e-16	2.219	p-val < 2e-16
s (age)	6.270	p-val = 0.0461	4.458	p-val = 0.469
s (educ)	6.050	p-val < 2e-16	1.000	p-val = 0.144
σ	70	(63.5, 77.3)		
ν	1.55	(1.52, 1.59)		
θ	0.386	(0.229, 0.527)		

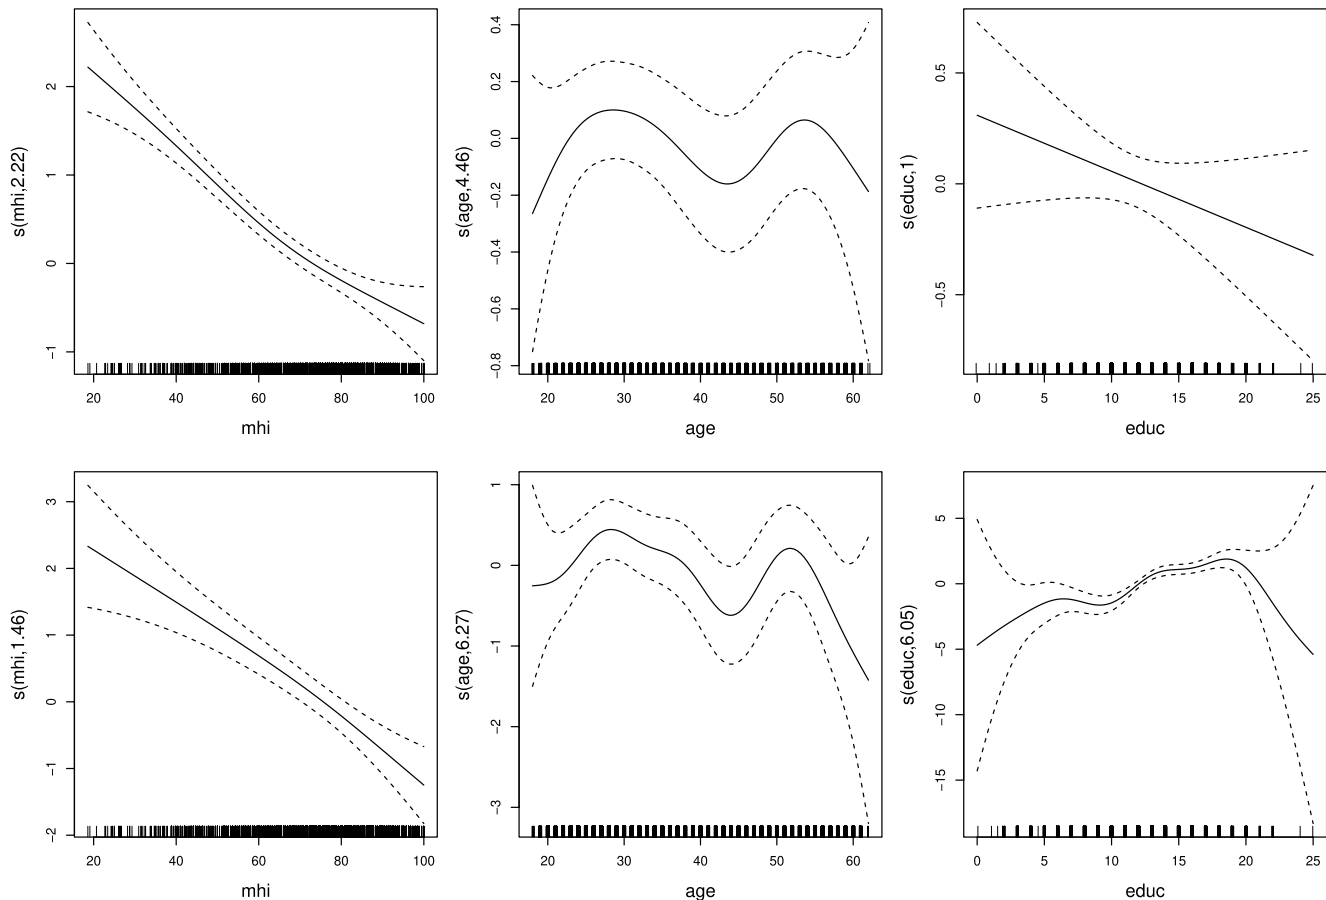


FIGURE 5 Estimated smooth effects of mhi, age and educ and associated 95% point-wise intervals obtained when fitting a Gaussian copula model with Bernoulli (with probit link) and Tweedie margins to data from the Rand Health Insurance Experiment (RHIE). The top plots refer to the probit equation whereas the bottom ones to the Tweedie equation. The number in brackets in the y-axis caption represents the effective degrees of freedom of the respective smooth curve. The jittered rug plot, at the bottom of each graph, shows the covariate values.




FIGURE 6 Estimated average partial effect (APE), with the corresponding 95% interval obtained by Bayesian posterior simulation, for different values of the cut-off point used for the mental health variable, namely $x \in \{30, 40, 50, 60, 70, 80, 90\}$.

spending and the probability of low mental health. An example would be the emergence of a mental health problem due to an adverse life event.

The methodology presented in this paper is fundamentally parametric and as such may suffer from the usual potential drawbacks resulting from model mis-specifications. Indeed, we provide evidence, by way of a simulation study, that the choice of copula and marginal distributions matters: mis-specification of either has detrimental effects on bias and variance of the parameter of interest.

Therefore, our modeling framework has the advantage of enabling model exploration via the various functional forms available in GJRM, which provides a wide set of copulae and marginal distributions, as well as much flexibility in the way regressor effects are specified. This helps researchers to explore several permutations of functional forms and regressor effects, and to select and test for the best model. Arguably, such exploration is consistent with the philosophy of non-parametric methods, in that it lets the data to point to meaningful model structures.

ACKNOWLEDGEMENT

We benefited from helpful comments by two anonymous reviewers. The first and third authors were supported by the Engineering and Physical Sciences Research Council [EP/T033061/1].

Open access funding provided by Universitat Zurich.

CONFLICT OF INTEREST STATEMENT

Drs. Marra, Fasiolo, Radice and Winkelmann have nothing to disclose.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in the National Archive of Computerized Data on Aging at <https://doi.org/10.3886/ICPSR06439.v1>, reference number ICPSR 6439.

ORCID

Rainer Winkelmann  <https://orcid.org/0000-0002-8686-5560>

REFERENCES

- Aron-Dine, A., Einav, L., & Finkelstein, A. (2013). The RAND health insurance experiment, three decades later. *The Journal of Economic Perspectives*, 27(1), 197–222. <https://doi.org/10.1257/jep.27.1.197>
- Brook, R. H., Ware, J. E., Rogers, W. H., Keeler, E. B., Davies, A. R., Donald, C. A., Goldberg, G. A., Lohr, K. N., Masthay, P. C., & Newhouse, J. P. (1983). Does free care improve adults' health? Results from a randomized controlled trial. *New England Journal of Medicine*, 309(23), 1426–1434. <https://doi.org/10.1056/nejm198312083092305>
- Chernozhukov, V., Fernandez-Val, I., & Melly, B. (2013). Inference on counterfactual distributions. *Econometrica*, 81(6), 2205–2268.
- Collett, D. (2002). *Modelling binary data* (2nd ed.). Chapman and Hall/CRC.
- Deb, P., & Trivedi, P. (2006). Specification and simulated likelihood estimation of a non-normal treatment-outcome model with selection: Application to health care utilization. *The Econometrics Journal*, 9(2), 307–331. <https://doi.org/10.1111/j.1368-423x.2006.00187.x>
- Deb, P., & Trivedi, P. K. (2002). The structure of demand for health care: Latent class versus two-part models. *Journal of Health Economics*, 21(4), 601–625. [https://doi.org/10.1016/s0167-6296\(02\)00008-5](https://doi.org/10.1016/s0167-6296(02)00008-5)
- Dunn, P. K., & Smyth, G. K. (1996). Randomized quantile residuals. *Journal of Computational & Graphical Statistics*, 5(3), 236–245. <https://doi.org/10.2307/1390802>
- Dunn, P. K., & Smyth, G. K. (2005). Series evaluation of Tweedie exponential dispersion model densities. *Statistics and Computing*, 15(4), 267–280. <https://doi.org/10.1007/s11222-005-4070-y>
- Farbmacher, H., Ihle, P., Schubert, I., Winter, J., & Wuppermann, A. (2017). Heterogeneous effects of a nonlinear price schedule for outpatient care. *Health Economics*, 26(10), 1234–1248. <https://doi.org/10.1002/hec.3395>
- Han, S., & Vytlačil, E. J. (2017). Identification in a generalization of bivariate probit models with dummy endogenous regressors. *Journal of Econometrics*, 199(1), 63–73. <https://doi.org/10.1016/j.jeconom.2017.04.001>
- Humphreys, B. R., McLeod, L., & Ruseski, J. E. (2014). Physical activity and health outcomes: Evidence from Canada. *Journal of Health Economics*, 23(1), 33–54. <https://doi.org/10.1002/hec.2900>
- Joe, H. (2014). *Dependence modeling with copulas*. CRC Press.
- Jrgensen, B. (1987). Exponential dispersion models. *Journal of the Royal Statistical Society B*, 49(2), 127–162. <https://doi.org/10.1111/j.2517-6161.1987.tb01685.x>
- Kurz, C. F. (2017). Tweedie distributions for fitting semicontinuous health care utilization cost data. *BMC Medical Research Methodology*, 17(1), 171. <https://doi.org/10.1186/s12874-017-0445-y>
- Manning, W. G., & Mullahy, J. (2001). Estimating log models: To transform or not to transform? *Journal of Health Economics*, 20(4), 461–494. [https://doi.org/10.1016/s0167-6296\(01\)00086-8](https://doi.org/10.1016/s0167-6296(01)00086-8)
- Manning, W. G., Newhouse, J. P., Duan, N., Keeler, E. B., & Leibowitz, A. (1987). Health insurance and the demand for medical care: Evidence from a randomized experiment. *The American Economic Review*, 77(3), 251–277.
- Manning, W. G., Wells, K. B., Buchanan, J. L., Keeler, E. B., Burciaga Valdez, R., & Newhouse, J. P. (1989). *Effects of mental health insurance: Evidence from the health insurance experiment*. RAND Corporation. Retrieved from <https://www.rand.org/pubs/reports/R3815.html>
- Manning, W. G., Wells, K. B., Duan, N., Newhouse, J. P., & Ware, J. E. (1986). How cost sharing affects the use of ambulatory mental health services. *The Journal of the American Medical Association*, 256(14), 1930–1934. <https://doi.org/10.1001/jama.1986.03380140100030>
- Marra, G., & Radice, R. (2022). GJRM: Generalized joint regression modeling. R package version 0.2-6.1. Retrieved from <https://cran.r-project.org/package=GJRM>
- Marra, G., Radice, R., & Zimmer, D. M. (2020). Estimating the binary endogenous effect of insurance on doctor visits by copula-based regression additive models. *Journal of the Royal Statistical Society: Series A C*, 69(4), 953–971. <https://doi.org/10.1111/rssc.12419>
- Mazurenko, O., Balio, C. P., Agarwal, R., Carroll, A. E., & Menachemi, N. (2018). The effects of Medicaid expansion under the ACA: A systematic review. *Health Affairs*, 37(6), 944–950. <https://doi.org/10.1377/hlthaff.2017.1491>
- Mullahy, J. (1998). Much ado about two: Reconsidering retransformation and the two-part model in health econometrics. *Journal of Health Economics*, 17(3), 247–281. [https://doi.org/10.1016/s0167-6296\(98\)00030-7](https://doi.org/10.1016/s0167-6296(98)00030-7)
- Nelsen, R. (2006). *An introduction to copulas* (2nd ed.). Springer.
- Nikoloulopoulos, A. K., & Karlis, D. (2010). Regression in a copula model for bivariate count data. *Journal of Applied Statistics*, 37(9), 1555–1568. <https://doi.org/10.1080/02664760903093591>
- OECD. (2014). Making mental health count: The social and economic costs of neglecting mental health care.
- R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Rhodes, J. H. (2018). Changes in the utilization of mental health care services and mental health at the onset of medicare. *The Journal of Mental Health Policy and Economics*, 21(1), 29–41.
- Rigby, R. A., & Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series A C*, 54(3), 507–554. <https://doi.org/10.1111/j.1467-9876.2005.00510.x>
- Rivers, D., & Vuong, Q. H. (1988). Limited information estimators and exogeneity tests for simultaneous probit models. *Journal of Econometrics*, 39(3), 347–366. [https://doi.org/10.1016/0304-4076\(88\)90063-2](https://doi.org/10.1016/0304-4076(88)90063-2)
- Schepsmeier, U., StoeberBrechmann, E. C., Graeler, B., Nagler, T., Erhardt, T., Almeida, C., Min, A., Czado, C., Hofmann, M., Killiches, M., Joe, H., & Vatter, T. (2019). VineCopula: Statistical inference of vine copulas. R package version 2.2.0. Retrieved from <https://cran.r-project.org/package=VineCopula>
- Sklar, A. (1973). Random variables, joint distributions, and copulas. *Kybernetika*, 9(6), 449–460.

- Smyth, G. K., & Jrgensen, B. (2002). Fitting tweedie's compound Poisson model to insurance claims data: Dispersion modeling. *Astin Bulletin*, 32(1), 143–157. <https://doi.org/10.2143/ast.32.1.1020>
- Terza, J., Basu, A., & Rathouz, P. (2008). Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling. *Journal of Health Economics*, 27(3), 531–543. <https://doi.org/10.1016/j.jhealeco.2007.09.009>
- Trivedi, P., & Zimmer, D. (2017). A note on identification of bivariate copulas for discrete count data. *Econometrics*, 5(1), 10. <https://doi.org/10.3390/econometrics5010010>
- United Nations. (2015). UN general assembly, transforming our world: The 2030 agenda for sustainable development. Retrieved from <https://www.refworld.org/docid/57b6e3e44.html>
- Waters, H. R. (1999). Measuring the impact of health insurance with a correction for selection bias - a case study of Ecuador. *Health Economics*, 8(5), 473–483. [https://doi.org/10.1002/\(sici\)1099-1050\(199908\)8:5<473::aid-hec453>3.0.co;2-c](https://doi.org/10.1002/(sici)1099-1050(199908)8:5<473::aid-hec453>3.0.co;2-c)
- Wood, S. N. (2017). *Generalized additive models: An introduction with R* (2nd ed.). Chapman and Hall/CRC.
- Wood, S. N., Pya, N., & Säfken, B. (2017). Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association*, 111(516), 1548–1563. <https://doi.org/10.1080/01621459.2016.1180986>
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. The MIT Press.

How to cite this article: Marra, G., Fasiolo, M., Radice, R., & Winkelmann, R. (2023). A flexible copula regression model with Bernoulli and Tweedie margins for estimating the effect of spending on mental health. *Health Economics*, 1–18. <https://doi.org/10.1002/hec.4668>

APPENDIX A: TWEEDIE CDF AND ITS DERIVATIVES

For $1 < \nu < 2$, a Tweedie variable Y can be written as the sum of N independent Gamma-distributed random variables W_1, \dots, W_N with shape $-\alpha$ and scale γ , while N follows a Poisson distribution with rate ω . In terms of the Tweedie parameters, we have

$$\omega = \frac{\mu^{2-\nu}}{\sigma(2-\nu)}, \quad \alpha = (2-\nu)/(1-\nu), \quad \gamma = \sigma(\nu-1)\mu^{\nu-1}.$$

Hence, the cdf of Y is

$$F(y) = P\left(\sum_{i=1}^N w_i < y\right) = \sum_{k=1}^{\infty} P_G\left(\sum_{i=1}^k w_i < y\right) P_P(N = k), \quad (6)$$

where the second equality holds due to the Law of Total Probability, P_G is equivalent to the cdf (F_G^k) of a Gamma distribution with parameters $-k\alpha$ and γ , while P_P is the probability mass function (f_P) of a Poisson with rate ω .

We approximate the infinite sum in Equation (6) by

$$F(y) \approx \hat{F}(y) = \sum_{k=k_{\min}}^{k_{\max}} F_G^k(y) f_P(k),$$

for some $k_{\max} \geq k_{\min}$ whose values are chosen as follows. Given that f_P is maximal at $k = \lfloor \omega \rfloor$ and then monotonically decreases as k moves away from the mode, a reasonable approach is to choose k_{\min} and k_{\max} such that $f_P(k) < \epsilon f_P(\lfloor \omega \rfloor)$ for $k < k_{\min}$ or $k > k_{\max}$, for some small ϵ . If we do so, it is clear that a very pessimistic upper bound on the approximation error is

$$|F(y) - \hat{F}(y)| < P_P(N < k_{\min} \text{ or } N > k_{\max}) = 1 - P_P(k_{\min} < N < k_{\max}),$$

which is easy to compute. Of course, k_{\min} and k_{\max} are not known in advance, but we can initialize k to $\lfloor \omega \rfloor$ and then increase k until $f_P(k) < \epsilon f_P(\lfloor \omega \rfloor)$, which leads to k_{\max} . Then k is set to the Poisson mode and decreased until k_{\min} is found or $k = 0$.

First and second derivatives of the (approximate) Tweedie cdf, $\hat{F}(y)$, w.r.t. $\theta = \{\mu, \sigma, \nu\}$ are obtained by firstly calculating the derivatives w.r.t. $\psi = \{\alpha, \gamma, \omega\}$ and then using the Jacobian of the transformation to convert them to the θ -based parametrization. The gradient of $\hat{F}(y)$ and the diagonal entries of the Hessian w. r.t ψ are

$$\frac{\partial^j \hat{F}}{\partial \alpha^j} \approx \sum_{k=k_{\min}}^{k_{\max}} \frac{\partial^j F_G^k}{\partial \alpha^j} f_P, \quad \frac{\partial^j \hat{F}}{\partial \gamma^j} \approx \sum_{k=k_{\min}}^{k_{\max}} \frac{\partial^j F_G^k}{\partial \gamma^j} f_P, \quad \frac{\partial^j \hat{F}}{\partial \omega^j} \approx \sum_{k=k_{\min}}^{k_{\max}} F_G^k \frac{\partial^j f_P}{\partial \omega^j},$$

for $j = 1$ or 2 , while the non-diagonal Hessian elements are

$$\frac{\partial^2 \hat{F}}{\partial \alpha \partial \gamma} \approx \sum_{k=k_{\min}}^{k_{\max}} \frac{\partial^2 F_G^k}{\partial \alpha \partial \gamma} f_P, \quad \frac{\partial^2 \hat{F}}{\partial \alpha \partial \omega} \approx \sum_{k=k_{\min}}^{k_{\max}} \frac{\partial F_G^k}{\partial \alpha} \frac{\partial f_P}{\partial \omega}, \quad \frac{\partial^2 \hat{F}}{\partial \gamma \partial \omega} \approx \sum_{k=k_{\min}}^{k_{\max}} \frac{\partial F_G^k}{\partial \gamma} \frac{\partial f_P}{\partial \omega}.$$

The derivatives w.r.t. γ and ν are given by

$$\begin{aligned} \frac{\partial F_G^k}{\partial \gamma} &= -\frac{1}{\Gamma(-k\alpha)} \frac{y}{\gamma^2} \left(\frac{y}{\gamma}\right)^{-k\alpha-1} e^{-\frac{y}{\gamma}}, \quad \frac{\partial^2 F_G^k}{\partial \gamma \partial \alpha} = k \frac{\partial F_G^k}{\partial \gamma} \left\{ \frac{\Gamma^{(1)}(-k\alpha)}{\Gamma(-k\alpha)} - \log \frac{y}{\gamma} \right\}, \\ \frac{\partial^2 F_G^k}{\partial \gamma^2} &= \frac{y \exp\left\{-\frac{y}{\gamma} - (k\alpha + 1) \log\left(\frac{y}{\gamma}\right) - \log \Gamma(-k\alpha)\right\} \{\gamma(1 - k\alpha) - y\}}{\gamma^4}, \\ \frac{\partial f_P}{\partial \omega} &= f_P(k) \left(\frac{k}{\omega} - 1\right), \quad \frac{\partial^2 f_P}{\partial \omega^2} = \frac{\partial f_P}{\partial \omega} \left(\frac{k}{\omega} - 1\right) - \frac{k}{\omega^2} f_P(k), \end{aligned}$$

where Γ is the gamma function and $\Gamma^{(1)}(-k\alpha)/\Gamma(-k\alpha)$ is the digamma function. Derivatives w.r.t. α require computing the derivative of the lower incomplete gamma function w.r.t. its first argument. To our best knowledge, no numerical routine is currently available to approximate this quantity efficiently and stably, hence we use finite differences to approximate (mixed) derivatives of \hat{F}_G w.r.t. α . The gradient of \hat{F} w.r.t. θ is obtained by $\nabla_{\theta} \hat{F} = \mathbf{J} \nabla_{\psi} \hat{F}$, where

$$\mathbf{J} = \begin{bmatrix} \frac{\partial \omega}{\partial \mu} & \frac{\partial \alpha}{\partial \mu} & \frac{\partial \gamma}{\partial \mu} \\ \frac{\partial \omega}{\partial \sigma} & \frac{\partial \alpha}{\partial \sigma} & \frac{\partial \gamma}{\partial \sigma} \\ \frac{\partial \omega}{\partial \nu} & \frac{\partial \alpha}{\partial \nu} & \frac{\partial \gamma}{\partial \nu} \end{bmatrix} = \begin{bmatrix} \frac{2-\nu}{\mu} \omega & 0 & \frac{\nu-1}{\mu} \gamma \\ -\frac{\omega}{\sigma} & 0 & \frac{\gamma}{\sigma} \\ \frac{\mu^{2-\nu} \{1 + (\nu-2) \log \mu\}}{(\nu-2)^2 \sigma} & \frac{1}{(\nu-1)^2} & \gamma \left(\log \mu + \frac{1}{\nu-1} \right) \end{bmatrix}$$

is the Jacobian of the transformation. The Hessian is then obtained by

$$\nabla'_{\theta} \nabla_{\theta} \hat{F} = \mathbf{J} \nabla'_{\psi} \nabla_{\psi} \hat{F} \mathbf{J}^T + \left[\frac{\partial \mathbf{J}}{\partial \mu} \nabla_{\psi} \hat{F}, \frac{\partial \mathbf{J}}{\partial \sigma} \nabla_{\psi} \hat{F}, \frac{\partial \mathbf{J}}{\partial \nu} \nabla_{\psi} \hat{F} \right],$$

where

$$\begin{aligned} \frac{\partial \mathbf{J}}{\partial \mu} &= \begin{bmatrix} \mathbf{J}_{11} \left(\frac{1-\nu}{\mu} \right) & 0 & \mathbf{J}_{13} \left(\frac{\nu-2}{\mu} \right) \\ -\frac{\mathbf{J}_{11}}{\sigma} & 0 & \frac{\mathbf{J}_{13}}{\sigma} \\ \frac{2-\nu}{\mu} \mathbf{J}_{31} + \frac{1}{\mu} \frac{\mu^{2-\nu}}{(\nu-2)\sigma} & 0 & \frac{\gamma}{\mu} + \mathbf{J}_{13} \left(\log \mu + \frac{1}{\nu-1} \right) \end{bmatrix}, \quad \frac{\partial \mathbf{J}}{\partial \sigma} = \begin{bmatrix} \frac{2-\nu}{\mu} \mathbf{J}_{21} & 0 & \frac{\nu-1}{\mu} \mathbf{J}_{23} \\ -2 \frac{\mathbf{J}_{21}}{\sigma} & 0 & 0 \\ -\frac{\mathbf{J}_{31}}{\sigma} & 0 & \mathbf{J}_{23} \left(\log \mu + \frac{1}{\nu-1} \right) \end{bmatrix}, \\ \frac{\partial \mathbf{J}}{\partial \nu} &= \begin{bmatrix} \frac{2-\nu}{\mu} \mathbf{J}_{31} - \frac{\omega}{\mu} & 0 & \frac{\nu-1}{\mu} \mathbf{J}_{33} + \frac{\gamma}{\mu} \\ -\frac{\mathbf{J}_{31}}{\sigma} & 0 & \frac{\mathbf{J}_{33}}{\sigma} \\ -\mathbf{J}_{31} \left(\log \mu + \frac{2}{\nu-2} \right) + \frac{\mu^{2-\nu} \log \mu}{(\nu-2)^2 \sigma} & -\frac{2}{(\nu-1)^3} & -\gamma \frac{1}{(\nu-1)^2} + \mathbf{J}_{33} \left(\log \mu + \frac{1}{\nu-1} \right) \end{bmatrix}, \end{aligned}$$

with \mathbf{J}_{ij} indicating the element in row i and column j of the Jacobian.

APPENDIX B: SOME FITTING DETAILS

Let us assume that a random sample $\{(y_{1i}, y_{2i}, \mathbf{z}_i)\}_{i=1}^n$ is available, then the log-likelihood function, $\ell(\delta)$, can be obtained by taking the logarithm of $f_{12}(y_1, y_2)$ defined in Equation (2) and creating the indicator variables corresponding to the four possible combinations of the responses. The parameters are defined as $\pi_i = g_{\pi}^{-1}(\eta_{\pi_i})$, $\mu_i = g_{\mu}^{-1}(\eta_{\mu_i})$, $\sigma_i = g_{\sigma}^{-1}(\eta_{\sigma_i})$, $\nu_i = g_{\nu}^{-1}(\eta_{\nu_i})$ and $\theta_i = g_{\theta}^{-1}(\eta_{\theta_i})$, and δ contains the coefficient vectors associated with η_{π_i} , η_{μ_i} , η_{σ_i} , η_{ν_i} and η_{θ_i} , that is $\delta = (\beta'_{\pi}, \beta'_{\mu}, \beta'_{\sigma}, \beta'_{\nu}, \beta'_{\theta})'$.

Because of the modeling flexibility offered by the additive predictors in the model, parameter estimation is achieved by maximizing the penalized log-likelihood

$$\ell_p(\delta) = \ell(\delta) - \frac{1}{2} \delta' \mathbf{S}_{\lambda} \delta, \quad (7)$$

where $\mathbf{S}_\lambda = \text{diag}(\lambda_\pi \mathbf{D}_\pi, \lambda_\mu \mathbf{D}_\mu, \lambda_\sigma \mathbf{D}_\sigma, \lambda_\nu \mathbf{D}_\nu, \lambda_\theta \mathbf{D}_\theta)$, each smoothing parameter vector in \mathbf{S}_λ contains all the smoothing parameters related to the corresponding \mathbf{D} component, and $\lambda = (\lambda'_\pi, \lambda'_\mu, \lambda'_\sigma, \lambda'_\nu, \lambda'_\theta)'$. Estimation of δ and λ is based on an extension of the efficient and stable trust region algorithm with integrated automatic multiple smoothing parameter selection discussed in Marra et al. (2020), which uses analytical derivative information of $\ell_p(\delta)$. We would like to mention that, while the implementation of the proposed model exploited the infrastructure of GJRM as well as several of its internal functions, extending GJRM to accommodate a Tweedie marginal required a great deal of programming work.

Inferential results are derived using known theory for general penalized likelihood-based models. Specifically, at convergence, reliable intervals for any linear function of δ are obtained using the Bayesian large sample approximation $\delta \stackrel{a}{\sim} N(\hat{\delta}, (\mathbf{H}(\hat{\delta}) + \mathbf{S}_\lambda)^{-1})$, where $\mathbf{H}(\hat{\delta})$ is the observed information matrix (Hessian of the negative log likelihood) at $\hat{\delta}$. Intervals for non-linear functions of δ can be conveniently obtained via posterior simulation (e.g., Marra et al., 2020), whereas p -values for all the terms in the model can be obtained by using the results in Wood (2017) which are based on $\mathbf{H}_p(\hat{\delta})^{-1}$. Under some classical model assumptions, it can be proved that $\hat{\delta} - \delta^\circ = O_P(n^{-1/2})$ as $n \rightarrow \infty$, where δ° denotes the “true” parameter vector. For more information, the reader can consult the on-line supplementary material of Marra et al. (2020) which provides general asymptotic arguments that can also be applied to the current context.

Fitting the model in R

The proposed copula model can be employed via the `gjrm()` function in the R package GJRM (Marra & Radice, 2022). An example of the syntax is

```
fl <- list(y1 ~ y2 + s(z1) + z3,
          y2 ~ z1 + s(z2) + z3,
            ~ z1 + z3,
            ~ s(z1) + s(z2) + z3
            ~ z1 + z3)
mo <- gjrm(fl, margins = c("probit", "TW"), data = md, BivD = "C0", Model = "B")
```

where `fl` is a list containing five equations: the first equation is for parameter π of the Bernoulli distribution of the binary outcome `y1` with `probit` link function (`logit` and `cloglog` are also allowed for); the second, third and fourth equations are for parameters μ , σ and ν of the Tweedie distribution used to model `y2`; the fifth equation is for the copula dependence parameter θ . `z1` and `z2` are continuous covariates and `z3` is a binary covariate. Argument `BivD` specifies the copula function (Clayton in this case), `Model = "B"` implies that a bivariate model is employed and `md` is a data frame. Symbol `s()` refers to the smooth function mentioned in Section 2. Default is `bs = "tp"` (penalized low rank thin plate spline) with `k = 10` (number of basis functions) and `m = 2` (order of derivatives). However, argument `bs` can also be set to, for example, `cr` (penalized cubic regression spline), `ps` (P-spline) and `mrf` (Markov random field), to name but a few. Note that, for example, for uni-dimensional smooth functions of continuous covariates, the specific choice of spline definition will not have an impact on the estimated curves. Furthermore, the default value of `k = 10` (which can be increased if desired) is arbitrary although it generally offers enough modeling flexibility in applications. Functions such as `AIC()`, `summary()` and `predict()` can be employed in the usual manner. Function `post.check()` will produce, for the Tweedie margin, a histogram and normal Q-Q plot of modified normalized quantile residuals constructed as follows. For a continuous distribution, $F_2(y_{2i} | \mu_i, \sigma_i, \nu_i)$ is uniform. Under correct specification, therefore, the sample values $F_2(y_{2i} | \hat{\mu}_i, \hat{\sigma}_i, \hat{\nu}_i)$, $i = 1, \dots, n$, should be approximately uniform, and the transformed values $\Phi^{-1}(F_2(y_{2i} | \hat{\mu}_i, \hat{\sigma}_i, \hat{\nu}_i))$, where $\Phi^{-1}(\cdot)$ is the quantile function of a standard normal distribution, approximately standard normal. To account for the probability mass at zero, we employ an adjustment based on the idea of sampling uniform variates with bounds given by 0 and the upper probability of those observations (Dunn & Smyth, 1996). Residual analysis for the binary margin is not informative (e.g., Collett, 2002). Here, a sensitivity analysis based on different link functions can be carried out; experience suggests that the model fit will not be significantly affected by this choice.