

# Bayesian Semiparametric Covariate Informed Multivariate Density Deconvolution

Abhra Sarkar

abhra.sarkar@utexas.edu

Department of Statistics and Data Sciences, The University of Texas at Austin  
2317 Speedway D9800, Austin, TX 78712-1823, USA

## Abstract

Estimating the marginal and joint densities of the long-term average intakes of different dietary components is an important problem in nutritional epidemiology. Since these variables cannot be directly measured, data are usually collected in the form of 24-hour recalls of the intakes. The problem of estimating the density of the latent long-term average intakes from their observed but error contaminated recalls then becomes a problem of multivariate deconvolution of densities. The underlying densities could potentially vary with the subjects' demographic characteristics such as sex, ethnicity, age, etc. The problem of density deconvolution in the presence of associated precisely measured covariates has, however, never been considered before, not even in the univariate setting. We present a flexible Bayesian semiparametric approach to covariate informed multivariate deconvolution. Building on recent advances on copula deconvolution and conditional tensor factorization techniques, our proposed method not only allows the joint and the marginal densities to vary flexibly with the associated predictors but also allows automatic selection of the most influential predictors. Importantly, the method also allows the density of interest and the density of the measurement errors to vary with potentially different sets of predictors. We design Markov chain Monte Carlo algorithms that enable efficient posterior inference, appropriately accommodating uncertainty in all aspects of our analysis. The empirical efficacy of the proposed method is illustrated through simulation experiments. Its practical utility is demonstrated in the afore-described nutritional epidemiology applications in estimating covariate adjusted long term intakes of different dietary components. An important by-product of the approach is a solution to covariate informed ordinary multivariate density estimation. Supplementary materials include substantive additional details and R codes are also available online.

**Some Key Words:** Copula, Covariates, Multivariate density regression, Multivariate density deconvolution, Measurement error, Nutritional epidemiology, Tensor factorization.

**Short/Running Title:** Covariate Informed Multivariate Deconvolution

# 1 Introduction

The distribution of the dietary intakes can provide answers to important questions such as what proportion of the population consume certain dietary components above, between or below certain amounts etc. The last question is particularly important as it relates to the proportion of the population that are deficient in certain dietary components. Estimating the long-term average intakes of different dietary components  $\mathbf{x}$  and their marginal and joint distributions is thus a fundamentally important problem in nutritional epidemiology.

By the very nature of the problem,  $\mathbf{x}$  can never be observed directly. Data are thus often collected in the form of 24-hour recalls of the intakes. Treating the recalls  $\mathbf{w}$ , shown in Table 1, to be surrogates for the latent  $\mathbf{x}$  contaminated with additive measurement errors  $\mathbf{u}$  generated as  $\mathbf{w} = \mathbf{x} + \mathbf{u}$ , the problem of estimating the joint and marginal distributions of  $\mathbf{x}$  from the recalls  $\mathbf{w}$  then becomes a problem of multivariate deconvolution of densities.

Dietary intakes may potentially vary with additional precisely measured demographic covariates  $\mathbf{c}$  such as sex, ethnicity and age. Women, for example, consume practically all dietary components in lesser amounts compared to men, on average. To our knowledge, however, the problem of deconvolution in the presence of covariates has never been considered in the literature, not even in the univariate setting, not at least in a statistically principled manner. This article attempts to address this gap, developing a novel Bayesian semiparametric approach that not only allows robust estimation of the density of  $\mathbf{x}$  as it varies with  $\mathbf{c}$  while also letting the density of the measurement errors  $\mathbf{u}$  to depend flexibly on both  $\mathbf{x}$  and  $\mathbf{c}$  but also additionally selects the most important predictors influencing the distributions of  $\mathbf{x}$  and  $\mathbf{u}$  from the set of all available predictors  $\mathbf{c}$ .

We adopt the following generic notation for marginal, joint and conditional densities, respectively. For random vectors  $\mathbf{s}$  and  $\mathbf{t}$ , we denote the marginal density of  $\mathbf{s}$ , the joint density of  $(\mathbf{s}, \mathbf{t})$ , and the conditional density of  $\mathbf{s}$  given  $\mathbf{t}$ , by the generic notation  $f_{\mathbf{s}}, f_{\mathbf{s}, \mathbf{t}}$  and  $f_{\mathbf{s}|\mathbf{t}}$ , respectively. Likewise, for univariate random variables  $s$  and  $t$ , the corresponding densities are denoted by  $f_s, f_{s,t}$  and  $f_{s|t}$ , respectively. To avoid introducing more notation, with some abuse, barring few exceptions, for any random variable  $s$  or vector  $\mathbf{s}$ , their specific values would also be denoted by the same notation, i.e.,  $s$  and  $\mathbf{s}$ .

**The EATS Data Set:** The Eating at America’s Table Study (EATS) (Subar *et al.*, 2001) is a large scale epidemiological study conducted by the National Cancer Institute in which  $i = 1, \dots, n = 965$  participants were interviewed  $j = 1, \dots, m_i = 4$  times over the course of a year and, for many different dietary components  $\ell$ , their 24-hour dietary recalls  $w_{\ell,i,j}$  were recorded. Error free demographic covariates  $\mathbf{c}_i = (c_{1,i}, c_{2,i}, c_{3,i})^T \equiv (\text{sex}, \text{ethnicity}, \text{age})^T$  are additionally available for each individual  $i$ .

The long term average intakes may vary between different combinations of the levels of the predictors. The left panels of Figure 2, for example, show the histograms of the subject-specific means  $\bar{w}_{\ell,i} = \sum_{j=1}^{m_i} w_{\ell,i,j}/4$  for three different minerals, namely iron, magnesium and sodium, separately for men and women but superimposed on each other. The consumptions for men tend to be higher on average and also have much heavier right tails compared

Subject	Sex	Ethn	Age	24-hour recalls							
				Dietary Component 1				Dietary Component 2			
1	$c_{1,1}$	$c_{2,1}$	$c_{3,1}$	$w_{1,1,1}$	$w_{1,1,2}$	$w_{1,1,3}$	$w_{1,1,4}$	$w_{2,1,1}$	$w_{2,1,2}$	$w_{2,1,3}$	$w_{2,1,4}$
2	$c_{1,2}$	$c_{2,2}$	$c_{3,2}$	$w_{1,2,1}$	$w_{1,2,2}$	$w_{1,2,3}$	$w_{1,2,4}$	$w_{2,2,1}$	$w_{2,2,2}$	$w_{2,2,3}$	$w_{2,2,4}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
n	$c_{1,n}$	$c_{2,n}$	$c_{3,n}$	$w_{1,n,1}$	$w_{1,n,2}$	$w_{1,n,3}$	$w_{1,n,4}$	$w_{2,n,1}$	$w_{2,n,2}$	$w_{2,n,3}$	$w_{2,n,4}$

Table 1: The EATS data set showing the recalls for two regularly consumed dietary components and associated subject-specific predictors. Here  $w_{\ell,i,j}$  is the reported intake for the  $j^{\text{th}}$  recall of the  $i^{\text{th}}$  individual for the  $\ell^{\text{th}}$  dietary component;  $c_{h,i}$  is the value of the  $h^{\text{th}}$  predictor for the  $i^{\text{th}}$  individual.




Figure 1: Observed distributions of the demographic predictors in the EATS data set.

to women for all three dietary components. The right upper panels of Figure 2 show the histograms of ‘measurement error residuals’  $w_{\ell,i,j} - \bar{w}_{\ell,i}$  for men and women. The histograms are all right skewed and the histograms for women are slightly more concentrated around zero compared to men. The right lower panels of Figure 2 show  $\bar{w}_{\ell,i}$  vs the subject-specific variances  $s_{w,\ell,i}^2 = \sum_{j=1}^{m_i} (w_{\ell,i,j} - \bar{w}_{\ell,i})^2 / 3$  for the 24-hour recalls, providing crude estimates of the conditional variances  $\text{var}(u_{\ell,i,j} | x_{\ell,i})$ , suggesting strongly that  $\text{var}(u | x)$  increases as  $x$  increases for both men and women, although the patterns may not be significantly different between the two gender categories. Not all demographic variables may actually be important. Figure 3 summarizes similar exploratory analysis but for the predictor race, specifically the groups ‘whites’ and ‘blacks’. Unlike the two gender categories, in this case however, the consumptions do not seem to vary significantly between the two levels. Comparison between race groups ‘whites’ and ‘missing’, presented in Figure S.2 in the supplementary material, may indicate stark differences in consumption patterns at a quick glance but this may just be an artifact of the sparse representation of the ‘missing’ group (5 subjects only) in the EATS data set. Treating the subjects with missing race labels to come from a separate specific racial group is certainly a bit ad-hoc but will be instructive in illustrating the robustness of our proposed approach to the presence of small outlying groups in the data. Overall, these exploratory analyses illustrate the need for sophisticated density deconvolution methods that can accommodate the available demographic covariates and can also formally assess their statistical importance in influencing the long-term average consumptions.

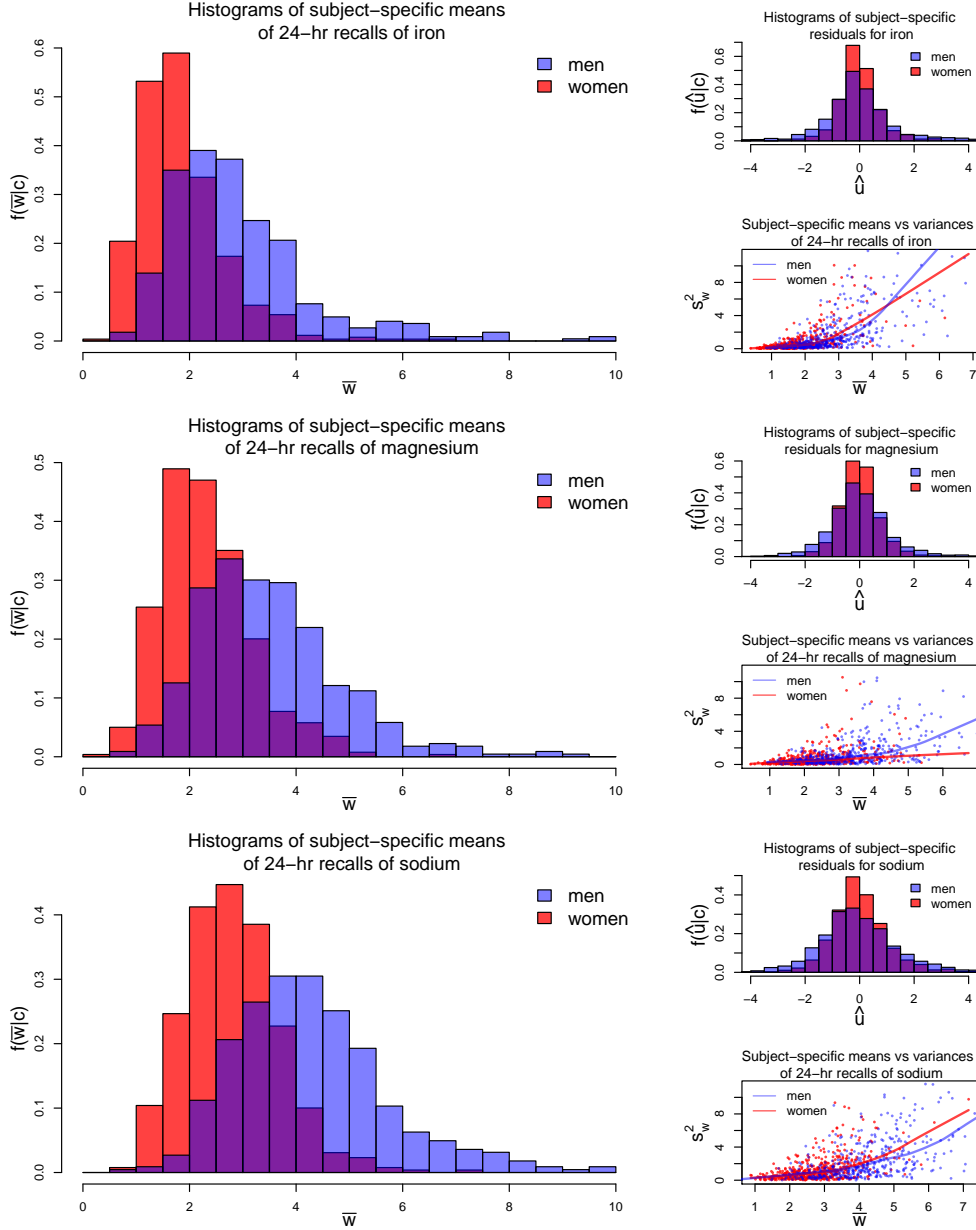


Figure 2: Exploratory plots. Left panels: histograms of subject-specific means  $\bar{w}_{\ell,i}$ , crude estimates of  $x_{\ell,i}$ ; right upper panels: histograms of ‘residuals’  $\hat{u}_{\ell,i,j} = (w_{\ell,i,j} - \bar{w}_{\ell,i})$ , crude estimates of  $u_{\ell,i,j}$ ; right lower panels: subject-specific means  $\bar{w}_{\ell,i}$  vs variances  $s_{w,\ell,i}^2$ , crude estimates of  $\text{var}(u_{\ell,i,j} | x_{\ell,i})$ , superimposed with lowess fits.

**Existing Methods:** The literature on univariate density deconvolution, in which context we denote the variable of interest by  $x$  and the measurement errors by  $u$ , and the surrogates by  $w$ , is massive. The classical literature, reviews of which can be found in Carroll *et al.* (2006) and Buonaccorsi (2010), mostly focused on the additive model  $w = x + u$  subject to  $\mathbb{E}(u) = 0$  with restrictive assumptions, such as known  $f_u$ , homoscedasticity of  $u$ ,




Figure 3: Exploratory plots. Left panels: histograms of subject-specific means  $\bar{w}_{\ell,i}$ , crude estimates of  $x_{\ell,i}$ ; right upper panels: histograms of ‘residuals’  $\hat{u}_{\ell,i,j} = (w_{\ell,i,j} - \bar{w}_{\ell,i})$ , crude estimates of  $u_{\ell,i,j}$ ; right lower panels: subject-specific means  $\bar{w}_{\ell,i}$  vs variances  $s_{w,\ell,i}^2$ , crude estimates of  $\text{var}(u_{\ell,i,j} | x_{\ell,i})$ , superimposed with loess fits.

independence of  $u$  from  $x$ , etc. These assumptions are often highly unrealistic, especially in nutritional epidemiology applications.

Recent works by Staudenmayer *et al.* (2008), Su *et al.* (2020), Sarkar *et al.* (2014, 2018, 2021) have shown that Bayesian hierarchical frameworks and associated computational machinery can provide powerful tools for solving complex deconvolution problems under more

realistic scenarios, including when the errors  $u$  can be conditionally heteroscedastic. In their seminal work, Staudenmayer *et al.* (2008) considered the model  $w = x + u$  with  $(u | x) \sim \text{Normal}\{0, s^2(x)\}$ , utilizing mixtures of B-splines to estimate  $f_x$  as well the conditional variability  $\text{var}(u | x) = s^2(x)$ . Sarkar *et al.* (2014) relaxed the assumption of normality of  $u$ , employing flexible mixtures of normals (Escobar and West, 1995; Frühwirth-Schnatter, 2006) to model both  $f_x$  and  $f_{u|x}$ . Sarkar *et al.* (2018) extended the methods to multivariate settings  $\mathbf{w} = \mathbf{x} + \mathbf{u}$  subject to  $\mathbb{E}(\mathbf{u} | \mathbf{x}) = \mathbf{0}$ , modeling  $f_{\mathbf{x}}$  and  $f_{\mathbf{u}|\mathbf{x}}$  using mixtures of multivariate normals. Sarkar *et al.* (2021) adopted a complimentary approach, modeling the marginals  $f_{x_\ell}$  and  $f_{u_\ell|x_\ell}$  first and then building the joint distributions  $f_{\mathbf{x}}$  and  $f_{\mathbf{u}|\mathbf{x}}$  by modeling the dependence structures separately using Gaussian copulas.

To the best of our knowledge, however, the problem of deconvoluting  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\mathbf{u}|\mathbf{x},\mathbf{c}}$  in the presence of precisely measured covariates  $\mathbf{c}$  from surrogates generated as  $\mathbf{w} = \mathbf{x} + \mathbf{u}$  subject to  $\mathbb{E}(\mathbf{u} | \mathbf{x}, \mathbf{c}) = \mathbf{0}$  has never been considered in the literature, not even in the univariate setting. The only practical solution we can mention in this context is the multi-stage pseudo-Bayesian approach of Zhang *et al.* (2011), where component-wise Box-Cox transformed (Box and Cox, 1964) recalls were assumed to follow a linear mixed model, comprising a subject specific random effect component and a covariate dependent linear fixed effects component with no interaction terms and an error component. The error and the random effects components were then both modeled using single component multivariate normal distributions. Multivariate normal priors were also assumed for the fixed effects coefficients. Estimates of the long-term intakes were then obtained via individual transformations back to the original scale. The density of interest is then obtained by applying a separate off-the-shelf kernel density estimation method on the estimated intakes. As shown in Sarkar *et al.* (2014), Box-Cox transformations for surrogate observations have severe limitations, including almost never being able to produce transformed surrogates that conform to normality, homoscedasticity, and independence. Single component multivariate normal models are thus often highly inadequate for the densities even after transformations (Sarkar *et al.*, 2021).

**Our Proposed Approach:** In this article, we develop a Bayesian semiparametric approach to covariate dependent multivariate density deconvolution, carefully accommodating the prominent features of nutritional epidemiology data sets, including conditional heteroscedasticity, departures from normality, etc. Our proposed approach not only allows  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\mathbf{u}|\mathbf{x},\mathbf{c}}$  to vary flexibly with the predictors  $\mathbf{c}$  but also allows us to determine which predictors among  $\mathbf{c}$  are the most influential ones for  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\mathbf{u}|\mathbf{x},\mathbf{c}}$ , including accommodating the possibility that the sets of influential predictors can be different for  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\mathbf{u}|\mathbf{x},\mathbf{c}}$ .

Following Staudenmayer *et al.* (2008) and Sarkar *et al.* (2014, 2018, 2021), we begin with the assumption that the measurement errors  $u_\ell$  decompose into a variance function  $v_\ell$  that explains their conditional heteroscedasticity and a scaled error component  $\epsilon_\ell$  that captures their general distributional shapes and other properties. Building on Sarkar *et al.* (2021), we model the joint densities using a copula based approach with the marginal densities  $f_{x_\ell}$  and  $f_{\epsilon_\ell}$  and the variance functions  $v_\ell$  characterized as flexible mixtures of dictionary

functions shared between all univariate components and all predictor level combinations. Unlike previous approaches, however, we now allow the mixture probabilities to vary with the predictors. A predictor is thus considered important if the mixture probabilities vary significantly between its levels, thereby significantly altering the densities. Viewing these mixture probabilities as a conditional probability tensor and relying on tensor factorization techniques (Yang and Dunson, 2016), we then parameterize the mixture probabilities themselves as mixtures of ‘core’ probability kernels with mixture weights depending on the level combinations of the predictors. The parameterization allows explicit identification of the set of important predictors while also implicitly capturing complex higher order interactions between them in a parsimonious manner. The elimination of the redundant predictors and the implicit modeling of the interactions among the important ones lead to a significant two fold reduction in the effective number parameters required to flexibly characterize the mixture probabilities. The daunting challenge of implementing a tensor factorization model separately for each dietary component is avoided via a simple innovation of treating the component labels to comprise the levels of an additional categorical predictor. We assign sparsity inducing priors that favor such lower dimensional representations. We assign a hierarchical Dirichlet prior on the core probability kernels, encouraging the model to shrink further towards lower dimensional structures by borrowing strength across these components as well. We develop a Markov chain Monte Carlo (MCMC) algorithm to approximately sample from the posterior. Importantly, our proposed method allows the density of interest  $f_{\mathbf{x}|\mathbf{c}}$  and the density of the scaled measurement errors  $f_{\epsilon|\mathbf{c}}$  to vary with potentially different sets of covariates. Applied to our motivating EATS data set, the proposed method estimates the distributions of long-term consumptions of different dietary components for different level combinations of the predictors, while also selecting the important predictors, providing novel insights into how the distributions of the intakes as well as the distributions of the associated measurement errors vary with the available subject specific demographic covariates.

**Outline of the Article:** The rest of the article is organized as follows. Section 2 details the proposed Bayesian hierarchical framework. Section 3 presents results of our proposed method applied to the motivating nutritional epidemiology problems. Section S.6 presents the results of some synthetic experiments. Section 4 concludes with a discussion. A brief review of copula and conditional tensor factorization models, the Markov chain Monte Carlo (MCMC) algorithm we used to sample from the posterior, results of synthetic numerical experiments, and some additional results are included in the supplementary material.

## 2 Deconvolution Models

We are interested in estimating the unknown joint density of a  $d$ -dimensional continuous random vector  $\mathbf{x}$  in the presence of associated  $p$ -dimensional categorical covariate  $\mathbf{c}$ , the  $h^{th}$  component  $c_h$  taking  $d_h$  different categorical values  $\{1, \dots, d_h\}$ . There are  $i = 1, \dots, n$  subjects. The covariates  $\mathbf{c}_i$  are precisely measured for each subject  $i$ . For  $\mathbf{x}_i$ , however,

only replicated proxies  $\mathbf{w}_{i,j}$  contaminated with measurement errors  $\mathbf{u}_{i,j}$  are available for  $j = 1, \dots, m_i$  for each subject  $i$ . The density of  $\mathbf{x}$  as well as the density of  $\mathbf{u}$  may both potentially vary with  $\mathbf{c}$ . The replicates are assumed to be generated by the model

$$\mathbf{w}_{i,j} = \mathbf{x}_i + \mathbf{u}_{i,j} \quad \text{subject to} \quad \mathbb{E}(\mathbf{u}_{i,j} | \mathbf{x}_i, \mathbf{c}_i) = \mathbf{0}. \quad (1)$$

To accommodate conditional heteroscedasticity in the measurement errors, adapting ideas from Sarkar *et al.* (2018), we let

$$\begin{aligned} (\mathbf{u}_{i,j} | \mathbf{x}_i) &= \mathbf{S}(\mathbf{x}_i)\boldsymbol{\epsilon}_{i,j}, \quad \text{with} \quad \mathbb{E}(\boldsymbol{\epsilon}_{i,j} | \mathbf{c}_i) = \mathbf{0}, \\ \text{and} \quad \mathbf{S}(\mathbf{x}_i) &= \text{diag}\{s_1(x_{1,i}), \dots, s_d(x_{d,i})\}. \end{aligned}$$

The model implies that  $\text{cov}(\mathbf{u}_{i,j} | \mathbf{x}_i, \mathbf{c}_i) = \mathbf{S}(\mathbf{x}_i) \text{cov}(\boldsymbol{\epsilon}_{i,j} | \mathbf{c}_i) \mathbf{S}(\mathbf{x}_i)$  and marginally  $\text{var}(u_{\ell,i,j} | x_{\ell,i}, \mathbf{c}_i) = s_\ell^2(x_{\ell,i})\text{var}(\epsilon_{\ell,i,j} | \mathbf{c}_i)$ . Other features of the predictor dependent distribution of  $\mathbf{u}$ , including its shape and correlation structure, are derived from  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$ .




Figure 4: Graphical model depicting the data generating processes considered in classical density deconvolution methods; in our previous research in Sarkar *et al.* (2014, 2018, 2021); and our proposed model described in Section 2. The unfilled and the shaded nodes depict latent and observable variables, respectively. Subject and replicate subscripts ( $i$  and  $j$ , respectively) are suppressed for clarity.

As discussed in detail in Sarkar *et al.* (2018), the above model arises naturally for conditionally heteroscedastic multivariate measurement errors. Additionally, the model also automatically accommodates multiplicative measurement errors: Suppressing the covariates  $\mathbf{c}_i$  and setting  $s(x_{\ell,i}) = x_{\ell,i}$  and  $\epsilon_{\ell,i,j} = (\tilde{u}_{\ell,i,j} - 1)$ , we have  $w_{\ell,i,j} = x_{\ell,i} + s(x_{\ell,i})\epsilon_{\ell,i,j} = x_{\ell,i} + x_{\ell,i}(\tilde{u}_{\ell,i,j} - 1) = x_{\ell,i}\tilde{u}_{\ell,i,j}$  with  $\tilde{u}_{\ell,i,j}$  independent of  $x_{\ell,i}$  and  $\mathbb{E}(\tilde{u}_{\ell,i,j}) = 1$ .

Importantly, in our formulation, the covariates  $\mathbf{c}$  may influence not only the density of  $\mathbf{x}$  but also the density of the scaled errors  $\boldsymbol{\epsilon}$  (Figure 4). The actual sets of influential predictors may be proper subsets of  $\mathbf{c}$  and are allowed to be different for  $\mathbf{x}$  and  $\boldsymbol{\epsilon}$ .

It is possible that the variance function components  $v_\ell(x)$  also vary with  $\mathbf{c}$ . Exploratory analysis, however, suggest that the functions  $\text{var}(u_\ell | x_\ell, \mathbf{c})$  are very similar for different values of  $\mathbf{c}$ . Since  $\text{var}(u_\ell | x_\ell, \mathbf{c})$  are already allowed to vary flexibly with  $\mathbf{c}$  through  $\text{var}(\boldsymbol{\epsilon}_\ell | \mathbf{c})$ , it thus becomes difficult to separate the influence of  $\mathbf{c}$  on  $v_\ell(x)$ . For nutritional epidemiology applications, our proposed model seems to provide a sufficient compromise.

Cast in a Bayesian hierarchical framework, the problem reduces to one of flexibly modeling  $f_{\mathbf{x}|\mathbf{c}}$ ,  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$  and  $v_\ell(x_\ell)$  while also selecting the set of most influential covariates for  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$ . The methodology developed below addresses these daunting statistical challenges. In



what follows, the component, subject and replicate subscripts  $\ell, i, j$  are often omitted and assumed to be implicitly understood to keep the notation simple.

## 2.1 Modeling the Density $f_{\mathbf{x}|\mathbf{c}}$

The joint density  $f_{\mathbf{x}|\mathbf{c}}$  is specified using a Gaussian copula density model

$$f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} | \mathbf{c}) = |\mathbf{R}_{\mathbf{x}}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \mathbf{y}_{\mathbf{x}}^{\text{T}} (\mathbf{R}_{\mathbf{x}}^{-1} - \mathbf{I}_d) \mathbf{y}_{\mathbf{x}} \right\} \prod_{\ell=1}^d f_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c}),$$

where  $F_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c}) = \Phi(y_{x,\ell})$ ,  $F_{x,\ell|\mathbf{c}}$  being the cdf corresponding to  $f_{x,\ell|\mathbf{c}}$ ;  $\mathbf{y}_{\mathbf{x}} = (y_{x,1}, \dots, y_{x,d})^{\text{T}}$ ;  $\Phi(\cdot)$  denotes the cdf of a standard normal distribution; and  $\mathbf{R}_{\mathbf{x}}$  is the correlation matrix between the  $d$  components of  $\mathbf{x}$  for all  $\mathbf{c}$ . The Gaussian copula maps  $\mathbf{x}$  to  $\mathbf{y}_{\mathbf{x}}$  such that  $\mathbf{y}_{\mathbf{x}} \sim \text{MVN}_d(\mathbf{0}, \mathbf{R}_{\mathbf{x}})$ , which allows the dependence relationships between the components of  $\mathbf{x}$  be conveniently modeled separately from their marginals  $f_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c})$ .

As in Sarkar *et al.* (2021), we model the marginal densities using mixtures of truncated normal distributions with atoms shared across the different dimensions. To model the influence of the associated observed covariates, the mixture probabilities are now, however, allowed to vary flexibly with the covariates. Specifically, we let

$$f_{x,\ell|\mathbf{c}}(x_{\ell} | c_1, \dots, c_p) = \sum_{k=1}^{k_x} P_{x,\ell|\mathbf{c}}(k | c_1, \dots, c_p) \text{TN}(x_{\ell} | \mu_{x,k}, \sigma_{x,k}^2, [A, B]),$$

where  $\text{TN}(x | \mu, \sigma^2, [A, B])$ 's are truncated normal mixture kernels with location  $\mu$ , scale  $\sigma$  and support restricted to the interval  $[A, B]$ ;  $P_{x,\ell|\mathbf{c}}(k | c_1, \dots, c_p)$ 's are the associated predictor dependent mixture probabilities. The corresponding cdfs  $F_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c})$  are thus given by

$$F_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c}) = \sum_{k=1}^{k_x} P_{x,\ell|\mathbf{c}}(k | c_1, \dots, c_p) \int_A^{x_{\ell}} \text{TN}(x | \mu_{x,k}, \sigma_{x,k}^2, [A, B]) dx.$$

Sharing the mixture components across different predictor combinations here allows efficient estimation of the atoms  $\{(\mu_{x,k}, \sigma_{x,k}^2)\}_{k=1}^{k_x}$  borrowing information across these combinations. This way, since the dependence of the densities on the associated predictors is modeled entirely through the mixture probabilities, a predictor will be important if the mixture probabilities vary significantly between its levels.

Mixtures of truncated normal kernels are just as flexible as mixtures of normals but also make the support of the densities  $f_{x,\ell|\mathbf{c}}(x_{\ell} | \mathbf{c})$  consistent with the support of  $s^2(x)$  which we model shortly in Section 2.3 using mixtures of B-splines which by construction are finitely supported, here on the interval  $[A, B]$ . The choices of the truncation limits  $A$  and  $B$  are discussed in Section S.3 in the supplementary material.

Modeling the mixture probabilities  $P_{x,\ell|\mathbf{c}}(k | c_1, \dots, c_p)$  separately for each dimension  $\ell$  would, however, be an extremely challenging task. We solve this issue using a simple trick - by including the component label itself as a separate categorical predictor  $c_0 \in \{1, \dots, d_0\}$ . Here  $d_0$  clearly just equals  $d$ , the dimension of  $\mathbf{x}$ , but is additionally introduced to be consistent with the notation  $d_h$  denoting the number of categories of  $c_h$ . Unlike the other  $c_h$ 's which take a single value for each individual  $i$  (e.g., sex), the predictor  $c_0$  however

takes each value in  $\{1, \dots, d_0\}$  for each  $i$  depending on which component we are looking at. More explicitly, we have  $c_{0,\ell,i} = \ell$  for all  $(\ell, i)$ . With  $c_0$  defined in this manner and  $f_{x,\ell|\mathbf{c}}(x | c_1, \dots, c_p) = f_{x|\mathbf{c}}(x | c_0 = \ell, c_1, \dots, c_p)$ , where, with some abuse in notation, we let  $\mathbf{c}$  include  $c_0$  as  $\mathbf{c} = (c_0, c_1, \dots, c_p)^\top$ , we model the marginal densities as

$$f_{x|\mathbf{c}}(x | c_0, c_1, \dots, c_p) = \sum_{k=1}^{k_x} P_{x|\mathbf{c}}(k | c_0, c_1, \dots, c_p) \text{TN}(x | \mu_{x,k}, \sigma_{x,k}^2, [A, B]). \quad (2)$$

The mixture kernels are now shared not just between the associated external predictors  $c_1, \dots, c_p$  but also across different dimensions  $c_0 = \ell$ .




Figure 5: Pictorial representation of the factorization of a conditional probability tensor  $\mathbf{P}$  with 3 categorical covariates with core tensor  $\boldsymbol{\lambda}$  and mode matrices  $\boldsymbol{\pi}^{(j)}, j = 1, 2, 3$ .

Modeling the conditional mixture probabilities  $P_{x|\mathbf{c}}(k | c_0, c_1, \dots, c_p)$  is still a daunting challenge. Being indexed by the  $\prod_{h=0}^p d_h$  different possible values of the predictors  $c_0, c_1, \dots, c_p$ , a completely unrestricted model for  $P_{x|\mathbf{c}}(k | c_0, c_1, \dots, c_p)$  would involve a total number of  $(k_x - 1) \prod_{h=0}^p d_h$  parameters which increases exponentially and becomes too large to be estimated efficiently with datasets of the sizes typically encountered in practice. The issue is further significantly complicated not only by the fact that the mixture component labels associated with the  $x_{\ell,i}$ 's are latent but that the  $x_{\ell,i}$ 's are also measured with error. For our motivating nutritional epidemiology application, for instance, for a  $d_0 = 3$  dimensional problem with sex ( $d_1 = 2$ ), ethnicity ( $d_2 = 5$ ) and age ( $d_3 = 6$ ) as the associated predictors and  $k_x = 30$  components in the mixture, the total number of parameters becomes  $\approx 30 \times 3 \times 2 \times 5 \times 6 = 5400$ . Without imposing additional model structure, it is practically impossible to estimate this many parameters based on the available error contaminated data.

To this end, we look toward higher order singular value decomposition (HOSVD) inspired conditional tensor factorization techniques that have been greatly successful in flexibly yet efficiently modeling high-dimensional conditional probabilities of the type  $P_{y|\mathbf{c}}(k | c_1, \dots, c_p)$  in measurement error free settings, where  $y$  is a categorical response taking values in the set  $\{1, \dots, k_y\}$  and  $c_1, \dots, c_p$  are associated categorical predictors (Yang and Dunson, 2016). Structuring the transition probabilities  $P_{y|\mathbf{c}}(y | c_1, \dots, c_p)$  as a  $k_y \times d_1 \times \dots \times d_p$  dimensional  $(p + 1)$ -way tensor, they considered the following HOSVD-type factorization

$$P_{y|\mathbf{c}}(y | c_1, \dots, c_p) = \sum_{k_1=1}^{k_{y,1}} \cdots \sum_{k_p=1}^{k_{y,p}} \lambda_{k_1, \dots, k_p}(y) \prod_{h=1}^p \pi_{h, c_h}(k_h).$$

See Figure 5. Here  $1 \leq k_{y,h} \leq d_h$  for all  $j$  and the parameters  $\lambda_{k_1, \dots, k_p}(y)$  and  $\pi_{h, c_h}(k_h)$  are all non-negative and satisfy (a)  $\sum_{y=1}^{k_y} \lambda_{k_1, \dots, k_p}(y) = 1$ , for each combination  $(k_1, \dots, k_p)$ , and (b)  $\sum_{k_h=1}^{k_{y,h}} \pi_{h, c_h}(k_h) = 1$ , for each pair  $(h, c_h)$ . Yang and Dunson (2016) further showed that any conditional probability tensor can be represented as (3) with the parameters satisfying the constraints (a) and (b).

Importantly, when  $k_h = 1$ ,  $\pi_{h, c_h}(1) = 1$  and  $P_{y|\mathbf{c}}(y | c_1, \dots, c_p)$  does not vary with  $c_h$ . The variable  $k_h$  thus determines the inclusion of the  $h^{\text{th}}$  predictor  $c_h$  in the model. When  $k_{y,h} \geq 2$ ,  $c_h$  is an important predictor of  $y$ , and when  $k_{y,h} = 1$ , it does not have any influence on  $y$ . The variable  $k_h$  also determines the number of latent classes for the  $h^{\text{th}}$  predictor  $c_h$ . The number of parameters in such a factorization is given by  $(k_y - 1) \prod_{h=1}^p k_{y,h} + k_y \sum_{h=1}^p (k_{y,h} - 1)$ , which will be much smaller than the number of parameters  $(k_y - 1) \prod_{h=1}^p d_h$  required to specify a full model, if  $\prod_{h=1}^p k_{y,h} \ll \prod_{h=1}^p d_h$ .

Building on these ideas to model the conditional probabilities  $P_{x|\mathbf{c}}(k | c_0, c_1, \dots, c_p)$  in our setting, we let

$$P_{x|\mathbf{c}}(k | c_0, c_1, \dots, c_p) = \sum_{k_0=1}^{k_{x,0}} \sum_{k_1=1}^{k_{x,1}} \cdots \sum_{k_p=1}^{k_{x,p}} \lambda_{x, k_0, k_1, \dots, k_p}(k) \prod_{h=0}^p \pi_{x, h, c_h}(k_h). \quad (3)$$

Additionally, we restrict the probabilities  $\pi_{x, h, c_h}(k_h)$ 's to satisfy  $\pi_{x, h, c_h}(k_h) = 1$  for one  $k_h$  and 0 otherwise, allowing each  $c_h$  to be associated with exactly one latent cluster  $k_h$ , simplifying the model structure and thereby facilitating posterior computation and model interpretability while also maintaining full model flexibility.

Introducing latent variables  $z_{x, \ell, i}$ ,  $z_{x, 0, \ell, i} = z_{x, 0}(c_{0, \ell, i})$ , and  $z_{x, h, \ell, i} = z_{x, h}(c_{h, i})$  for  $h = 1, \dots, p$ , we can rewrite the model as

$$\begin{aligned} (x_{\ell, i} | z_{x, \ell, i} = k) &\sim \text{TN}(x | \mu_k, \sigma_k^2, [A, B]), \\ (z_{x, \ell, i} | z_{x, h, \ell, i} = k_h, h = 0, \dots, p) &\sim \text{Mult}\{\lambda_{x, k_0, k_1, \dots, k_p}(1), \dots, \lambda_{x, k_0, k_1, \dots, k_p}(k_x)\}. \end{aligned}$$

Model specification for the marginal densities of  $\mathbf{x}$  is completed by assigning priors to the model parameters. For the mixture kernels  $\lambda_{x, k_0, k_1, \dots, k_p}$ , we let

$$\lambda_{x, k_0, k_1, \dots, k_p} \sim \text{Dir}\{\alpha_x \lambda_{x, 0}(1), \dots, \alpha_x \lambda_{x, 0}(k_x)\}, \quad \lambda_{x, 0} \sim \text{Dir}(\alpha_{x, 0}, \dots, \alpha_{x, 0}).$$

For the variable selection parameters  $k_{x, h}$ 's, we assign exponentially decaying priors as

$$p_0(k_{x, h}) \propto \exp(-\varphi_x k_{x, h}), \quad h = 0, \dots, p.$$

Large values of  $k_{x, h}$ 's are thus penalized, favoring sparsity. For the mixture atoms, we let

$$(\mu_{x, k}, \sigma_{x, k}^2) \sim \text{TN}(\mu_{x, 0}, \sigma_{x, 0}^2, [A, B]) \times \text{Inv-Ga}(a_{x, \sigma^2}, b_{x, \sigma^2}, [A_{x, \sigma^2}, B_{x, \sigma^2}]),$$

where  $\text{Inv-Ga}(a, b, [A, B])$  denotes an Inverse-Gamma distribution with shape  $a$ , rate  $b$ , restricted to the interval  $[A, B]$ .

We have not imposed any strict identifiability constraints on the mixture components as we are only interested (a) in estimating the overall shapes of the marginal densities, which are robust to overfitting and invariant to label switching, and (b) in selecting the influential predictors, which are determined by the predictors' influences on the overall distribution of the latent  $z_{x,\ell,i}$ 's. Overfitting could be an issue for the latter problem – two mixture components can be close enough to be considered practically the same but two different levels of a covariate  $c_h$  may differently prefer one component to the other, spuriously inferring  $c_h$  to have an important effect on the marginal densities. Extensive numerical experiments, however, suggest that such a situation almost never really happens in practice. Due to the sparsity inducing properties of the priors for the mixture models, in steady states of our MCMC based implementation, the mixture components generally get well separated and the redundant components become near-empty in the sense that practically zero probabilities will get assigned to such components. We are not invoking any notion of a true number of latent components here, but are rather interested in a relatively sparse data adaptive mixture model representation that well approximates the overall shapes of the marginal densities and allows inference about the influences of the predictors on them.

Next, we consider the problem of modeling  $\mathbf{R}_x$ . The correlation structure between the  $\mathbf{x}_i$ 's may certainly vary with the associated predictors  $\mathbf{c}_i$ 's. The problem of modeling such dependence is, however, an extremely difficult one even in the absence of measurement errors and only gets an order of magnitude more difficult when the  $\mathbf{x}_i$ 's are all measured with error. Empirical explorations also do not seem to suggest any real effect here. Practical benefits of accommodating such effects in our model would thus be limited at best, outweighed by the additional computational burden introduced. In this article, we thus assume the correlation structures to remain fixed across all predictor combinations.

We use a model based on spherical coordinate representation of Cholesky factorizations used before in Sarkar *et al.* (2021); Zhang *et al.* (2011) that allows the involved parameters to be treated separately of each other, simplifying posterior computation while guaranteeing the resulting matrix to always be a valid correlation matrix. To keep this article self-contained, we describe the model below. We drop the subscript  $\mathbf{x}$  to keep the notation clean. With  $\mathbf{R} = \mathbf{V}\mathbf{V}^T$ , where  $\mathbf{V}^{d \times d}$  is a lower triangular matrix, we have

$$\mathbf{V} = \begin{pmatrix} v_{1,1} & 0 & \dots & 0 \\ v_{2,1} & v_{2,2} & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ v_{d,1} & v_{d,2} & \dots & v_{d,d} \end{pmatrix}.$$

We have  $r_{\ell,\ell'} = \sum_{k=1}^{\ell} v_{\ell,k}v_{\ell',k}$  for all  $\ell \leq \ell'$ . The restriction that  $\mathbf{R}$  is a correlation matrix then implies  $\sum_{k=1}^{\ell} v_{\ell,k}^2 = 1$  for all  $\ell = 1, \dots, d$ . The restrictions are satisfied by the following

parameterization

$$\begin{aligned}
v_{1,1} &= 1, \\
v_{2,1} &= b_1, \quad v_{2,2} = \sqrt{1 - b_1^2}, \\
v_{3,1} &= b_2 \sin \theta_1, \quad v_{3,2} = b_2 \cos \theta_1, \quad v_{3,3} = \sqrt{1 - b_2^2}, \\
v_{\ell,1} &= b_{\ell-1} \sin \theta_{i_1(\ell)}, \\
v_{\ell,k} &= b_{\ell-1} \cos \theta_{i_1(\ell)} \cos \theta_{i_1(\ell)+1} \dots \cos \theta_{i_1(\ell)+k-2} \sin \theta_{i_1(\ell)+k-1}, \\
&\hspace{15em} \text{for } k = 2, \dots, (\ell - 2), \\
v_{\ell,\ell-1} &= b_{\ell-1} \cos \theta_{i_1(\ell)} \cos \theta_{i_1(\ell)+1} \dots \cos \theta_{i_2(\ell)-1} \cos \theta_{i_2(\ell)}, \quad v_{\ell,\ell} = \sqrt{1 - b_{\ell-1}^2},
\end{aligned}$$

where  $\ell = 4, \dots, d$ ,  $i_1(\ell) = 1 + \{1 + \dots + (\ell - 3)\} = (\ell^2 - 5\ell + 8)/2$  and  $i_2(\ell) = i_1(\ell) + (\ell - 3) = (\ell^2 - 3\ell + 2)/2$ ,  $|b_t| \leq 1$ ,  $t = 1, \dots, (d - 1)$ ,  $|\theta_s| \leq \pi$ ,  $s = 1, \dots, i_2(d)$ . The total number of parameters is  $\{1 + 2 + \dots + (d - 1)\} = d(d - 1)/2$ . We have  $|\mathbf{R}| = |\mathbf{V}|^2 = \prod_{\ell=2}^d v_{\ell,\ell}^2 = \prod_{\ell=1}^{d-1} (1 - b_\ell^2)$ . The model for  $\mathbf{R}$  is completed by assigning uniform priors on  $b_t$ 's and  $\theta_s$ 's

$$b_t \sim \text{Unif}(-1, 1), \quad \theta_s \sim \text{Unif}(-\pi, \pi).$$

Here  $\text{Unif}(a, b)$  denotes a uniform distribution with support  $(a, b)$ .

## 2.2 Modeling the Density $f_{\boldsymbol{\epsilon}|\mathbf{c}}$

As in Section 2.1, we use a Gaussian copula model to specify the density  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$  but the model now has to satisfy mean zero constraints. Specifically, we let

$$\begin{aligned}
f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon} | \mathbf{c}) &= |\mathbf{R}\boldsymbol{\epsilon}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \mathbf{y}\boldsymbol{\epsilon}^T (\mathbf{R}\boldsymbol{\epsilon}^{-1} - \mathbf{I}_d) \mathbf{y}\boldsymbol{\epsilon} \right\} \prod_{\ell=1}^d f_{\boldsymbol{\epsilon},\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c}), \\
&\text{subject to } \mathbb{E}_{f_{\boldsymbol{\epsilon},\ell|\mathbf{c}}}(\epsilon_\ell | \mathbf{c}) = 0, \quad \text{for } \ell = 1, \dots, d.
\end{aligned}$$

Here  $F_{\boldsymbol{\epsilon},\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c}) = \Phi(y_{\boldsymbol{\epsilon},\ell})$  for all  $\ell$  and all  $\mathbf{c}$  with  $F_{\boldsymbol{\epsilon},\ell|\mathbf{c}}$  being the cdf corresponding to  $f_{\boldsymbol{\epsilon},\ell|\mathbf{c}}$ ;  $\mathbf{y}\boldsymbol{\epsilon} = (y_{\boldsymbol{\epsilon},1}, \dots, y_{\boldsymbol{\epsilon},d})^T$ ; and  $\mathbf{R}\boldsymbol{\epsilon}$  is the correlation matrix between the error components.

Following Section 2.1 again, we use predictor dependent mixture models with shared atoms to model the marginal densities  $f_{\boldsymbol{\epsilon},\ell|\mathbf{c}}(\epsilon | c_1, \dots, c_p) = f_{\boldsymbol{\epsilon}|\mathbf{c}}(\epsilon | c_0 = \ell, c_1, \dots, c_p)$  as

$$\begin{aligned}
f_{\boldsymbol{\epsilon}|\mathbf{c}}(\epsilon | c_0, c_1, \dots, c_p) &= \sum_{k=1}^{k_\epsilon} P_{\boldsymbol{\epsilon}|\mathbf{c}}(k | c_0, c_1, \dots, c_p) f_{c\boldsymbol{\epsilon}}(\epsilon | p_{\boldsymbol{\epsilon},k}, \mu_{\boldsymbol{\epsilon},k}, \sigma_{\boldsymbol{\epsilon},k,1}^2, \sigma_{\boldsymbol{\epsilon},k,2}^2), \\
P_{\boldsymbol{\epsilon}|\mathbf{c}}(k | c_0, c_1, \dots, c_p) &= \sum_{k_0=1}^{k_{\boldsymbol{\epsilon},0}} \sum_{k_1=1}^{k_{\boldsymbol{\epsilon},1}} \dots \sum_{k_p=1}^{k_{\boldsymbol{\epsilon},p}} \lambda_{\boldsymbol{\epsilon},k_0,k_1,\dots,k_p}(k) \prod_{h=0}^p \pi_{\boldsymbol{\epsilon},h,c_h}(k_h).
\end{aligned} \tag{4}$$

Here  $f_{c\boldsymbol{\epsilon}}(\epsilon | p, \mu, \sigma_1^2, \sigma_2^2) = \{p \text{ Normal}(\epsilon | \mu_1, \sigma_1^2) + (1 - p) \text{ Normal}(\epsilon | \mu_2, \sigma_2^2)\}$ , with  $\mu_1 = c_1\mu$ ,  $\mu_2 = c_2\mu$ ,  $c_1 = (1 - p)/\{p^2 + (1 - p)^2\}^{1/2}$ ,  $c_2 = -p/\{p^2 + (1 - p)^2\}^{1/2}$ . The zero mean constraint on the errors is satisfied, since  $p\mu_1 + (1 - p)\mu_2 = \{pc_1 + (1 - p)c_2\}\mu = 0$ . Normal densities are included as special cases with  $(p, \mu) = (0.5, 0)$  or  $(0, 0)$  or  $(1, 0)$ . The mixture atoms  $\{(p_{\boldsymbol{\epsilon},k}, \mu_{\boldsymbol{\epsilon},k}, \sigma_{\boldsymbol{\epsilon},k,1}^2, \sigma_{\boldsymbol{\epsilon},k,2}^2)\}_{k=1}^{k_\epsilon}$  are again shared between different predictor combinations to facilitate dimension reduction and borrowing of information.

As in the case of  $\mathbf{x}$ , we use a parsimonious conditional tensor factorization based model for the predictor dependent mixture probabilities  $P_{\epsilon|\mathbf{c}}(k | c_0, c_1, \dots, c_p)$  as

$$P_{\epsilon|\mathbf{c}}(k | c_0, c_1, \dots, c_p) = \sum_{k_0=1}^{k_{\epsilon,0}} \sum_{k_1=1}^{k_{\epsilon,1}} \cdots \sum_{k_p=1}^{k_{\epsilon,p}} \lambda_{\epsilon,k_0,k_1,\dots,k_p}(k) \prod_{h=0}^p \pi_{\epsilon,h,c_h}(k_h),$$

where the parameters satisfy  $\sum_{k=1}^{k_{\epsilon}} \lambda_{\epsilon,k_0,k_1,\dots,k_p}(k) = 1$  for all  $(k_0, k_1, \dots, k_p)$  and  $\pi_{\epsilon,h,c_h}(k_h) = 1$  for one  $k_h$  and 0 otherwise.

Introducing latent variables  $z_{\epsilon,\ell,i,j}$ ,  $z_{\epsilon,0,\ell,i,j} = z_{\epsilon,0}(c_{0,\ell,i})$ , and  $z_{\epsilon,h,\ell,i,j} = z_{\epsilon,h}(c_{h,i})$  for  $h = 1, \dots, p$ , we can rewrite the model as

$$\begin{aligned} (\epsilon_{\ell,i,j} | z_{\epsilon,\ell,i,j} = k) &\sim f_{\mathcal{C}\epsilon}(\epsilon_{\ell,i,j} | p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2), \\ (z_{\epsilon,\ell,i,j} | z_{\epsilon,h,\ell,i,j} = k_h, h = 0, \dots, p) &\sim \text{Mult}\{\lambda_{\epsilon,k_0,k_1,\dots,k_p}(1), \dots, \lambda_{\epsilon,k_0,k_1,\dots,k_p}(k_{\epsilon})\}. \end{aligned}$$

We assume hierarchical Dirichlet priors for  $\boldsymbol{\lambda}_{\epsilon,k_0,k_1,\dots,k_p}$  as

$$\boldsymbol{\lambda}_{\epsilon,k_0,k_1,\dots,k_p} \sim \text{Dir}\{\alpha_{\epsilon}\lambda_{\epsilon,0}(1), \dots, \alpha_{\epsilon}\lambda_{\epsilon,0}(k_{\epsilon})\}, \quad \boldsymbol{\lambda}_{\epsilon,0} \sim \text{Dir}(\alpha_{\epsilon,0}, \dots, \alpha_{\epsilon,0}).$$

For the variable selection parameters  $k_{\epsilon,h}$ 's, we assign exponentially decaying priors as

$$p_0(k_{\epsilon,h}) \propto \exp(-\varphi_{\epsilon}k_{\epsilon,h}), \quad h = 0, \dots, p.$$

We assume non-informative priors for  $(p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2)$  as

$$\begin{aligned} (p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2) &\sim p_0(p_{\epsilon,k}) p_0(\mu_{\epsilon,k}) p_0(\sigma_{\epsilon,k,1}^2) p_0(\sigma_{\epsilon,k,2}^2) \\ &= \text{Unif}(0, 1) \text{Normal}(0, \sigma_{\epsilon,\mu}^2) \text{Inv-Ga}(a_{\epsilon}, b_{\epsilon}) \text{Inv-Ga}(a_{\epsilon}, b_{\epsilon}), \end{aligned}$$

where  $\text{Unif}(\ell, u)$  denotes a uniform distribution on the interval  $[\ell, u]$ .

As in the case of  $\mathbf{R}_{\mathbf{x}}$ , we assume  $\mathbf{R}_{\boldsymbol{\epsilon}}$  is independent of  $\mathbf{c}$  and let  $\mathbf{R}_{\boldsymbol{\epsilon}}^{d \times d} = ((r_{\boldsymbol{\epsilon},\ell,\ell'})) = \mathbf{V}_{\boldsymbol{\epsilon}} \mathbf{V}_{\boldsymbol{\epsilon}}^{\text{T}}$  and parameterize the elements of  $\mathbf{V}_{\boldsymbol{\epsilon}}$  using spherical coordinates. We assign uniform priors on  $b_{\boldsymbol{\epsilon},t}, t = 1, \dots, d-1$  and  $\theta_{\boldsymbol{\epsilon},s}, s = 1, \dots, i_2(d)$

$$b_{\boldsymbol{\epsilon},t} \sim \text{Unif}(-1, 1), \quad \theta_{\boldsymbol{\epsilon},s} \sim \text{Unif}(-\pi, \pi).$$

## 2.3 Modeling the Variance Functions $v_{\ell}(x_{\ell})$

We model the variance functions  $v_{\ell}(x) = s_{\ell}^2(x)$  as flexible mixtures of B-splines

$$\begin{aligned} v_{\ell}(x) &= s_{\ell}^2(x) = \sum_{j=1}^J b_j(x) \exp(\vartheta_{\ell,j}) = \mathbf{B}_J(x) \exp(\boldsymbol{\vartheta}_{\ell}), \\ \boldsymbol{\vartheta}_{\ell} &\sim \text{MVN}_J\{\mathbf{0}, (\boldsymbol{\Sigma}_{\vartheta,0}^{-1} + \mathbf{P}_{\vartheta,0}/\sigma_{\vartheta,0}^2)^{-1}\}, \end{aligned}$$

where  $\boldsymbol{\vartheta}_{\ell} = (\vartheta_{\ell,1}, \dots, \vartheta_{\ell,J})^{\text{T}}$  are spline coefficients,  $\text{MVN}_J(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  denotes a  $J$  dimensional multivariate normal distribution with mean  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ . We choose  $\mathbf{P}_{\vartheta,0} = \mathbf{D}_{\vartheta,0}^{\text{T}} \mathbf{D}_{\vartheta,0}$ , where the  $(J-2) \times J$  matrix  $\mathbf{D}_{\vartheta,0}$  is such that  $\mathbf{D}_{\vartheta,0} \boldsymbol{\vartheta}_{\ell}$  computes the second order differences in  $\boldsymbol{\vartheta}_{\ell}$ . The model thus penalizes  $\sum_j (\nabla^2 \vartheta_{\ell,j})^2 = \boldsymbol{\vartheta}_{\ell}^{\text{T}} \mathbf{P}_{\vartheta,0} \boldsymbol{\vartheta}_{\ell}$ , the sum of squares of second order differences in  $\boldsymbol{\vartheta}_{\ell}$  (Eilers and Marx, 1996). The variance parameter  $\sigma_{\vartheta,0}^2$  models the smoothness of the variance functions, smaller  $\sigma_{\vartheta,0}^2$  inducing smoother functions.

The methodology proposed here builds on a few diverse topics. A high-level overview of the different model components is presented as a box-summary in the supplementary materials for easy reference. Brief reviews of a few these topics are also presented in the supplementary materials for easy reference – conditional copula models in Section S.1, conditional tensor factorization in Section S.2, and mixtures with shared atoms in Section S.3.

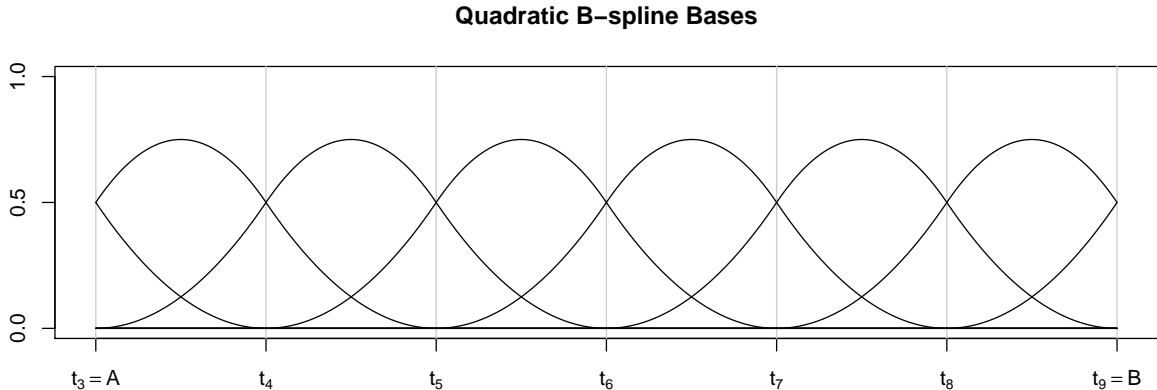


Figure 6: Plot of 9 quadratic ( $q = 2$ ) B-splines on  $[A, B]$  defined using 11 knot points that divide  $[A, B]$  into  $K = 6$  equal subintervals.

### 3 Applications in Nutritional Epidemiology

In this section, we discuss the results of our method applied to the EATS data set. Specifically, we consider the problem of estimating the distributions of long-term average daily intakes of iron, magnesium and sodium.

Figure 8 shows the estimated inclusion probabilities of different predictors in the models for  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$ . We recall that a predictor  $c_h$  is considered important if its levels form at least two clusters, that is,  $k_h \geq 2$ . Our MCMC based implementation produces estimates of posterior distribution of the  $k_h$ 's, accommodating uncertainty in variable election. Using a median probability rule (Barbieri and Berger, 2004), that is, selecting predictors with at least 50% posterior probability of being included in the model, the set of significant predictors for the density of main interest  $f_{\mathbf{x}|\mathbf{c}}$  is found to comprise the dimension labels ( $c_0$ ) and sex ( $c_1$ ). For  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$ , however, none of the potential predictors were found to be significant. The significance of gender is consistent with common knowledge that the men on average consume more than women, as was also clearly seen the exploratory analysis of Figure 2. We must not immediately extend the non-significance of age and ethnicity to the entire population and claim that long term dietary intakes do not vary with these covariates at all. Based on the finite size EATS data set, however, there is insufficient evidence to claim otherwise. The error distributions  $f_{\boldsymbol{\epsilon},\ell|\mathbf{c}}$  all collapsing together, while not immediately apparent from the




Figure 7: Graphical model depicting the dependency structure in the Bayesian hierarchical covariate informed density deconvolution model described in Section 2. The unfilled and the shaded nodes signify latent and observable variables, respectively. Subject and replicate subscripts ( $i$  and  $j$ , respectively) are suppressed to keep the notation clean.

histograms of the ‘residuals’ in Figure 2, is consistent with them having very similar right skewed shapes observed in separate univariate analyses (not shown here).

Figures 9 and 10 illustrate how the redundant mixture components become near-empty after reaching steady states in our MCMC based implementation. While we started with twenty mixture components, only six are finally being used for modeling the densities  $f_{x,\ell|\mathbf{c}}$ . Likewise, while we again started with twenty mixture components, only three are finally being used for modeling the densities  $f_{\epsilon,\ell|\mathbf{c}}$ .

Figures 9 and 10 additionally show how the mixture component specific parameters get shared across different dietary components and predictor combinations in our model and how the associated mixture probabilities vary across these combinations. For the densities  $f_{x,\ell|\mathbf{c}}$ , the mixture probabilities clearly vary significantly between men and women as well as between different dietary components, hence these variables were selected as important by our method. For the densities  $f_{\epsilon,\ell|\mathbf{c}}$ , on the other hand, the mixture probabilities are very similar between men and women as well as between different dietary components, hence these covariates were adjudged non-significant by our method.

Figure 11 shows the estimated densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  superimposed over histograms of the corresponding estimated  $x_\ell$ ’s obtained by our method. Figure S.7 in the supplementary materials shows the estimated joint densities  $f_{\mathbf{x},\ell_1,\ell_2|\mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  obtained by our method in






Figure 8: Results for the EATS data set showing the estimated probabilities of different numbers of clusters of the associated predictors’ levels being included in the model. The left panel shows the results for modeling the densities  $f_{\mathbf{x}|\mathbf{c}}$ . The right panel shows the results for modeling the densities  $f_{\epsilon|\mathbf{c}}$ . At the median 0.5 probability level, in the left panel, the dietary component labels and the sex of the subjects are important predictors for modeling the densities  $f_{\mathbf{x}|\mathbf{c}}$ , whereas in the right panel, none of the predictors are found to be important for modeling the densities  $f_{\epsilon|\mathbf{c}}$ .

the off-diagonal panels. The results suggest the model to provide a good fit for the EATS data, including especially being able to capture the heavily skewed consumption distributions for men with heavy right tails. In comparison, the distributions for women look more symmetric and have much lighter tails.

We compare our results with estimates produced by the method by Zhang *et al.* (2011). Zhang *et al.* (2011), strictly speaking, is not a principled deconvolution approach but rather a multi-stage pseudo-Bayesian mixed model approach. They use Box-Cox transformations (Box and Cox, 1964) of the recalls  $w_{\ell,i,j}$  separately for each component. The rest of the analysis is done conditional on the estimated Box-Cox parameters, assuming the transformed variables to follow a linear mixed model, a subject specific random effect component and a covariate dependent linear fixed effects component with no interaction terms and an independent error component. All covariates are included in the model as there is also no mechanism to select the important ones. The random effects components and the errors are modeled using multivariate normal distributions. Multivariate normal priors are also assumed for the fixed effects regression coefficients. Estimates of the long-term intakes are then obtained via individual transformations back to the original scale. As shown in Sarkar *et al.* (2014, 2021), Box-Cox transformations for surrogate observations have severe limitations, including almost never being able to produce transformed surrogates that conform to the assumed parametric assumptions, including normality, homoscedasticity, and independence of the errors. Single component multivariate normal models are thus highly inadequate for the densities even after transformations. Estimates of the marginal densities in the original




Figure 9: Results for the EATS data set. These results correspond to the final MCMC iteration but are representative of other iterations in steady state. The left panel shows the component specific parameters  $(\mu_{x,k}, \sigma_{x,k}^2)$  for the six mixture components that were actually used to model the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . The right panel shows the associated ‘empirical’ mixture probabilities  $\hat{p}_x(k | c_0, c_1, \dots, c_p) = \sum_{i=1}^n 1\{z_{x,\ell,i} = k, c_{0,\ell,i} = c_0, c_{1,\ell,i} = c_1, \dots, c_{p,\ell,i} = c_p\} / n$  for men and women and for the three dietary components, namely iron, magnesium and sodium, from left to right. Results for different combinations of dietary component and gender are shown here as they are the only predictors found important for  $\mathbf{x}$ . The mixture probabilities vary significantly between different predictor combinations.




Figure 10: Results for the EATS data set. These results correspond to the final MCMC iteration but are representative of other iterations in steady state. The left panel shows the component specific parameters  $(p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2)$  for the three mixture components used to model the densities  $f_{\epsilon,\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c})$ . The right panel shows the associated ‘empirical’ mixture probabilities  $\hat{p}_\epsilon(k | c_0, c_1, \dots, c_p) = \sum_{i=1}^n \sum_{j=1}^{m_i} 1\{z_{\epsilon,\ell,i,j} = k, c_{0,\ell,i} = c_0, c_{1,\ell,i} = c_1, \dots, c_{p,\ell,i} = c_p\} / \sum_{i=1}^n m_i$  for men and women and for the three dietary components, namely iron, magnesium and sodium, in that order, from left to right. Results for different combinations of dietary component and gender are shown here for illustrative purposes even though none of the predictors were found important for  $\epsilon$ . The mixture probabilities do not vary significantly between different predictor combinations.




Figure 11: Results for the EATS data set obtained by our method. From top to bottom, the left panels show the estimated conditional densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  of iron, magnesium and sodium, respectively. The right panels show the associated conditional distributions of the scaled errors  $f_{\epsilon,\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c})$ . Results for different dietary component and gender combinations are shown here as they are the only predictors found important for modeling the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . Also shown are the associated variance functions  $v_\ell(x_\ell) = s_\ell^2(x_\ell)$ .

scale are thus obtained not by applying what the model actually implies but by applying a univariate kernel density estimation method on the estimated intakes in the observed scale thereby mitigating the highly restrictive effects of the inherent parametric assumptions.

For clarity, we summarize the estimates obtained by the method of Zhang *et al.* (2011)

separately in Figure S.8, moved to the supplementary materials for space limitations. The shapes of the estimated densities are in general agreement with those produced by our method. The fixed effects regression coefficient estimates are presented in Table S.1 in the supplementary material. With no mechanism to select the important predictors, all covariates are included in the model. Taking the exclusion of zero from a 90% central credible interval to be a (somewhat ad-hoc) post-processing rule to determine the significance of the associated predictor, we can eliminate some, but a good number of coefficients associated with race and age still remain included in the model. Exploratory analysis (Figure 3) suggests these effects may still be spurious - a result of the presence of the ‘missing’ group which is of very small size with only five subjects but includes replicates that look very different from the remaining groups (Figure S.2 in the supplementary material).

The posterior of our proposed Bayesian hierarchical method did not include age as an important predictor and only included race as important in a small percentage of the MCMC iterations. Borrowing information across different predictor groups, it is very robust to the presence of small outlying groups such as the ‘missing’ race.

## 4 Discussion

In this article, we considered the problem of multivariate density deconvolution in the presence of categorical predictors. The problem is important in nutritional epidemiology for estimating long-term intakes of regularly consumed dietary components in the presence of associated demographic variables age, sex and ethnicity. We developed a copula based deconvolution approach that focuses on the marginals first and then models the dependence among the components to build the joint densities. Our proposed method not only allows the densities to vary flexibly with the associated predictors but also allows automatic selection of the most influential predictors. Importantly, our proposed method also allows the sets of predictors influencing the density of interest and the density of the measurement errors to potentially be different. Applied to our motivating nutritional epidemiology data set, we found gender to be an important predictor for the density of long term average intakes of different dietary components.

The applicability of the methodology developed here for covariate informed multivariate densities is not restricted to deconvolution problems but the different model components can be adapted to other important problems in statistics as well. For instance, the methodology developed in Section 2.1 for modeling  $f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} | \mathbf{c})$  can be straightforwardly applied to the problem of ordinary multivariate density estimation without measurement errors in the presence of associated potentially high-dimensional precisely measured covariates. Likewise, the methodology developed in Section 2.2 for modeling  $f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon} | \mathbf{c})$  can be straightforwardly applied to modeling covariate dependent regression errors in the presence of associated potentially high-dimensional precisely measured covariates. Section S.8 in the supplementary materials provides additional brief discussions and some simulations evaluating the perfor-

mance of our method for ordinary density estimation problems. More rigorous expositions of these problems will be pursued elsewhere.

The methodology developed here is semiparametric in nature, where some model components are highly flexible while some others are highly parametric. At a first glance, the parametric choices may be perceived as restrictive. Deconvolution problems are, however, well known to be extremely difficult ones and methods that work for measurement error free settings may not always work for measurement error problems. For example, it was shown in Sarkar *et al.* (2014) that methods that could be adapted to allow all aspects of the error distributions to vary flexibly with covariates, e.g., Chung and Dunson (2009), are not numerically feasible for measurement errors even for moderately large data sets like the EATS, and a multiplicative structural assumption  $u = s(x)\epsilon$ , as considered in our article, although in theory more restrictive, is in fact a highly efficient practical choice. Likewise, the assumption of covariate independent Gaussian copula can, in principle, be relaxed to include covariates as well as other copula classes. In practice, however, these problems are extremely challenging even in measurement error free scenarios (dos Santos Silva and Lopes, 2008). The Gaussian copula is easy to understand, interpret, and implement and hence is an effective practical choice for deconvolution problems.

The method of course has other important limits. The trick used here to include the component labels as the levels of a categorical covariate allowed us to greatly simplify the tensor decomposition computations but also restricted each component to be influenced by the same set of important covariates. An important direction for future research is to relax this restriction to allow different sets of important predictors for different dietary components using more flexible partition models. Our previous work in Sarkar *et al.* (2021) also showed that mixtures of truncated normals do not work well for zero-inflated recall data for episodically consumed dietary components but requires other modeling strategies to accommodate the sharp boundaries of the densities encountered in such problems. Adaptations for episodic components, however, forms a crucial step forward toward a more sophisticated framework for estimating the Healthy Eating Index (HEI, <https://www.fns.usda.gov/resource/healthy-eating-index-hei>), a performance measure developed by the US Department of Agriculture (USDA) to assess and promote healthy diets (Guenther *et al.*, 2008; Krebs-Smith *et al.*, 2018), forming another important direction for future research.

## Supplementary Materials

The supplementary materials details the choice of hyper-parameters and the MCMC algorithm used to sample from the posterior. R programs implementing the deconvolution methods developed in this article are included as separate files in the supplementary material. The EATS data analyzed in Section 3 can be accessed from National Cancer Institute by arranging a Material Transfer Agreement. A simulated data set, simulated according to one of the designs described in Section S.6, and a ‘readme’ file providing additional details are also included in the supplementary material.

## Acknowledgments

We thank the University of Texas Advanced Computing Center (TACC) for providing computing resources that contributed to the research reported here.

## References

- Barbieri, M. M. and Berger, J. O. (2004). Optimal predictive model selection. *The annals of statistics*, **32**, 870–897.
- Box, G. E. and Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B*, **26**, 211–252.
- Buonaccorsi, J. P. (2010). *Measurement Error : Models, Methods, and Applications*. Chapman & Hall/CRC interdisciplinary statistics series. CRC Press, Boca Raton.
- Carroll, R. J., Ruppert, D., Stefanski, L. A., and Crainiceanu, C. M. (2006). *Measurement Error in Nonlinear Models: A Modern Perspective, Second Edition*. Chapman and Hall, Boca Raton.
- Chung, Y. and Dunson, D. B. (2009). Nonparametric Bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association*, **104**, 1646–1660.
- dos Santos Silva, R. and Lopes, H. F. (2008). Copula, marginal distributions and model selection: a Bayesian note. *Statistics and Computing*, **18**, 313–320.
- Eilers, P. H. C. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical Science*, **11**, 89–121.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, **90**, 577–588.
- Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. Springer, New York.
- Guenther, P. M., Reedy, J., and Krebs-Smith, S. M. (2008). Development of the Healthy Eating Index-2005. *Journal of the American Dietetic Association*, **108**, 1896–1901.
- Krebs-Smith, S. M., Pannucci, T. E., Subar, A. F., Kirkpatrick, S. I., Lerman, J. L., Tooze, J. A., Wilson, M. M., and Reedy, J. (2018). Update of the Healthy Eating Index: HEI-2015. *Journal of the Academy of Nutrition and Dietetics*, **118**, 1591–1602.
- Sarkar, A., Mallick, B. K., Staudenmayer, J., Pati, D., and Carroll, R. J. (2014). Bayesian semiparametric density deconvolution in the presence of conditionally heteroscedastic measurement errors. *Journal of Computational and Graphical Statistics*, **24**, 1101–1125.
- Sarkar, A., Pati, D., Chakraborty, A., Mallick, B. K., and Carroll, R. J. (2018). Bayesian semiparametric multivariate density deconvolution. *Journal of the American Statistical Association*, **113**, 401–416.

- Sarkar, A., Pati, D., Mallick, B. K., and Carroll, R. J. (2021). Bayesian copula density deconvolution for zero-inflated data in nutritional epidemiology. *Journal of the American Statistical Association*, **116**, 1075–1087.
- Staudenmayer, J., Ruppert, D., and Buonaccorsi, J. R. (2008). Density estimation in the presence of heteroscedastic measurement error. *Journal of the American Statistical Association*, **103**, 726–736.
- Su, Y., Bhattacharya, A., Zhang, Y., Chatterjee, N., and Carroll, R. J. (2020). Non-parametric Bayesian deconvolution of a symmetric unimodal density. arXiv preprint arXiv:2002.07255.
- Subar, A. F., Thompson, F. E., Kipnis, V., Midthune, D., Hurwitz, P., McNutt, S., McIntosh, A., and Rosenfeld, S. (2001). Comparative validation of the Block, Willett, and National Cancer Institute food frequency questionnaires - The Eating at America's Table Study. *American Journal of Epidemiology*, **154**, 1089–1099.
- Yang, Y. and Dunson, D. B. (2016). Bayesian conditional tensor factorization for high-dimensional classification. *Journal of the American Statistical Association*, **111**, 656–669.
- Zhang, S., Midthune, D., Guenther, P. M., Krebs-Smith, S. M., Kipnis, V., Dodd, K. W., Buckman, D. W., Tooze, J. A., Freedman, L., and Carroll, R. J. (2011). A new multivariate measurement error model with zero-inflated dietary data, and its application to dietary assessment. *Annals of Applied Statistics*, **5**, 1456–1487.

---

Supplementary Materials for

**Bayesian Semiparametric  
Covariate Informed Multivariate Density  
Deconvolution**

Abhra Sarkar

abhra.sarkar@utexas.edu

Department of Statistics and Data Sciences, The University of Texas at Austin  
2317 Speedway D9800, Austin, TX 78712-1823, USA

Supplementary materials discuss copula models where the marginal densities vary with associated precisely measured covariates, a brief review of tensor factorization models for easy reference, the choice of hyper-parameters and details of the MCMC algorithm we designed to sample from the posterior. Supplementary materials also present some additional figures summarizing the analysis of the EATS data set. Separate files additionally include a synthetic data set simulated according to one of the designs described in Section S.6, a ‘readme’ file providing additional details of this data set, and R programs implementing the deconvolution method developed in this article.



Method Overview

The methodology proposed here builds on many diverse topics. A high-level overview of the different model components is presented here for easy reference.

1. **Model:**

$$\begin{aligned} \mathbf{w}_{i,j} &= \mathbf{x}_i + \mathbf{u}_{i,j}, & \mathbf{u}_{i,j} &= \mathbf{S}(\mathbf{x}_i)\boldsymbol{\epsilon}_{i,j}, \\ (\mathbf{x}_i \mid \mathbf{c}_i) &\sim f_{\mathbf{x}|\mathbf{c}}(\mathbf{x}_i \mid \mathbf{c}_i), & (\boldsymbol{\epsilon}_{i,j} \mid \mathbf{c}_i) &\sim f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon}_{i,j} \mid \mathbf{c}_i) \text{ with } \mathbb{E}(\boldsymbol{\epsilon}_{i,j} \mid \mathbf{c}_i) = \mathbf{0}, \\ \mathbf{S}(\mathbf{x}_i) &= \text{diag}\{s_1(x_{1,i}), \dots, s_d(x_{d,i})\}. \end{aligned}$$

(a) **Model  $f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} \mid \mathbf{c})$  using copula (Section 2.1 in the main paper).**

$$f_{\mathbf{x}|\mathbf{c}}(\mathbf{x}_i \mid \mathbf{c}_i) = c(\mathbf{x}_i) \prod_{\ell} f_{x,\ell|\mathbf{c}}(x_{\ell,i} \mid \mathbf{c}_i).$$

- i. Model the marginals  $f_{x,\ell|\mathbf{c}}$  using mixtures of truncated normals with covariate informed mixture probabilities with atoms shared between different components  $\ell$ .
- ii. Model the covariate informed mixture probabilities using conditional tensor factorization that also allows identification of the important covariates.
- iii. Model the dependence function  $c(\mathbf{x})$  using a Gaussian copula and its correlation matrix using a polar coordinate parametrization.

(b) **Model  $f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon} \mid \mathbf{c})$  using copula (Section 2.2 in the main paper).**

$$f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon}_{i,j} \mid \mathbf{c}_i) = c(\boldsymbol{\epsilon}_{i,j}) \prod_{\ell} f_{\epsilon,\ell|\mathbf{c}}(\epsilon_{\ell,i,j} \mid \mathbf{c}_i).$$

- i. Model the marginals  $f_{\epsilon,\ell|\mathbf{c}}$  using mixtures of two-component normals centered to have mean zero with covariate informed mixture probabilities with atoms shared between different components  $\ell$ .
- ii. Model the covariate informed mixture probabilities using conditional tensor factorization that also allows identification of the important covariates.
- iii. Model the dependence function  $c(\boldsymbol{\epsilon})$  using a Gaussian copula and its correlation matrix using a polar coordinate parametrization.

(c) **Model  $\mathbf{S}(\mathbf{x})$  using B-splines (Section 2.3 in the main paper).** Model each component  $s_{\ell}^2(x_{\ell,i})$  using mixtures of B-splines with large second order differences between the adjacent coefficients penalized to induce smoothness.

2. **Prior (in relevant sections in the main paper).** Priors, assigned to different model parameters, are described alongside the model components. The choice of prior hyper-parameters is discussed in Section S.5 in the supplementary materials.
3. **Posterior (Section S.5 in the supplementary materials).** The joint posterior is complex and cannot be computed in closed form nor can be easily sampled from. Inference is based on samples from the posterior using an MCMC algorithm.

## S.1 Conditional Copula Models

The literature on copula models is enormous. See, for example, Nelsen (2007); Joe (2015); Shemyakin and Kniazev (2017) and the references therein.

A function  $\mathbb{C}(\mathbf{u}) = \mathbb{C}(u_1, \dots, u_d) : [0, 1]^d \rightarrow [0, 1]$  is called a copula if  $\mathbb{C}(\mathbf{u})$  is a continuous cumulative distribution function (cdf) on  $[0, 1]^d$  such that each marginal is a uniform cdf on  $[0, 1]$ . That is, for any  $\mathbf{u} \in [0, 1]^d$ ,  $\mathbb{C}(\mathbf{u}) = \mathbb{C}(u_1, \dots, u_d) = \Pr(u_1 \leq u_1, \dots, u_d \leq u_d)$  with  $\mathbb{C}(1, \dots, 1, u_i, 1, \dots, 1) = \Pr(u_i \leq u_i) = u_i, i = 1, \dots, d$ . If  $\{x_i\}_{i=1}^d$  are absolutely continuous random variables having marginal cdf  $\{z_i(x_i)\}_{i=1}^d$  and marginal probability density functions (pdf)  $\{z_i(x_i)\}_{i=1}^d$ , joint cdf  $H(x_1, \dots, x_d)$  and joint pdf  $h(x_1, \dots, x_d)$ , then a copula  $\mathbb{C}$  can be defined in terms of  $H$  as  $\mathbb{C}(u_1, \dots, u_d) = H(x_1, \dots, x_d)$  where  $u_i = z_i(x_i), i = 1, \dots, d$ . It follows that  $h(x_1, \dots, x_d) = c(u_1, \dots, u_d) \prod_{i=1}^d z_i(x_i)$ , where  $c(u_1, \dots, u_d) = \partial^d \mathbb{C}(u_1, \dots, u_d) / (\partial u_1 \dots \partial u_d)$ . This defines a copula density  $c(\mathbf{u})$  in terms of the joint and marginal pdfs of  $\{x_i\}_{i=1}^d$  as

$$c(u_1, \dots, u_d) = h(x_1, \dots, x_d) / \prod_{i=1}^d z_i(x_i). \quad (\text{S.1})$$

Conversely, if  $\{v_i\}_{i=1}^d$  are continuous random variables having fixed marginal cdfs  $\{F_i(v_i | c_1, \dots, c_p)\}_{i=1}^d$ , then their joint cdf  $F(v_1, \dots, v_d | c_1, \dots, c_p)$ , with a dependence structure introduced through a copula  $\mathbb{C}$ , can be defined as

$$F(v_1, \dots, v_d | c_1, \dots, c_p) = \mathbb{C}\{F_1(v_1 | c_1, \dots, c_p), \dots, F_d(v_d | c_1, \dots, c_p)\} = \mathbb{C}(u_1, \dots, u_d) \quad (\text{S.2})$$

where  $u_i = F_i(v_i | c_1, \dots, c_p), i = 1, \dots, d$ . If  $\{v_i\}_{i=1}^d$  have marginal densities  $\{f_i(v_i | c_1, \dots, c_p)\}_{i=1}^d$ , then from (S.2) it follows that the joint density  $f(v_1, v_2, \dots, v_d | c_1, \dots, c_p)$  is given by

$$f(v_1, \dots, v_d | c_1, \dots, c_p) = c(u_1, \dots, u_d) \prod_{i=1}^d f_i(v_i | c_1, \dots, c_p). \quad (\text{S.3})$$

With  $F_i(v_i | c_1, \dots, c_p) = u_i = z_i(x_i), i = 1, \dots, d$ , substitution of the copula density (S.1) into (S.3) gives

$$\begin{aligned} f(v_1, \dots, v_d | c_1, \dots, c_p) &= c(u_1, \dots, u_d) \prod_{i=1}^d f_i(v_i | c_1, \dots, c_p) \\ &= \left\{ \frac{h(x_1, \dots, x_d)}{\prod_{i=1}^d z_i(x_i)} \right\} \prod_{i=1}^d f_i(v_i | c_1, \dots, c_p). \end{aligned} \quad (\text{S.4})$$

Equation (S.2) can be used to define flexible multivariate dependence structure using standard known multivariate densities (Sklar, 1959). Let  $\text{MVN}_d(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  denote a  $p$ -variate normal distribution with mean vector  $\boldsymbol{\mu}$  and positive semidefinite covariance matrix  $\boldsymbol{\Sigma}$ . An important case is  $\mathbf{x} = (x_1, \dots, x_d)^T \sim \text{MVN}_d(\mathbf{0}, \mathbf{R})$ , where  $\mathbf{R}$  is a correlation matrix. In this case,  $\mathbb{C}(u_1, \dots, u_d | \mathbf{R}) = \Phi_d\{\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d) | \mathbf{R}\}$ , where  $\Phi(x) = \Pr\{X \leq x | X \sim \text{Normal}(0, 1)\}$  and  $\Phi_d(x_1, \dots, x_d | \mathbf{R}) = \Pr\{x_1 \leq x_1, \dots, x_d \leq x_d | \mathbf{x} \sim \text{MVN}_d(\mathbf{0}, \mathbf{R})\}$ . If  $\mathbf{x} \sim N_d(\mathbf{0}, \boldsymbol{\Sigma})$ , where  $\boldsymbol{\Sigma} = ((\sigma_{i,j}))$  is a covariance matrix with  $\sigma_{ii} = \sigma_i^2$ , then defining  $\boldsymbol{\Lambda} = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$  and  $\mathbf{y} = \boldsymbol{\Lambda}^{-\frac{1}{2}} \mathbf{x}$  and noting that  $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{1/2} \mathbf{R} \boldsymbol{\Lambda}^{1/2}$ , we have

$$\begin{aligned} c(u_1, \dots, u_d) &= \text{MVN}_d(\mathbf{x} | \mathbf{0}, \boldsymbol{\Sigma}) / \text{MVN}_d(\mathbf{x} | \mathbf{0}, \boldsymbol{\Lambda}) = |\boldsymbol{\Lambda}|^{1/2} |\boldsymbol{\Sigma}|^{-1/2} \exp\{-\mathbf{x}^T (\boldsymbol{\Sigma}^{-1} - \boldsymbol{\Lambda}^{-1}) \mathbf{x} / 2\} \\ &= |\mathbf{R}|^{-1/2} \exp\{-\mathbf{y}^T (\mathbf{R}^{-1} - \mathbf{I}_d) \mathbf{y} / 2\} = \text{MVN}_d(\mathbf{y} | \mathbf{0}, \mathbf{R}) / \text{MVN}_d(\mathbf{y} | \mathbf{0}, \mathbf{I}_d). \end{aligned}$$

Sticking to the standard normal case, a flexible dependence structure between random variables  $\{v_i\}_{i=1}^d$  with given marginals  $\{F_i(v_i | c_1, \dots, c_p)\}_{i=1}^d$  may thus be obtained assuming a Gaussian distribution on the latent random variables  $\{y_i\}_{i=1}^d$  obtained through the transformations  $F_i(v_i | c_1, \dots, c_p) = u_i = \Phi(y_i), i = 1, \dots, d$ . The joint density of  $\mathbf{v} = (v_1, \dots, v_d | c_1, \dots, c_p)^T$  is then given by

$$f(v_1, \dots, v_d | c_1, \dots, c_p) = c(u_1, \dots, u_d) \prod_{i=1}^d f_i(v_i | c_1, \dots, c_p) = \frac{\text{MVN}_d(\mathbf{y} | \mathbf{0}, \mathbf{R})}{\text{MVN}_d(\mathbf{y} | \mathbf{0}, \mathbf{I}_d)} \prod_{i=1}^d f_i(v_i | c_1, \dots, c_p).$$

We have

$$\begin{aligned} & \Pr(v_1 \leq v_1, \dots, v_d \leq v_d | c_1, \dots, c_p) \\ &= \Pr[y_1 \leq \Phi^{-1}\{F_1(v_1 | c_1, \dots, c_p)\}, \dots, y_d \leq \Phi^{-1}\{F_d(v_d | c_1, \dots, c_p)\} | \mathbf{Y} \sim \text{MVN}_d(\mathbf{0}, \mathbf{R})]. \end{aligned}$$

For  $q \leq d$ , with  $(y_1, \dots, y_q)^T \sim \text{MVN}_q(\mathbf{0}, \mathbf{R}_q)$ , we then have

$$\begin{aligned} & \Pr(v_1 \leq v_1, \dots, v_q \leq v_q | c_1, \dots, c_p) \\ &= \Pr[y_1 \leq \Phi^{-1}\{F_1(v_1 | c_1, \dots, c_p)\}, \dots, y_q \leq \Phi^{-1}\{F_q(v_q | c_1, \dots, c_p)\} | \mathbf{y}_{1:q} \sim \text{MVN}_q(\mathbf{0}, \mathbf{R}_q)], \end{aligned}$$

implying that the density of  $(v_1, \dots, v_q)$  will be

$$\begin{aligned} f(v_1, \dots, v_q | c_1, \dots, c_p) &= c(u_1, \dots, u_q) \prod_{i=1}^q f_i(v_i | c_1, \dots, c_p) \\ &= \frac{\text{MVN}_q(\mathbf{y}_{1:q} | \mathbf{0}, \mathbf{R}_q)}{\text{MVN}_q(\mathbf{y}_{1:q} | \mathbf{0}, \mathbf{I}_q)} \prod_{i=1}^q f_i(v_i | c_1, \dots, c_p). \end{aligned}$$

## S.2 Conditional Tensor Factorization Models

There is a vast literature on tensor factorization techniques, the two most popular approaches being the parallel factor analysis (PARAFAC) and the higher order singular value decomposition (HOSVD). The PARAFAC approach Harshman (1970) decomposes a  $d_1 \times \dots \times d_p$  dimensional tensor  $\mathbf{M} = \{m_{c_1, \dots, c_p}\}$  as the sum of rank one tensors as

$$m_{c_1, \dots, c_p} = \sum_{k=1}^{k_m} g_k \prod_{h=1}^p u_{h, c_h}(k). \tag{S.5}$$


In contrast, the HOSVD approach, proposed by Tucker (1966) for three way tensors and extended to multi-way tensors of arbitrary orders by De Lathauwer *et al.* (2000), would factorize  $\mathbf{M}$  as

$$m_{c_1, \dots, c_p} = \sum_{k_1=1}^{k_{m,1}} \dots \sum_{k_p=1}^{k_{m,p}} g_{k_1, \dots, k_p} \prod_{h=1}^p u_{h, c_h}(k_h), \tag{S.6}$$


where  $\mathbf{G} = \{g_{h_1, \dots, h_p}\}$ , a  $k_{m,1} \times \dots \times k_{m,p}$ -dimensional core tensor, captures interactions between the different components and the  $k_{m,h} \times d_h$  dimensional mode matrices  $\mathbf{U}_h = \{u_{h, c_h}(k_h)\}$  consist of the corresponding component specific weights. See Figure S.1.

HOSVD achieves better data compression and requires fewer components compared to the PARAFAC model which can be obtained as a special case of HOSVD with the core tensor  $\mathbf{G}$  restricted to being diagonal.

Yang and Dunson (2016) leveraged these ideas while regressing a categorical response



(a) PARAFAC decomposition



(b) HOSVD/Tucker decomposition

Figure S.1: Pictorial representation of tensor factorizations of a three dimensional tensor  $\mathbf{M}$  (a) parallel factor decomposition with vectors (b) higher order singular value decomposition with core tensor  $\mathbf{G}$  and mode matrices  $\mathbf{U}^{(h)}$ ,  $h = 1, 2, 3$ .

variable  $y \in \{1, \dots, k_y\}$  on a set of categorical predictors  $c_h \in \{1, \dots, d_h\}$ ,  $h = 1, \dots, p$ . Structuring the conditional probabilities  $P_{y|\mathbf{c}}(y | c_h, h = 1, \dots, p)$  as the elements of a  $k_y \times d_1 \times \dots \times d_p$  dimensional tensor, they proposed the following HOSVD-type factorization

$$P_{y|\mathbf{c}}(y | c_h, h = 1, \dots, p) = \sum_{k_1=1}^{k_{y,1}} \dots \sum_{k_p=1}^{k_{y,p}} \lambda_{k_1 \dots k_p}(y) \prod_{j=1}^p \pi_{h,c_h}(k_h), \quad (\text{S.7})$$

where  $1 \leq k_{y,h} \leq d_h$  for  $h = 1, \dots, p$  and the parameters  $\lambda_{k_1, \dots, k_p}(y)$  and  $\pi_{h,c_h}(k_h)$  are all non-negative and satisfy the constraints (a)  $\sum_{y=1}^{d_0} \lambda_{k_1 \dots k_p}(y) = 1$  for each combination  $(k_1, \dots, k_p)$ , and (b)  $\sum_{k_h=1}^{k_{y,h}} \pi_{h,c_h}(k_h) = 1$  for each pair  $(h, c_h)$ . See Figure 5 in the main paper.

Yang and Dunson (2016) established that any conditional probability tensor can be represented as (S.7), with the parameters satisfying the constraints (a) and (b). When  $k_{y,h} = 1$ ,  $\pi_1^{(h)}(c_h) = 1$  and  $P_{y|\mathbf{c}}(y | c_h, h = 1, \dots, p)$  does not vary with  $c_h$ . The number of parameters involved in the factorization is given by  $(d_0 - 1) \prod_{h=1}^p k_{y,h} + d_0 \sum_{h=1}^p (k_{y,h} - 1)$ , which is much less than the number of parameters  $(d_0 - 1) \prod_{h=1}^p d_h$  required to specify a fully parametrized model, if  $\prod_{h=1}^p k_{y,h} \ll \prod_{h=1}^p d_h$ .

### S.3 Copula Mixture Models with Shared Atoms

A mixture model with mixture probabilities  $p_k$  and mixture kernels  $\mathcal{K}(\cdot | \boldsymbol{\theta}_k)$  parametrized by atoms  $\boldsymbol{\theta}_k$  is specified as

$$f_x(x) = \sum_{k=1}^{k_x} p_k \mathcal{K}(x | \boldsymbol{\theta}_k).$$

A Bayesian framework then assigns priors on the mixture probabilities  $(p_1, \dots, p_{k_x})$  and the atoms  $\{\boldsymbol{\theta}_k\}_{k=1}^{k_x}$ . The number of mixture components  $k_x$  can be finite (Frühwirth-Schnatter, 2006, etc.) or be a-priori set at  $\infty$  in which case the number of ‘expressed’ components can be inferred from the data (Escobar and West, 1995, etc.). Aside from providing a flexible representation of the density  $f_x$ , such models also implicitly induce a clustering of the observations  $x_i, i = 1, \dots, n$ , the  $x_i$ ’s associated with the  $k^{\text{th}}$  mixture component assumed to belong to the  $k^{\text{th}}$  cluster.

For grouped data  $x_{\ell,i}, \ell = 1, \dots, d, i = 1, \dots, n_\ell$ , with a focus on shared clustering between groups, common atoms mixture models that keep the atoms fixed between the groups but allow the associated mixture probabilities to vary can be specified as

$$f_{x,\ell}(x) = \sum_{k=1}^{k_x} p_{\ell,k} \mathcal{K}(x | \boldsymbol{\theta}_k).$$

A number of highly sophisticated Bayesian nonparametric hierarchical priors have been developed in the literature that induce such models at the observation level while may or may not allow additional clustering of the entire group distributions  $f_{x,\ell}, \ell = 1, \dots, d$ . See, for example, Beraha *et al.* (2021); Denti *et al.* (2021) and the references cited therein.

In deconvolution context, where the main focus is not so much on clustering but instead on obtaining a flexible representation of the joint density of a multivariate continuous random variable  $\mathbf{x} = (x_1, \dots, x_d)^{\text{T}}$ , a similar idea with finite mixture models for the marginals and a Gaussian copula for the dependence structure was independently developed in Sarkar *et al.* (2021), where the different component dimensions  $\ell = 1, \dots, d$  played the roles of different groups and the dependence between the components was separately modeled using a Gaussian copula. Specifically, they let

$$f_{\mathbf{x}}(\mathbf{x}) = |\mathbf{R}_{\mathbf{x}}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \mathbf{y}_{\mathbf{x}}^{\text{T}} (\mathbf{R}_{\mathbf{x}}^{-1} - \mathbf{I}_d) \mathbf{y}_{\mathbf{x}} \right\} \prod_{\ell=1}^d f_{x,\ell}(x_\ell),$$

where  $F_{x,\ell}(x_\ell) = \Phi(y_{x,\ell})$  for all  $\ell$ ,  $F_{x,\ell}$  being the cdf corresponding to  $f_{x,\ell}$ ;  $\mathbf{y}_{\mathbf{x}} = (y_{x,1}, \dots, y_{x,d})^{\text{T}}$ ;  $\Phi(\cdot)$  denotes the cdf of a standard normal distribution; and  $\mathbf{R}_{\mathbf{x}}$  is the correlation matrix between the  $d$  components of  $\mathbf{x}$  and the marginals were modeled as mixtures of truncated normals with shared atoms  $\{(\mu_k, \sigma_k^2)\}_{k=1}^{k_x}$  as

$$f_{x,\ell}(x) = \sum_{k=1}^{k_x} p_{\ell,k} \text{TN}(x | \mu_k, \sigma_k^2, [A, B]).$$

A related model where the marginals were modeled using infinite mixtures of normals, each having its own set of atoms, and were linked using a Bernstein polynomial copula had previously appeared in Burda and Prokhorov (2014).

## S.4 Additional Exploratory Figures




Figure S.2: Exploratory plots for iron, magnesium and sodium (from top to bottom). Left panels: histograms of subject-specific means  $\bar{w}_{\ell,i}$ , crude estimates of  $x_{\ell,i}$ ; right upper panels: histograms of ‘residuals’  $\hat{u}_{\ell,i,j} = (w_{\ell,i,j} - \bar{w}_{\ell,i})$ , crude estimates of  $u_{\ell,i,j}$ ; right lower panels: subject-specific means  $\bar{w}_{\ell,i}$  vs variances  $s_{w,\ell,i}^2$ , crude estimates of  $\text{var}(u_{\ell,i,j} | x_{\ell,i})$ , superimposed with lowess fits.

## S.5 Hyper-parameter Choices and Posterior Computation

Samples from the posterior can be drawn using the MCMC algorithm described below. In what follows,  $\zeta$  denotes a generic variable that collects the data as well as all parameters of the model, including the sampled values of  $\mathbf{x}_{1:n}$  and  $\boldsymbol{\epsilon}_{1:N}$ , that are not explicitly mentioned. Also, the generic notation  $p_0$  is sometimes used for specifying priors and hyper-priors.

We now discuss our choices for the prior hyper-parameters and the initial values of the MCMC sampler. Carefully chosen starting values facilitate convergence of our sampler. The starting values of some of the parameters for the multivariate problem are determined by first running samplers for the covariate independent univariate model of Sarkar *et al.* (2014). We describe the hyper-parameter choices and the initial values for the sampler for these univariate models first. Unless otherwise mentioned, the prior hyper-parameter choices for similar model components for the multivariate model remain the same as that used for the univariate marginal models. We only detail the sampling steps for the multivariate method. The steps for the univariate method are detailed in Sarkar *et al.* (2014).

To make the recalls for all the components to be unit free and have a shared support, we transformed the recalls as  $w_{\ell,i,j} = 20 \times \frac{w_{\ell,i,j}}{\max\{w_{\ell,i,j}\}}$ . The latent  $x_{\ell,i}$ 's can then be assumed to lie in  $[A, B] = [0, 10]$ , greatly simplifying model specification and hyper-parameter selection. As opposed to the non-linear Box-Cox transformations used in the previous literature, including Zhang *et al.* (2011), which often result in loss of information and introduce bias, we only make linear scale transformations here that preserve all features of the original data points.

For the univariate samplers for the marginal components, we used the subject-specific sample means  $\bar{w}_{\ell,1:n}$  as the starting values for  $x_{\ell,1:n}$ . The appropriate number of mixture components in a mixture model depends on the flexibility of the component mixture kernels as well as on specific demands of the particular application at hand. With appropriately chosen mixture kernels, univariate mixture models with 5-10 components have often been found to be sufficiently flexible. Detailed guidelines on selecting the number of mixture components for the specific context of deconvolution problems can be found in Section S.1 and S.6 in the supplementary materials of Sarkar *et al.* (2018). Based on such guidelines, we used  $J_\ell = 12$  equidistant knot points for the B-splines supported on  $[A, B] = [0, 10]$  for modeling the variance functions, and  $k_{x,\ell} = 10$  truncated normals for modeling the densities. We also allowed  $k_{\epsilon,\ell} = 10$  mixture components for the mixtures modeling the densities of the scaled errors. For the Dirichlet prior hyper-parameters, we set  $\alpha_x = 1/k_x$ ,  $\alpha_\epsilon = 1/k_\epsilon$ . The hyper-parameters for the smoothness inducing parameters are set to be mildly informative as  $a_\vartheta = 10$ ,  $b_\vartheta = 1$ . Introducing latent mixture component allocation variables  $\mathbf{z}_{x,1:d,1:n}$ ,  $\mathbf{z}_{\epsilon,1:d,1:N}$  and  $\mathbf{z}^2_{\epsilon,1:d,1:N}$ , we can write

$$\begin{aligned} (x_{\ell,i} \mid z_{x,\ell,i} = k, \mu_{x,\ell,k}, \sigma_{x,\ell,k}^2) &\sim \text{TN}(x_{\ell,i} \mid \mu_{x,\ell,k}, \sigma_{x,\ell,k}^2, [A, B]), \quad \text{and} \\ (\epsilon_{\ell,i,j} \mid z_{\epsilon,\ell,i,j} = k, z^2_{\epsilon,\ell,i,j} = t, \mu_{\epsilon,\ell,k,t}, \sigma_{\epsilon,\ell,k,t}^2) &\sim \text{Normal}(\epsilon_{\ell,i,j} \mid \mu_{\epsilon,\ell,k,t}, \sigma_{\epsilon,\ell,k,t}^2), \quad \ell = 1, \dots, d. \end{aligned}$$

The mixture labels  $z_{x,\ell,i}$ 's, and the component specific parameters  $\mu_{x,\ell,k}$ 's and  $\sigma_{x,\ell,k}$ 's are initialized by fitting a  $k$ -means algorithm with  $k = k_x$ . The parameters of the distribution of scaled errors are initialized at values that correspond to the special standard normal case. The initial values of the smoothness inducing parameters are set at  $\sigma_{\vartheta,\ell}^2 = \sigma_{\xi,\ell}^2 = \sigma_{\epsilon,\ell}^2 = 0.1$ . The associated mixture labels  $z_{\epsilon,\ell,i,j}$ 's are thus all initialized at  $z_{\epsilon,\ell,i,j} = 1$ . The initial values of  $\boldsymbol{\vartheta}_\ell$ 's are obtained by maximizing

$$\ell(\boldsymbol{\vartheta}_\ell \mid \sigma_{\vartheta,\ell}^2, \bar{\mathbf{w}}_{\ell,1:n}) = -\frac{\boldsymbol{\vartheta}_\ell^\top \mathbf{P}_\ell \boldsymbol{\vartheta}_\ell}{2\sigma_{\vartheta,\ell}^2} - \sum_{i=1}^n \frac{1}{2s_\ell^2(\bar{w}_{\ell,i}, \boldsymbol{\vartheta}_\ell)} \sum_{j=1}^{m_i} (w_{\ell,i,j} - \bar{w}_{\ell,i})^2$$

with respect to  $\boldsymbol{\vartheta}_\ell$ .

We now discuss how we set the initial values of the sampler for the multivariate method. The starting values of the  $w_{\ell,i,j}$ 's,  $x_{\ell,i}$ 's,  $u_{\ell,i,j}$ 's,  $\boldsymbol{\vartheta}_\ell$ 's were all set at the corresponding estimates returned by the univariate samplers. We set the number of shared atoms of the mixture models for the densities  $f_{x,\ell}$  and  $f_{\epsilon,\ell}$  at  $k_x = k_\epsilon = \max\{5d, 20\}$ . We set  $\alpha_{x,\ell} = \alpha_{\epsilon,\ell} = 1$ . The atoms of the mixtures of truncated normals for the marginal densities  $f_{x,\ell}$  of the regular components are shared, so are the atoms of the mixture models for the univariate marginals  $f_{\epsilon,\ell}$  of the scaled errors, and hence these parameters could not be initialized directly using the univariate model output. We initialized these parameters by iteratively sampling them from their posterior full conditionals 100 times, keeping the estimated  $x_{\ell,i}$ 's fixed. Adopting a similar strategy, we initialized the parameters specifying the densities  $f_{\epsilon,\ell}$  of the scaled errors by iteratively sampling them from their posterior full conditionals 100 times, keeping the estimated errors  $u_{\ell,i,j}$ 's fixed. Finally, the parameters specifying  $\mathbf{R}_x$  and  $\mathbf{R}_\epsilon$  were set at values that correspond to the special case  $\mathbf{R}_x = \mathbf{R}_\epsilon = \mathbf{I}_d$ .

We set the variable selection parameters  $k_{x,0} = d$  and  $k_{x,h} = 1$  for all  $h = 1, \dots, p$ . Likewise,  $k_{\epsilon,0} = d$  and  $k_{\epsilon,h} = 1$  for all  $h = 1, \dots, p$ . So only the component label is considered important initially. The values of the associated latent variables are set accordingly. Specifically, we let  $z_{x,0,\ell,i} = \ell$  for all  $i = 1, \dots, n$  and  $z_{x,h,\ell,i} = 1$  for all  $h = 1, \dots, p; \ell = 1, \dots, d; i = 1, \dots, n$ . Likewise,  $z_{\epsilon,0,\ell,i,j} = \ell$  for all  $i = 1, \dots, n; j = 1, \dots, m_i$  and  $z_{\epsilon,h,\ell,i,j} = 1$  for all  $h = 1, \dots, p; \ell = 1, \dots, d; i = 1, \dots, n; j = 1, \dots, m_i$ .

In our sampler for the multivariate problem, we first update the parameters specifying the different marginal densities using a pseudo-likelihood that ignores the contribution of the copula. The parameters characterizing the copula and the latent  $\mathbf{x}_i$ 's are then updated using the exact likelihood function conditionally on the parameters obtained in the first step. We then update the parameters of the marginal densities again and so forth. A more appealing approach would have been to perform joint estimation of the marginal distributions and the copula functions. Joint estimation algorithms, most involving carefully designed Metropolis-Hastings (M-H) moves, have been proposed in much simpler settings in Pitt *et al.* (2006); Wu *et al.* (2014, 2015) etc. Designing such moves for our complex deconvolution problem is a daunting task. Importantly, the results of dos Santos Silva and Lopes (2008) suggest that two-stage approaches often perform just as good as joint estimation procedures, validating their use for practical reasons.

Our sampler for the multivariate model iterates between the following steps.

**1. Updating the parameters specifying  $f_{x|\mathbf{c}}$ :** We have

$$\begin{aligned} \Pr(\boldsymbol{\lambda}_{x,k_0,k_1,\dots,k_p} \mid \boldsymbol{\zeta}) &= \text{Dir}\{\alpha_x \lambda_{x,0}(1) + n_{x,k_0,k_1,\dots,k_p}(1), \dots, \alpha_x \lambda_{x,0}(k_x) + n_{x,k_0,k_1,\dots,k_p}(k_x)\}, \\ P_{x|\mathbf{c}}(z_{x,\ell,i} = k \mid z_{x,0,\ell,i} = k_0, z_{x,1,\ell,i} = k_1, \dots, z_{x,p,\ell,i} = k_p, \boldsymbol{\zeta}) \\ &\propto \lambda_{x,k_0,k_1,\dots,k_p}(k) \times \text{TN}(x_{\ell,i} \mid \mu_{x,k}, \sigma_{x,k}^2, [A, B]), \end{aligned}$$



where  $n_{x,k_0,k_1,\dots,k_p}(k) = \sum_{\ell=1}^d \sum_{i=1}^n 1(z_{x,\ell,i} = k, z_{x,0,\ell,i} = k_0, z_{x,1,\ell,i} = k_1, \dots, z_{x,p,\ell,i} = k_p)$ . To update  $\lambda_{x,0}$ , mimicking ideas presented in Sarkar and Dunson (2020), for  $s = 1, \dots, n_{x,k_0,k_1,\dots,k_p}(k)$ , we first sample an auxiliary variable  $\omega_s$  as

$$\omega_s \mid \zeta \sim \text{Bernoulli} \left\{ \frac{\alpha_x \lambda_{x,0}(k)}{s - 1 + \alpha_x \lambda_{x,0}(k)} \right\}.$$

We then set  $m_{x,k_0,k_1,\dots,k_p}(k) = \sum_s \omega_s$ ,  $m_{x,0}(k) = \sum_{(k_0,k_1,\dots,k_p)} m_{x,k_0,k_1,\dots,k_p}(k)$ . Finally, we sample  $\lambda_{x,0}$  as

$$\lambda_{x,0} \mid \zeta \sim \text{Dir}\{\alpha_{x,0}/k_x + m_{x,0}(1), \dots, \alpha_{x,0}/k_x + m_{x,0}(k_x)\}.$$

The full conditionals of  $\mu_{x,k}$  and  $\sigma_{x,k}^2$  are given by

$$\begin{aligned} \Pr(\mu_{x,k} \mid \zeta) &\propto p_0(\mu_{x,k}) \times \prod_{\ell=1}^d \prod_{\{i: z_{x,\ell,i}=k\}} \text{TN}(x_{\ell,i} \mid \mu_{x,k}, \sigma_{x,k}^2, [A, B]), \\ \Pr(\sigma_{x,k}^2 \mid \zeta) &\propto p_0(\sigma_{x,k}^2) \times \prod_{\ell=1}^d \prod_{\{i: z_{x,\ell,i}=k\}} \text{TN}(x_{\ell,i} \mid \mu_{x,k}, \sigma_{x,k}^2, [A, B]). \end{aligned}$$

These parameters are updated by Metropolis-Hastings (MH) steps with the proposals  $q(\mu_{x,k} \rightarrow \mu_{x,k,new}) = \text{Normal}(\mu_{x,k,new} \mid \mu_{x,k}, \sigma_{x,\ell,\mu}^2)$  and  $q(\sigma_{x,k}^2 \rightarrow \sigma_{x,k,new}^2) = \text{TN}(\sigma_{x,k,new}^2 \mid \sigma_{x,k}^2, \sigma_{x,\ell,\sigma}^2, [\max\{0, \sigma_{x,k}^2 - 1\}, \sigma_{x,k}^2 + 1])$ , respectively.

**2. Updating the covariate selection parameters  $\mathbf{k}_x$  and  $\mathbf{z}_x$ :** The current values of  $\mathbf{k}_x = (k_{x,0}, k_{x,1}, \dots, k_{x,p})^\top$  and  $\mathbf{z}_{x,\ell,i} = (z_{x,0,\ell,i}, z_{x,1,\ell,i}, \dots, z_{x,p,\ell,i})^\top$  induce a partition of the  $d_h$  levels of  $c_h$  into  $k_{x,h}$  clusters  $\{\mathcal{C}_{x,h,r} : r = 1, \dots, k_{x,h}\}$  with each cluster  $\mathcal{C}_{x,h,r}$  corresponding to the latent class  $z_{x,h} = r$ . Integrating out  $\lambda_{x,k_0,k_1,\dots,k_p}$ , conditional on the cluster configurations  $\mathcal{C}_x = \{\mathcal{C}_{x,h,r} : h = 0, \dots, p, r = 1, \dots, k_{x,h}\}$ , we have

$$\Pr(\mathbf{z}_x \mid \mathcal{C}_x, \zeta) = \prod_{(k_0,k_1,\dots,k_p)} \frac{\beta\{\alpha_x \lambda_{x,0}(1) + n_{x,k_0,k_1,\dots,k_p}(1), \dots, \alpha_x \lambda_{x,0}(k_x) + n_{x,k_0,k_1,\dots,k_p}(k_x)\}}{\beta\{\alpha_x \lambda_{x,0}(1), \dots, \alpha_x \lambda_{x,0}(k_x)\}}, \quad (\text{S.8})$$

where  $\beta(a_1, \dots, a_r) = \prod_{j=1}^r \Gamma(a_j) / \Gamma(a_1 + \dots + a_r)$  is the Beta function. Given the current values of  $k_{x,h}$  and the current clusters  $\mathcal{C}_x = \{\mathcal{C}_{x,h,r} : h = 0, \dots, p, r = 1, \dots, k_{x,h}\}$ , we do the following for  $h = 0, \dots, p$ . If  $k_{x,h} < c_h$ , we propose to increase  $k_{x,h}$  to  $(k_{x,h} + 1)$ . If  $k_{x,h} > 1$ , we propose to decrease  $k_{x,h}$  to  $(k_{x,h} - 1)$ . For  $1 < k_{x,h} < c_h$ , the moves are proposed with equal probabilities. For  $k_{x,h} = 1$ , the increase move is selected with probability 1. For  $k_{x,h} = c_h$ , the decrease move is selected with probability 1. If an increase move is proposed, we randomly split a cluster into two. If a decrease move is proposed, we randomly merge two clusters into a single one. We accept the proposed moves with acceptance rates based on the marginal likelihood (S.8). Finally, we set  $\mathbf{z}_x$  to be the cluster allocation variables determined by the updated cluster mappings.

**3. Updating the parameters specifying  $f_{\epsilon|\mathbf{c}}$ :** We have

$$\begin{aligned} \Pr(\lambda_{\epsilon,k_0,k_1,\dots,k_p} \mid \zeta) &= \text{Dir}\{\alpha_\epsilon \lambda_{\epsilon,0}(1) + n_{\epsilon,k_0,k_1,\dots,k_p}(1), \dots, \alpha_\epsilon \lambda_{\epsilon,0}(k_\epsilon) + n_{\epsilon,k_0,k_1,\dots,k_p}(k_\epsilon)\}, \\ P_{\epsilon|\mathbf{c}}(z_{\epsilon,\ell,i,j} = k \mid z_{\epsilon,0,\ell,i,j} = k_0, z_{\epsilon,1,\ell,i,j} = k_1, \dots, z_{\epsilon,p,\ell,i,j} = k_p, \zeta) \\ &\propto \lambda_{\epsilon,k_0,k_1,\dots,k_p}(k) \times f_{w_\ell|x_\ell}(w_{\ell,i,j} \mid p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2, \zeta), \end{aligned}$$

where  $n_{\epsilon, k_0, k_1, \dots, k_p}(k) = \sum_{\ell=1}^d \sum_{i=1}^n \sum_{j=1}^{m_i} 1(z_{\epsilon, \ell, i, j} = k, z_{\epsilon, 0, \ell, i, j} = k_0, z_{\epsilon, 1, \ell, i, j} = k_1, \dots, z_{\epsilon, p, \ell, i, j} = k_p)$ . To update  $\lambda_{\epsilon, 0}$ , following the same ideas for sampling  $\lambda_{x, 0}$  above, for  $s = 1, \dots, n_{\epsilon, k_0, k_1, \dots, k_p}(k)$ , we first sample an auxiliary variable  $\omega_s$  as

$$\omega_s \mid \zeta \sim \text{Bernoulli} \left\{ \frac{\alpha_{\epsilon} \lambda_{\epsilon, 0}(k)}{s - 1 + \alpha_{\epsilon} \lambda_{\epsilon, 0}(k)} \right\}.$$

We set  $m_{\epsilon, k_0, k_1, \dots, k_p}(k) = \sum_s \omega_s$ ,  $m_{\epsilon, 0}(k) = \sum_{(k_0, k_1, \dots, k_p)} m_{\epsilon, k_0, k_1, \dots, k_p}(k)$ , and sample  $\lambda_{\epsilon, 0}$  as

$$\lambda_{\epsilon, 0} \mid \zeta \sim \text{Dir}\{\alpha_{\epsilon, 0}/k_{\epsilon} + m_{\epsilon, 0}(1), \dots, \alpha_{\epsilon, 0}/k_{\epsilon} + m_{\epsilon, 0}(k_{\epsilon})\}.$$

We propose a new  $(p_{\epsilon, k}, \mu_{\epsilon, k}, \sigma_{\epsilon, k, 1}^2, \sigma_{\epsilon, k, 2}^2)$  with the proposal  $q\{\theta_{\epsilon, k} = (p_{\epsilon, k}, \mu_{\epsilon, k}, \sigma_{\epsilon, k, 1}^2, \sigma_{\epsilon, k, 2}^2) \rightarrow (p_{\epsilon, k, \text{new}}, \mu_{\epsilon, k, \text{new}}, \sigma_{\epsilon, k, 1, \text{new}}^2, \sigma_{\epsilon, k, 2, \text{new}}^2) = \theta_{\epsilon, k, \text{new}}\} = \text{TN}(p_{\epsilon, k, \text{new}} \mid p_{\epsilon, k}, \sigma_{p, \epsilon}^2, [0, 1]) \times \text{Normal}(\mu_{\epsilon, k, \text{new}} \mid \mu_{\epsilon, k}, \sigma_{\epsilon, \mu}^2) \times \text{TN}(\sigma_{\epsilon, k, 1, \text{new}}^2 \mid \sigma_{\epsilon, k, 1}^2, \sigma_{\epsilon, \sigma}^2, [0, \infty)) \times \text{TN}(\sigma_{\epsilon, k, 2, \text{new}}^2 \mid \sigma_{\epsilon, k, 2}^2, \sigma_{\epsilon, \sigma}^2, [0, \infty))$ . We update  $\theta_k$  to the proposed value  $\theta_{k, \text{new}}$  with probability

$$\min \left\{ 1, \frac{q(\theta_{\epsilon, k, \text{new}} \rightarrow \theta_{\epsilon, k}) \prod_{\ell=q+1}^{2q+p} \prod_{\{i, j: c_{\epsilon, \ell, i, j} = k\}} f_{w_{\ell} | \tilde{x}_{\ell}}(w_{\ell, i, j} \mid \theta_{\epsilon, k, \text{new}}, \zeta) p_0(\theta_{\epsilon, k, \text{new}})}{q(\theta_{\epsilon, k} \rightarrow \theta_{\epsilon, k, \text{new}}) \prod_{\ell=q+1}^{2q+p} \prod_{\{i, j: c_{\epsilon, \ell, i, j} = k\}} f_{w_{\ell} | \tilde{x}_{\ell}}(w_{\ell, i, j} \mid \theta_{\epsilon, k}, \zeta) p_0(\theta_{\epsilon, k})} \right\}.$$

**4. Updating the covariate selection parameters  $\mathbf{k}_{\epsilon}$  and  $\mathbf{z}_{\epsilon}$ :** The values of  $\mathbf{k}_{\epsilon} = (k_{\epsilon, 0}, k_{\epsilon, 1}, \dots, k_{\epsilon, p})^T$  and  $\mathbf{z}_{\epsilon, \ell, i, j} = (z_{\epsilon, 0, \ell, i, j}, z_{\epsilon, 1, \ell, i, j}, \dots, z_{\epsilon, p, \ell, i, j})^T$  are updated mimicking the same strategy used to update  $\mathbf{k}_x$  and  $\mathbf{z}_x$ .

**5. Updating the parameters specifying  $v_{\ell}$ :** The full conditional of each  $\vartheta_{\ell}$  is

$$\Pr(\vartheta_{\ell} \mid \mathbf{w}_{\ell, 1:N}, \zeta) \propto p_0(\vartheta_{\ell}) \times \prod_{i, j} f_{w_{\ell} | x_{\ell}}(w_{\ell, i, j} \mid x_{\ell, i}, \vartheta_{\ell}, \zeta).$$

We use M-H sampler with random walk proposal  $q(\vartheta_{\ell} \rightarrow \vartheta_{\ell, \text{new}}) = \text{MVN}(\vartheta_{\ell, \text{new}} \mid \vartheta_{\ell}, \Sigma_{\vartheta, \ell})$ .

**6. Updating the values of  $\mathbf{x}$ :** The full conditionals for  $\mathbf{x}_i$  are given by

$$\begin{aligned} (\mathbf{x}_i \mid \mathbf{c}, \zeta) &\propto f_{\mathbf{x} | \mathbf{c}}(\mathbf{x}_i \mid \mathbf{c}, \zeta) \times \prod_{j=1}^{m_i} f_{\mathbf{w} | \mathbf{x}}(\mathbf{w}_{i, j} \mid \mathbf{x}_i, \zeta) \\ &= |\mathbf{R}_{\mathbf{x}}|^{-1/2} \exp \left\{ -\frac{1}{2} \mathbf{y}_{\mathbf{x}, i}^T (\mathbf{R}_{\mathbf{x}}^{-1} - \mathbf{I}_d) \mathbf{y}_{\mathbf{x}, i} \right\} \prod_{\ell=1}^d f_{x_{\ell} | \mathbf{c}}(x_{\ell, i} \mid \mathbf{c}, \zeta) \\ &\quad \times \prod_{j=1}^{m_i} \left[ |\mathbf{R}_{\epsilon}|^{-1/2} \exp \left\{ -\frac{1}{2} \mathbf{y}_{\epsilon, i, j}^T (\mathbf{R}_{\epsilon}^{-1} - \mathbf{I}_d) \mathbf{y}_{\epsilon, i, j} \right\} \prod_{\ell=1}^d f_{w_{\ell} | x_{\ell}}(w_{\ell, i, j} \mid x_{\ell, i}, \zeta) \right], \end{aligned}$$

where  $F_{x, \ell}(x_{\ell, i} \mid \zeta) = \Phi(y_{x, \ell, i})$  and  $F_{\epsilon, \ell}\{(w_{\ell, i, j} - x_{\ell, i})/s_{\ell}(x_{\ell, i}) \mid \zeta\} = \Phi(y_{\epsilon, \ell, i, j})$ . The full conditionals do not have closed forms. MH steps with independent truncated normal proposals for each component are used within the Gibbs sampler.

**7. Updating the parameters specifying the copula:** We have  $F_{x, \ell}(x_{\ell, i} \mid \zeta) = \Phi(y_{x, \ell, i})$  for all  $i = 1, \dots, n$  and  $\ell = 1, \dots, d$ . Conditionally on the parameters specifying the marginals,  $\mathbf{y}_{\mathbf{x}, 1:d, 1:n}$  are thus known quantities. We plug-in these values and use that  $(\mathbf{y}_{\mathbf{x}, i} \mid \mathbf{R}_{\mathbf{x}}) \sim \text{MVN}_d(\mathbf{0}, \mathbf{R}_{\mathbf{x}})$  to update  $\mathbf{R}_{\mathbf{x}}$ . The full conditionals of the parameters specifying  $\mathbf{R}_{\mathbf{x}}$  do not have closed forms. We use M-H steps to update these parameters.

(a) For  $t = 1, \dots, (d-1)$ , we discretized the values of  $b_{x, t}$  to the set  $\{-0.99 + 2 \times 0.99(m-1)/(M-1)\}$ , where  $m = 1, \dots, M$  and we chose  $M = 41$ . A new value  $b_{x, t, \text{new}}$  is proposed at random from the set comprising the current value of  $b_{x, t}$  and its two neighbors. The

proposed value is accepted with probability  $\min\{1, a(b_{x,t,new})/a(b_{x,t})\}$ , where

$$a(b_{x,t}) = (1 - b_{x,t}^2)^{-n/2} \times \exp \left\{ -(1/2) \sum_{i=1}^n \sum_{j=1}^{m_i} \mathbf{y}_{\mathbf{x},i,j}^T \{ \boldsymbol{\Sigma}_{\mathbf{x}}(b_{x,t}, \boldsymbol{\zeta}) \}^{-1} \mathbf{y}_{\mathbf{x},i,j} \right\}.$$

(b) For  $s = 1, \dots, (d-1)(d-2)/2$ , we discretized the values of  $\theta_{x,s}$  to the set  $\{-3.14 + 2 \times 3.14(m-1)/(M-1)\}$ , where  $m = 1, \dots, M$  and  $M = 41$ . A new value  $\theta_{s,new}$  is proposed at random from the set comprising the current value and its two neighbors. The proposed value is accepted with probability  $\min\{1, a(\theta_{x,s,new})/a(\theta_{x,s})\}$ , where

$$a(\theta_{x,s}) = \exp \left\{ -(1/2) \sum_{i=1}^n \sum_{j=1}^{m_i} \mathbf{y}_{\mathbf{x},i,j}^T \{ \boldsymbol{\Sigma}_{\mathbf{x}}(\theta_{x,s}, \boldsymbol{\zeta}) \}^{-1} \mathbf{y}_{\mathbf{x},i,j} \right\}.$$

The parameters specifying  $\mathbf{R}_{\boldsymbol{\epsilon}}$  are updated in a similar fashion.

With carefully chosen initial values and proposal densities for the MH steps, we were able to achieve quick convergence for the MCMC samplers. For our proposed method, 5,000 MCMC iterations were run in each case with the initial 3,000 iterations discarded as burn-in. The remaining samples were further thinned by a thinning interval of 5. We programmed in R. With  $n = 1000$  subjects and  $m_i = 3$  proxies for each subject, on an ordinary desktop, 5,000 MCMC iterations required approximately 3 hours to run.

## S.6 Simulation Studies

We focus here on comparisons with our main competitor, the method of Zhang *et al.* (2011). Simulation scenarios to perform these comparisons were designed as follows.

We mimic some aspects of the real data set analyzed in Section 3 and the design in Sarkar *et al.* (2021) as closely as possible while modifying some others to illustrate the flexibility and efficiency of the proposed method. We chose  $n = 965$ ,  $m_i = 3$  replicates per subject, and  $d = 3$  dimensional  $\mathbf{x}$ . While our method scales quite well to much higher dimensional problems, with 3 total components the results can be conveniently graphically summarized.

We maintained the same distribution of the covariates as in the EATS data (Figure 1).

To generate the true  $x_{\ell,i}$ 's for  $\ell = 1, \dots, d$ , we (a) first sampled  $\mathbf{x}_i^\Delta \sim \text{MVN}_d(\mathbf{0}, \mathbf{R}_{\mathbf{x}})$ , (b) then, set  $\mathbf{x}_i^{\Delta\Delta} = \Phi(\mathbf{x}_i^\Delta)$ , (c) finally, set  $x_{\ell,i} = F_{TN,mix}^{-1}(x_{\ell,i}^{\Delta\Delta} \mid \boldsymbol{\pi}_{x,\ell}, \boldsymbol{\mu}_{x,\ell}, \boldsymbol{\sigma}_{x,\ell}^2, x_{\ell,L}, x_{\ell,U})$ , where  $F_{TN,mix}(x \mid \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, x_L, x_U) = \sum_{k=1}^k \pi_k F_{TN}(x \mid \mu_k, \sigma_k^2, x_L, x_U)$ . The marginal distributions are thus mixtures of truncated normal distributions and hence can take widely varying shapes (Figure S.6) while the correlation between different components is  $\mathbf{R}_{\mathbf{x}}$ . We set

$$\mathbf{R}_{\mathbf{x}} = \begin{pmatrix} 1 & 0.7 & 0.7^2 \\ & 1 & 0.7 \\ & & 1 \end{pmatrix}, \quad \boldsymbol{\mu}_{\mathbf{x}} = \begin{pmatrix} \boldsymbol{\mu}_{x,1}^T \\ \boldsymbol{\mu}_{x,2}^T \\ \boldsymbol{\mu}_{x,3}^T \end{pmatrix} = \begin{pmatrix} 1.5 & 1.5 & 3.0 & 5.0 \\ 2.0 & 2.0 & 4.0 & 5.0 \\ 2.0 & 3.0 & 4.0 & 5.0 \end{pmatrix},$$

$$x_{\ell,L} = 0, \quad x_{\ell,U} = 10 \text{ for all } \ell, \text{ and } \sigma_{x,\ell,k}^2 = 0.75^2 \text{ for all } \ell, k.$$

We assume, as we have seen in the case of the real data application, that only  $c_1 = \text{gender}$  ( $c_1 = 1 \equiv \text{man}$ ,  $c_1 = 2 \equiv \text{woman}$ ) is an important predictor for  $f_{x,\ell|\mathbf{c}}$ . We set

$$\begin{pmatrix} \Pr(1 | c_0, 1, c_2, \dots, c_p) \\ \Pr(2 | c_0, 1, c_2, \dots, c_p) \\ \Pr(3 | c_0, 1, c_2, \dots, c_p) \\ \Pr(4 | c_0, 1, c_2, \dots, c_p) \end{pmatrix} = \begin{pmatrix} 0.10 \\ 0.40 \\ 0.20 \\ 0.30 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \Pr(1 | c_0, 2, c_2, \dots, c_p) \\ \Pr(2 | c_0, 2, c_2, \dots, c_p) \\ \Pr(3 | c_0, 2, c_2, \dots, c_p) \\ \Pr(4 | c_0, 2, c_2, \dots, c_p) \end{pmatrix} = \begin{pmatrix} 0.40 \\ 0.40 \\ 0.00 \\ 0.20 \end{pmatrix}$$

for all  $(c_0, c_2, \dots, c_p)$ . The densities  $f_{x,\ell|\mathbf{c}}$ 's also vary between different dimensions  $c_0 = \ell$ . This becomes clearer from a rearrangement of the distinct values of the components of  $\boldsymbol{\mu}_x$  and associated mixture probabilities for different combinations of  $(c_0, c_1)$  in Table S.1.

Dim	Sex	$\mu_{x,k}$				
		1.5	2.0	3.0	4.0	5.0
Associated Probabilities						
1	M	0.50	0.00	0.20	0.00	0.30
	W	0.80	0.00	0.00	0.00	0.20
2	M	0.00	0.50	0.00	0.20	0.30
	W	0.00	0.80	0.00	0.00	0.20
3	M	0.00	0.10	0.40	0.20	0.30
	W	0.00	0.40	0.40	0.00	0.20

Table S.1: Distinct values of the components of  $\boldsymbol{\mu}_x$  and associated mixture probabilities for different combinations of  $(c_0, c_1) = (\text{dimension, sex})$  used in the simulation design.

We used a similar procedure to simulate the true scaled errors  $\epsilon_{\ell,i,j}$ 's,  $\ell = 1, \dots, d$ . Following Sarkar *et al.* (2021), we (a) first sampled  $\boldsymbol{\epsilon}_{i,j}^\Delta \sim \text{MVN}_d(\mathbf{0}, \mathbf{R}_\epsilon)$ , (b) then, set  $\boldsymbol{\epsilon}_{i,j}^{\Delta\Delta} = \Phi(\boldsymbol{\epsilon}_i^\Delta)$ , (c) finally, set  $\epsilon_{\ell,i,j} = F_{\epsilon,\ell,mix,scaled}^{-1}(\epsilon_{\ell,i,j}^{\Delta\Delta} | \boldsymbol{\pi}_{\epsilon,\ell}, \boldsymbol{\theta}_{\epsilon,\ell})$ . Here for  $\ell = 1, 2$ ,  $F_{\epsilon,\ell,mix,scaled}$  is a scaled version of  $F_{\epsilon,\ell,mix}(\epsilon | \boldsymbol{\pi}_{\epsilon,\ell}, \boldsymbol{\theta}_{\epsilon,\ell}) = \sum_{k=1}^{k_{\epsilon,\ell}} \pi_{\epsilon,\ell,k} F_{cc}(\epsilon | p_{\epsilon,\ell,k}, \mu_{\epsilon,\ell,k}, \sigma_{\epsilon,\ell,k,1}^2, \sigma_{\epsilon,\ell,k,2}^2)$ . with  $\boldsymbol{\theta}_{\epsilon,\ell} = \{(p_{\epsilon,\ell,k}, \mu_{\epsilon,\ell,k}, \sigma_{\epsilon,\ell,k,1}^2, \sigma_{\epsilon,\ell,k,2}^2)\}_{k=1}^{k_{\epsilon,\ell}}$ . And, for  $\ell = 3$ ,  $F_{\epsilon,\ell,mix,scaled}$  is a scaled version of  $F_{\epsilon,\ell,mix}(\epsilon | \boldsymbol{\pi}_{\epsilon,\ell}, \boldsymbol{\theta}_{\epsilon,\ell}) = \sum_{k=1}^{k_{\epsilon,\ell}} \pi_{\epsilon,\ell,k} F_{\text{Laplace}}(\epsilon | m_{\epsilon,\ell,k}, b_{\epsilon,\ell,k})$  with  $\boldsymbol{\theta}_{\epsilon,\ell} = \{(m_{\epsilon,\ell,k}, b_{\epsilon,\ell,k})\}_{k=1}^{k_{\epsilon,\ell}}$ , adjusted to have mean zero and variance 1. This way, the marginal distributions can take widely varying shapes (Figure S.6) while the marginal correlation between different components is  $\mathbf{R}_\epsilon$ . In this case, we set

$$\mathbf{R}_\epsilon = \begin{pmatrix} 1 & 0.5 & 0.5^2 \\ & 1 & 0.5 \\ & & 1 \end{pmatrix}, \quad \boldsymbol{\pi}_{\epsilon,\ell} = \begin{pmatrix} 0.25 \\ 0.50 \\ 0.25 \end{pmatrix} \quad \text{for all } \ell,$$

$$\boldsymbol{\theta}_\epsilon = \begin{pmatrix} \boldsymbol{\theta}_{\epsilon,1}^\top \\ \boldsymbol{\theta}_{\epsilon,2}^\top \\ \boldsymbol{\theta}_{\epsilon,3}^\top \end{pmatrix} = \begin{pmatrix} (0.4, 2, 2, 1) & (0.4, 2, 2, 1) & (0.4, 2, 2, 1) \\ (0.5, 0, 0.25, 0.25) & (0.5, 0, 0.25, 0.25) & (0.5, 0, 5, 5) \\ (0, 1) & (0, 1) & (0, 1) \end{pmatrix}.$$

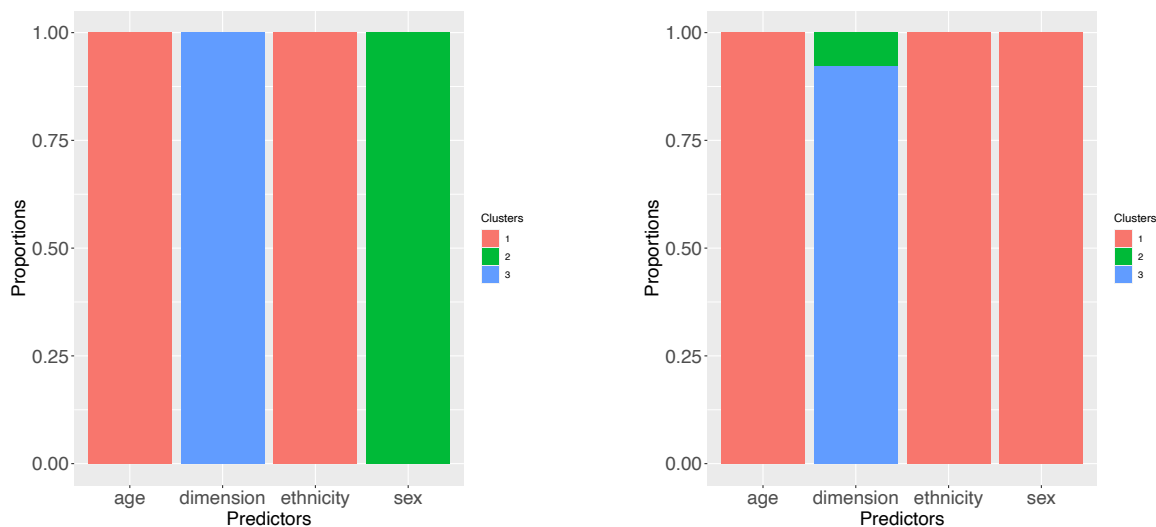


Figure S.3: Results for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs, showing the estimated probabilities of different numbers of clusters of the associated predictors’ levels being included in the model. The left panel shows the results for modeling the densities  $f_{x|c}$ . The right panel shows the results for modeling the densities  $f_{\epsilon|c}$ . At the median 0.5 probability level, in the left panel, the component labels and a binary covariate proxying for sex of the subjects are important predictors for modeling the densities  $f_{x|c}$ , whereas in the right panel, only the component labels are important for modeling the densities  $f_{\epsilon|c}$ . The results are consistent with the true simulation scenario.

The representations with  $k_{\epsilon,\ell} = 3$  components above are more than what are really needed to describe the particular assumed truths – we are effectively using a single component mixture of two-component scaled normals for  $f_{\epsilon,1}$  producing a bimodal error distribution, a two component (0.75, 0.25) mixture of two-component scaled normals for  $f_{\epsilon,2}$  producing a unimodal but heavier tailed error distribution, and finally a single component scaled Laplace for  $f_{\epsilon,3}$  producing a unimodal error distribution with a spike at zero (Figure S.6). It is clear, however, that such 3-component models are capable of generating a very wide variety of shapes, including multimodality heavy-tails etc., for the error distributions. Unlike the real data application, we thus chose the marginal distributions of the measurement errors associated with the three components to all be different.

Finally, we set  $w_{\ell,i,j} = x_{\ell,i} + u_{\ell,i,j}$  with  $u_{\ell,i,j} = s_{\ell}(x_{\ell,i})\epsilon_{\ell,i,j}$  and  $s_{\ell}(x_{\ell}) = x_{\ell}/3$  for each  $\ell$ .

Component	Sex	Median ISE $\times 1000$	
		Zhang, et al. (2011)	Our Proposed Method
1	M	6.08	<b>2.99</b>
	W	4.69	<b>2.75</b>
2	M	18.66	<b>4.92</b>
	W	15.61	<b>2.52</b>
3	M	10.06	<b>5.88</b>
	W	18.84	<b>3.70</b>

Table S.2: Median integrated squared error (MISE) performance of the covariate informed density deconvolution method developed in this article compared with the methods of Zhang *et al.* (2011). Here M and W are abbreviations for ‘men’ and ‘women’, respectively.




Figure S.4: Results for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs. These results correspond to the final MCMC iteration but are representative of other iterations in steady state. The left panel shows the component specific parameters  $(\mu_{x,k}, \sigma_{x,k}^2)$  for the six mixture components that were actually used to model the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . The right panel shows the associated ‘empirical’ mixture probabilities  $\hat{p}_x(k | c_0, c_1, \dots, c_p) = \sum_{i=1}^n 1\{z_{x,\ell,i} = k, c_{0,\ell,i} = c_0, c_{1,\ell,i} = c_1, \dots, c_{p,\ell,i} = c_p\}/n$  for ‘men’ and ‘women’ and for the three components, from left to right. Results for different combinations of component and ‘gender’ are shown here as they are the only predictors important for  $\mathbf{x}$ . Consistent with the simulation truth, the mixture probabilities vary significantly between these predictor combinations.




Figure S.5: Results for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs. These results correspond to the final MCMC iteration but are representative of other iterations in steady state. The left panel shows the component specific parameters  $(p_{\epsilon,k}, \mu_{\epsilon,k}, \sigma_{\epsilon,k,1}^2, \sigma_{\epsilon,k,2}^2)$  for the three mixture components used to model the densities  $f_{\epsilon,\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c})$ . The right panel shows the associated ‘empirical’ mixture probabilities  $\hat{p}_\epsilon(k | c_0, c_1, \dots, c_p) = \sum_{i=1}^n \sum_{j=1}^{m_i} 1\{z_{\epsilon,\ell,i,j} = k, c_{0,\ell,i} = c_0, c_{1,\ell,i} = c_1, \dots, c_{p,\ell,i} = c_p\} / \sum_{i=1}^n m_i$  for ‘men’ and ‘women’ and for the three components, from left to right. Results for different combinations of component and ‘gender’ are shown here. Consistent with the simulation truth, the mixture probabilities vary significantly between different components but not between the gender categories.

The integrated squared error (ISE) of estimation of  $f_{x,\ell}$  by  $\hat{f}_{x,\ell}$  is defined as  $ISE =$




Figure S.6: Results for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs. From top to bottom, the left panels show the estimated conditional densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  obtained by our method and the corresponding truths. The right panels show the associated conditional distributions of the scaled errors  $f_{\epsilon,\ell|\mathbf{c}}(\epsilon_\ell | \mathbf{c})$  obtained by our method and the corresponding truths. Results for different component and ‘gender’ combinations are shown here as they are the only predictors important for modeling the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . Also shown are the associated variance functions  $v_\ell(x_\ell) = s_\ell^2(x_\ell)$ , estimated by our method and the corresponding truths.

$\int \{f_{x,\ell}(x_\ell) - \hat{f}_{x,\ell}(x_\ell)\}^2 dx_\ell$ . A Monte Carlo estimate of ISE for the  $b^{\text{th}}$  simulated data set is given by  $ISE_{est} = \sum_{m=1}^M \{f_{x,\ell}(x_m) - \hat{f}_{x,\ell}^{(b)}(x_m)\}^2 \Delta_m$ , where  $\{x_m\}_{m=1}^M$  are grid points covering  $[A, B]$ , the support of each  $x_\ell$ . Table S.2 reports the median ISEs (MISEs) for estimating the

trivariate joint densities and the univariate marginals obtained by our method, compared with the method of Zhang *et al.* (2011). The MISEs reported here are all based on  $B = 100$  simulated data sets. As Table S.2 shows, our method outperforms Zhang *et al.* (2011) in all cases, often significantly. Our method also produces detailed information about the distributions of the intakes as well as the distributions of the associated measurement errors, including specifically which predictors are mostly important in influencing these distributions. Additional graphical summaries for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs are provided below.

Figure S.3 shows the estimated inclusion probabilities of different predictors in the models for  $f_{\mathbf{x}|\mathbf{c}}$  and  $f_{\epsilon|\mathbf{c}}$ . Consistent with the simulation truth, the set of significant predictors for the density of main interest  $f_{\mathbf{x}|\mathbf{c}}$  is found to comprise the dimension labels ( $c_0$ ) and ‘gender’ ( $c_1$ ), and, for  $f_{\epsilon|\mathbf{c}}$ , the set of significant predictors comprises only the dimension labels ( $c_0$ ).

Figures S.4 and S.5 show how the redundant mixture components become empty after reaching steady states of the MCMC algorithm. They also show how the mixture component specific parameters get shared across different components and predictor combinations. For the densities  $f_{x,\ell|\mathbf{c}}$ , the mixture probabilities vary significantly between ‘men’ and ‘women’ as well as between different dietary components. For the densities  $f_{\epsilon,\ell|\mathbf{c}}$ , on the other hand, the mixture probabilities vary between different dietary components but are very similar between ‘men’ and ‘women’. These results are all consistent with the simulation truth.

Figure S.6 shows the estimated densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  superimposed over histograms of the corresponding estimated  $x_\ell$ ’s obtained by our method. Figure S.9 and S.10 repeat the univariate densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  in the diagonal panels but also show the estimated joint densities  $f_{\mathbf{x},\ell_1,\ell_2|\mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  obtained by our method in the off-diagonal panels separately for men and women, respectively. The results suggest the model to provide a good fit for the simulated data, including being able to capture skewness, multimodality and heavy tails.

Results obtained by the method of Zhang *et al.* (2011) in numerical experiments are discussed in Section S.8 of the supplementary material. The results summarized in the supplementary materials are for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs across all predictor combinations. The fixed effects regression coefficient estimates in the transformed scale for the data set are presented in Table S.1 in the supplementary material. With no mechanism to select the important predictors, all covariates are included in the model. Interestingly, however, a 90% central credible interval based post-processing rule to determine the significance of the associated predictor still produces a good number of spurious significant coefficients - which we believe is an artifact of the bias introduced in the analysis due to non-linear transformations, strong parametric assumptions, exclusion of interaction effects, etc. Figure S.11 in the supplementary materials shows the estimated univariate densities which, as in the case of the real data application, are in general agreement with the estimates produced by our method but are much worse compared to ours in terms of finer details. To reiterate, (a) the main reason the highly restrictive parametric assumptions of Zhang *et al.* (2011) do not get reflected as much in the final density estimates is because these estimates are obtained by a separate kernel density estimation approach applied to the estimated intakes in the original scale; (b) additionally, the approach of Zhang *et al.* (2011) does not provide any insight into the distribution of the measurement errors, including their shapes, how they may be influenced by the associated predictors, etc.



## S.7 Additional Results for the EATS Data Set

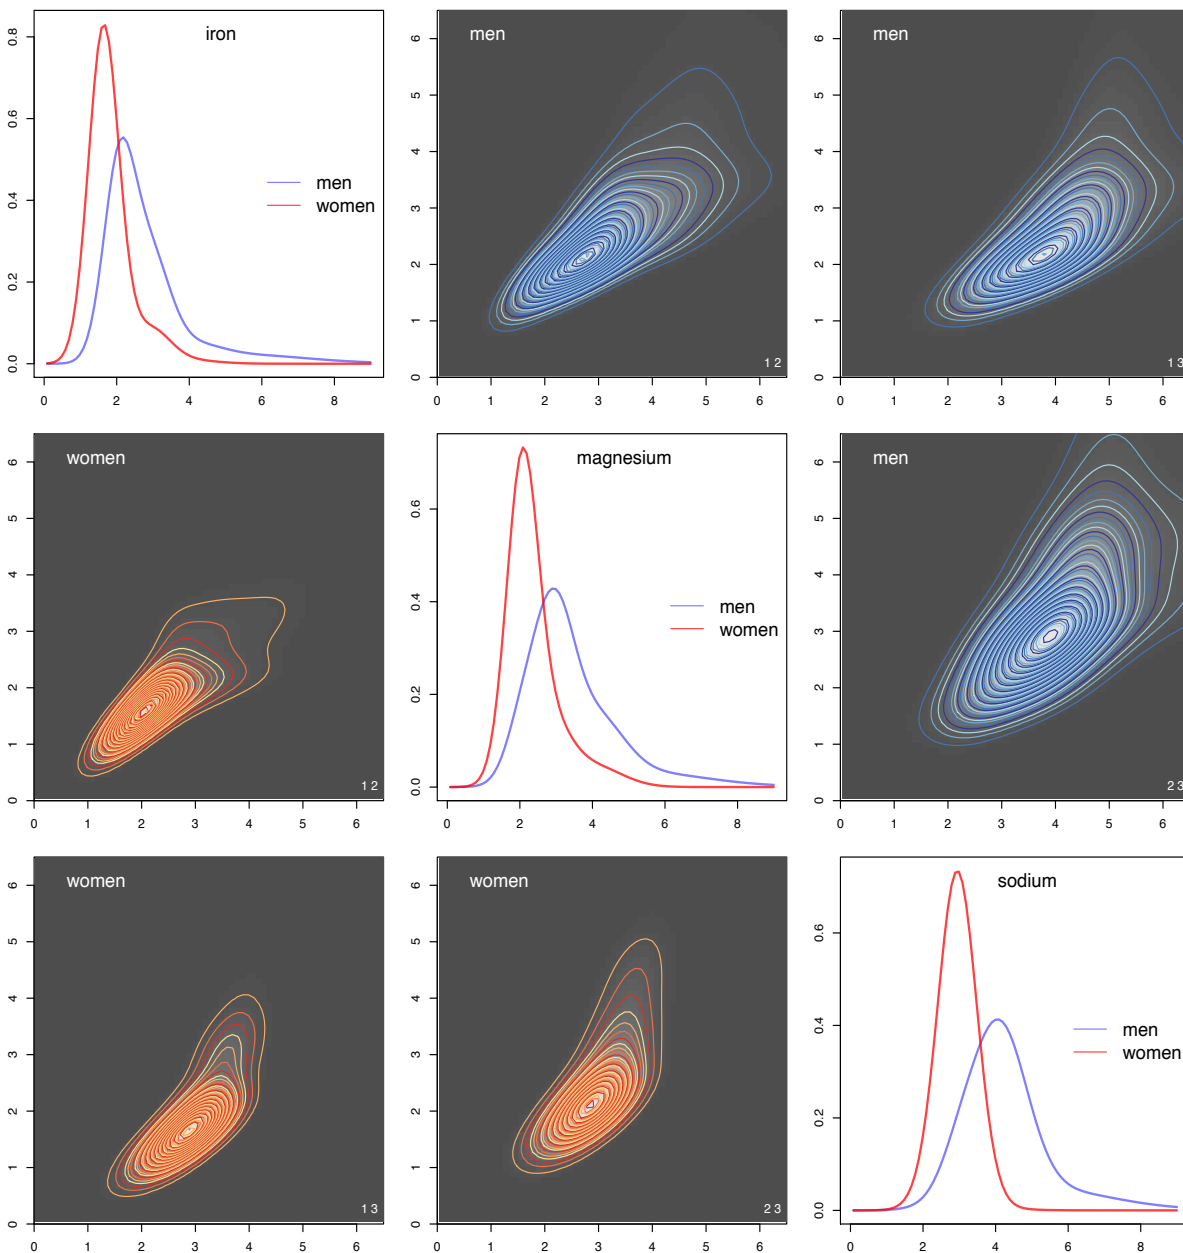


Figure S.7: Results for the EATS data set obtained by our method. Results for different dietary component and gender combinations are shown here as they are the only predictors found important for modeling the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . The diagonal panels show the estimated conditional densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  of iron, magnesium and sodium, respectively. The off-diagonal panels show the estimated joint densities  $f_{\mathbf{x},\ell_1,\ell_2|\mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$ . Results for men are shown in blue, for women are shown in red. Axis labels are suppressed to allow more space for the individual panels.




Figure S.8: Results for the EATS data set produced by the method of Zhang *et al.* (2011). From top to bottom, the left panels show the estimated densities of iron, magnesium and sodium, respectively. Results for men are shown in blue, for women are shown in red.

	Iron	Magnesium	Sodium
Intercept	<b>-1.6949</b> (-2.4356, -0.9409)	<b>-2.0106</b> (-2.8563, -1.1225)	<b>-2.1528</b> (-2.8852, -1.4361)
Woman	<b>1.0028</b> (0.9180, 1.0966)	<b>1.0608</b> (0.9516, 1.1653)	<b>1.1138</b> (1.0292, 1.2034)
Race 1	<b>0.9754</b> (0.3134, 1.6014)	<b>1.1526</b> (0.3603, 1.9438)	<b>0.8766</b> (0.2316, 1.5447)
Race 2	0.6977 (-0.0177, 1.3439)	0.6774 (-0.1347, 1.5296)	<b>0.8989</b> (0.2633, 1.5974)
Race 3	<b>0.8953</b> (0.2107, 1.5648)	<b>1.1218</b> (0.3038, 1.9787)	<b>0.7712</b> (0.0820, 1.4438)
Race 4	<b>0.9703</b> (0.2442, 1.6625)	<b>1.1668</b> (0.3320, 1.9847)	<b>0.7059</b> (0.0196, 1.3802)
Age 1	0.4554 (-0.0302, 0.9439)	0.4151 (-0.0810, 0.9439)	<b>1.0209</b> (0.5497, 1.4854)
Age 2	0.3299 (-0.1165, 0.8084)	0.4570 (-0.0109, 0.9666)	<b>0.8306</b> (0.3454, 1.2670)
Age 3	0.2595 (-0.2279, 0.7423)	0.3936 (-0.0954, 0.8919)	<b>0.7061</b> (0.2275, 1.1361)
Age 4	0.1847 (-0.2870, 0.6661)	0.3623 (-0.1358, 0.8610)	<b>0.6949</b> (0.2022, 1.1429)
Age 5	0.1954 (-0.3022, 0.7023)	0.4332 (-0.0778, 0.9654)	<b>0.6903</b> (0.1878, 1.1354)

Table S.1: Results for the EATS data set produced by the method of Zhang *et al.* (2011). Posterior means of the fixed effects regression effects and associated 90% credible intervals. The estimates with credible intervals not including zero are highlighted in bold.

## S.8 Additional Results for the Simulation Experiments

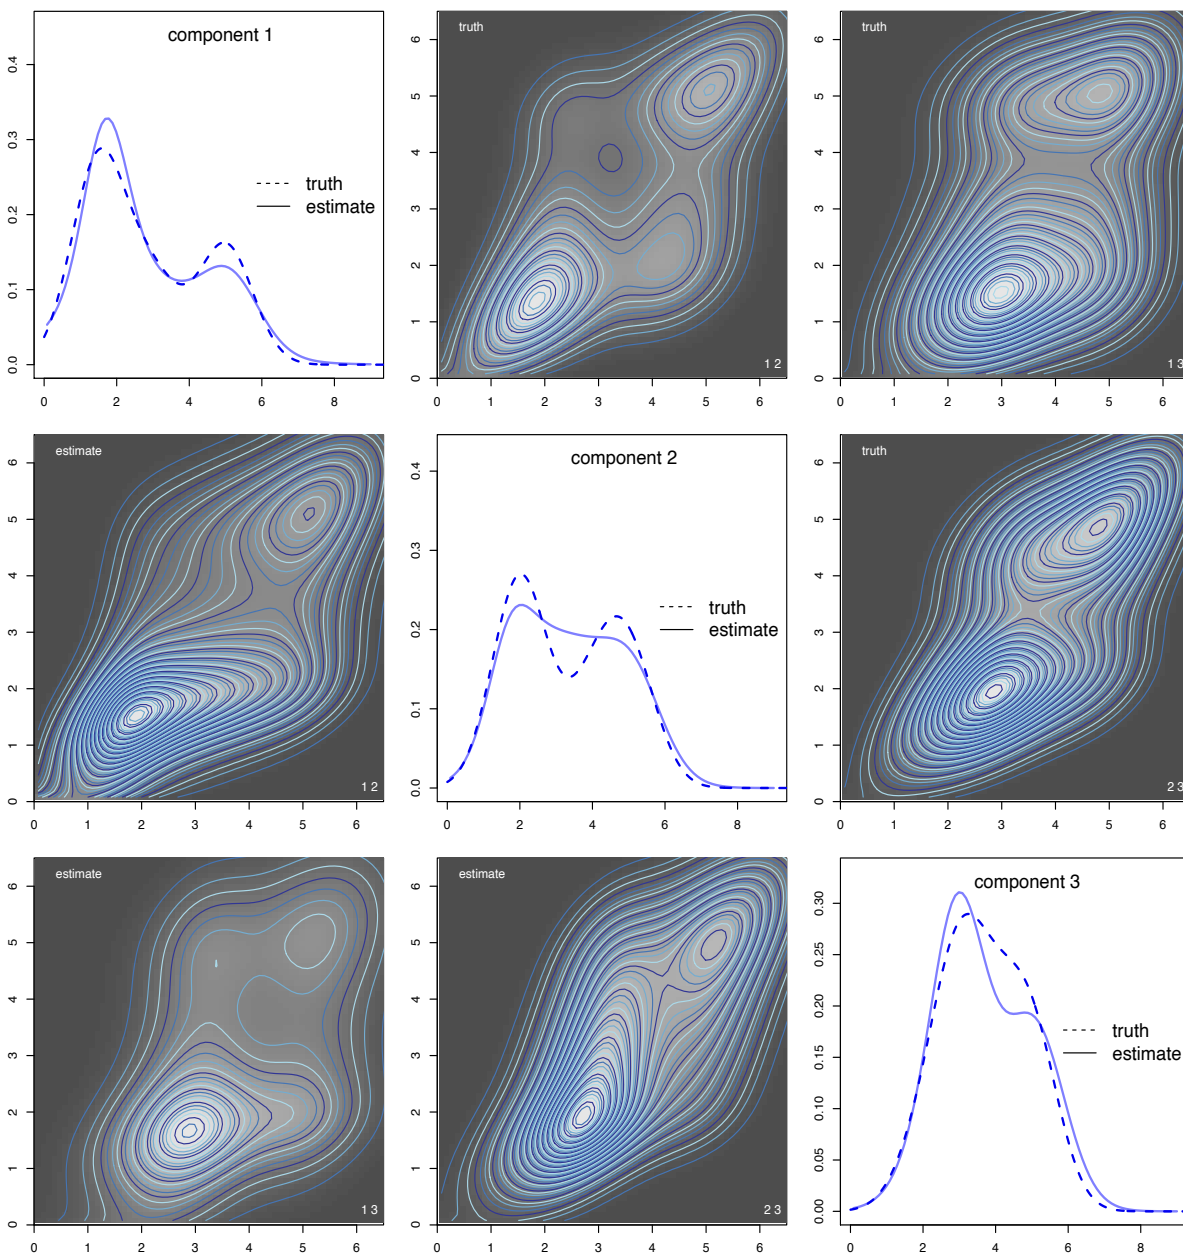


Figure S.9: Results for the synthetic data set corresponding to the 25th percentile of the average ISEs. Results for ‘men’ are shown here. The diagonal panels show one dimensional marginal densities and the corresponding estimates produced by our method. The off-diagonal panels show the contour plots of the true two-dimensional densities  $f_{\mathbf{x}, \ell_1, \ell_2 | \mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  and the corresponding estimates obtained by our method. Axis labels are suppressed to allow more space for the individual panels.

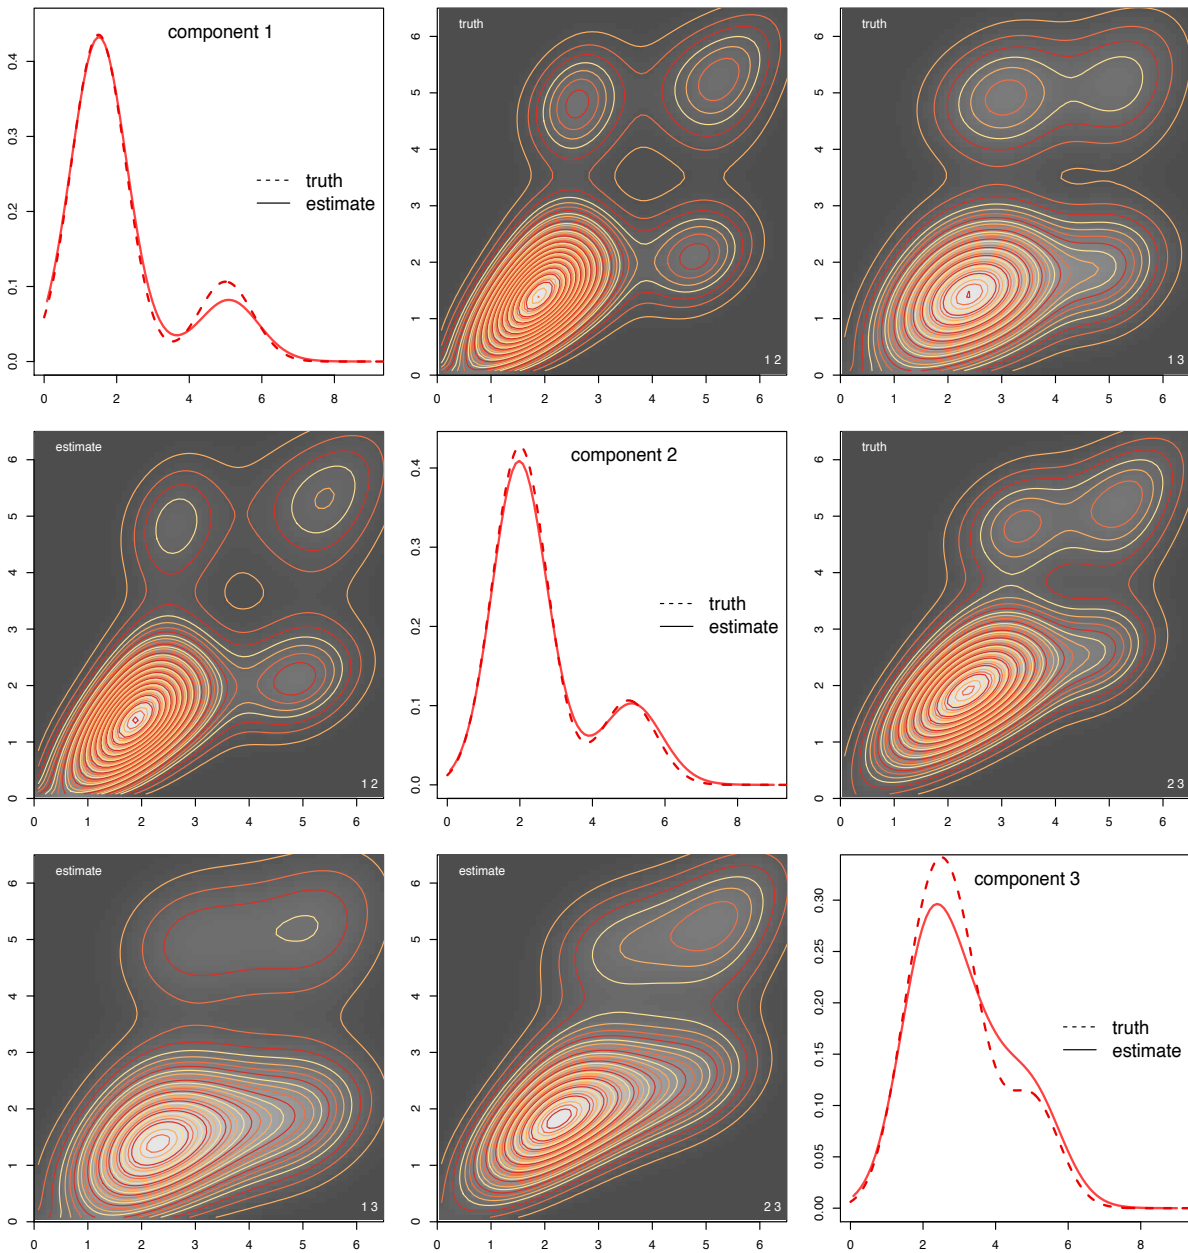


Figure S.10: Results for the synthetic data set corresponding to the 25th percentile of the average ISEs. Results for ‘women’ are shown here. The diagonal panels show one dimensional marginal densities and the corresponding estimates produced by our method. The off-diagonal panels show the contour plots of the true two-dimensional densities  $f_{\mathbf{x}, \ell_1, \ell_2 | \mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  and the corresponding estimates obtained by our method. Axis labels are suppressed to allow more space for the individual panels.




Figure S.11: Results for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs produced by the method of Zhang *et al.* (2011). The panels show the estimated conditional densities  $f_{x,\ell|c}(x_\ell | \mathbf{c})$  obtained by the method of Zhang *et al.* (2011) and the corresponding truths. Results for different component and gender combinations are shown here as they are the only predictors important for modeling the densities  $f_{x,\ell|c}(x_\ell | \mathbf{c})$ .

	Component 1	Component 2	Component 3
Intercept	-0.1166 (-1.2525, 0.9613)	<b>-0.8776</b> (-1.7734, -0.0620)	-0.1128 (-0.8528, 0.6691)
Woman	<b>0.6092</b> (0.4852, 0.7252)	<b>0.5062</b> (0.3980, 0.6159)	<b>0.5023</b> (0.3960, 0.6003)
Race 1	0.2858 (-0.6955, 1.2389)	0.6613 (-0.0601, 1.4242)	0.1167 (-0.5912, 0.8163)
Race 2	0.3730 (-0.6361, 1.4630)	<b>0.7778</b> (0.0377, 1.5489)	0.3307 (-0.4025, 1.0383)
Race 3	0.2222 (-0.8744, 1.2838)	<b>0.9318</b> (0.1707, 1.7002)	0.3871 (-0.3557, 1.0959)
Race 4	0.7873 (-0.3523, 1.7383)	<b>0.8241</b> (0.0159, 1.6462)	0.2596 (-0.4826, 0.9952)
Age 1	-0.5335 (-1.2414, 0.1857)	-0.0646 (-0.6731, 0.5503)	-0.2986 (-0.8749, 0.2256)
Age 2	-0.4836 (-1.1630, 0.2302)	-0.0865 (-0.6755, 0.4994)	-0.2780 (-0.8378, 0.2431)
Age 3	-0.5214 (-1.1801, 0.1904)	-0.0637 (-0.6242, 0.5048)	-0.2779 (-0.8315, 0.2064)
Age 4	-0.2687 (-0.9604, 0.4463)	0.1233 (-0.4948, 0.7186)	-0.2016 (-0.7821, 0.3007)
Age 5	-0.5619 (-1.3525, 0.1676)	-0.0868 (-0.6703, 0.4953)	-0.2880 (-0.8680, 0.2599)

Table S.1: Results for the synthetic data set corresponding to the 25th percentile of the average ISEs produced by the method of Zhang *et al.* (2011). Posterior means of the fixed effects regression effects and associated 90% credible intervals. The estimates with credible intervals not including zero are highlighted in bold.

## S.9 Multivariate Density Regression

The applicability of the methodology developed herein for modeling covariate informed multivariate densities is not restricted exclusively to deconvolution problems but the different model components can be adapted to other important statistics problems as well.

For example, the methodology developed in Section 2.1 in the main paper for modeling  $f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} | \mathbf{c})$  can be straightforwardly applied to the (order-of-magnitude simpler) problem of ordinary multivariate density estimation without measurement errors in the presence of associated potentially high-dimensional precisely measured covariates. Likewise, the methodology developed in Section 2.2 in the main paper for modeling  $f_{\boldsymbol{\epsilon}|\mathbf{c}}(\boldsymbol{\epsilon} | \mathbf{c})$  can be straightforwardly applied to modeling covariate dependent regression errors in the presence of associated potentially high-dimensional precisely measured covariates.

As a matter of illustration, we discuss here the problem of covariate dependent density estimation using the methodology developed in Section 2.1 in the main paper. More rigorous exposition of this problem as well as the problem of modeling covariate dependent multivariate regression errors may be pursued separately elsewhere.

Specifically, we have precisely measured observations  $\mathbf{x}_i = (x_{1,i}, \dots, x_{d,i})^T$  and associated precisely measured covariates  $\mathbf{c}_i = (c_{1,i}, \dots, c_{p,i})^T$  for  $n$  observational units  $i = 1, \dots, n$ , and the goal is to estimate  $f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} | \mathbf{c})$ .

The statistics literature on univariate density estimation is enormous; and the problem of covariate density estimation, sometimes referred to as density regression, has also received attention (see, e.g., Payne *et al.*, 2020, and the references therein). The literature on multivariate density estimation, in contrast, is small and the literature on multivariate density regression almost non-existent. There does exist a body of works on conditional copula estimation. See, for example, Gijbels *et al.* (2011) and the subsequent works citing this paper. The main focus here is often on modeling the conditional dependence patterns. Additionally, most of these works are restricted to simple bivariate settings and/or single covariates. An optimal transport based approach was recently considered in Tabak *et al.* (2020). Variational Bayes inference for a mixture model with multivariate normal component kernels with their means and the mixture probabilities varying with associated covariates has been considered in Dao and Tran (2021). Copula based models with flexible marginals have previously been shown to outperform mixtures of multivariate kernels in realistic scenarios (Sarkar *et al.*, 2021). Additionally, it is not clear how scalable these methods are for high-dimensional problems – high-dimensional  $\mathbf{x}$  as well as high-dimensional  $\mathbf{c}$ . Our method, on the other hand, can efficiently accommodate high-dimensional  $\mathbf{x}$ 's by sparse shared atoms mixture models as well as high-dimensional  $\mathbf{c}$ 's by sparse conditional tensor factorization techniques.

We now report the numerical performance of our method for ordinary density estimation problems in simulation experiments. To our knowledge, software for the alternatives cited above are not publicly available, we restrict ourselves to evaluating the numerical performance of our method alone. We simulated from the same true  $f_{\mathbf{x}|\mathbf{c}}(\mathbf{x} | \mathbf{c})$  that we had considered for the deconvolution problem in Section S.6 in the main paper. As in the case of the deconvolution problem, we set  $n = 965$ , and  $d = 3$ , and the same number  $p = 3$  and distribution of covariates as seen in Figure 1 in the main paper. Other details can be found there in the main paper.

Table S.2 reports the median ISEs (MISEs) for estimating the trivariate joint densities and the univariate marginals obtained by our method. The MISEs reported here are based



Component	Sex	Median ISE $\times 1000$
		Our Model in Section 2.1
1	M	<b>0.112</b>
	W	<b>0.129</b>
2	M	<b>0.081</b>
	W	<b>0.108</b>
3	M	<b>0.046</b>
	W	<b>0.033</b>

Table S.1: Median integrated squared error (MISE) performance of the model developed in Section 2.1 in the main paper applied to covariate informed DENSITY ESTIMATION problem. Here M and W are abbreviations for ‘men’ and ‘women’, respectively.

on  $B = 100$  simulated data sets. As expected, the MISEs are now an order-of-magnitude smaller than the corresponding MISEs for the deconvolution problem reported in Table S.2 in the main paper.

Figure S.12 shows the estimated inclusion probabilities of different predictors in the models for  $f_{\mathbf{x}|\mathbf{c}}$ . Consistent with the simulation truth, the set of significant predictors for the density of main interest  $f_{\mathbf{x}|\mathbf{c}}$  is found to comprise the dimension labels ( $c_0$ ) and ‘gender’ ( $c_1$ ), and, for  $f_{\boldsymbol{\epsilon}|\mathbf{c}}$ , the set of significant predictors comprises only the dimension labels ( $c_0$ ).

Figures S.13 shows how the mixture component specific parameters get shared across different components and predictor combinations. Consistent with the simulation truth, for the densities  $f_{x,\ell|\mathbf{c}}$ , the mixture probabilities vary significantly between ‘men’ and ‘women’ as well as between different components. Figure S.14 shows the estimated densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  obtained by our method superimposed over histograms of the corresponding  $x_\ell$ ’s. Figure S.15 and S.16 repeat the univariate densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  in the diagonal panels and also show the estimated joint densities  $f_{\mathbf{x},\ell_1,\ell_2|\mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  obtained by our method in the off-diagonal panels separately for ‘men’ and ‘women’, respectively. These graphical summaries suggest the model to provide an excellent fit for the simulated data.

As also discussed in the concluding section of the main paper, our construction implies each component depends on the same of important covariates which can be restrictive in real world applications. The problem can be addressed by constructing more flexible partition structures. In the absence of measurement errors, the copula function can potentially also be modeled more flexibly, including possibly allowing it to vary with the covariates. The problem of modeling flexible regression errors can be similarly addressed. These research directions are being pursued separately elsewhere.

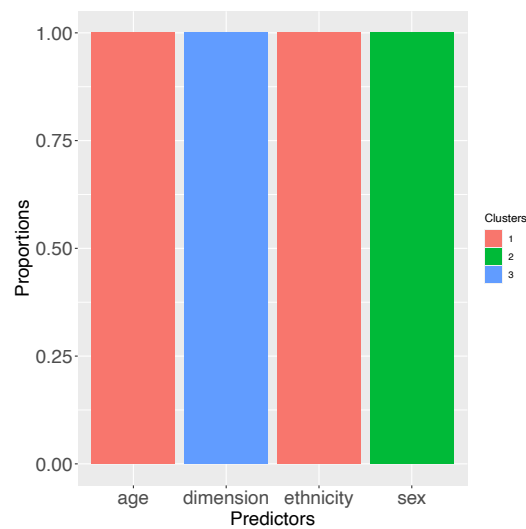


Figure S.12: Results for DENSITY ESTIMATION for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs, showing the estimated probabilities of different numbers of clusters of the associated predictors' levels being included in the model for the densities  $f_{\mathbf{x}|\mathbf{c}}$ . At the median 0.5 probability level, the component labels and the sex of the subjects are important predictors for modeling the densities  $f_{\mathbf{x}|\mathbf{c}}$ . The results are consistent with the true simulation scenario.




Figure S.13: Results for DENSITY ESTIMATION for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs. These results correspond to the final MCMC iteration but are representative of other iterations in steady state. The left panel shows the component specific parameters  $(\mu_{x,k}, \sigma_{x,k}^2)$  for the eleven mixture components that were actually used to model the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ . The right panel shows the associated ‘empirical’ mixture probabilities  $\hat{p}_x(k | c_0, c_1, \dots, c_p) = \sum_{i=1}^n 1\{z_{x,\ell,i} = k, c_{0,\ell,i} = c_0, c_{1,\ell,i} = c_1, \dots, c_{p,\ell,i} = c_p\}/n$  for ‘men’ and ‘women’ and for the three components, from left to right. Results for different combinations of component and ‘gender’ are shown here as they are the only predictors important for  $\mathbf{x}$ . The mixture probabilities vary significantly between these predictor combinations.




Figure S.14: Results for DENSITY ESTIMATION for the synthetic data set corresponding to the 25<sup>th</sup> percentile of the average ISEs. From top to bottom, the left panels show the estimated conditional densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$  obtained by our method and the corresponding truths. Results for different component and ‘gender’ combinations are shown here as they are the only predictors important for modeling the densities  $f_{x,\ell|\mathbf{c}}(x_\ell | \mathbf{c})$ .

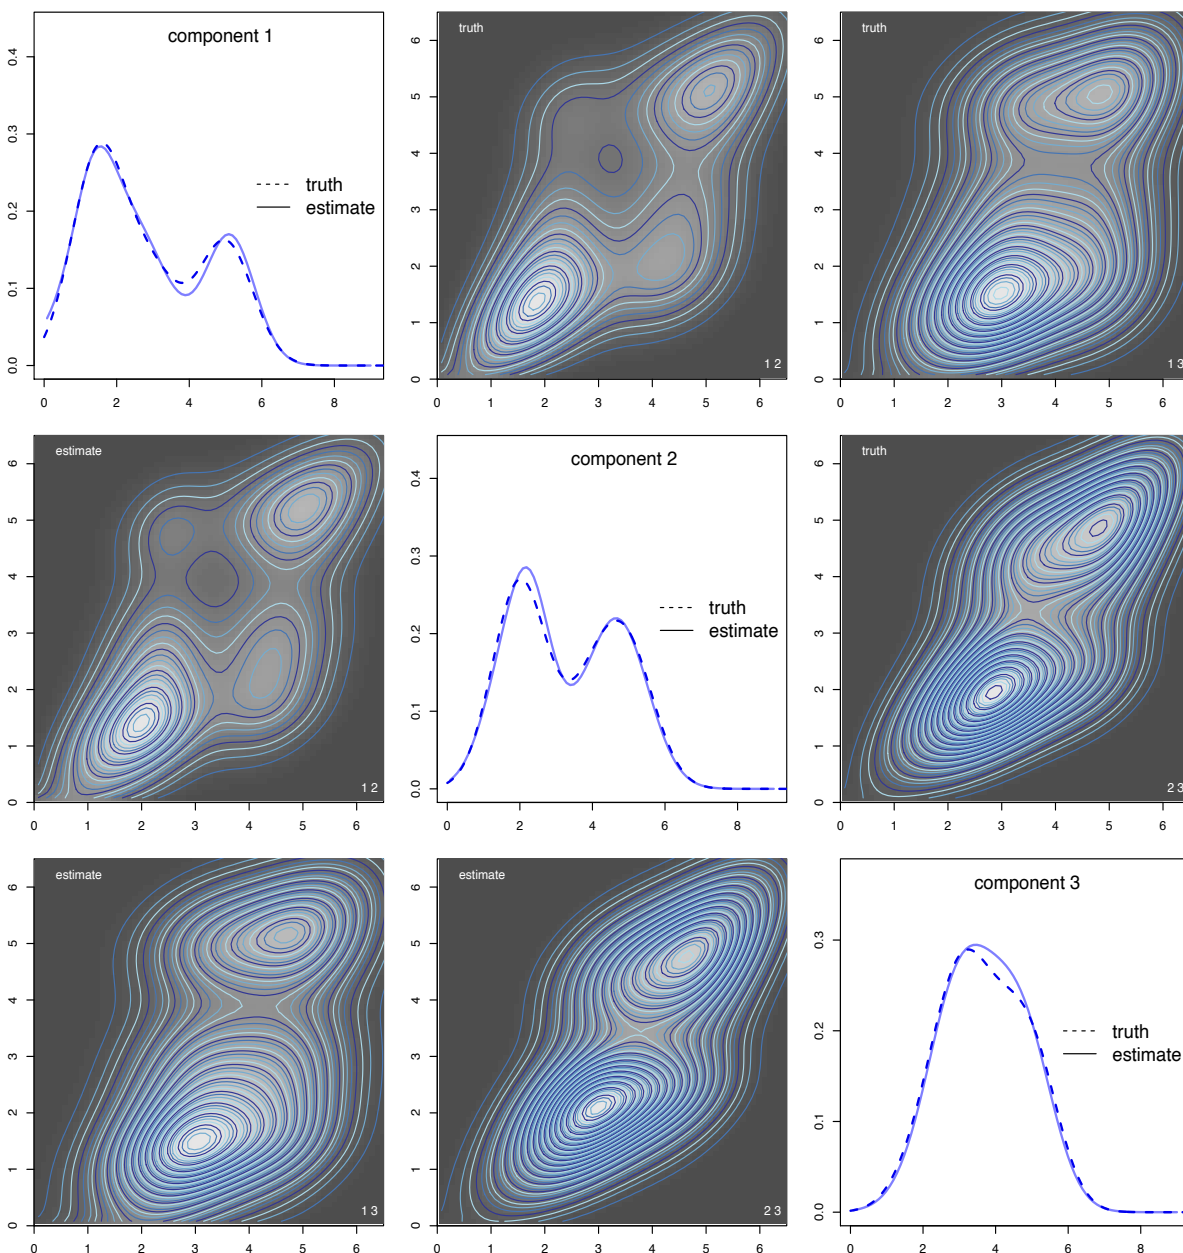


Figure S.15: Results for DENSITY ESTIMATION for the synthetic data set corresponding to the 25th percentile of the average ISEs. Results for ‘men’ are shown here. The diagonal panels show one dimensional marginal densities and the corresponding estimates produced by our method. The off-diagonal panels show the contour plots of the true two-dimensional densities  $f_{\mathbf{x}, \ell_1, \ell_2 | \mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  and the corresponding estimates obtained by our method. Axis labels are suppressed to allow more space for the individual panels.




Figure S.16: Results for DENSITY ESTIMATION for the synthetic data set corresponding to the 25th percentile of the average ISEs. Results for ‘women’ are shown here. The diagonal panels show one dimensional marginal densities and the corresponding estimates produced by our method. The off-diagonal panels show the contour plots of the true two-dimensional densities  $f_{\mathbf{x}, \ell_1, \ell_2 | \mathbf{c}}(x_{\ell_1}, x_{\ell_2} | \mathbf{c})$  and the corresponding estimates obtained by our method. Axis labels are suppressed to allow more space for the individual panels.

## References

- Beraha, M., Guglielmi, A., and Quintana, F. A. (2021). The semi-hierarchical dirichlet process and its application to clustering homogeneous distributions. *Bayesian Analysis*, **16**, 1187–1219.
- Burda, M. and Prokhorov, A. (2014). Copula based factorization in Bayesian multivariate infinite mixture models. *Journal of Multivariate Analysis*, **127**, 200–213.
- Dao, T. and Tran, M.-N. (2021). Flexible multivariate regression density estimation. *Communications in Statistics-Theory and Methods*, **50**, 4703–4717.
- De Lathauwer, L., De Moore, B., and Vandewalle, J. (2000). A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications*, **21**, 1253–1278.
- Denti, F., Camerlenghi, F., Guindani, M., and Mira, A. (2021). A common atom model for the Bayesian nonparametric analysis of nested data. *Journal of the American Statistical Association*, (just-accepted), 1–22.
- dos Santos Silva, R. and Lopes, H. F. (2008). Copula, marginal distributions and model selection: a Bayesian note. *Statistics and Computing*, **18**, 313–320.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, **90**, 577–588.
- Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. Springer, New York.
- Gijbels, I., Veraverbeke, N., and Omelka, M. (2011). Conditional copulas, association measures and their applications. *Computational Statistics & Data Analysis*, **55**, 1919–1932.
- Harshman, R. (1970). Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis. *UCLA working papers in phonetics*, **16**, 1–84.
- Joe, H. (2015). *Dependence Modeling with Copulas*. CRC Press, Boca Raton.
- Nelsen, R. B. (2007). *An Introduction to Copulas*. Springer Science & Business Media, New York.
- Payne, R. D., Guha, N., Ding, Y., and Mallick, B. K. (2020). A conditional density estimation partition model using logistic Gaussian processes. *Biometrika*, **107**, 173–190.
- Pitt, M., Chan, D., and Kohn, R. (2006). Efficient Bayesian inference for Gaussian copula regression models. *Biometrika*, **93**, 537–554.
- Sarkar, A. and Dunson, D. (2020). Bayesian semiparametric higher order hidden markov models. arXiv preprint.
- Sarkar, A., Pati, D., Mallick, B. K., and Carroll, R. J. (2021). Bayesian copula density deconvolution for zero-inflated data in nutritiona epidemiology. *Journal of the American Statistical Association*, **116**, 1075–1087.

- Shemyakin, A. and Kniazev, A. (2017). *Introduction to Bayesian Estimation and Copula Models of Dependence*. John Wiley & Sons, Hoboken.
- Sklar, A. (1959). Fonctions de repartition a n dimensions et leurs marges. *Publications de l'Institut de Statistique de L'Universite de Paris*, **8**, 229–231.
- Tabak, E. G., Trigila, G., and Zhao, W. (2020). Conditional density estimation and simulation through optimal transport. *Machine Learning*, **109**, 665–688.
- Tucker, L. (1966). Some mathematical notes on three-mode factor analysis. *Psychometrika*, **31**, 273–282.
- Wu, J., Wang, X., and Walker, S. G. (2014). Bayesian nonparametric inference for a multivariate copula function. *Methodology and Computing in Applied Probability*, **16**, 747–763.
- Wu, J., Wang, X., and Walker, S. G. (2015). Bayesian nonparametric estimation of a copula. *Journal of Statistical Computation and Simulation*, **85**, 103–116.
- Yang, Y. and Dunson, D. B. (2016). Bayesian conditional tensor factorization for high-dimensional classification. *Journal of the American Statistical Association*, **111**, 656–669.