



<b>Title</b>	Inferring food intake from multiple biomarkers using a latent variable model
<b>Authors(s)</b>	Brennan, Lorraine, D'Angelo, Silvia, Gormley, Isobel Claire
<b>Publication date</b>	2021-12
<b>Publication information</b>	Brennan, Lorraine, Silvia D'Angelo, and Isobel Claire Gormley. "Inferring Food Intake from Multiple Biomarkers Using a Latent Variable Model." Institute of Mathematical Statistics, December 2021. <a href="https://doi.org/10.1214/21-AOAS1478">https://doi.org/10.1214/21-AOAS1478</a> .
<b>Publisher</b>	Institute of Mathematical Statistics
<b>Item record/more information</b>	<a href="http://hdl.handle.net/10197/13012">http://hdl.handle.net/10197/13012</a>
<b>Publisher's version (DOI)</b>	10.1214/21-AOAS1478

Downloaded 2026-05-02 01:16:16

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd\_oa)



© Some rights reserved. For more information

# INFERRING FOOD INTAKE FROM MULTIPLE BIOMARKERS USING A LATENT VARIABLE MODEL

Silvia D'Angelo<sup>1</sup>, Lorraine Brennan<sup>2</sup> and Isobel Claire Gormley<sup>1</sup>

<sup>1</sup>School of Mathematics and Statistics, Insight Centre for Data Analytics,  
University College Dublin, Belfield, Dublin 4.

<sup>2</sup>School of Agriculture and Food Science, Conway Institute,  
University College Dublin, Belfield, Dublin 4.

May 7, 2021

## Abstract

Metabolomic based approaches have gained much attention in recent years due to their promising potential to deliver objective tools for assessment of food intake. In particular, multiple biomarkers have emerged for single foods. However, there is a lack of statistical tools available for combining multiple biomarkers to quantitatively infer food intake. Furthermore, there is a paucity of approaches for estimating the uncertainty around biomarker-based inferred intake.

Here, to estimate the relationship between multiple metabolomic biomarkers and food intake in an intervention study conducted under the A-DIET research programme, a latent variable model, multiMarker, is proposed. The multiMarker model integrates factor analytic and mixture of experts models: the observed biomarker values are related to intake which is described as a continuous latent variable which follows a flexible mixture of experts model with Gaussian components. The multiMarker model also facilitates inference on the latent intake when only biomarker data are subsequently observed. A Bayesian hierarchical modelling framework provides flexibility to adapt to different biomarker distributions and facilitates inference of the latent intake along with its associated uncertainty.

Simulation studies are conducted to assess the performance of the multiMarker model, prior to its application to the motivating application of quantifying apple intake.

**Keywords:** Latent variable models, Factor analysis, Mixture of experts, Metabolomics, Ordinal regression, Biomarkers, Food intake

## 1 Introduction

Dietary biomarkers have emerged in recent years as objective measures of food intake ([Baldrick et al. \(2011\)](#)). A dietary biomarker is a small molecule called a metabolite that can provide information on the level of intake of a food ([Gao et al. 2017](#)). By definition a dietary biomarker increases with food intake and is validated against a

defined set of criteria (Dragsted et al. 2018). In recent years, many biomarkers have emerged for a range of foods. The importance of such biomarkers stems from the fact that classical dietary assessment approaches rely on self-reported data which can be subjective and biased, and such issues are well documented in the literature (e.g. Bingham (2002), Kipnis et al. (2002), Subar et al. (2003), Lloyd et al. (2019), Siddique et al. (2019)). As a consequence, dietary biomarkers have emerged as a potential objective tool to aid food intake assessment. As the biomarker field progressed and the analytical tools improved, the number of biomarkers identified as potential biomarkers of intake has increased. Indeed, there are now multiple biomarkers for individual foods. However, there is a paucity of statistical tools for modelling the relationship between multiple biomarkers and food intake. The work to date has employed biomarkers as panels to classify intake into categories, e.g. consumers and non-consumers. For example, Garcia-Aloy et al. (2017) propose a panel of biomarkers for cocoa intake, which is employed to estimate cocoa consumption or non-consumption through a forward stepwise logistic regression model. Rothwell et al. (2014) propose a partial least squares discriminant analysis (PLS-DA) approach to distinguish between low and high consumption of coffee in a sample of individuals, using a panel of three coffee-specific biomarkers. Recent work by Vázquez-Manjarrez et al. (2019) on discovery and validation of banana intake biomarkers also employs a PLS-DA approach to detect low, medium and high consumption of banana. The proposed panel of biomarkers proved to be effective in distinguishing between low and high consumers, while medium consumers were difficult to separate from the two extremal groups. Gürdeniz et al. (2016) present a PLS-DA approach to detect beer intake (consumers versus non-consumers) using a panel of aggregated biomarkers. While the aforementioned approaches avail of a panel of biomarkers, they provide only a categorical quantification of intake, not an estimation of the quantity of intake, nor its associated uncertainty.

The A-DIET research programme ([www.ucdnutrimarkers.com/a-diet](http://www.ucdnutrimarkers.com/a-diet)) aims to identify new metabolomic biomarkers of dietary intake, and here provides the motivating context. Data from a panel of four novel metabolomic biomarkers were collected from an intervention study, where a group of participants consumed three different food quantities (of apple), over a three week period. The question of interest and associated statistical challenge is twofold: given the intervention study data, the aim is to estimate the relationship between the biomarkers and food intake. At a later stage, when biomarker data only have been collected, the objective is to infer intake from the panel of biomarkers alone, and to provide the associated uncertainty.

Here, to estimate the relationship between multiple biomarkers and food intake we have developed the “multiMarker” model. MultiMarker relies on a factor analytic latent variable construct (Knott & Bartholomew 1999) to capture the relationship between the panel of observed biomarkers and the unobserved intake. The distribution of the latent factor (i.e. the unobserved food intake) is expressed through a flexible mixture model. Moreover, in order to “refine” inferred intakes, multiMarker avails of a mixture of experts framework (Jacobs et al. (1991), Gormley & Frühwirth-Schnatter (2019)), such that the biomarker data first informs the inferred level of intake (relative to the most likely food quantity), with the more refined inferred intakes then resulting from the factor analytic aspect of multiMarker. The model is developed in a Bayesian framework, naturally allowing for uncertainty quantification and therefore providing more informed quantification of intake.

Traditional factor analysis models assume that the correlation structure between a collection of observed variables can be represented in terms of a linear combination of a lower number of latent variables, the factors (Knott & Bartholomew 1999). The latent variable(s) are often assumed to follow a Gaussian distribution, and

many extensions have been proposed. Allowing an infinite number of factors to facilitate greater modelling flexibility has been proposed by [Bhattacharya & Dunson \(2011\)](#), and [Murphy et al. \(2020\)](#) extend the framework to an infinite mixture of infinite factor analysers. When in the presence of an heterogeneous population, [Montanari & Viroli \(2010\)](#) proposed an heteroscedastic factor mixture analysis model, where factors are distributed according to a mixture of multivariate Gaussian distributions. Relaxations of the Gaussianity assumption for the factors have been considered by many authors, see for example [McLachlan et al. \(2007\)](#), [Murray et al. \(2014\)](#) and [Lin et al. \(2016\)](#). An extension of the factor mixture analysis approach to multivariate binary response data has been considered by [Cagnone & Viroli \(2012\)](#). [Galimberti et al. \(2009\)](#) propose an approach to dimensionality reduction in factor mixture analysis models, where factor loadings are shrunk through a penalized likelihood approach. Robustification of factor analysis models was addressed by [Pison et al. \(2003\)](#), who introduced a method to address outliers. A further issue with factor analytic models is non-identifiability, see for example [Lopes & West \(2004\)](#), [Ročková & George \(2016\)](#) and [Frühwirth-Schnatter & Lopes \(2018\)](#).

While much factor analytic research has focused on two common issues, the number of factors to employ and the distribution to adopt, attempting to model the scale of the latent factors has received little attention. Typically the latent factors are perceived as instrumental tools to achieve a lower-dimensional representation of the data at hand. On the contrary, motivated by the application, our approach is focused on a single latent variable that is a proxy for the latent intake. Information on the scale of this latent variable is available and necessary in order to provide practically useful quantification of intake. While the availability of such information also ensures the multiMarker model is identifiable, unlike general factor analytic models, it adds complexity to inferring the latent variable which is no longer a simple instrument. Modelling the latent factor via a mixture model allows for a flexible framework, but introduces the issue of properly modelling mixture weights. Indeed, mixture components have the role of locating different regions in the intake range for the latent variable, where the order of such regions is relevant and should not be ignored. To this end, we embed the model for the latent variable in a mixture of experts framework, where the weights are informed by the observed biomarkers. Further, when modelling the weights we directly account for the ordinal nature of the food quantities via an ordinal regression model ([Agresti 1999](#)), employing the robust Cauchit link function ([Morgan & Smith \(1992\)](#)). Given subsequently observed biomarker data only, inference on the latent intake and its associated uncertainty is available through the latent variable’s posterior predictive distribution.

In what follows, [Section 2](#) details the motivating application of inferring food (specifically apple) intake. [Section 3](#) outlines the multiMarker model, while [Section 4](#) develops an efficient Metropolis-within-Gibbs sampling strategy for Bayesian inference. [Section 5](#) provides details of thorough simulation studies exploring the performance of the proposed model across a series of realistic settings. The multiMarker model is applied in the motivating context of inferring apple intake in [Section 6](#), with the concluding [Section 7](#) discussing the application outcomes, the multiMarker model and possible extensions. An R ([R Core Team 2020](#)) package, multiMarker ([D’Angelo et al. 2020](#)), is freely available to facilitate widespread use of the method, and with which all results presented herein were produced.

## 2 Modelling the relationship between food intake and multiple biomarkers: apple intake as an example

In the present work we develop a model to estimate the relationship between apple intake and  $P = 4$  urinary biomarkers, identified using an untargeted metabolomics approach (McNamara et al. 2020). The four urinary biomarkers are: Xylose, Epicatechin Sulfate, [(4-{3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-4,6-dihydroxy-2-methoxyphenyl}-2-methylbut-2-en-1-yl)oxy] sulfonic acid, and Glucodistylin. Throughout the paper we will refer to the third biomarker as  $(4 - 3 - [2 - (2, 4 - dihydroxyphenyl) - 2 - oxoethyl] - DHMPMB - SA)$ .

The data were collected as part of an intervention study, where a group of 32 participants consumed different quantities of apple daily, over a three week intervention period. The intervention study was conducted as part of the A-DIET research programme ([www.ucdnutrimarkers.com/a-diet](http://www.ucdnutrimarkers.com/a-diet)), which aims to identify new metabolomic biomarkers of dietary intake. Data were available following consumption of  $D = 3$  apple quantities: 50, 100 and 300 grams. The participants were fed the different food quantities in each of three intervention weeks, where each intervention week was followed by a resting week. Out of the 32 participants, 29 consumed 50 grams of apple, while respectively 28 and 29 ate 100 and 300 grams. Throughout the paper we treat the total of  $n = 86$  recorded intakes as independent observations and this assumption is assessed in Section 6.

The original values for Epicatechin Sulfate,  $(4 - 3 - [2 - (2, 4 - dihydroxyphenyl) - 2 - oxoethyl] - DHMPMB - SA)$  and Glucodistylin caused computational instability (most values were larger than  $10^7$ ) and consequently were scaled (see Section 3.1 of the Supplementary Material). However, the transformation did not alter the correlations between the four biomarkers. A visualization of the data is given in Figure 1. The first three urinary biomarkers have similar median values for the first two apple quantities, and all biomarkers are highly variable for the 300 grams apple quantity. The boxplots corresponding to different apple quantities partially overlap, indicating that such quantities cannot be easily discriminated.

Given the intervention study data, we estimate the relationship between the panel of biomarkers and apple intake. We then infer the apple consumption (in grams) of new observations, for which only the metabolomic biomarker measurements are available.

## 3 The multiMarker model

To quantify food (in this case apple) intake from a panel of biomarkers we propose the multiMarker model which integrates factor analytic and mixture of experts models. The model estimates the relationship between the multiple biomarkers and food intake and facilitates inference on the latent intake when only biomarker data are subsequently observed.

### 3.1 Model specification

Consider a biomarker matrix  $Y$  of dimension  $n \times P$ , storing  $P$  different biomarkers measurements on a set of  $n$  independent observations. The  $D$  food quantities considered in the intervention study are denoted by  $\mathbf{X} = \{X_1, \dots, X_d, \dots, X_D\}$ . Elements in  $\mathbf{X}$  are ordered such that  $X_d < X_{d+1}$ . We assume that the biomarker




Figure 1: Biomarker levels following consumption of different quantities of apple. Boxplots for the  $n = 86$  observations on the four biomarkers, for the three apple quantities (50, 100 and 300 grams).

measurements are related to an unobserved continuous intake value, for which the food quantities are proxies, leading to the following factor analytic model:

$$y_{ip} = \alpha_p + \beta_p z_i + \epsilon_{ip}, \quad \forall \quad i = 1, \dots, n, \quad p = 1, \dots, P, \quad (1)$$

where the (one dimensional) latent variable  $z_i$  denotes the latent intake of observation  $i$ . The  $\alpha_p$  and  $\beta_p$  parameters characterize, respectively, the intercept and the scaling effect for the  $p^{\text{th}}$  biomarker. We assume a truncated Gaussian prior distribution,  $\alpha_p \sim \mathcal{N}_{[0, \infty)}(\mu_\alpha, \sigma_\alpha^2)$ , for the intercept parameters. The non-negative assumption is required as biomarker levels cannot be negative. The  $p^{\text{th}}$  scaling parameter  $\beta_p$  captures the effect of an increment in consumption of a given food on the observed level of biomarker  $p$ . As validated biomarkers are known to be reflective of intake, we assume  $\beta_p > 0, \forall p = 1, \dots, P$  and place a truncated Gaussian prior distribution on  $\beta_p$  i.e.  $\beta_p \sim \mathcal{N}_{(0, \infty)}(\mu_\beta, \sigma_\beta^2)$ .

The error term  $\epsilon_{ip}$  is the variability associated with the  $p^{\text{th}}$  biomarker. As is common in factor analytic models, we assume that  $\epsilon_{ip} \sim \mathcal{N}(0, \sigma_p^2)$ , where  $\sigma_p^2$  serves as a proxy for the precision of the  $p^{\text{th}}$  biomarker. A precise biomarker will have a value of  $\epsilon_{ip}$  close to zero but may take positive or negative values and thus  $\epsilon_{ip} \in (-\infty, \infty)$ . However, to guarantee non-negative biomarker values, the further assumption  $y_{ip} \sim \mathcal{N}_{[0, \infty)}(\alpha_p + \beta_p z_i, \sigma_p^2)$  is required. Here, differently from standard factor analytic models, the scale of the latent variable plays a central role, and its accurate, quantitative recovery is a central requirement of the motivating application.

The likelihood function conditional on  $\mathbf{z} = (z_1, \dots, z_n)^T$  is

$$\ell(\mathbf{Y}|\alpha, \beta, \mathbf{z}, \Sigma) = C + \sum_{p=1}^P \sum_{i=1}^n \left( -\frac{1}{2} \log(2\pi\sigma_p^2) - \frac{1}{2\sigma_p^2} (y_{ip} - \alpha_p - \beta_p z_i)^2 \right) \quad (2)$$

where  $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_P^2)$  and  $C$  is a constant.

## 3.2 Modelling the latent intakes

Although the quantities of interest  $\mathbf{z} = (z_1, \dots, z_n)^T$  cannot be directly measured, food quantities from the intervention study are available and are employed to inform the distribution of the latent intakes. While the specific food quantity consumed by each intervention study observation is known, here this information is not directly modelled, but is used to inform hyperparameter settings and starting values when estimating the multiMarker model (see Supplementary Material, Section 1.1). This approach ensures that the same structural model is used when both estimating the relationship between the biomarkers and intake from the intervention study observations' data, and when inferring intake for new observations for which biomarker data only are available.

To specify a prior distribution for the latent intake, we assume that the true but latent intakes of the intervention study's observations will be distributed around the food quantities used in the study. Thus we assume a mixture of experts model with  $D$  truncated Gaussian components as the prior distribution for the latent intake

$$z_i \sim \sum_{d=1}^D \pi_{id} \mathcal{N}_{[0, \infty)}(X_d, \theta_d^2), \quad \pi_i = \{\pi_{i1}, \dots, \pi_{iD}\}, \quad (3)$$

for  $i = 1, \dots, n$  and  $d = 1, \dots, D$ . The use of truncated distributions follows naturally from the definition of intake, which is non-negative. The  $d^{\text{th}}$  component in the mixture of experts model is centred at the  $d^{\text{th}}$  food quantity  $X_d$ . The variances  $\Theta = \{\theta_1^2, \dots, \theta_D^2\}$  represent intake variability around the intervention study's food quantities; in such intervention studies the range of food quantities considered is carefully chosen to realistically reflect feasible intake amounts. Finally, the observation-specific mixture weights are assumed to be a function of the observed biomarker values,  $\pi_{id} = \pi_{id}(y_i)$ , ensuring that estimation of the relationship between the biomarkers and food intake, and the subsequent inference of food intake from biomarkers alone, are derived from the same model, philosophically and structurally.

As is standard in mixture models, an unobserved component allocation label  $c_i$  is introduced for observation  $i$  and  $\mathbf{c} = \{c_1, \dots, c_n\}$ . These latent auxiliary variables indicate the food quantity/component to which each observation belongs, depending on the corresponding biomarker measurements. Formally,  $c_i \sim \text{Multinomial}\{1, \pi_i = (\pi_{i1}, \dots, \pi_{iD})\}$ . For more details on  $\mathbf{c}$  see Section 3.2.1. The mixture model in (3) can then be written according to its complete data representation:

$$z_i | c_i \sim \prod_{d=1}^D \left\{ \mathcal{N}_{[0, \infty)}(X_d, \theta_d^2) \right\}^{[c_i=d]} \quad (4)$$

where  $[ \cdot ]$  is the Iverson bracket ( $[A]$  returns 1 if the proposition  $A$  is true, and 0 otherwise).

### 3.2.1 Modelling observation-specific weights

Given the inherent ordering of the food quantities  $X_1, \dots, X_D$  in the intervention study, and hence the ordering in the mixture components, we employ an ordinal regression model with Cauchit link function to model the observation-specific weights. The Cauchit link function is more robust than the standard logit or probit link

functions and thus is suited to the often highly variable biomarker data. Specifically, we assume:

$$\begin{aligned}
p(\pi \mid \gamma, \eta, \mathbf{c}, \mathbf{Y}, \mathbf{X}) &= \prod_{i=1}^n \prod_{d=1}^D \pi_{id}^{[c_i=d]} = \prod_{i=1}^n \prod_{d=1}^D \pi_{id}(y_i \mid \gamma_d, \eta)^{[c_i=d]} \\
&= \prod_{i=1}^n \prod_{d=1}^D \left[ \Pr(x_i \leq X_d \mid \gamma_d, \eta, y_i) - \Pr(x_i \leq X_{d-1} \mid \gamma_{d-1}, \eta, y_i) \right]^{[c_i=d]}
\end{aligned} \tag{5}$$

where  $\Pr(x_i \leq X_d \mid \gamma_d, \eta, y_i) = \frac{1}{\pi} \left[ \arctan \left( \frac{1}{2}(\gamma_d + \eta y_i) \right) + \frac{\pi}{2} \right]$  is the probability that observation  $i$ 's consumed quantity  $x_i$  is less than or equal to the food quantity  $X_d$ , given its  $P$  biomarker measurements  $y_i$ , expressed through a Cauchit link. The vector  $\gamma = \{-\infty = \gamma_0, \gamma_1, \dots, \gamma_d, \dots, \gamma_D = \infty\}$  contains food quantity-specific intercepts. The  $\eta$  parameter is a  $P$  dimensional scaling coefficient for the biomarkers, expressing the contribution of each to the determination of the observation-specific mixture weights.

## 4 An MCMC algorithm for parameter estimation

We adopt a Bayesian approach to estimate the hierarchical model's parameters, implemented through a Metropolis within Gibbs Markov chain Monte Carlo (MCMC) algorithm. Hyperprior distributions are assumed on the prior parameters with the corresponding hyperparameter values fixed based on the data at hand, following an empirical Bayes approach. Hyperparameter specifications are reported in the Supplementary Material (Section 1.1). Figure 2 provides a graphical representation of the multiMarker model, with the steps of the MCMC algorithm detailed below.




Figure 2: Hierarchical structure of the multiMarker model. Shaded circles represent the data. Parameters and latent variables are indicated with transparent circles and hyperparameters are indicated with no circles.

### 4.1 Biomarker-specific regression parameters

For the  $p^{\text{th}}$  biomarker, the truncated Gaussian prior distributions of both  $\alpha_p$  and  $\beta_p$  are combined with the likelihood function leading to the full conditional distributions of  $\alpha_p$  and  $\beta_p$ :

$$\alpha_p \sim \mathcal{N}_{[0, \infty)}(\mu_{\alpha_p}^*, \sigma_{\alpha_p}^{2*}); \quad \beta_p \sim \mathcal{N}_{(0, \infty)}(\mu_{\beta_p}^*, \sigma_{\beta_p}^{2*})$$

where

$$\sigma_{\alpha_p}^{2*} = \frac{\sigma_p^2 \sigma_\alpha^2}{n\sigma_\alpha^2 + \sigma_p^2}; \quad \mu_{\alpha_p}^* = \sigma_{\alpha_p}^{2*} \left[ \frac{\sum_{i=1}^n (y_{ip} - \beta_p z_i)}{\sigma_p^2} + \frac{\mu_\alpha}{\sigma_\alpha^2} \right];$$

$$\sigma_{\beta_p}^{2*} = \frac{\sigma_p^2 \sigma_\beta^2}{\sigma_\beta^2 \sum_{i=1}^n z_i^2 + \sigma_p^2}; \quad \mu_{\beta_p}^* = \sigma_{\beta_p}^{2*} \left[ \frac{\sum_{i=1}^n z_i (y_{ip} - \alpha_p)}{\sigma_p^2} + \frac{\mu_\beta}{\sigma_\beta^2} \right]$$

## 4.2 Error term's variance parameter

Assuming an inverse gamma prior distribution for the  $p^{\text{th}}$  variance term  $\sigma_p^2$ , with shape parameter  $\nu_{P1}$  and scale parameter  $\nu_{P2}$ , leads to the following full conditional distribution:

$$\sigma_p^2 \sim \text{Inv}\Gamma\left(\nu_{P1}^* = \frac{n}{2} + \nu_{P1}, \quad \nu_{P2}^* = \frac{1}{2} \sum_{i=1}^n (y_{ip} - \alpha_p - \beta_p z_i)^2 + \nu_{P2}\right)$$

## 4.3 Nuisance parameters

To allow for the uncertainty in the parameters of the prior distributions of  $\alpha_p$  and  $\beta_p$ , hyperprior distributions are specified for  $\mu_\alpha$ ,  $\mu_\beta$  and  $\sigma_\beta^2$ . Specifically,  $\mu_\alpha \sim \mathcal{N}_{[0,\infty)}(m_\alpha, \tau_\alpha \sigma_\alpha^2)$ ,  $\mu_\beta \sim \mathcal{N}_{[0,\infty)}(m_\beta, \tau_\beta \sigma_\beta^2)$ , and  $\sigma_\beta^2 \sim \text{Inv}\Gamma(\nu_{\beta1}, \nu_{\beta2})$ , leading to the following full conditionals:

$$\mu_g \sim \mathcal{N}_{[0,\infty)}(\mu_g^*, \sigma_g^{2*}), \quad \sigma_\beta^2 \sim \text{Inv}\Gamma(\nu_{\beta1}^*, \nu_{\beta2}^*),$$

with

$$\sigma_g^{2*} = \frac{\tau_g \sigma_g^2}{\tau_g P + 1}, \quad \mu_g^* = \sigma_g^{2*} \left[ \frac{\tau_g \sum_{p=1}^P g_p + m_g}{\tau_g \sigma_g^2} \right],$$

$$\nu_{\beta1}^* = \frac{P + 1 + 2\nu_{\beta1}}{2}, \quad \nu_{\beta2}^* = \nu_{\beta2} + \frac{\tau_\beta \sum_{p=1}^P (\beta_p - \mu_\beta)^2 + (\mu_\beta - m_\beta)^2}{2\tau_\beta},$$

where  $g = \alpha, \beta$ . The set of hyperparameters  $(m_\alpha, m_\beta, \tau_\alpha, \tau_\beta, \nu_{\beta1}, \nu_{\beta2})$  are fixed in advance, some of which are informed by the observation-specific food quantity information from the intervention study. Practical guidelines are outlined in the Supplementary Material (Section 1.1). Finally, we fix the value of  $\sigma_\alpha^2 = 1$ , for identifiability.

## 4.4 Observation-specific weight parameters

The prior distributions for the  $\gamma = \{\gamma_0, \gamma_1, \dots, \gamma_d\}$  parameters represent the constrained characteristic of ordinal data, that is:  $X_1 < \dots < X_d < \dots < X_D$ . These constraints correspond to the following in terms of model parameters:  $\gamma_0 < \dots < \gamma_d < \dots < \gamma_D$ . Thus, the prior distribution for  $\gamma_d$  is taken to be  $\gamma_d \sim \mathcal{N}_{(m_{\gamma_{d-1}}, m_{\gamma_{d+1}})}(m_{\gamma_d}, \kappa)$ , for  $d = 1, \dots, D - 1$ . Biomarkers' intercepts'  $\eta = \{\eta_1, \dots, \eta_P\}$  prior distribution is  $\eta_p \sim \mathcal{N}(m_{\eta_p}, \kappa)$ , for  $p = 1, \dots, P$ .

Both  $\gamma$  and  $\eta$  parameters are updated via a Metropolis Hastings step using random walk proposal distributions. Hyperparameters  $m_\gamma$  and  $m_\eta$ , where  $m_\gamma = \{m_{\gamma_0}, m_{\gamma_1}, \dots, m_{\gamma_{D-1}}\}$  and  $m_\eta = \{m_{\eta_1}, \dots, m_{\eta_P}\}$ , are set via their corresponding estimates from the ordinal regression of the known observation-specific component allocation labels on biomarker measurements from the intervention study. More details on the hyperparameters, and on  $\gamma$  and  $\eta$  initialization are deferred to the Supplementary Material (Section 1.2).

Given current  $\mathbf{z}$ ,  $\theta_d^2$ ,  $\gamma$  and  $\eta$  values, weights are updated using (5), together with the unobserved observation-specific component allocation labels:  $c_i = \arg \max_d \frac{\pi_{id} \mathcal{N}_{[0,\infty)}(z_i | X_d, \theta_d^2)}{\sum_{d=1}^D \pi_{id} \mathcal{N}_{[0,\infty)}(z_i | X_d, \theta_d^2)}$ .

## 4.5 Latent intakes

As is common in latent variable models, following (4), we sample the  $i^{\text{th}}$  observation's latent intake value from  $z_i \mid c_i = d, \dots \sim \mathcal{N}_{[0, \infty)}(\mu_{id}^*, \theta_{id}^{2*})$ , with

$$\theta_{id}^{2*} = \left( \sum_{p=1}^P \frac{\beta_p^2 \theta_d^2 + \sigma_p^2 / P}{\theta_d^2 \sigma_p^2} \right)^{-1}, \quad \mu_{id}^* = \sigma_{id}^{2*} \left[ \sum_{p=1}^P \frac{\beta_p (y_{ip} - \alpha_p)}{\sigma_p^2} + \frac{X_d}{\theta_d^2} \right].$$

Further, assuming an inverse gamma prior distribution on the components' variance parameters,  $\theta_d^2 \sim \text{Inv}\Gamma(\nu_{z1}, \nu_{z2})$ , the full conditional distribution is  $\theta_d^2 \sim \text{Inv}\Gamma(\nu_{z1}^*, \nu_{z2}^*)$ , with

$$\nu_{z1}^* = \frac{n_d}{2} + \nu_{z1}, \quad \nu_{z2}^* = \nu_{z2} + \frac{\sum_{i=1}^n [c_i = d] (z_i - x_d)^2}{2},$$

where  $n_d = \sum_{i=1}^n [c_i = d]$ .

## 4.6 Latent intakes' posterior predictive distribution

The model defined by (2) and (4) facilitates estimation of the relationship between biomarkers and latent intake, taking advantage of the information on food quantities used in an intervention study. We are then interested in inferring latent intake values having observed only biomarker measurements for a new group of  $n^*$  observations. Such inference is available through the latent variable's posterior predictive distribution. Denoting by  $z_j^*$  the  $j^{\text{th}}$  new observation's latent intake, and by  $y_j^* = (y_{j1}^*, \dots, y_{jP}^*)^T$  their corresponding biomarker measurements, following (2), (4) and (5), the  $j^{\text{th}}$  latent intake's sampling distribution is:

$$\begin{aligned} p(z_j^* \mid y_j^*, \Omega) &= \mathcal{N}_{[0, \infty)}(\mu_{zj}, \sigma_{zj}^2) \sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty)}(X_d, \theta_d^2) \\ &= \sum_{d=1}^D \pi_{jd} (y_j^* \mid \gamma_d, \eta) \mathcal{N}_{[0, \infty)} \left[ \frac{\mu_{zj} \theta_d^2 + X_d \sigma_{zj}^2}{\sigma_{zj}^2 + \theta_d^2}, \left( \frac{1}{\theta_d^2} + \frac{1}{\sigma_{zj}^2} \right)^{-1} \right] \end{aligned} \quad (6)$$

where for brevity  $\sigma_{zj}^2 = \left( \sum_{p=1}^P \frac{\beta_p^2}{\sigma_p^2} \right)^{-1}$ ,  $\mu_{zj} = \sigma_{zj}^2 \left( \sum_{p=1}^P \frac{\beta_p (y_{jp}^* - \alpha_p)}{\sigma_p^2} \right)$  and  $\Omega = \{\alpha, \beta, \Sigma, \mathbf{X}, \Theta, \eta, \gamma\}$ .

As in Section 3.2, let  $c_j^*$  be the unobserved  $j^{\text{th}}$  observation-specific component allocation label. Hence (6) can be written in its complete data representation:

$$p(z_j^* \mid y_j^*, c_j^*, \Omega) = \prod_{d=1}^D \left[ \mathcal{N}_{[0, \infty)} \left( \frac{\mu_{zj} \theta_d^2 + X_d \sigma_{zj}^2}{\sigma_{zj}^2 + \theta_d^2}, \left( \frac{1}{\theta_d^2} + \frac{1}{\sigma_{zj}^2} \right)^{-1} \right) \right]^{[c_j^* = d]} \quad (7)$$

Thus, the posterior predictive distribution of latent intake  $z_j^*$ , given the observed biomarker data  $y_j^*$  is:

$$p(z_j^* \mid y_j^*, c_j^*) \propto \int p(z_j^* \mid y_j^*, c_j^*, \Omega) p(\Omega \mid \mathbf{Y}) d\Omega \quad (8)$$

Note that the conditional distribution in equation 8 depends on both the intervention study's biomarker data  $\mathbf{Y}$  and new observation  $j$ 's biomarker measurements  $y_j^*$ . Hence, intuitively, to infer latent intake values for  $n^*$  new observations for which biomarker data only are available, one needs to have first estimated the multiMarker model on some intervention study data. Moreover, the  $P$  biomarkers measured for the  $n^*$  new observations must be the same as the ones used in the intervention study.

Given  $P$  biomarker measurements  $y_j^*$ , latent intake values  $z_j^*$  can be sampled from the posterior predictive (8) in two steps. First, relevant model parameters,  $\Omega$ , are sampled from their estimated full conditional and

proposal distributions, given observed  $y_j^*$  values and current  $z_j^*$ ,  $c_j^*$  values, for  $j = 1 \dots, n^*$ . Then, given the updated parameter values  $\Omega$ , observation-specific weights  $\pi_{jd}$  are computed together with the unobserved component-allocation labels  $c_j^* = \arg \max_d \frac{\pi_{jd} \mathcal{N}_{[0, \infty)}(z_j^* | X_d, \theta_d^2)}{\sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty)}(z_j^* | X_d, \theta_d^2)}$ , for  $d = 1 \dots, D$  and  $j = 1, \dots, n^*$ . Finally, the latent intake value  $z_j^*$  is sampled from (7), conditioning on the current  $c_j^*$  value and model parameters  $\Omega$ . The choice of sampling latent intakes from (7) rather than from (6) is driven by the fact that the latter implies averaging across the whole mixture distribution to draw  $z_j^*$ , leading to inferred intake values close to the food quantities' sample mean  $\bar{X}_d = \frac{1}{D} \sum_d X_d$ . This leads to poorly inferred  $z_j^*$  values, particularly for observations that truly belong to components located near small or large food quantities. Thus (7), which conditions on the current component allocation labels, is employed as the sampling distribution for the latent intakes in the MCMC algorithm.

## 5 Simulation studies

Several simulation scenarios have been constructed to explore the performance of the proposed approach. The term “training data” refers to simulated intervention study data employed to estimate the multiMarker model while “test data” refers to simulated biomarker only data, from which intake is to be inferred. To depict common experimental settings, the number of biomarkers considered is  $P = 4$ . To represent a variety of real-world scenarios, different food quantity values and sample sizes have been considered. Training datasets' sample sizes are  $n = \{30, 60, 99, 150\}$ , representing a low sample size scenario to a larger sample size one. The test datasets have  $n^* = \lfloor 0.4 \times n \rfloor$  observations.

Three different specifications are considered for the  $\alpha_p$  and  $\beta_p$  parameters, to represent different types of biomarker measures. Also, across the  $P$  biomarkers, three different specifications are considered for  $\sigma_p^2$ , to explore the impact of increasing biomarker variability: a small (*scenario 1*)(S.1), mixed (*scenario 2*)(S.2) or large (*scenario 3*)(S.3) range of variance values. Finally, the mixture components' parameters  $\{\mathbf{X}, \theta_1^2, \dots, \theta_D^2\}$  employed reflect real-world intervention study scenarios: food quantities with stable increments (“stable increments”), food quantities with increasing increments (“increasing increments”), and food quantities with decreasing increments (“decreasing increments”). Two settings are implemented for the components variances  $\theta_d^2$ , corresponding to low and high variability. Observation-specific component allocation labels  $\mathbf{c}$  are sampled at random from the set  $\{1, \dots, D\}$ , leading to similar, yet not equal, proportions of observations being assigned to the food quantities in the training datasets. Combinations of these settings (20 simulated datasets for each) are used to investigate different aspects of performance in three simulation studies detailed below. Figure 3 illustrates two simulated training datasets with low (*scenario 1*) and high (*scenario 3*) biomarker variability levels.

For the MCMC algorithm (Section 4), 30000 iterations were run, both when estimating the model from a training dataset and when inferring intake from a test set's biomarker data only, with the first 6000 iterations being discarded. Further details on the simulation settings, hyperparameter specifications and initializations are deferred to the Supplementary Material (Section 2).

In the absence of comparable approaches for intake quantification in the presence of multiple biomarker data, we compare our results with those obtained using Bayesian linear regression (BLR) and partial least squares (PLS) regression. In BLR, the food quantities are regressed on a linear combination of the biomarkers in

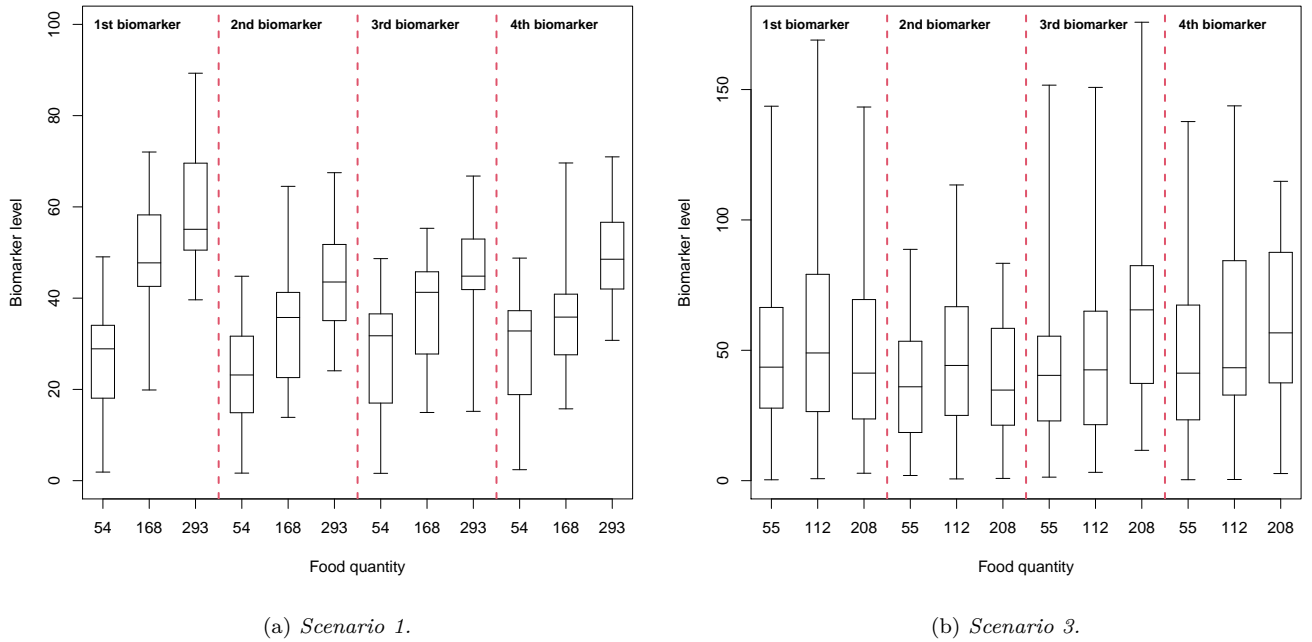


Figure 3: Simulation study I. Examples of simulated biomarker data ( $P = 4, D = 3$ ) under (a) the small (Scenario 1) and (b) large (Scenario 3) biomarkers' variability scenarios. Here,  $n = 99$  with (a) "stable increments" and (b) "increasing increments".

the training dataset with the resulting parameters used to infer intake in the test dataset. In PLS, the food quantities are regressed on a dimensionally reduced representation of the biomarker data in the training dataset.

## 5.1 Simulation study I: intake quantification under varying biomarker variability

Simulation study I assesses performance under increasing noise levels in the  $P$  biomarkers. In the case where  $D = 3$ , we compare results from the three biomarker variability scenarios. Table 1 (Simulation Study I columns) reports the median absolute errors between true and estimated intakes from the training data and true and inferred intakes from the test data.

Overall, for estimated intakes, absolute error values are quite low, with median values ranging from 3g in scenario 1 (S.1)(small biomarker variability) to 9g in scenario 3 (S.3)(large biomarker variability). Notably, the within-scenario error variability is low, and in all cases is commensurate with the median error. Moreover, the absolute errors are relatively stable across increasing noise levels  $\sigma_p^2$ , suggesting robustness of the approach to biomarker variability. Notably the relevant parameters for the multiMarker model have been recovered within a 95% credible interval (see Section 2.4 of the Supplementary Material). Table 1 also shows that the estimated absolute errors under BLR and PLS are higher and more uncertain than those obtained using the multiMarker approach. Also, the BLR and PLS errors increase with increasing biomarker variability  $\sigma_p^2$ .

In terms of inferring intake from biomarker data alone, the multiMarker approach performs well, but with larger errors in scenario 3 (S.3); however, the data simulated under scenario 3 are extreme and unlikely to be seen in real applications (see Figure 3). The BLR errors for these test datasets are much larger and more variable.

Under PLS, the errors are comparable with the proposed approach, however, as is borne out in the application (see Section 6), error values under PLS tend to be relatively low as the inferred intakes tend to be close to the food quantities’ sample mean. While such mean-tendency yields good results in terms of median absolute error values, it corresponds to a lack of precision in the inferred intakes, with more extreme intakes being inferred in the middle of the intake range.

Finally, while a small number of food quantities (e.g.  $D = 3$ ) is practical and cost efficient, a larger  $D$  could increase the coverage of the realistic intake range. Thus, we have investigated the impact of a larger  $D$  value ( $D = 6$ ) on estimation of the multiMarker model and on inferring intake from biomarker data alone. The results and a comparison with the findings from simulation study I are reported in the Supplementary Material (Section 2.1). In general, the performance of multiMarker in both cases ( $D = 3$  or  $D = 6$ ) is comparable, suggesting that the benefit of employing a larger number of food quantities is minimal, and that modelling the latent intake range with a mixture distribution is a flexible and cost efficient approach.

Table 1: Simulation studies’ results. Median (95% CI width) absolute error values (in grams) computed between true and estimated (E) or inferred (I) latent intakes. The values are reported for the three simulation studies (I,II,III) and the three biomarkers’ variability scenarios (S.1,S.2,S.3). Results for the multiMarker (MM) model are reported, as well as those from Bayesian linear regression (BLR) and PLS (PLS) regression.

Model		Simulation Study								
		I			II			III		
		S.1	S.2	S.3	S.1	S.2	S.3	S.1	S.2	S.3
MM	E	3(7)	4(8)	6(18)	3(8)	4(8)	7(62)	6(35)	7(44)	11(62)
	I	4(9)	4(33)	9(59)	5(25)	7(22)	22(77)	8(49)	9(56)	27(74)
BLR	E	77(200)	62(136)	62(136)	61(181)	113(222)	80(118)	66(227)	98(269)	87(197)
	I	76(216)	111(263)	64(137)	62(184)	112(224)	88(108)	67(231)	112(298)	101(138)
PLS	E	10(24)	20(39)	35(60)	9(30)	22(33)	39(56)	26(72)	43(80)	64(65)
	I	10(26)	21(41)	37(69)	10(37)	23(43)	41(67)	31(77)	46(87)	70(80)

## 5.2 Simulation study II: intake quantification under discrepancies between training and test data generation

Due to the controlled environment, observations in an intervention study will have intake values which are similar to the administered food quantities. Thus the associated latent intakes will exhibit low variability around the intake mixture components’ means (i.e. low  $\Theta$  values). However, this is unlikely to be the case when new biomarker measurements for the  $n^*$  test observations are considered. Hence we explore the performance of multiMarker when the variability of the intakes in the test dataset is larger than that of the intakes in the training data.

We simulated data as in simulation study I, but where  $\theta_d^2 \forall d = 1, \dots, D$  are sampled from inverse gamma

distributions with expected values of 6 and 12 in training and test sets respectively (in the small variances setting), and with expected values of 12 and 24 in training and test sets respectively (in the large variances setting). Table 1 (Simulation Study II columns) presents the absolute errors between estimated and inferred latent intakes and the true values. Here, median error values are low and similar to those computed under simulation study I, for scenarios 1 (S.1) and 2 (S.2). Scenario 3 (S.3) presents a relatively larger median error value, and associated 95% CI width, between inferred and true latent intakes (22). However, the multiMarker approach still performs well, given the range of the simulated intakes and the performance of the competitors under the same scenario. This is due to the flexibility of the model for the latent intakes, and multiMarker’s factor analytic aspect, that allow refinement of the inferred latent intakes from biomarker data only to values that potentially differ from those of the intervention study. The BLR and PLS results overall yield similar conclusions to those resulting from simulation study I.

Three additional settings were explored in which there were discrepancies between how training and test data were generated (see Supplementary Material, Section 2.2). In the first additional setting different food quantity values, sampled in comparable and realistic ranges, are considered to generate training and test datasets. Cases in which the number of food quantities in the test data is  $D^* \neq D$  with either  $D^* = D + 1$  or  $D^* = D - 1$  were considered. In the second additional setting, latent intakes in the test data were generated from a diffuse uniform distribution on  $(0, 350)$ , to assess robustness to deviations from the mixture of Gaussians assumption in (3). The third additional setting explored the impact of differently-sized mixture components in the test data. Overall, results (Table 1 in the Supplementary Material) are similar to those obtained under simulation studies I and II (see Table 1), with the exception of larger median absolute errors under variability scenario 3 in the additional settings. However the median error values (and corresponding 95% CIs) are smaller than those of PLS and BLR.

### 5.3 Simulation study III: intake quantification under model misspecification

The multiMarker model assumes a linear relationship between latent intake and observed biomarkers. To assess performance of multiMarker in the presence of model misspecification, we simulated data according to the settings of simulation study I, but using a non-linear relationship between intake and biomarkers:

$$y_{ip} = \alpha_p + \beta_p z_i^2 + \epsilon_{ip}, \quad \forall \quad i = 1, \dots, n, \quad p = 1, \dots, P \quad (9)$$

Here the scaling coefficients  $\beta_p$  have been re-scaled to 0.1% of the values used in simulation study I, to obtain similar simulated biomarker data ranges.

Table 1 (Simulation Study III columns) reports the absolute errors between estimated and inferred latent intakes and the true values. Here, error values are low and similar to those computed under simulation study I. Indeed, median error values are always lower than 27g in the test data, indicating that true intakes can still be recovered quite well, even when the underlying model is misspecified. This is due to the role of the mixture distribution on the latent intake, which anchors the intakes around the intervention study’s food quantities. However, model parameters are hardly ever recovered and often overestimated, indicating that the model misspecification is absorbed by their estimates. The price of biased parameters estimates is a relatively low one to pay in this context, as our goal is to obtain reliable inference on intake. Further, the range of errors’ confidence intervals is larger in simulation study III than in simulation study I, but comparable to that of simulation study

II. Finally, comparison with the BLR and PLS results yields similar conclusions to those presented in simulation studies I and II.

## 6 Application of multiMarker to biomarkers associated with apple intake

The motivation for multiMarker is estimation of the relationship between four biomarkers and apple intake from an intervention study, with a view to subsequently inferring intake when only the four biomarkers are available. MultiMarker was employed to estimate the relationship between the biomarkers and apple intake using 30000 MCMC iterations, discarding the first 6000, on a Intel(R) Core(TM) i7-8565U@1.80GHz computer. Model hyperparameters were fixed as outlined in the Supplementary Material (Section 1.1).

To assess the ability of multiMarker to infer apple intake, we use leave-one-out cross validation. Thus,  $n = 86$  multiMarker models were fitted, each with a different set of  $(n - 1)$  observations. In each case, the intake for the left out observation was then inferred using only its biomarker data. The MCMC algorithms were run several times and demonstrated no substantial difference in the parameter estimates and inferred latent intakes. Details of expected sample size values are available in the Supplementary Material (Section 3.3). The computation time was approximately 3.51 minutes per estimated model, and 2.67 seconds per data set when inferring intake.

Results from the leave-one-out cross validation procedure are shown in Figure 4. Differences between the known consumed apple quantities and median inferred intake values are reported, together with corresponding 95% credible intervals. Overall, there is good agreement between consumed apple quantities and inferred intakes. Indeed, median inferred values are concentrated around the true quantities and generally the latent intake posterior predictive distributions include the true apple quantity values. Only a few observations have median inferred intakes far from their true values. In general, such “miss-inferred” intakes have larger associated uncertainty, presented through their wider 95% credible intervals. The median absolute difference between inferred intake and true consumed food quantity is 15.29 grams for the 50 grams apple quantity group, 35.22 grams for the 100 grams quantity and 9.19 grams for the 300 grams quantity. These errors correspond to, respectively, 31%, 35% and 3% of the true apple quantity.

To further visualize the results, Figure 5 presents inferred posterior predictive intake distributions, for 6 observations, two for each apple quantity. Plots on the left side of Figure 5 correspond to accurately inferred intakes, while those on the right are poorly inferred. When intakes are correctly inferred, the range of the posterior predictive distribution does not incorporate apple quantities far from the truth (e.g. observations 13, 32 and 82). For some of these observations, when the corresponding true intake was either 50 or 100 grams, the inferred distribution is bimodal (e.g. observations 13, 32). However, for all of them the associated 95% credible interval is always relatively narrow and the mode closer to the true intake value is consistently higher than the other one. Among “miss-inferred” intakes, those whose true quantity was 300 grams tend to have notably underestimated inferred intakes (e.g. observation 73). This is due to the ordinal approach in the modelling of the observation-specific weights, which is more robust to overestimation than to underestimation. Similar posterior predictive distributions to those presented in Figure 5 were obtained for the remaining observations.

For comparative purposes, BLR and PLS regression were also used to infer apple intake via leave-one-out

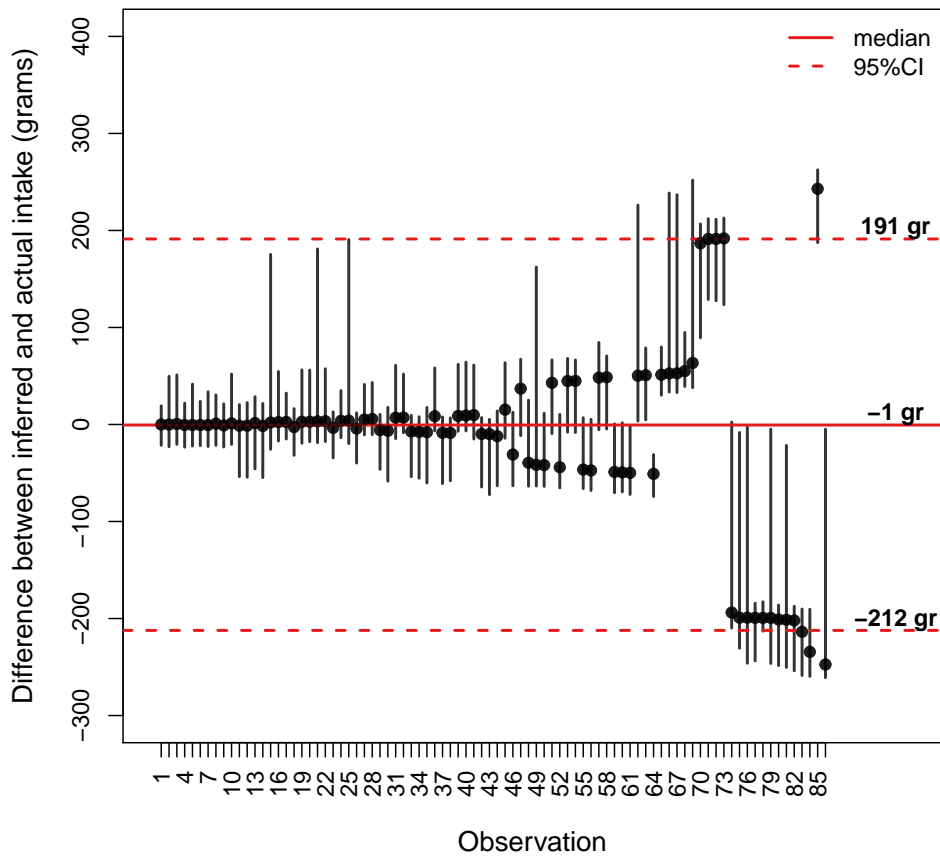


Figure 4: Performance of the multiMarker model when inferring apple intake (in grams) from four biomarkers. Each dot represents the difference between posterior median predicted intake and known consumed apple quantity, with vertical lines illustrating the 95% credible interval. Horizontal lines represent overall median and 95% credible intervals for the differences, with the corresponding numerical values reported in bold.

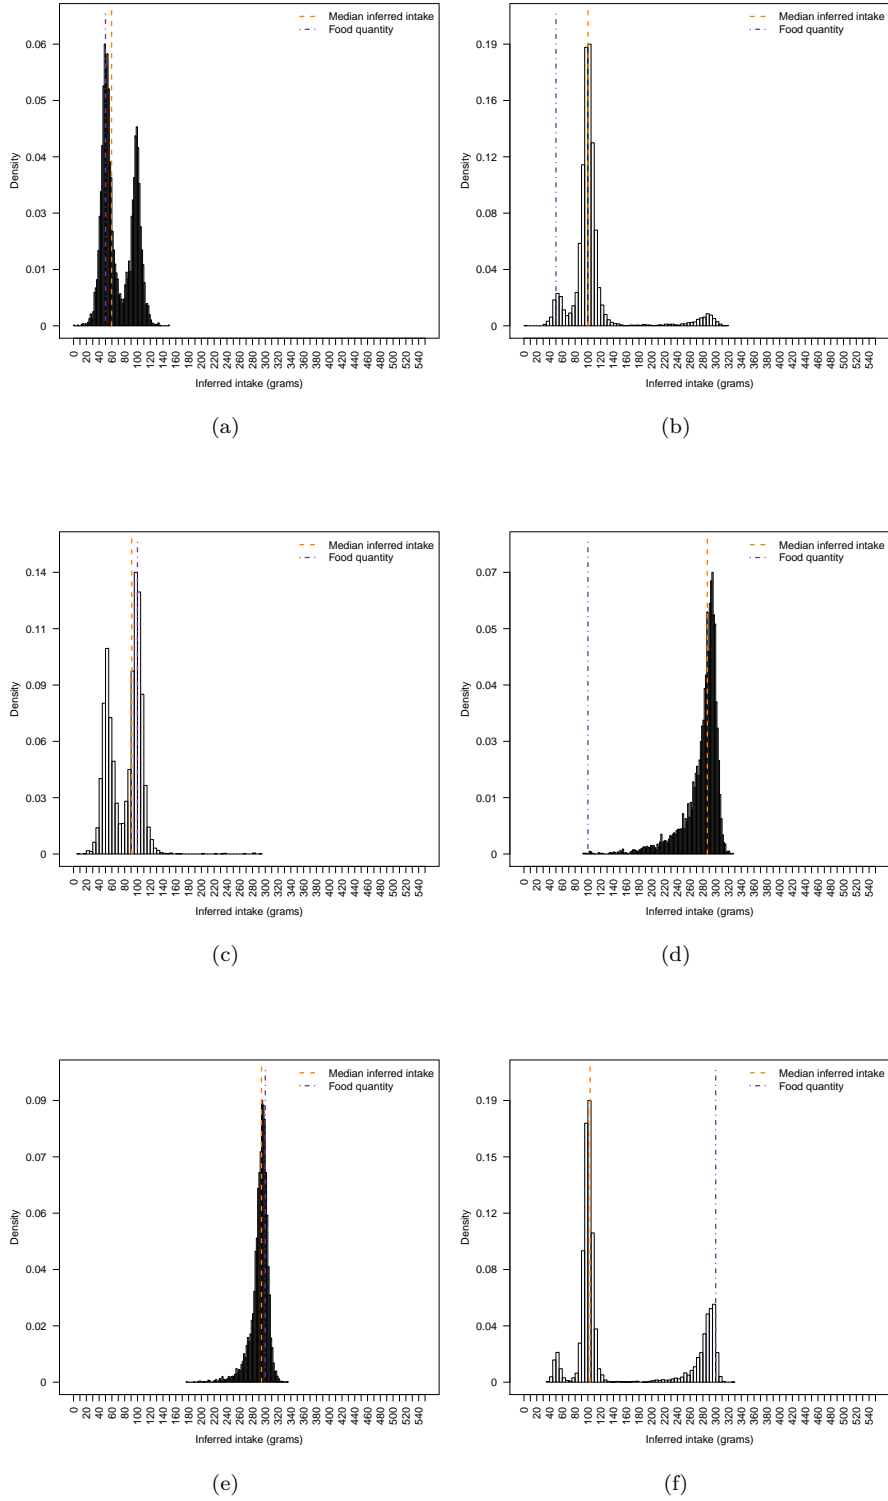


Figure 5: Posterior predictive distributions of apple intake (in grams) for 6 observations. Orange and purple lines denote median inferred intake and known apple quantity, respectively.

cross validation. Similar to the observed performance in the simulation studies, both methods proved unable to precisely infer intake, with most predictions around the average of the three apple quantities ( $\approx 150$  grams). Furthermore, 95% credible intervals obtained using BLR are more than twice the size of those obtained under multiMarker, indicating low reliability of the inferred values. Further, the uncertainty quantification under PLS regression was not meaningful, as the constructed 95% confidence intervals via cross validation yielded very narrow ranges. Further details on these results are deferred to the Supplementary Material (Section 3.2).

To assess whether treating the  $n = 86$  observations as independent was appropriate, given that they resulted from a set of 32 intervention study participants, we estimated the multiMarker model using slightly modified versions of the dataset. Specifically, we constructed modified versions of the original dataset such that the same participant appeared only once i.e. each participant is present in only one apple quantity, which was selected at random. This procedure was repeated 100 times, and each time the multiMarker model was fitted to the modified dataset. Table 2 presents the posterior median parameter estimates obtained from fitting the model to the original data and to the modified datasets. For the  $\alpha_p, \beta_p, \sigma_p$  parameters, dimensions 1 to 4 correspond to the four biomarkers, while for the  $\theta_d$  parameter the dimensions correspond to the three apple quantities. The median estimates for the first three parameters ( $\alpha_p, \beta_p, \sigma_p$ ) are similar under the original and modified data although the similarity is weaker for the variances  $\sigma_p$ . Also, posterior medians of these parameters from the original data are always contained in the corresponding 95% credible intervals inferred from the modified datasets. These results suggest the inference is robust to our assumption of independence between the  $n = 86$  observations. As for the  $\theta_d$  parameters, the estimated component variances are much larger in the original data than of the modified data. This finding is in line with our expectations, as the  $\theta_d$  parameters are a function of the sample size (see Section 4.5). Indeed, as discussed in simulation study II, component variances are nuisance parameters, as their role is instrumental to adapting the inferred latent intake distribution to the data analysed, but their value is of no interest.

Table 2: Posterior median (95% CI width) parameter values inferred from the original apple data, and average posterior median (average 95% CI width) parameter values inferred from the modified apple datasets.

Data	Dimension	Parameter			
		$\alpha_p$	$\beta_p$	$\sigma_p$	$\theta_d$
Original	1	0.206 (0.267)	0.003 (0.002)	0.353 (0.117)	5.546 (6.514)
	2	0.489 (0.215)	0.005 (0.002)	0.271 (0.104)	7.546 (14.038)
	3	0.612 (0.514)	0.007 (0.004)	0.677 (0.234)	98.989 (71.595)
	4	0.614 (0.323)	0.008 (0.004)	0.382 (0.174)	-
Modified	1	0.214 (0.433)	0.003 (0.003)	0.597 (0.157)	1.806 (1.176)
	2	0.515 (0.416)	0.005 (0.004)	0.600 (0.264)	2.032 (1.773)
	3	0.517 (0.755)	0.009 (0.007)	0.746 (0.258)	8.407 (10.358)
	4	0.662 (0.635)	0.008 (0.006)	0.676 (0.226)	-

## 7 Discussion

Motivated by the need to estimate the relationship between metabolomic biomarkers and food intake to allow inference on intake from biomarker data alone, we have introduced the general and flexible multiMarker model. MultiMarker builds upon two classical regression models, multiple linear regression and ordinal regression, combining them in the wider frameworks of factor analysis and mixture of experts models. The multiMarker model facilitates estimation of the relationship between multiple biomarkers and food intake from intervention studies. Further, it allows one to subsequently infer intake in a quantitative manner when only biomarker data are available. The multiMarker model advances on current approaches, which focus on inferring categorical levels of intake, by providing more detailed inference. Moreover, as multiMarker is embedded in a Bayesian framework, uncertainty quantification is readily available, which is an important requirement in the motivating application domain.

The multiMarker model and inferential algorithm were assessed in an extensive simulation study, where a large variety of biomarker data have been generated to check model performances under, among other things, different levels of biomarker or intake variability. Furthermore, performances under model misspecification have been verified. In all of the cases the framework performed well, as intake values and their range were recovered with low error, even in a model misspecification context.

The multiMarker model was successfully employed to estimate the relationship between biomarkers and apple intake and to subsequently infer apple intake from biomarker data alone, leveraging data collected under the motivating A-DIET research programme. Leave-one-out cross validation results showed that generally inferred intakes were concentrated around the true apple quantities. Comparison with existing regression models showed that multiMarker was able not only to provide useful uncertainty quantification, but also much more reliable inference on apple intake.

The proposed framework allows quantification of any unobserved quantity of interest for which prior information on its scale is observed (here, food quantities) and for which proxies can be measured (here, biomarkers). Furthermore, the observed variables' scales are not relevant to recover that of the latent variable. Such a feature is appealing for researchers, as biomarkers for the same intake could correspond to quite different measurement ranges, due to their sources (blood, urine, etc), the instrument used for collection and so on.

Although multiMarker has been motivated by the specifications of the A-DIET study, the model is general and designed to be accessible to non-statisticians. In multiMarker, the extra layer of complexity introduced by the latent variable and associated model parameters bestows the advantage of inferring refined, quantitative intakes and allows the model to adapt to different scenarios in the data. Modelling the weights of the mixture distribution using a mixture of experts framework guides quantitative inference, informing the latter on which part of the intake range is to be considered when inferring intake. Unlike previous approaches, we explicitly model the ordinal nature of the food quantities. However, multiMarker could easily be modified to suit different application contexts, and different specifications could be employed to represent different data settings. For example, should the ordinal regression feature be inappropriate for a particular setting, multinomial regression can easily be incorporated as a substitute. Further, some of the application-driven restrictions imposed here, for example bounded prior distributions for some parameters, could be easily lifted to adapt the proposed model to

different contexts without resulting in any additional complications from either a model definition or estimation point of view.

As a key requirement of the motivating application was to quantify uncertainty, particularly when inferring intakes from biomarker data alone, inference was carried out via a Metropolis within Gibbs MCMC algorithm. Alternative inferential procedures could easily be considered and may be particularly fruitful in alternative specifications of multiMarker that use different distributional assumptions, leading to more complex and computationally expensive posterior distributions. The MCMC approach may also have limitations should the number of biomarkers or observations increase. A natural alternative in such a setting is a variational inference approach (Blei et al. 2017), however it can also be sensitive to initialization and only guarantees convergence to a local optimum. Further, uncertainty quantification is not automatic under a variational framework, although Chen et al. (2018) demonstrate that a bootstrap procedure is still valid, albeit for variational parameters which may or may not coincide with the true parameters of interest.

Although designed for a particular problem, the quantification of apple intake from a panel of urinary biomarkers, the multiMarker approach has general applicability outside of nutrition. Indeed, the model could have applicability in any scenario where multiple outcomes are associated with an unobserved variable of interest, such as in toxicology or social science studies.

Possible extensions to the multiMarker model are many and varied, including explicitly modelling the repeated biomarker measurements in the same group of participants. Individual heterogeneity could be accounted for by introducing random effects in the model (see for example Muthén (1989)). Further, the introduction of subject-specific covariates in the latent intake mixture weights, to allow a more flexible model fit, would be feasible. Thus there is much potential in the use of latent variable models to infer food intake.

**Acknowledgments.** Supported by a research grant from the European Research Council (ERC)(647783).

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number SFI/12/RC/2289\_P2. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission'

We would like to acknowledge the help of the following people with respect to the Apple Biomarker data: Aoife E. McNamara, Cassandra Collins, Pedapati S. C. Sri Harsha and Diana Gonzalez-Pena.

## References

- Agresti, A. (1999), 'Modelling ordered categorical data: recent advances and future challenges', *Statistics in medicine* **18**, 2191–2207.
- Baldrick, F., Woodside, J., Elborn, J., Young, I. & McKinley, M. (2011), 'Biomarkers of fruit and vegetable intake in human intervention studies: a systematic review', *Critical Reviews in Food Science and Nutrition* **51**, 795–815.
- Bhattacharya, A. & Dunson, D. (2011), 'Sparse Bayesian infinite factor models', *Biometrika* **98**(2), 291–306.
- Bingham, S. (2002), 'Biomarkers in nutritional epidemiology', *Public Health Nutrition* **5**(6(A)), 821–827.

- Blei, D. M., Kucukelbir, A. & McAuliffe, J. (2017), ‘Variational inference: a review for statisticians’, *Journal of the American Statistical Association* **112**(518), 859–877.
- Cagnone, S. & Viroli, C. (2012), ‘A factor mixture analysis model for multivariate binary data’, *Statistical Modelling* **12**(3), 257–277.
- Chen, Y., Wang, Y. & Erosheva, E. (2018), ‘On the use of bootstrap with variational inference: Theory, interpretation, and a two-sample test example’, *Annals of Applied Statistics* **12**(2), 846–876.
- D’Angelo, S., Brennan, L. & Gormley, I. (2020), *multiMarker: Latent Variable Model to Infer Food Intake from Multiple Biomarkers*. R package version 1.0.1.  
**URL:** <https://CRAN.R-project.org/package=multiMarker>
- Dragsted, L. et al. (2018), ‘Validation of biomarkers of food intake—critical assessment of candidate biomarkers’, *Genes and Nutrition* **13**(14).
- Frühwirth-Schnatter, S. & Lopes, H. (2018), ‘Sparse Bayesian factor analysis when the number of factors is unknown’, *arXiv:1804.04231*.
- Galimberti, G., Montanari, A. & Viroli, C. (2009), ‘Penalized factor mixture analysis for variable selection in clustered data’, *Computational Statistics and Data Analysis* **53**(12), 4301–4310.
- Gao, Q. et al. (2017), ‘A scheme for a flexible classification of dietary and health biomarkers’, *Genes and Nutrition* **12**(34).
- Garcia-Aloy, M., Rabassa, M., Casas-Agustench, P., Hidalgo-Liberona, N., Llorach, R. & Andres-Lacueva, C. (2017), ‘Novel strategies for improving dietary exposure assessment: Multiple-data fusion is a more accurate measure than the traditional single-biomarker approach’, *Trends in Food Science and Technology* **69**(B), 220–229.
- Gormley, I. & Frühwirth-Schnatter, S. (2019), *Handbook of mixture analysis. Chapter 12: mixture of experts models*, Chapman and Hall/CRC.
- Gürdeniz, G. et al. (2016), ‘Detecting beer intake by unique metabolite patterns’, *Journal of Proteome Research* **15**(12), 4544–4556.
- Jacobs, R., Jordan, M., Nowlan, S. & Hinton, G. (1991), ‘Adaptive mixtures of local experts’, *Neural Computation* **3**, 79–87.
- Kipnis, V. et al. (2002), ‘Bias in dietary-report instruments and its implications for nutritional epidemiology’, *Public Health Nutrition* **5**(6(A)), 915–923.
- Knott, M. & Bartholomew, D. (1999), *Latent variable models and factor analysis. Number 7.*, Edward Arnold, London, second edition.
- Lin, T., McLachlan, G. & Lee, S. (2016), ‘Extending mixtures of factor models using the restricted multivariate skew-normal distribution’, *Journal of Multivariate Analysis* **143**, 398–413.

- Lloyd, A., Willis, N., Wilson, T., Zubair, H., Chambers, E., Garcia-Perez, I., Xie, L., Taillart, K., Beckmann, M., Mathers, J. & Draper, J. (2019), ‘Addressing the pitfalls when designing intervention studies to discover and validate biomarkers of habitual dietary intake’, *Metabolomics* **15**.
- Lopes, H. & West, M. (2004), ‘Bayesian model assessment in factor analysis’, *Statistica Sinica* pp. 41—67.
- McLachlan, G., Bean, R. & Jones, L. (2007), ‘Extension of the mixture of factor analyzers model to incorporate the multivariate t-distribution’, *Computational Statistics and Data Analysis* **51**(11), 5327–5338.
- McNamara, A., Collins, C., Sri Harsha, P., González-Peña, D., Gibbons, H., McNulty, B., Nugent, A., Walton, J., Flynn, A. & Brennan, L. (2020), ‘Metabolomic-based approach to identify biomarkers of apple intake’, *Molecular Nutrition and Food Research* .
- Montanari, A. & Viroli, C. (2010), ‘Heteroscedastic factor mixture analysis’, *Statistical Modelling* **10**(4), 441—460.
- Morgan, B. & Smith, D. (1992), ‘A note on wadley’s problem with overdispersion’, *Applied Statistics* **41**(2), 287–497.
- Murphy, K., Viroli, C. & Gormley, I. (2020), ‘Infinite mixtures of infinite factor analysers’, *Bayesian Analysis* .
- Murray, P., Browne, R. & McNicholas, P. (2014), ‘Mixtures of skew-t factor analyzers’, *Computational Statistics and Data Analysis* **77**, 326–335.
- Muthén, B. O. (1989), ‘Latent variable modeling in heterogeneous populations’, *Psychometrika* **54**(4), 557—585.
- Pison, G., Rousseeuw, P., Filzmoser, P. & Croux, C. (2003), ‘Robust factor analysis’, *Journal of Multivariate Analysis* **84**, 145—172.
- R Core Team (2020), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.  
**URL:** <https://www.R-project.org/>
- Rothwell, J. et al. (2014), ‘New biomarkers of coffee consumption identified by the non-targeted metabolomic profiling of cohort study subjects’, *PLOS One* .
- Ročková, V. & George, E. (2016), ‘Fast Bayesian factor analysis via automatic rotations to sparsity’, *Journal of the American Statistical Association* **111**, 1608—1622.
- Siddique, J., Daniels, M., Carroll, R., Raghunathan, T., Stuart, E. & Freedman, L. (2019), ‘Measurement error correction and sensitivity analysis in longitudinal dietary intervention studies using an external validation study’, *Biometrics* **75**(3), 927–937.
- Subar, A., Kipnis, V., Troiano, R., Midthune, D., Schoeller, D., Bingham, S., Sharbaugh, C., Trabulsi, J., Runswick, S., Ballard-Barbash, R., Sunshine, J. & Schatzkin, A. (2003), ‘Using intake biomarkers to evaluate the extent of dietary misreporting in a large sample of adults: The open study’, *American Journal of Epidemiology* **158**(11), 1–13.

Vázquez-Manjarrez, N. et al. (2019), 'Discovery and validation of banana intake biomarkers using untargeted metabolomics in human intervention and cross-sectional studies', *The Journal of Nutrition* **149**(10), 1701–1713.

## A Posterior distribution and MCMC algorithm

Considering the multiMarker model's likelihood, prior and hyperprior distributions, the posterior distribution is:

$$\begin{aligned}
P(\Omega \mid \mathbf{Y}, \mathbf{X}) = & L(\mathbf{Y} \mid \alpha, \beta, \mathbf{z}, \Sigma) p(\alpha \mid \mu_\alpha, \sigma_\alpha^2) p(\mu_\alpha \mid m_\alpha, \tau_\alpha, \sigma_\alpha^2) p(\beta \mid \mu_\beta, \sigma_\beta^2) p(\mu_\beta \mid m_\beta, \tau_\beta, \sigma_\beta^2) \\
& p(\sigma_\beta^2 \mid \nu_{\beta 1}, \nu_{\beta 2}) p(\Sigma \mid \nu_{p1}, \nu_{p2}) p(\mathbf{z} \mid \pi, \mathbf{X}, \Theta) p(\Theta \mid \nu_{z1}, \nu_{z2}) \\
& p(\pi \mid \gamma, \eta, \mathbf{Y}, \mathbf{c}, \mathbf{X}) p(\gamma \mid m_\gamma, \kappa) p(\eta \mid m_\eta, \kappa)
\end{aligned} \tag{10}$$

where for brevity  $\Omega$  denotes the set of model parameters.

### A.1 Hyperparameter settings

Given an observed dataset, hyperparameter values are fixed automatically according to the following procedures. The overall means  $m_\alpha$  and  $m_\beta$  are fixed, respectively, as the estimated intercept and slope coefficient of the multiple linear regression defined using biomarkers as response variable and food quantity values as predictor. Variances' hyperparameters are fixed differently according to the parameter to which they refer:  $(\nu_{\beta 1}, \nu_{\beta 2}) = (2, 3)$  and  $(\nu_{p1}, \nu_{p2}) = (1, 3)$ . The  $(\nu_{z1}, \nu_{z2})$  latent intake variances' hyperparameters are  $D$ -dimensional vectors, with  $\nu_{z1d} = \frac{D-d+1}{2}$ , and  $\nu_{z2d} = n$ ,  $\forall d$ . Regarding the  $\alpha$  and  $\beta$  vector of parameters, their values are initialized solving the following system of equations:

$$\sum_{(x_i=X_d)} (y_{ip}) \approx \alpha_p + \beta_p x_d \quad \text{for } p = 1, \dots, P \quad \text{and } d = 1, \dots, D.$$

Biomarkers' error variances are initialized exploiting the definition of estimated error variances under the factor analytic model, adjusted for the extra variability brought in by the latent intakes prior distribution:  $\Sigma = \widehat{V(\mathbf{Y})} - \frac{1}{D} \beta \beta^T$ . Last, mixture components' variance parameters are initialized with the following values:

$$\sigma_d^2 = \frac{1}{P} \sum_{p=1}^P \frac{\widehat{V(Y_{pd})} - \sigma_\alpha^2 - \sigma_p^2}{\sigma_\beta^2}$$

where  $\widehat{V(Y_{pd})} = \widehat{var}(\mathbf{1}(x_i = X_d) y_{ip})$ . Last, denoting with  $\{x_1, \dots, x_i, \dots, x_n\}$  the food quantities known to be consumed by the  $n$  observations in an intervention study, latent intake values are initialized as  $z_i \sim \mathcal{N}_{(0, \infty)}(x_i, \theta_d)$ , where  $\theta_d^2 = 5^2$ ,  $\forall d$ , is the starting value for the components' variance parameters.

### A.2 Components' weights

The prior distributions for the  $\gamma = \{\gamma_0, \gamma_1, \dots, \gamma_D\}$  weights' parameters are defined to represent the constrained characteristic of ordinal data, that is:  $X_1 < \dots < X_d < \dots < X_D$ . These constraints correspond to the following in terms of model parameters:  $\gamma_0 < \dots < \gamma_d < \dots < \gamma_D$ . A natural choice for the corresponding prior distributions is the following:

$$\gamma_d \sim \mathcal{N}_{(m_{\gamma_{d-1}}, m_{\gamma_{d+1}})}(m_{\gamma_d}, \kappa), \quad \text{for } d = 1, \dots, D-1$$

Further, biomarkers' intercepts  $\eta = \{\eta_1, \dots, \eta_P\}$  have been given the following prior distribution:  $\eta_p \sim \mathcal{N}(m_{\eta_p}, \kappa)$ ,  $p = 1, \dots, P$ . Hyperparameters  $\{m_{\eta_p}\}_{p=1}^P$  and  $\{m_{\gamma_d}\}_{d=1}^{D-1}$  are set via their corresponding estimates from the ordinal regression of the known observation-specific component allocation labels on biomarker measurements from

the intervention study. Such estimates are obtained with the *ordinalNet* package, available on CRAN (<https://CRAN.R-project.org/package=ordinalNet>). Both  $\eta$  and  $\gamma$  parameters are initialized with their hyperparameter values ( $\{m_{\eta_p}\}_{p=1}^P, \{m_{\gamma_d}\}_{d=1}^{D-1}$ ) and are updated with a random walk Metropolis Hastings step inside the MCMC algorithm. Variance hyperparameter  $\kappa$  is fixed to some value  $\kappa \leq 2$ . Small variations around this threshold have been tested and did not produce any substantial difference in the  $\eta$  and  $\gamma$  parameters estimates.

### A.3 Latent intakes posterior predictive distribution

In Section 4.2 we have introduced a sampling distribution for the latent intakes, to be used when deriving the latent intake posterior predictive distribution. This distribution was derived as the product of two terms, the first being the log-likelihood of the model, expressed as a function of the latent intakes:

$$\begin{aligned}
\ell(y_j^* | \alpha, \beta, z_j^*, \Sigma) &= p(z_j^* | \alpha, \beta, y_j^*, \Sigma) \propto \sum_{p=1}^P p(z_j^* | \alpha_p, \beta_p, y_j^*, \sigma_p^2) \\
&= \sum_{p=1}^P \left( -\frac{1}{2} \log(2\pi\sigma_p^2) - \frac{1}{2\sigma_p^2} \left( z_j^* \beta_p - (y_{jp}^* - \alpha_p) \right)^2 \right) \\
&= \sum_{p=1}^P \left( -\frac{1}{2} \log(2\pi\sigma_p^2) - \frac{\beta_p^2}{2\sigma_p^2} \left( z_j^* - \frac{y_{jp}^* - \alpha_p}{\beta_p} \right)^2 \right) \\
&\propto \sum_{p=1}^P \left( -\frac{\beta_p^2}{2\sigma_p^2} \left( z_j^* - \frac{y_{jp}^* - \alpha_p}{\beta_p} \right)^2 \right) \\
&\propto -\frac{1}{2} z_j^{*2} \left( \sum_{p=1}^P \frac{\beta_p^2}{\sigma_p^2} \right) + z_j^* \left( \sum_{p=1}^P \frac{\beta_p (y_{jp}^* - \alpha_p)}{\sigma_p^2} \right) = -\frac{1}{2\sigma_z^2} z_j^{*2} + \frac{\mu_z}{\sigma_z^2} z_j^*,
\end{aligned}$$

The second term is the mixture distribution presented in Section 3.2. The product of these two terms can be expressed as a mixture distribution:

$$\begin{aligned}
p(z_j^* | \Omega) &= \mathcal{N}_{[0, \infty]}(\mu_z, \sigma_z^2) \sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty]}(X_d, \theta_d^2) = \sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty]}(\mu_z, \sigma_z^2) \mathcal{N}_{[0, \infty]}(X_d, \theta_d^2) \\
&= \sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty]} \left( \frac{\mu_z \theta_d^2 + X_d \sigma_z^2}{\theta_d^2 + \sigma_z^2}, \left( \frac{1}{\theta_d^2} + \frac{1}{\sigma_z^2} \right)^{-1} \right) = \sum_{d=1}^D \pi_{jd} \mathcal{N}_{[0, \infty]}(\mu_{zd}, \sigma_{zd}^2)
\end{aligned} \tag{11}$$

where  $\Omega = \{\mu_{z1}, \dots, \mu_{zD}, \sigma_{zq}^2, \dots, \sigma_{zD}^2\}$ .

## B Additional simulation study details

In the simulation studies described in the paper, biomarkers' intercepts and scale coefficients ( $\alpha, \beta$ ) are sampled from their prior distributions with hyperparameters:

$$\{\mu_\alpha, \mu_\beta, \sigma_{\alpha^2}, \sigma_\beta^2\} = \{(1, 0.01, 1, 0.01), (20, 0.1, 4, 0.1), (100, 1, 14, 1)\}$$

These hyperparameter specifications correspond to small, medium and large range biomarker values, and are used to represent different types of biomarker measures (as for example measurements coming from different instruments or measurements non normalised by osmolality). Regarding the error standard deviations terms  $\sigma_p^2$ , these values are sampled from inverse gamma distributions with expected values dependent on the biomarkers'

range considered:  $(1^2, 3^2, 15^2)$  for small variances and  $(3^2, 20^2, 100^2)$  for large ones, respectively in the small, mixed and large biomarkers’ range frameworks. Values for  $X_d$  are sampled from  $D$  zero-truncated Gaussian distributions with means  $\mu_{X_d}$  ranging in between 30 and 320, with  $\mu_{X_d} < \mu_{X_{d+1}}$ . Three different settings are explored to represent “food quantities with stable increments”, “food quantities with increasing increments”, and “food quantities with decreasing increments”. In the first setting (“stable increments”), means  $\mu_{X_d}$  are equispaced, that is  $d(\mu_{X_{d-1}}, \mu_{X_d}) = d(\mu_{X_d}, \mu_{X_{d+1}})$ . Instead, in the second and third settings we have that  $d(\mu_{X_{d-1}}, \mu_{X_d}) < d(\mu_{X_d}, \mu_{X_{d+1}})$  and  $d(\mu_{X_{d-1}}, \mu_{X_d}) > d(\mu_{X_d}, \mu_{X_{d+1}})$ , respectively. In simulation studies I and III, values for  $\theta_d^2$  are sampled from inverse gamma distributions with either expected value of  $8^2$  (small variances) or of  $16^2$  (large variances). In simulation study II, values for  $\theta_d^2$  are sampled from Inverse Gamma distributions with

- Small variances setting: expected value of  $6^2$  in the intervention study data and of  $12^2$  in the biomarker only data;
- Large variances setting: expected value of  $12^2$  in the intervention study data and of  $24^2$  in the biomarker only data.

## B.1 Simulation study I: comparison between $D = 3$ and $D = 6$

Table 3 (first column block) reports the absolute errors computed between estimated and inferred latent intakes and the truth when  $D = 6$ . As observed in Simulation Study I in the paper, the errors under BLR are large and uncertain when compared to the multiMarker and PLS approaches which perform similarly when the number of food quantities is large. The multiMarker approach performs similarly when  $D = 3$  or  $D = 6$ , suggesting that the benefit of employing of a larger number of food quantities is minimal.

## B.2 Simulation study II: discrepancies between training and test data generation

### B.2.1 Varying $\mathbf{X}$ values

This simulation study considers the case in which different food quantity values, sampled in comparable ranges, are used to generate train and test data. In particular, latent intakes used to generate train data have been simulated from a mixture distribution with  $D = 3$  components, food quantity vector  $\mathbf{X}$ , and a vector  $\Theta$  of component variances. Instead, to generate test data, we either sample latent intake values from:

- a new vector of food quantities  $\mathbf{X}^*$ , of length  $D^* = D$ , whose range is similar to that of  $\mathbf{X}$ ,
- a new vector of food quantities  $\mathbf{X}^*$ , of length  $D^* = D - 1$ , whose range is similar to that of  $\mathbf{X}$ ,
- a new vector of food quantities  $\mathbf{X}^*$ , of length  $D^* = D + 1$ , whose range is similar to that of  $\mathbf{X}$ .

As in Section B, food quantity values  $\mathbf{X}^*$  are sampled from  $D^*$  zero-truncated Gaussian distributions with means  $\mu_{X_d^*}$  ranging in between 30 and 320, with  $\mu_{X_d^*} < \mu_{X_{d+1}^*}$ , and such that  $X_d \neq X_d^*, \forall d, d^*$ . Further,  $\mathbf{X}^*$  have been simulated from all three increments settings: “food quantities with stable increments”, “food quantities with increasing increments”, and “food quantities with decreasing increments”.

For comparability purposes, all other parameters and hyperparameters settings are set as in Simulation study I (Section 5.1 of the paper). Table 3 (second column block) reports the absolute errors computed between estimated

and inferred latent intakes and simulated values under “varying  $\mathbf{X}$ ” scenarios. In general, our approach is able to recover quite well “true” latent intakes under this type of mis-specification of the latent intake distribution. Indeed, absolute error values computed between true and estimated intakes under this setting are comparable with those of Simulation Study I (Section 5.1), that is to the case where the latent intake distribution is correctly specified.

### B.2.2 Diffuse intake distribution in the test data

This simulation study considers the case in which intakes in the test data are generated from a single diffuse distribution, to assess the impact of deviations from the mixture of Gaussians assumption (see Section 3.2 of the paper) on performance. Here, latent intakes used to generate train data are simulated from a mixture distribution with  $D = 3$  components (see Section B), while latent intakes in the test data are sampled from a Uniform distribution in  $(0, 350)$ . For all other parameters and hyperparameter settings, see Section 5.1 of the paper. Table 3 (third column block “Uniform intakes”) reports the absolute errors between estimated and inferred latent intakes and simulated values. In general, absolute error values are in line with those obtained with other settings (see Table 1 of the paper), indicating that intakes sampled from a single diffuse distribution (here, a uniform distribution) can still be well recovered using a mixture of Gaussians construct.

### B.2.3 Unbalanced components in the test data

This simulation study considers the case in which observations in the test data have intake values which are concentrated in a specific sub-region of the intake range explored in the intervention study. This corresponds to assuming  $n_d^* \gg n_k^*$ , for  $k \in \{1, \dots, D\}, k \neq d$ , where  $n_d^*$  is the number of observations in the  $d^{\text{th}}$  component, and  $n_k^*$  the number of observations allocated in a component different from  $d$ . To assess multiMarker’s performance under this setting, intakes are simulated intakes in the train data from a “balanced” mixture distribution with  $D = 3$  components (see Section 5.1 of the paper and Section B). However, intakes in the test data have been generated as follows:

1. A component index  $d$  is sampled at random from  $\{1, \dots, D\} = \{1, 2, 3\}$ .
2. Given  $d$ ,  $n_d^*$  observations are simulated from the  $d^{\text{th}}$  component, where  $n_d^*$  is sampled at random in the range  $(0.7n^*, 0.8n^*)$ .
3.  $n_{-d}^* = \{n_k^*, n_l^*\}$  observations, where  $l, k$  are the indexes for the two remaining components, are simulated from the corresponding components. The  $n_{-d}^*$  values are sampled at random in  $(0, n^* - n_d^*)$ , and normalized so that  $(n_k^* + n_l^*) = (n^* - n_d^*)$ .

All other parameters and hyperparameters are specified as in Section 5.1 of the paper. Table 3 (fourth column block “Unbalanced components”) reports the absolute errors computed between estimated and inferred latent intakes and simulated values. Results are in agreement with those obtained in all previous scenarios (see Table 1 of the paper) and the other column blocks in Table 3 for a comparison. This suggests that the ability to recover intakes in the test data is not impacted by how the intakes are spread across the intake range.

Table 3: Simulation studies’ results. Median (95% CI width) absolute error values (in grams) computed between true and estimated (E) or inferred (I) latent intakes. The values are reported for the cases  $D = 6$ , “Varying  $\mathbf{X}$ ”, “Uniform intakes” and “Unbalanced components”, and for the three biomarkers’ variability scenarios (S.1,S.2,S.3), introduced in Section 5 of the manuscript. Results for the multiMarker (MM) model are reported, as well as those from Bayesian linear regression (BLR) and PLS (PLS) regression.

Model		$D = 6$			Varying $\mathbf{X}$			Uniform intakes			Unbalanced components		
		S.1	S.2	S.3	S.1	S.2	S.3	S.1	S.2	S.3	S.1	S.2	S.3
MM	E	4(11)	5(16)	14(15)	4(9)	6(52)	12(15)	4(7)	5(8)	10(48)	4(8)	5(8)	9(38)
	I	4(15)	6(47)	33(59)	9(47)	10(51)	40(62)	10(21)	22(38)	37(74)	5(12)	19(23)	27(66)
BLR	E	73(207)	94(234)	69(138)	66(145)	81(60)	76(70)	93(263)	114(231)	72(109)	108(206)	142(228)	98(117)
	I	86(172)	126(282)	72(145)	68(150)	87(114)	81(296)	88(256)	110(245)	96(187)	96(187)	128(275)	96(169)
PLS	E	8(19)	21(46)	34(49)	19(60)	46(70)	14(15)	10(15)	23(30)	43(61)	10(25)	20(31)	38(65)
	I	8(26)	22(51)	38(58)	10(65)	19(67)	49(83)	11(20)	25(45)	47(74)	11(29)	22(44)	39(81)

### B.3 Simulation study: examples of simulated data

Figure 6 shows an example of simulated latent intakes in train and test data, under Simulation study II. In Figure 6 (a), an example of simulated latent intakes in simulation study II is reported, with intakes from the train data and test data depicted in black and grey, respectively. Simulated food quantity values are also reported, to show how in the test data increased  $\theta_d^2$  values,  $d = 1, \dots, D$ , lead to a flatter latent intake density. Indeed, latent intakes are no longer clearly concentrated around the food quantities, but are more widely spread across the intake range (roughly 0 to 310). Figure 6 (b) reports a comparison between simulated biomarker data under simulation studies I and III. Simulated biomarkers are in the medium range scenario, with food quantities presenting “stable increments” and under variability scenario S.1.

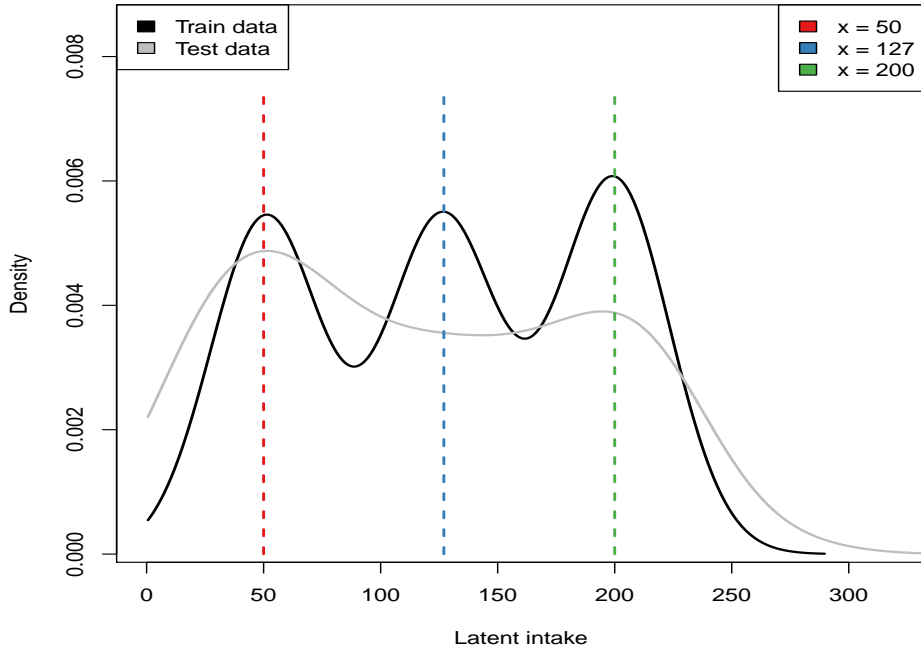
### B.4 Estimated model parameters

Figures 7, 8 and 9 show the median and 95% credible interval estimates, respectively for the  $\alpha_p$ ,  $\beta_p$  and  $\sigma_p$  parameters,  $\forall p$ , under Simulation study I. Medians and 95% credible intervals are computed using the subsets of simulated intervention study datasets (and corresponding parameter estimates) that shared the same true simulated value of the parameter of interest. True simulated parameter values are reported for a comparison, and results are reported separately for the different error variability scenarios and biomarker data ranges.

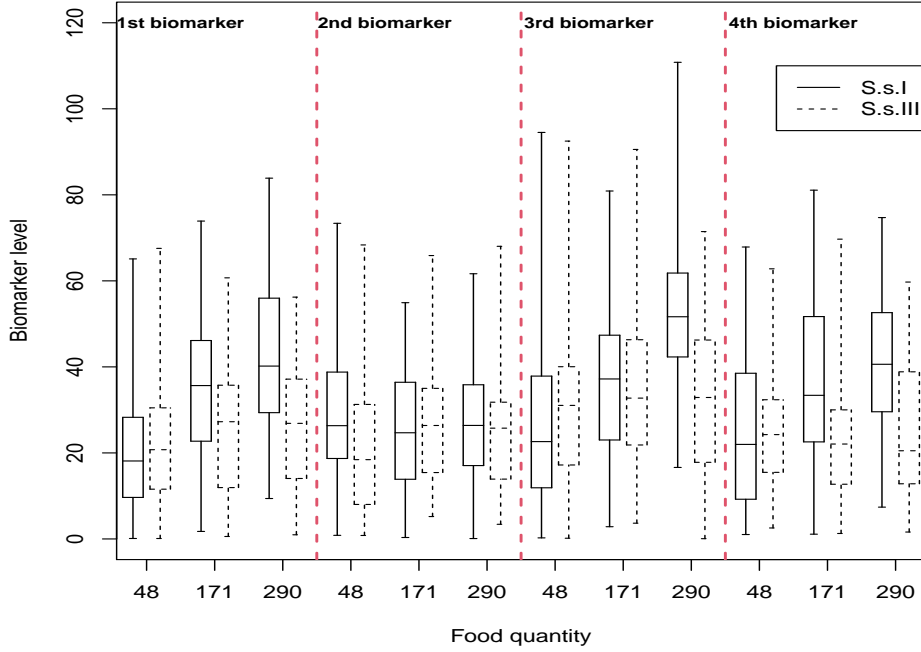
## C Apple intake data

### C.1 Data scaling

The original values for Epicatechin Sulfate,  $(4-3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-DHMPMB-SA)$  and Glucodistylin biomarkers caused computational instability (most values were larger than  $10^7$ ) and



(a) *Simulation study II. Examples of simulated latent intakes in train and test data.*



(b) *Simulation studies I and III. Examples of simulated biomarker data.*

Figure 6: *Simulation study. Example of simulated latent intakes (study II) (a) and biomarker data (studies I, III) (b).*

consequently were scaled. Given a biomarker  $p$ , the corresponding scaled values  $\tilde{y}_{ip}$  are computed as follows:

$$\tilde{y}_{ip} = \frac{y_{ip} - \bar{y}_p}{sd(y_p)} + 2 \left| \min_{i=1, \dots, n} \left( \frac{y_{ip} - \bar{y}_p}{sd(y_p)} \right) \right|$$

where  $y_{ip}$  are the original measurements,  $i = 1, \dots, n$ . The original measurements mean and standard deviation are denoted with  $\bar{y}_p$  and  $sd(y_p)$ , respectively. The transformation did not alter the correlations between the four

biomarkers.

## C.2 Comparison to BLR and PLS regression


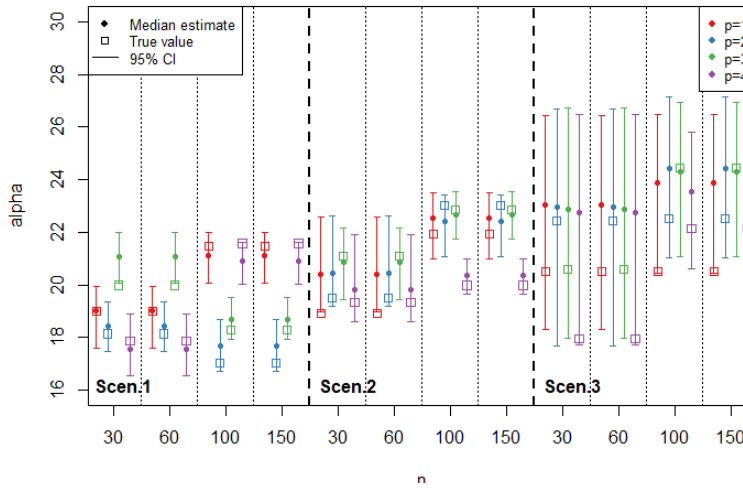

Figure 10 reports plots that are analogous to those in Figure 4 of the paper, obtained using either BLR or PLS regression, for a comparison.

## C.3 MCMC diagnostics

Table 4 reports a summary of the ESS (Expected Sample Size) values for the model parameters (30000 MCMC iterations considered). Figures 11 and 12 report, respectively, the estimated posterior distributions for the  $\alpha_p$  and  $\beta_p$  parameters,  $p = 1, \dots, P$ , in the apple intake data. In addition, Figures 13 and 14 report, respectively, the trace plots for the estimated  $\alpha_p$  and  $\beta_p$  parameters,  $p = 1, \dots, P$ , in the apple intake data.




Table 4: Apple data. Summary of the ESS (Expected Sample Size) values for the model parameters (30000 MCMC iterations considered).

	Parameter									
	$\mu_\alpha$	$\mu_\beta$	$\sigma_\beta^2$	$\alpha$	$\beta$	$\sigma_p^2$	$\theta_d^2$	$z$	$\gamma$	$\eta$
min	21673	17584	13399	1162	830	724	906	450	220	2014
median	24162	22262	21464	3308	1192	7213	2170	1159	870	2219
Max	27129	25890	24935	6607	2941	22843	15717	29174	1119	2312




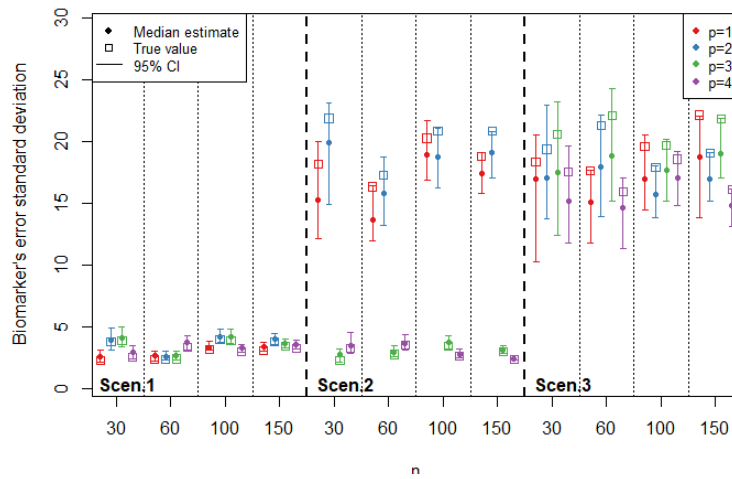

(c) Large biomarker data range.

Figure 7: Intercept parameters  $(\alpha_1, \dots, \alpha_P)$ . Dots, lines and squares represent, respectively, median estimates, 95% credible interval estimates and true values. Four different colors represent parameter referred to different biomarkers. Results are reported split by sample size and variability scenario.




(c) Large biomarker data range.

Figure 8: Scale coefficient parameters ( $\beta_1, \dots, \beta_P$ ). Dots, lines and squares represent, respectively, median estimates, 95% credible interval estimates and true values. Four different colors represent parameter referred to different biomarkers. Results are reported split by sample size and variability scenario.




(c) Large biomarker data range.

Figure 9: Biomarker's error standard deviations ( $\sigma_1, \dots, \sigma_P$ ). Dots, lines and squares represent, respectively, median estimates, 95% credible interval estimates and true values. Four different colors represent parameter referred to different biomarkers. Results are reported split by sample size and variability scenario.




(a)




(b)


Figure 10: Inferred intakes (in grams), obtained using leave one out cross validation and (a) fitting the PLS regression model, and (b) fitting the Bayesian linear regression model. These plots are analogous to that in Figure 4 of the paper.




(a) *Xylose.*



(b) *Epicatechin Sulfate.*




(c) *(4-3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-DHMPMB-SA).*




(d) *Glucodistylin.*


Figure 11: Apple data. Estimated posterior distributions for the  $\alpha_p$  parameters,  $p = 1, \dots, P$ .




(a) *Xylose.*



(b) *Epicatechin Sulfate.*




(c) *(4-3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-DHMPMB-SA).*




(d) *Glucodistylin.*

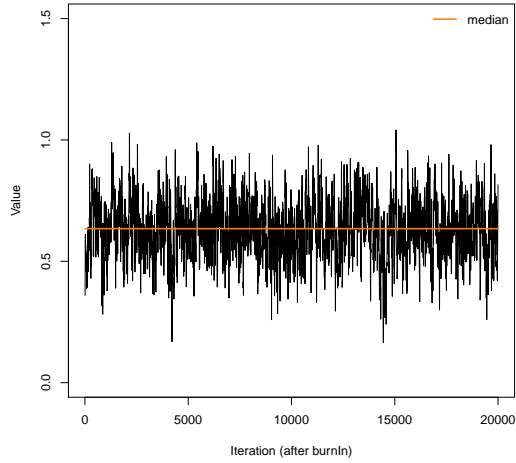
Figure 12: Apple data. Estimated posterior distributions for the  $\beta_p$  parameters,  $p = 1, \dots, P$ .




(a) *Xylose.*



(b) *Epicatechin Sulfate.*




(c) *(4-3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-DHMPMB-SA).*

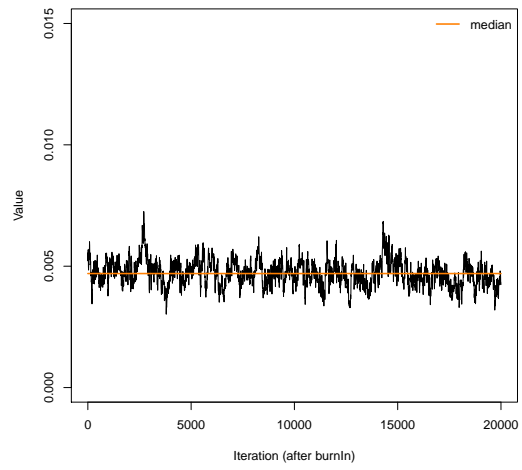


(d) *Glucodistylin.*


Figure 13: Apple data. Trace plots for the estimated  $\alpha_p$  parameters,  $p = 1, \dots, P$ .




(a) *Xylose.*



(b) *Epicatechin Sulfate.*



(c) *(4-3-[2-(2,4-dihydroxyphenyl)-2-oxoethyl]-DHMPMB-SA).*



(d) *Glucodistylin.*

Figure 14: Apple data. Trace plots for the estimated  $\beta_p$  parameters,  $p = 1, \dots, P$ .