A Bayesian partial membership model for multiple exposures with uncertain group memberships

[Link to publication record in Ulster University Research Portal](#)

Document Version
Author Accepted version

A Bayesian partial membership model for multiple exposures with uncertain group memberships

December 22, 2022

Abstract

We present a Bayesian partial membership model that estimates the associations between an outcome, a small number of latent variables, and multiple observed exposures where the number of latent variables is specified *a priori*. We assign one observed exposure as the sentinel marker for each latent variable. The model allows non-sentinel exposures to have complete membership in one latent group, or partial membership across two or more latent groups. MCMC sampling is used to determine latent group partial memberships for the non-sentinel exposures, and estimate all model parameters. We compare the performance of our model to competing approaches in a simulation study, and apply our model to inflammatory marker data measured in a large mother-child cohort of the Seychelles Child Development Study (SCDS). In simulations, our model estimated model parameters with little bias, adequate coverage, and tighter credible intervals compared to competing approaches. Under our partial membership model with two latent groups, SCDS inflammatory marker classifications generally aligned with the scientific literature. Incorporating additional SCDS inflammatory markers and more latent groups produced similar groupings of markers that also aligned with the literature. Associations between covariates and birth weight were similar across latent variable models, and were consistent with earlier work in this SCDS cohort.

Keywords: Immune response; Inflammation; Latent variables; Markov chain Monte Carlo; Multiple exposures; Seychelles Child Development Study.

1 Introduction

Several challenges may arise when modeling the association between multiple observed exposures and an outcome of interest. When the observed exposures are correlated, standard approaches such as ordinary least squares (OLS) regression are complicated by collinearity and variance inflation (Carrico et al., 2015). Several alternatives to OLS regression have been developed, including weighted quantile sum (WQS) regression (Carrico et al., 2015), grouped WQS regression (Wheeler et al., 2021), Bayesian Kernel Machine Regression (Bobb et al., 2014), and structural equation models (SEMs) (Sánchez et al., 2005). SEMs and variations on SEMs (Jedidi et al., 1997; Lee and Song, 2003; Muthén and Asparouhov, 2012) have been proposed as flexible approaches for modeling several predictors that may be correlated. While SEMs may consist of only observed variables, SEMs involving latent variables have been utilized in a wide variety of applications including nutritional epidemiology (Keller, 2006) and phthalate metabolites (Weuve et al., 2006).

SEMs allow associations between covariates, observed exposures, latent variables, and outcomes, however, in most cases relationships must be established *a priori*. For some applications, there may be interest in estimating a subset of these associations, particularly the assignment of observed exposures to latent variables. Researchers may wish to compare exposure assignments that are hypothesized *a priori* to exposure assignments that are estimated from data. Alignment between model-based exposure assignments and literature-based expectations may validate experimental data and strengthen scientific conclusions. Disagreement may alert researchers to potential data quality concerns, or lead to novel hypotheses regarding exposure relationships.

Uncertainty in exposure assignments may arise when studying the roles of maternal T-helper type 1 (Th1) cytokines and T-helper type 2 (Th2) cytokines in relation to an outcome, such as birth weight. Th1 cytokines are generally associated with pro-inflammatory responses, while Th2 cytokines are associated with anti-inflammatory responses (Berger, 2000). Some cytokines (i.e. interferon (IFN)-γ) are routinely classified as "Th1," while other cytokines (i.e. interleukin (IL)-4) are routinely classified as "Th2." However, de-

termining cytokine classifications is not always straightforward. Certain cytokines may have the capacity to behave as either pro-inflammatory or anti-inflammatory, depending on the specific context (Cavaillon, 2001). For example, IL-6 has been noted to have both pro- and anti-inflammatory properties (Scheller et al., 2011). Separate non-cytokine factors (e.g. chemokines), may also influence the Th1/Th2 paradigm (Kidd, 2003).

To address this uncertainty, Zavez et al. (2020) developed a latent Bayesian model for multiple observed exposures that allows exposure-specific slopes and variances, and additionally estimates the latent class membership for each exposure. The model assigns each observed exposure to exactly one latent group, and estimates the magnitude of the association between each exposure and its assigned latent group. The application of that model to prenatal inflammatory marker data collected from the Seychelles Child Development Study (SCDS) revealed that several cytokine permutations may be plausible, possibly indicating that certain cytokines do not belong completely to one latent group. A more flexible model that permits observed exposures (e.g., cytokines in this example) to have partial membership across multiple latent groups may be more effective at modeling the true relationships among these prenatal Th1/Th2 inflammatory markers. A partial membership model may also be desirable when modeling a larger set of inflammatory markers. For example, IL-6 was not categorized as "Th1" or "Th2" by McSorley et al. (2018) or Yeates et al. (2020) because it is not consistently classified as either Th1 or Th2 in the literature. IL-6, however, has a "dual" role in Th1/Th2 differentiation in that it promotes Th2 response by inhibiting Th1 polarization (Diehl and Rincón, 2002), and exhibits both pro- and anti-inflammatory properties (Scheller et al., 2011).

For this reason, we propose a novel Bayesian approach that allows observed exposures to have either complete membership in one latent group, or partial membership across several latent groups. Under our proposed model, the number of latent groups must be specified *a priori*. In addition, one observed exposure must be selected as the sentinel exposure for each latent group to prevent label switching (Stephens, 2000) and ensure model identifiability. Each sentinel exposure is assumed to have complete membership in the one latent group to which it is assigned, while MCMC sampling is used to determine

the latent membership(s) for the non-sentinel exposures and estimate partial membership coefficients.

We first apply our model to the seven Th1/Th2 cytokines previously analyzed by Mc-Sorley et al. (2018), Yeates et al. (2020), and Zavez et al. (2020) in the Seychelles Child Development Study (SCDS) Nutrition Cohort 2: IFN-γ, IL-2, IL-1β, tumor necrosis factor (TNF)-α, IL-4, IL-5, and IL-10, as well as two latent variables: "Th1" and "Th2." We then expand our analysis to thirteen markers by additionally including IL-6, monocyte chemoattractant protein (MCP)-1 , thymus- and activation-regulated chemokine (TARC), soluble fms-like tyrosine kinase 1 (sFlt-1), vascular endothelial growth factor (VEGF)-D, and C-reactive protein (CRP). Many of these additional markers may be involved with fetal development during pregnancy. MCP-1 and TARC are both chemokines, which may influence the maternal-fetal interface (Du et al., 2014), while sFlt-1 and VEGF-D may influence placental development (Yeates et al., 2020). Independently, CRP was a significant predictor of birth weight in the NC2 cohort (Yeates et al., 2020). For the expanded set of exposures, we fit our partial membership model under three latent groups and also under four latent groups.

The rest of the paper is organized as follows. In Section 2, we introduce our model and MCMC sampling procedure. In Section 3, we evaluate the performance of our model in several related simulation studies. In Section 4, we fit our proposed partial membership model to the Seychelles inflammatory marker data under several different partial membership model conditions. We close with a discussion in Section 5.

2 Model

2.1 Notation and Models

We center and scale all observed quantities (exposures, covariates, and the outcome) to account for the possibility of measurements on different scales and eliminate the need for intercept parameters. We use the indices (i, j, k) where i indexes the subject, j indexes the observed exposure, and k indexes the latent group, $i = 1, ..., n, j = 1, ..., J, k = 1, ..., K,$

where $K < J$. To simplify notation, we omit indices that are not relevant to a given quantity.

For the ith subject, we let y_i denote the centered and scaled outcome, $x_{i,1}... \ x_{i,K}$ denote the K unobserved latent groups, and $w_{i,1}, ..., w_{i,J}$ denote the J centered and scaled observed exposures. We take a classical measurement error approach (Carroll et al., 2006), and model the observed exposures (w's) conditional on the latent quantities (x's). Our models for y_i, $x_{i,k}$, $k = 1, ..., K$, and $w_{i,j}, j = 1, ..., J$ are

$$y_i | \boldsymbol{x_i}, \boldsymbol{Z_{y,i}}, \boldsymbol{\beta_x}, \boldsymbol{\beta_z}, \sigma^2 \sim N(\boldsymbol{x_i\beta_x} + \boldsymbol{Z_{y,i}\beta_z}, \sigma^2)$$

$$x_{i,k} | \boldsymbol{Z_{x,i}}, \boldsymbol{\gamma_k}, \eta_k^2 \sim N(\boldsymbol{Z_{x,i}\gamma_k}, \eta_k^2)$$

$$w_{i,j} | \boldsymbol{x_i}, \boldsymbol{\lambda_j}, \tau_j^2 \sim N(\boldsymbol{x_i\lambda_j}, \tau_j^2)$$

where, for subject i, $\boldsymbol{x_i} = (x_{i,1}, ..., x_{i,K})$ is the K-dimensional vector of latent group variables, $\boldsymbol{Z_{y,i}} = (Z_{y_1,i}, Z_{y_2,i}, ..., Z_{y_p,i})$, is a p-dimensional vector of covariates associated with the outcome, and $\boldsymbol{Z_{x,i}} = (Z_{x_1,i}, Z_{x_2,i}, ..., Z_{x_q,i})$, a q-dimensional vector of covariates associated with the latent groups. In the model for y_i, $\boldsymbol{\beta_x}$ denotes the K-dimensional vector of coefficients for the K latent variables, and $\boldsymbol{\beta_y}$ denotes the p-dimensional vector of coefficients for the Z_y covariates. In the model for $x_{k,i}$, $\boldsymbol{\gamma_k}$ denotes the q-dimensional vector of coefficients for the Z_x covariates. In the model for $w_{i,j}$, $\boldsymbol{\lambda_j} = (\lambda_{j,1}, ..., \lambda_{j,k})^T$ denotes the K-dimensional vector of coefficients for the K latent variables. We use $\lambda_{j,k} \geq 0$ to denote the proportion of exposure j that is assigned to latent group k, where $k = 1, ..., K$ and $j = 1, ..., J$. For each exposure, we denote the $\lambda_{j,k}$ values with $\boldsymbol{\lambda_j} = (\lambda_{j,1}, ..., \lambda_{j,K})$ and require that $\sum_{k=1}^{K} \lambda_{j,k} = 1$. The $\lambda_{j,k}$ parameterization is discussed further in Section 2.2.

It is relatively straightforward to develop a Bayesian modeling framework in which each observed exposure has membership in every latent group (see Web Appendix A of the Supporting Information for details). However, allowing all non-sentinel exposures to belong to all latent groups might not be biologically plausible, and may result in very small slopes that are difficult to interpret. For this reason, we utilize a parameterization that

permits non-sentinel exposures to have either complete membership in one latent group or partial membership across two or more latent groups. Our model parameterization also discourages very small partial memberships, which increases interpretability.

2.2 Parameterization Details

Our $\lambda_{j,k}$ parameterization was introduced by Xiao et al. (2014) in the context of multiple outcomes research, where child developmental outcomes may have partial membership in one or more latent domains. We adapt Xiao et al. (2014)'s approach in order to estimate partial memberships for the non-sentinel exposures, and write $\lambda_{j,k}$ in the following way

$$\lambda_{j,k} = \frac{z_{j,k} \exp(v_{j,k})}{\sum_{k'=1}^{K} z_{j,k'} \exp(v_{j,k'})}, \; j = 1, ..., J; \; k = 1, ..., K \tag{1}$$

where $z_{j,k} = 1$ if exposure j has at least partial membership in the kth latent group, and $z_{j,k} = 0$ otherwise. In this notation, $\exp(v_{j,k})$ represents the relative membership weight for exposure j in latent group k. Reparameterizing each $\lambda_{j,k}$ in this way (i.e., by using $v_{j,k}$ and $z_{j,k}$) allows some $\lambda_{j,k}$ values to be zero. Modeling the $\lambda_{j,k}$ parameters directly with a Dirichlet prior would not permit the same level of flexibility (Xiao et al., 2014).

We fix a subset of the $\lambda_{j,k}$ parameters to ensure identifiability and prevent label switching. We select one exposure as the sentinel marker for each of the K latent groups and set the corresponding $\lambda_{j,k}$ parameter equal to one. For example, if exposure 2 is the sentinel exposure for the first latent group, $\lambda_{2,1} = 1$. The remaining $(K - 1)$ $\lambda_{2,k}$ parameters, $k = 2, ..., K$ for exposure 2 must be equal to 0.

For each non-sentinel exposure, we have $\boldsymbol{\lambda_j} = (\lambda_{j,1}, ..., \lambda_{j,K})$ where $\boldsymbol{z_j} = (z_{j,1}, ..., z_{j,K})$ is a binary vector, and $\boldsymbol{v_j} = (v_{j,1}, ..., v_{j,K})$ is a log-transformed weight vector. To ensure that the denominator in Equation 1 is always nonzero, we utilize prior information to make an *a priori* partial membership assignment (PMA) for each nonsentinel exposure by fixing one of it's $z_{j,k}$ parameters at 1. Fixing $z_{j,k} = 1$ for exposure j is equivalent to saying, "exposure j has at least partial membership in latent group k." For each exposure, we can denote the corresponding latent group by $k^*(j)$. For example, if we assume that

exposure 3 has at least partial membership in latent group 1, then $k^*(3) = 1$.

For technical reasons, it will also be necessary to fix one $v_{j,k}$ equal to 0 for each exposure. This is because $\boldsymbol{v_j} = (v_{j,1}, ..., v_{j,K})$ and $\boldsymbol{v'_j} = (v_{j,1} + C, ..., v_{j,K} + C)$ will produce the same $\boldsymbol{\lambda_j}$, $\forall C \in \mathsf{R}$. For each exposure, we follow the recommendation of Xiao et al. (2014) and fix the $v_{j,k}$ that corresponds to $k^*(j)$ equal to zero. Thus, we set $v_{j,k^*(j)} = 0$.

2.3 MCMC Sampling Details

We utilize an MCMC sampling scheme that is similar to the approach taken by Xiao et al. (2014). For all model parameters except $\boldsymbol{\lambda_j}$, $j = 1, ..., J$ we use Gibbs sampling to iteratively sample each parameter (or parameter vector) from its posterior, conditional on the values of the other model parameters. To obtain posterior samples for the set of $\boldsymbol{\lambda_j}$ parameters, we use a Metropolis-Hasting sampler to sample $\boldsymbol{v_j}$ and Gibbs sampling to sample $\boldsymbol{z_j}$. Posterior distributions and further sampling details are provided in Web Appendix B of the Supporting Information.

2.4 Model Priors

We assume the same priors as Zavez et al. (2020) for $\boldsymbol{\beta} = (\beta_x, \boldsymbol{\beta_z})$, $\boldsymbol{\gamma_k}$, σ^2, and η_k^2, $k = 1, ..., K$. We let $\boldsymbol{\beta} \sim N(\boldsymbol{\beta_0}, \Sigma_\beta)$ and $\boldsymbol{\gamma_k} \sim N(\boldsymbol{\gamma_0}, \Sigma_\gamma)$, $k = 1, ..., K$ where $\boldsymbol{\beta_0} = \boldsymbol{\gamma_0} = \boldsymbol{0}$ and $\Sigma_\beta = \Sigma_\gamma = \text{diag}(100)$. We assume that all variance parameters are independent and follow Inverse-Gamma (IG) distributions, thus, $\sigma^2 \sim IG(a_{\sigma^2}, b_{\sigma^2})$, $\tau_j^2 \sim IG(a_{\tau^2}, b_{\tau^2})$, $j = 1, ..., J$ and $\eta_j \sim IG(a_{\eta^2}, b_{\eta^2})$, $j = 1, ..., J$. For prior A, we fix the IG shape hyperparameters $a_{\sigma^2} = a_{\tau^2} = a_{\eta^2} = 0.05$ and the IG scale hyperparameters $b_{\sigma^2} = b_{\tau^2} = b_{\eta^2} = 0.01$. For prior B, we fix $a_{\sigma^2} = a_{\tau^2} = a_{\eta^2} = 0.50$ and $b_{\sigma^2} = b_{\tau^2} = b_{\eta^2} = 0.10$. Zavez et al. (2020) provides further justification for the prior parameter values in the context of this Seychelles cohort.

Like Xiao et al. (2014), we utilize a Bernoulli prior distribution for $z_{j,k}$ and a Normal mixture prior for $v_{j,k}$. Specifically, we assume that $p(z_{j,k}, v_{j,k}) = p(z_{j,k})p(v_{j,k}|z_{j,k})$, $p(z_{j,k}) \sim \text{Bernoulli}(p_{j,k,0})$ and $p(v_{j,k}|z_{j,k}) \sim z_{j,k}N(0, \sigma_v^2) + (1 - z_{j,k})\mathrm{N}^-(0, g\sigma_v^2)$, where $g > 1$ is a fixed value. In this notation, $\mathrm{N}^-(0, g\sigma_v^2)$ represents the distribution of $-|X|$ where

$X \sim N(0, g\sigma_v^2)$.

Normal mixture priors have been utilized in other fields (Yi et al., 2003), and allow for greater flexibility than standard Normal priors. For example, $z_{j,k} = 0$ implies that exposure j has no membership in latent group k. The corresponding $\exp(v_{j,k})$, which represents the relative membership of exposure j in the kth latent group, should be close to zero. This will be the case for large, negative values of $v_{j,k}$, which are typical under the $N^-(0, g\sigma_v^2)$ component of the mixture prior.

Utilizing σ_v^2 to govern the mixture prior variance allows for greater sampling flexibility than a fixed variance, and is the approach taken by Xiao et al. (2014). In sensitivity analyses, we evaluate the three values for g ($g = 5$, $g = 10$, $g = 15$) and three different conditions for σ_v^2 ($\sigma_v^2 \sim U(1, 4)$, $\sigma_v^2 = 1$, $\sigma_v^2 = 4$) considered by Xiao et al. (2014). While a fixed value for σ_v^2 may be a feasible approach for some applications, our simulation results suggest that choosing a σ_v^2 that is too large may adversely affect estimation of the $\lambda_{j,k}$ parameters (see Web Appendix D of the Supporting Information for details).

Use of a traditional "spike and slab" prior (Mitchell and Beauchamp, 1988), in which the coefficients of one component of the mixture distribution are exactly zero, is a reasonable alternative to the method we used. The results are likely to be similar to our method. In our method, $z_{j,k} = 0$ sets $\lambda_{j,k}$ to zero for that iteration. This has the same effect as taking $\lambda_{j,k} = 0$ for the spike at zero in the spike and slab prior at that iteration.

3 Simulations

3.1 Simulation Study Design

We simulated data under the following model structure with three latent groups ($K = 3$) and nine observed exposures ($J = 9$):

$$y_i \sim N(\boldsymbol{x_i}\boldsymbol{\beta_x} + \boldsymbol{Z_{y,i}}\boldsymbol{\beta_z}, \sigma^2) \text{ where } \boldsymbol{\beta_x} = (0.25, -0.15, 0.10); \boldsymbol{\beta_z} = (0.15, -0.10, 0.05)$$

$$x_{i,1} \sim N(\mathbf{Z}_{x,i}\boldsymbol{\gamma_1}, \eta_1^2) \text{ where } \boldsymbol{\gamma_1} = (0.07, 0.21)$$

$$x_{i,2} \sim N(\mathbf{Z}_{x,i}\boldsymbol{\gamma_2}, \eta_2^2) \text{ where } \boldsymbol{\gamma_2} = (0.15, 0.05)$$

$$x_{i,3} \sim N(\mathbf{Z}_{x,i}\boldsymbol{\gamma_3}, \eta_3^2) \text{ where } \boldsymbol{\gamma_3} = (0.28, 0.11)$$

$$w_{i,j} \sim N(\boldsymbol{x_i}\boldsymbol{\lambda_j}, \tau_j^2) \, j = 1, ..., 9; \, \boldsymbol{x_i} = (x_{i,1}, x_{i,2}, x_{i,3})$$

We use a sample size of 500 and simulate each dataset from a multivariate Normal distribution. All observed exposures ($\boldsymbol{w_j}$'s), covariates ($\mathbf{Z_y}; \mathbf{Z_x}$), and the simulated outcome ($\boldsymbol{y}$) are simulated to have mean zero and unit variance which defines the scale for the residual variances. For the latent groups, we fix $Var(\boldsymbol{x_k}) = 0.5$ and set the pairwise correlation between the $\boldsymbol{x_k}$'s equal to 0.15 to ensure valid models for each sentinel exposure (see Web Appendix C of the Supporting Information for further details) and produce correlations among the simulated $\boldsymbol{w_j}$'s that are representative of the correlations among the observed Th1/Th2 cytokines in the NC2 cohort.

We set $\boldsymbol{\lambda_1} = \boldsymbol{\lambda_2} = (1, 0, 0), \boldsymbol{\lambda_3} = (0.6, 0.4, 0), \boldsymbol{\lambda_4} = (0, 1, 0), \boldsymbol{\lambda_5} = (0.3, 0.4, 0.3), \boldsymbol{\lambda_6} = (0, 0.7, 0.3), \boldsymbol{\lambda_7} = (0, 0.5, 0.5), \boldsymbol{\lambda_8} = (0.3, 0, 0.7),$ and $\boldsymbol{\lambda_9} = (0, 0, 1)$. The first ($w_1$), fourth ($w_4$) and ninth ($w_9$) simulated observed exposures are assumed to be the sentinel markers for the first (x_1), second (x_2) and third (x_3) latent variables, respectively. The remaining non-sentinel exposures ($w_2, w_3, w_5, w_6, w_7, w_8$) are permitted to have partial membership across one or more latent group. Under this simulation design, the range of exposure correlations in the simulated data (range: -0.06 to 0.59) was similar to the correlation range observed among the NC2 cytokines (range: -0.10 to 0.65).

3.2 Simulation Results

We simulate 100 datasets and evaluate our partial membership (PM) model against the single membership (SM) model presented by Zavez et al. (2020) and two different SEM-based approaches. For each model parameter, we compute parameter bias, credible interval length, and coverage. For bias, we report the mean of $(\hat{\theta} - \theta)$, where $\hat{\theta}$ denotes the posterior mean of true parameter θ from which the data are generated. For interval length, we report the mean 95% equal-tailed posterior interval length, and for coverage, we report the proportion of these intervals that cover the true parameter value.

Under all three competing approaches, each exposure is assumed to belong to exactly one latent group. For the SM model, the assignment is estimated using MCMC sampling, while for the SEM models assignments are specified *a priori*. For the first SEM-based model, we assume that each measured exposure is assigned to the latent group in which it has majority membership. We refer to this model as "SEM Best" because it assumes that optimal *a priori* decisions are made regarding exposure group memberships. For the second SEM-based model, we assume that each measured exposure is assigned poorly, and to a latent group in which it does not have majority membership. We refer to this model as "SEM Worst." These three approaches provide insight into how a single membership model will perform when applied to data that are simulated under a partial membership model.

Our partial membership model requires the specification of g and σ_v^2. Detailed sensitivity analyses suggest that model performance can be enhanced by optimizing these parameters (see Web Tables 1-3 of the Supporting Information for details). For this simulation, we fix $g = 15$ and $\sigma_v^2 = 4$, which were shown to be relatively poor choices of g and σ_v^2. This helps to ensure that we are not optimizing our proposed model compared to the competing approaches. For each approach (PM, SM, SEM Worst, and SEM Best), we run the MCMC for 4 chains with 5,000 iterations per chain. The first 1,000 iterations of each chain are considered "burn-in" draws and are discarded, while the remaining 16,000 draws are used to estimate model parameters.

Performance results for select model parameters are presented in Table 1. It is difficult

to compare exposure model parameters (i.e., λ_j's, τ_j^2's) across the four models because our partial membership model places certain constraints on the λ_j's (e.g. $\sum_{k=1}^{\cdot K} \lambda_{j,k} = 1$) that are not enforced by the other approaches. For this reason, we focus on comparing outcome model parameters ($\boldsymbol{\beta}, \sigma^2$) and latent variable model parameters (γ's, η_k^2's) across the four different models.

Table 1: Simulation results with three latent groups under prior A (Beta(1,1); IG(0.05, 0.01)) for our partial membership model ("PM"), the single membership model ("SM"), the best-choice SEM model ("SEM-B") and worst-choice SEM model ("SEM-W") for select model parameters, where (β, σ^2) and (γ, η^2) are slopes and variances for ($y \mid x, Z_y$) and ($x \mid Z_x$), respectively. Reported values are averages over 100 simulated datasets. For PM, we fix $g = 15$ and $\sigma_v^2 = 4$.

		Bias				Interval Length				Coverage			
	Truth	PM	SM	SEM-B	SEM-W	PM	SM	SEM-B	SEM-W	PM	SM	SEM-B	SEM-W
β_{x_1}	0.250	-0.013	-0.021	-0.226	-0.502	0.291	0.362	1.527	3.564	0.950	0.940	0.960	0.860
β_{x_2}	-0.150	0.028	-0.014	0.018	0.656	0.314	0.719	1.116	5.168	0.950	0.990	0.960	0.730
β_{x_3}	0.100	0.004	0.049	-0.083	-0.045	0.317	0.446	1.486	1.589	0.950	0.950	0.970	0.970
$\beta_{Z_{y\,1}}$	0.150	-0.000	-0.000	-0.000	-0.000	0.171	0.171	0.171	0.172	0.940	0.940	0.940	0.950
$\beta_{Z_{y\,2}}$	-0.100	0.001	0.001	0.001	0.001	0.171	0.171	0.171	0.172	0.960	0.950	0.970	0.940
$\beta_{Z_{y\,3}}$	0.050	0.001	0.001	0.002	0.001	0.171	0.171	0.171	0.172	0.940	0.940	0.940	0.940
σ^2	0.918	0.001	-0.000	0.001	-0.011	0.237	0.242	0.240	0.267	1.000	1.000	1.000	1.000
$\gamma_{1,1}$	0.070	-0.003	0.008	-0.001	-0.013	0.148	0.144	0.150	0.128	0.940	0.900	0.940	0.950
$\gamma_{1,2}$	0.210	0.002	-0.006	-0.021	-0.179	0.147	0.154	0.186	0.086	0.960	0.950	0.960	0.010
$\gamma_{2,1}$	0.150	-0.001	0.019	0.010	-0.096	0.159	0.160	0.166	0.127	0.960	0.920	0.920	0.300
$\gamma_{2,2}$	0.050	-0.006	0.016	0.017	0.004	0.159	0.144	0.138	0.124	0.970	0.950	0.940	0.940
$\gamma_{3,1}$	0.280	0.001	-0.018	-0.028	-0.078	0.153	0.168	0.224	0.203	0.960	0.910	0.960	0.740
$\gamma_{3,2}$	0.110	0.009	0.025	0.014	0.017	0.153	0.146	0.168	0.152	0.980	0.940	0.960	0.910
η_1^2	0.451	0.027	-0.014	-0.041	-0.424	0.210	0.303	0.348	0.064	0.990	0.990	0.970	0.000
η_2^2	0.475	0.035	-0.183	-0.193	-0.451	0.330	0.437	0.299	0.059	0.970	0.540	0.350	0.000
η_3^2	0.409	0.022	-0.077	-0.097	-0.271	0.295	0.352	0.328	0.207	0.970	0.850	0.800	0.030

Overall, our partial membership model estimated model parameters with little bias and tighter intervals compared to the SM and SEM approaches, though all four methods performed similarly with regards to estimating the $\boldsymbol{\beta_z}$ parameters. The biases for β_{x_1} and β_{x_3} under the SM model were fairly small, but were nonetheless 2-10 times larger than the corresponding biases under the PM model. The bias for β_{x_2} under the SM model was half that of the PM model, but the length of the SM model credible interval was more than double that of the PM model. In addition, 20% of SM interval lengths were greater than 1.0, while 100% of PM intervals were smaller than 0.5 (See Web Appendix E for details). The SM model also performed worse than the PM model in estimating η_2^2. The bias for η_2^2 under the SM model was over 4 times larger than bias under the PM model, and SM model coverage for η_2^2 (54%) was much lower than the PM model coverage for

η_2^2 (97%). Thus, the PM model produced more precise intervals, and generally smaller biases, than the SM membership model. In addition, the high SM coverage estimates for some parameters may be a reflection of larger credible intervals, and should be interpreted with this in mind.

While SEM Best estimated most parameters with little bias and adequate coverage, the biases for two of the three $\boldsymbol{\beta_x}$ parameters (the latent group slopes) under SEM Best were unacceptably large. In addition, credible intervals for $\boldsymbol{\beta_x}$ were more than three times wider than those under the partial membership model.

3.3 Sensitivity Analyses

Traditional SEM-based approaches require that each observed exposure be completely assigned to one latent group when fitting the model, and require that the number of latent groups, K, be known. Our model allows partial membership of observed exposures to latent groups, while having three requirements: (1) K is known; (2) one exposure for each latent group is assigned *a priori* as the sentinel exposure with complete membership in that group; and (3) each non-sentinel exposure is assigned partial membership in one latent group *a priori* (see Section 2.2). The assignment of K sentinel exposures allows the remaining $J - K$ exposures to have full or partial membership in latent group(s) as estimated by our partial membership model. We now evaluate the sensitivity of our model to these requirements.

3.3.1 Sensitivity to Number of Latent Groups

Our partial membership model requires the specification of K, the number of latent groups. To evaluate model sensitivity to this parameter, we performed additional simulations for the $K = 3$ simulation described above, but assuming either too few ($K = 2$) or too many ($K = 4$) latent groups. For the $K = 2$, $K = 3$, and $K = 4$ models, the first observed exposure (w_1) and fourth observed exposure (w_4) were assumed to be the sentinel exposures for the first (x_1) and second (x_2) latent groups, respectively. For the $K = 3$ and $K = 4$ models, the ninth observed exposure (w_9) was assumed to be the sentinel

exposures for the third (x_3) latent group. For the $K = 4$ model, the seventh observed exposure (w_7) was assumed to be the sentinel exposures for the fourth (x_4) latent group.

Since exposure model constraints are sensitive to K (e.g., $\sum_{k=1}^{L \cdot K} \lambda_{j,k} = 1$), we compare outcome model parameters ($\boldsymbol{\beta}, \sigma^2$) and latent variable model parameters (γ's, η_k^2's) across the $K = 2$, $K = 3$, and $K = 4$ models in Table 2. Since the fourth latent group does not exist under the $K = 3$ simulation design, no true values, biases, or coverages are provided for parameters specific to the fourth latent group (i.e., β_{x_4}, $\gamma_{4,1}$, $\gamma_{4,2}$, and η_4^2).

Table 2: Simulation results with three latent groups under prior A (Beta(1,1); IG(0.05, 0.01)) for our partial membership model under three different K specifications: $K = 2$ (too few latent groups), $K = 3$ (the correct number of latent groups), and $K = 4$ (too many latent groups). (β, σ^2) and (γ, η^2) are slopes and variances for ($y \mid x, Z_y$) and ($x \mid Z_x$), respectively. Reported values are averages over 100 simulated datasets. For each model, we fix $g = 15$ and $\sigma_v^2 = 4$.

		Bias			Interval Length			Coverage		
	Truth	$K = 2$	$K = 3$	$K = 4$	$K = 2$	$K = 3$	$K = 4$	$K = 2$	$K = 3$	$K = 4$
β_{x_1}	0.250	-0.001	-0.013	-0.013	0.300	0.291	0.292	0.970	0.950	0.940
β_{x_2}	-0.150	0.120	0.028	0.050	0.415	0.314	0.304	0.770	0.950	0.910
β_{x_3}	0.100	-	0.004	-0.004	-	0.317	0.320	-	0.950	0.960
β_{x_4}	-	-	-	-	-	-	0.488	-	-	-
$\beta_{Z_{y1}}$	0.150	-0.001	-0.000	-0.000	0.171	0.171	0.171	0.940	0.940	0.930
$\beta_{Z_{y2}}$	-0.100	0.001	0.001	0.001	0.171	0.171	0.171	0.960	0.960	0.960
$\beta_{Z_{y3}}$	0.050	0.001	0.001	0.001	0.171	0.171	0.171	0.940	0.940	0.950
σ^2	0.918	0.013	0.001	-0.003	0.237	0.237	0.240	1.000	1.000	1.000
$\gamma_{1,1}$	0.070	0.001	-0.003	-0.004	0.148	0.148	0.149	0.940	0.940	0.950
$\gamma_{1,2}$	0.210	0.003	0.002	0.003	0.146	0.147	0.149	0.960	0.960	0.960
$\gamma_{2,1}$	0.150	0.062	-0.001	-0.002	0.111	0.159	0.169	0.380	0.960	0.960
$\gamma_{2,2}$	0.050	0.026	-0.006	-0.005	0.112	0.159	0.168	0.840	0.970	0.960
$\gamma_{3,1}$	0.280	-	0.001	0.002	-	0.153	0.162	-	0.960	0.940
$\gamma_{3,2}$	0.110	-	0.009	0.012	-	0.153	0.162	-	0.980	0.980
$\gamma_{4,1}$	-	-	-	-	-	-	0.166	-	-	-
$\gamma_{4,2}$	-	-	-	-	-	-	0.167	-	-	-
η_1^2	0.451	0.018	0.027	0.034	0.227	0.210	0.225	1.000	0.990	0.990
η_2^2	0.475	-0.243	0.035	0.167	0.115	0.330	0.538	0.000	0.970	0.830
η_3^2	0.409	-	0.022	0.142	-	0.295	0.512	-	0.970	0.900
η_4^2	-	-	-	-	-	-	0.688	-	-	-

The $K = 2$ model produced relatively large bias and poor coverage for the model parameters that are directly associated with the second latent group: β_{x_2}, $\gamma_{2,1}$, $\gamma_{2,2}$, and η_2^2. Compared to the $K = 3$ model, the $K = 4$ model produced relatively large intervals, larger bias, and weaker coverage for several of the variance parameters (i.e., η_2^2 and η_3^2). In addition, convergence diagnostics from the $K = 4$ model showed clear convergence issues

in traceplots, the Gelman-Rubin diagnostic ($\hat{R}$), and the effective sample size estimates (n_{eff}), especially for parameters associated with the non-existent group (Gelman et al., 1992; Plummer et al., 2006). Further details are provided in Web Appendix F of the Supporting Information.

3.3.2 Sensitivity to Choice of Sentinel Exposures

Our partial membership model also requires the specification of one sentinel exposure for each latent group. In the simulations considered earlier, the sentinel exposures selected (w_1, w_4, and w_9) had complete membership in the latent group to which they were assigned.

We now evaluate our model's performance under two additional sentinel exposure assignments. Under Sentinel-2, we select the exposures with the second greatest proportion of membership (w_2, w_6, and w_8) to be the sentinel markers. Under Sentinel-3, we select the exposures with the third greatest proportion of membership (w_3, w_5, and w_7) to be the sentinel markers. Parameter results for this sensitivity analysis are presented in Table 3.

Table 3: Simulation results with three latent groups under prior A (Beta(1,1); IG(0.05, 0.01)) for our partial membership model under three different sentinel marker arrangements: Sent-1 (w_1, w_4, and w_9 as the sentinel markers), Sent-2 (w_2, w_6, and w_8 as the sentinel markers), and Sent-3 (w_3, w_5, and w_7 as the sentinel markers). (β, σ^2) and (γ, η^2) are slopes and variances for ($y \mid x, Z_y$) and ($x \mid Z_x$), respectively. Reported values are averages over 100 simulated datasets. For each model, we fix $g = 15$ and $\sigma_v^2 = 4$.

		Bias			Interval Length			Coverage		
	Truth	Sent-1	Sent-2	Sent-3	Sent-1	Sent-2	Sent-3	Sent-1	Sent-2	Sent-3
β_{x_1}	0.250	-0.013	-0.041	0.011	0.291	0.291	0.359	0.950	0.930	0.970
β_{x_2}	-0.150	0.028	0.018	0.112	0.314	0.368	0.576	0.950	0.970	0.780
β_{x_3}	0.100	0.004	0.051	-0.091	0.317	0.372	0.478	0.950	0.930	0.800
$\beta_{Z_{y_1}}$	0.150	-0.000	-0.001	-0.001	0.171	0.171	0.171	0.940	0.960	0.940
$\beta_{Z_{y_2}}$	-0.100	0.001	0.001	0.001	0.171	0.171	0.171	0.960	0.960	0.950
$\beta_{Z_{y_3}}$	0.050	0.001	0.000	0.000	0.171	0.171	0.172	0.940	0.940	0.940
σ^2	0.918	0.001	0.001	-0.001	0.237	0.239	0.242	1.000	1.000	1.000
$\gamma_{1,1}$	0.070	-0.003	-0.008	0.010	0.148	0.152	0.135	0.940	0.950	0.950
$\gamma_{1,2}$	0.210	0.002	0.005	-0.017	0.147	0.152	0.135	0.960	0.940	0.940
$\gamma_{2,1}$	0.150	-0.001	0.026	0.024	0.159	0.140	0.158	0.960	0.880	0.860
$\gamma_{2,2}$	0.050	-0.006	0.003	0.051	0.159	0.139	0.152	0.970	0.950	0.780
$\gamma_{3,1}$	0.280	0.001	-0.032	-0.048	0.153	0.146	0.147	0.960	0.870	0.740
$\gamma_{3,2}$	0.110	0.009	0.030	-0.024	0.153	0.146	0.140	0.980	0.860	0.900
η_1^2	0.451	0.027	0.057	-0.074	0.210	0.245	0.185	0.990	0.890	0.560
η_2^2	0.475	0.035	-0.103	-0.159	0.330	0.182	0.322	0.970	0.470	0.350
η_3^2	0.409	0.022	-0.051	-0.105	0.295	0.241	0.216	0.970	0.820	0.440

Two of the three latent group slopes (β_{x_1} and β_{x_3}) showed larger bias under Sentinel-2 than under Sentinel-1, and bias for $\boldsymbol{\beta_x}$ was generally larger under Sentinel-3. Sentinel-3 also produced wider intervals and lower coverage for β_{x_2}, β_{x_3} and select γ parameters. Under Sentinel-2 and Sentinel-3, η^2 parameters were estimated with more bias, wider intervals, and poor coverage compared to Sentinel-1.

3.3.3 Sensitivity to Partial Membership Assignments (PMAs)

We first evaluate our model's performance under three different *a priori* partial membership assignments (PMAs). Under PMA-1, each non-sentinel observed exposure is assigned to have partial membership in the latent group in which it has majority membership. Under PMA-2, non-sentinel exposures are assigned to have partial membership in latent groups in which they have at least partial membership. Under PMA-3, non-sentinel exposures are assigned to have partial membership in latent groups in which they have no membership.

Parameter results for this sensitivity analysis are presented in Table 4 and Table 5. First, we present outcome model and latent variable model parameters in Table 4. Performance was fairly consistent across PMA-1 and PMA-2, while $\boldsymbol{\beta_x}$ credible intervals were slightly increased under PMA-3. For all three PMA conditions, parameters were estimated with very little bias and coverage probabilities were close to the 0.95 level.

Estimates for the $\lambda_{j,k}$'s under the different PMAs are presented in Table 5. Overall, there was little difference in performance across PMA-1 and PMA-2. Under PMA-3, parameter bias was slightly higher for some λ parameters. Compared to PMA-1 and PMA-2, credible intervals under PMA-3 were wider, but still considerably tighter than the SEM Best intervals presented previously, in Table 1. Under PMA-3, some λ coverage probabilities were exactly zero when a non-sentinel exposure is partially assigned to a latent group in which it has no membership. Partially assigning exposure j to latent group k requires exposure j to have non-zero membership in latent group k. As a result, it must be the case that $\lambda_{j,k} > 0$.

PMA misspecification may or may not affect the coverage of other $\lambda_{j,k}$ parameters.

Table 4: Simulation results with three latent groups under prior A (Beta(1,1); IG(0.05, 0.01)) for our partial membership model under three different partial membership assignments: PMA-1 (optimal latent assignments), PMA-2 (intermediate assignments), and PMA-3 (inappropriate assignments). (β, σ^2) and (γ, η^2) are slopes and variances for $(y \mid x, Z_y)$ and $(x \mid Z_x)$, respectively. Reported values are averages over 100 simulated datasets. For each model, we fix $g = 15$ and $\sigma_v^2 = 4$.

		Bias			Interval Length			Coverage		
	Truth	PMA-1	PMA-2	PMA-3	PMA-1	PMA-2	PMA-3	PMA-1	PMA-2	PMA-3
β_{x_1}	0.250	-0.013	-0.013	-0.013	0.291	0.291	0.310	0.950	0.940	0.960
β_{x_2}	-0.150	0.028	0.028	0.019	0.314	0.315	0.355	0.950	0.950	0.980
β_{x_3}	0.100	0.004	0.004	0.018	0.317	0.317	0.356	0.950	0.950	0.980
$\beta_{Z_{y\,1}}$	0.150	-0.000	-0.000	-0.000	0.171	0.171	0.171	0.940	0.960	0.940
$\beta_{Z_{y\,2}}$	-0.100	0.001	0.001	0.001	0.171	0.171	0.171	0.960	0.970	0.970
$\beta_{Z_{y\,3}}$	0.050	0.001	0.001	0.001	0.171	0.171	0.171	0.940	0.940	0.940
σ^2	0.918	0.001	0.001	0.001	0.237	0.237	0.238	1.000	1.000	1.000
$\gamma_{1,1}$	0.070	-0.003	-0.003	-0.003	0.148	0.148	0.149	0.940	0.940	0.940
$\gamma_{1,2}$	0.210	0.002	0.002	-0.002	0.147	0.147	0.149	0.960	0.960	0.970
$\gamma_{2,1}$	0.150	-0.001	-0.001	0.009	0.159	0.159	0.159	0.960	0.960	0.940
$\gamma_{2,2}$	0.050	-0.006	-0.006	-0.002	0.159	0.159	0.157	0.970	0.970	0.960
$\gamma_{3,1}$	0.280	0.001	0.001	-0.008	0.153	0.153	0.162	0.960	0.960	0.960
$\gamma_{3,2}$	0.110	0.009	0.009	0.013	0.153	0.153	0.154	0.980	0.980	0.980
η_1^2	0.451	0.027	0.028	0.014	0.210	0.210	0.281	0.990	0.990	1.000
η_2^2	0.475	0.035	0.036	-0.008	0.330	0.330	0.468	0.970	0.970	1.000
η_3^2	0.409	0.022	0.022	0.005	0.295	0.293	0.425	0.970	0.970	0.990

For example, first consider non-sentinel exposure 6, which has partial membership in latent groups 2 and 3, but no membership in latent group 1. If exposure 6 is incorrectly assigned to have partial membership in latent group 1, $\lambda_{6,1}$ cannot include the true value of 0, so the coverage for $\lambda_{6,1}$ is zero. However, the coverage for $\lambda_{6,2}$ and $\lambda_{6,3}$ do not appear to be reduced by this misspecification. Alternatively, consider non-sentinel exposure 2, which has complete membership in latent group 1. If exposure 2 is incorrectly assigned partial membership to latent group 3, the coverage for both $\lambda_{2,1}$ and $\lambda_{2,3}$ will be zero. Since $\lambda_{2,3} > 0$ and since $\sum_{k=1}^{3} \lambda_{2,k} = 1$, it must be the case that $\lambda_{2,1} < 1$. Therefore, the coverage for $\lambda_{2,1}$, which has a true value of 1, must be zero. The coverage for $\lambda_{2,3}$, which has a true value of 0, must also be zero.

Table 5: Non-sentinel λ estimates under prior A (Beta(1,1); IG(0.05, 0.01)) for our partial membership model with three latent groups under three different partial membership assignments (denoted by *) where $\lambda_{j,k}$ is the partial membership of the jth exposure for the kth latent group. Reported values are averages over 100 simulated datasets. For each model, we fix $g = 15$ and $\sigma_v^2 = 4$.

		Bias			Interval Length			Coverage		
	Truth	PMA-1	PMA-2	PMA-3	PMA-1	PMA-2	PMA-3	PMA-1	PMA-2	PMA-3
$\lambda_{2,1}$	1.000	-0.040*	-0.040*	-0.045	0.165*	0.166*	0.269	1.000*	1.000*	0.000
$\lambda_{2,2}$	0.000	0.021	0.021	0.007	0.116	0.118	0.071	1.000	1.000	1.000
$\lambda_{2,3}$	0.000	0.019	0.019	0.038*	0.114	0.114	0.236*	1.000	1.000	0.000*
$\lambda_{3,1}$	0.600	-0.012*	-0.013	0.032	0.237*	0.237	0.578	0.960*	0.960	0.960
$\lambda_{3,2}$	0.400	-0.009	-0.010*	-0.068	0.245	0.247*	0.530	0.960	0.970*	0.970
$\lambda_{3,3}$	0.000	0.021	0.023	0.036*	0.126	0.128	0.207*	1.000	1.000	0.000*
$\lambda_{5,1}$	0.300	0.008	0.005	-0.004	0.257	0.268	0.333	0.920	0.960	0.950
$\lambda_{5,2}$	0.400	0.000*	0.001	0.024	0.299*	0.316	0.448	0.910*	0.920	0.980
$\lambda_{5,3}$	0.300	-0.008	-0.007*	-0.021*	0.306	0.303*	0.408*	0.910	0.930*	0.950*
$\lambda_{6,1}$	0.000	0.028	0.030	0.039*	0.126	0.127	0.182*	0.980	0.990	0.000*
$\lambda_{6,2}$	0.700	-0.047*	-0.049	0.040	0.336*	0.330	0.606	0.940*	0.930	0.990
$\lambda_{6,3}$	0.300	0.019	0.019*	-0.079	0.322	0.318*	0.541	0.990	0.970*	0.960
$\lambda_{7,1}$	0.000	0.024	0.025	0.032*	0.126	0.125	0.168*	1.000	1.000	0.000*
$\lambda_{7,2}$	0.500	-0.011	-0.012*	0.034	0.309	0.310*	0.794	0.960	0.960*	0.980
$\lambda_{7,3}$	0.500	-0.013*	-0.013	-0.066	0.315*	0.310	0.781	0.950*	0.950	0.950
$\lambda_{8,1}$	0.300	0.006	0.005*	-0.080	0.251	0.250*	0.439	0.970	0.970*	0.990
$\lambda_{8,2}$	0.000	0.029	0.031	0.060*	0.147	0.150	0.252*	0.990	1.000	0.000*
$\lambda_{8,3}$	0.700	-0.034*	-0.036	0.021	0.297*	0.294	0.569	0.960*	0.960	0.990

4 Application: Seychelles Child Development Study

4.1 Study Background

The Seychelles Child Development Study (SCDS) is an on-going, multi-center research collaboration studying the association between prenatal mercury exposure from fish consumption and child neurodevelopment. We illustrate our model with prenatal inflammatory marker data collected from Nutrition Cohort 2 (NC2), the largest mother-child cohort ($n = 1518$ eligible mothers) in the SCDS. NC2 mothers were recruited between 2008 and 2011 on Mahé, the main island of Seychelles. For a detailed description of the NC2 cohort, including exclusion and inclusion criteria, see Strain et al. (2015).

Our partial membership model uses data from several observed exposures (i.e. inflammatory markers), covariates, and one outcome, and makes use of multiple latent variables. Since markers of maternal inflammation may influence birth outcomes (Yeates et al., 2020), we select birth weight as the model outcome for our work. We consider

two different sets of inflammatory markers. We first fit our partial membership model to the seven inflammatory markers that were classified as either Th1 (IFN-γ, TNF-α, IL-1β, IL-2) or Th2 (IL-5, IL-10, and IL-4) by McSorley et al. (2018) and Yeates et al. (2020), with two latent variables representing "Th1" and "Th2." This set of exposures was previously analyzed by Zavez et al. (2020) using a Bayesian single membership model with two latent variables. We then expand our analysis to include six additional NC2 inflammatory markers (IL-6, MCP-1, TARC, sFlt-1, VEGF-D, and CRP), with either three or four latent variables. In this larger model, the first and second latent groups again represent "Th1" and "Th2" respectively. The third latent group is meant to represent "Chemokines", a class of non-cytokine factors that may influence both the Th1/Th2 paradigm (Kidd, 2003) and birth weight. In the model with four latent variables, the fourth latent group is meant to represent "placental-fetal" circulatory markers, as some of these exposures (e.g., sFlt-1 and VEGF-D) may influence placental development (Yeates et al., 2020), and placental weight may be an important determinant of both birth weight and fetal growth (Roland et al., 2012).

Previous analyses of these NC2 inflammatory markers (McSorley et al. (2018), Yeates et al. (2020), and Zavez et al. (2020)) used the natural logarithm of each marker (after adding a constant of 1) to better satisfy regression assumptions. We apply the same transformation procedure here. Like the single membership model presented by Zavez et al. (2020), our partial membership model requires covariates associated with the latent variables (denoted with Z_x), and covariates associated with the outcome (denoted with Z_y). To facilitate comparisons between the single membership and partial membership approaches, we utilize the same covariates as Zavez et al. (2020). Specifically, we select maternal BMI, gestational age at time of blood draw, total n6 polyunsaturated fatty acids (PUFA), total n3 PUFA, maternal MeHg, and socioeconomic status as $Z_{x_1}, Z_{x_2}, Z_{x_3}, Z_{x_4}, Z_{x,5}$ and Z_{x_6}, respectively. For Z_y, we select child sex (equal to 1 if male; 0 if female), gestational age at birth, maternal BMI, maternal age, and socioeconomic status as $Z_{y_1}, Z_{y_2}, Z_{y_3}, Z_{y_4}$ and Z_{y_5}, respectively.

4.2 MCMC Details

We fit our model using two different *IG* priors for model variances (prior A: *IG*(0.01, 0.05); prior B: *IG*(0.1, 0.5)) and two different values of *g*. These *IG* priors were also used in Zavez et al. (2020). For all applications, we fix $\sigma^2_v = 1$, as this showed improved performance in our simulation study sensitivity analyses compared to different options for σ^2_v (see Web Tables 1-3 of the Supporting Information for more details). For each Seychelles application, we fit our partial membership model using four MCMC chains with 10,000 iterations per chain. We discard the first 1,000 iterations of each chain as "burn-in" draws, and utilize the remaining iterations for inference. To determine model convergence, we investigate traceplots, compute the Gelman-Rubin diagnostic ($\hat{R}$) and compute effective sample size estimates (n_{eff}) for each model parameter (Gelman et al., 1992; Plummer et al., 2006).

4.3 Seychelles Application for K = 2

We first fit our partial membership model to the Seychelles application with seven inflammatory markers and two latent groups. Like Zavez et al. (2020), we assign IFN-γ as the sentinel cytokine in the first (i.e., "Th1") latent group and IL-4 as the sentinel cytokine in the second (i.e., "Th2") latent group. Parameter estimates and posterior intervals for prior A and prior B (under $g = 5$ and $g = 10$) were similar and are presented in Web Table 6 and Web Table 7 of the Supporting Information.

Estimates and intervals were similar under $g = 5$ and $g = 10$. TNF-α (w_2), IL-1β (w_3) and IL-2 (w_4) had majority membership in the Th1 group with IFN-γ (w_1), while IL-5 (w_5) and IL-10 (w_6) had majority membership in the Th2 group with IL-4 (w_7). Neither latent Th1 nor latent Th2 was associated with child birth weight. Several Z_y covariates (child sex, gestational age at birth, maternal BMI, and maternal age) were positively associated with birth weight. Total n6 PUFA was positively associated with both latent Th1 and latent Th2, while maternal MeHg was negatively associated with both latent Th1 and latent Th2. Socioeconomic status was negatively associated with latent Th2.

Zavez et al. (2020) analyzed the same dataset under the single membership model

(also with two latent groups). In the single membership model, there was also no association between child birth weight and either latent Th1 or Th2. Covariate associations were preserved, with child sex, gestational age at birth, maternal BMI, and maternal age also having positive associations with child birth weight in the single membership model. Associations involving total n6 PUFA, maternal MeHg, and socioeconomic status were also similar. With regards to exposure classifications, the dominant membership (either Th1 or Th2) in the partial membership model generally aligned with the single membership model classification. There was only one difference in that IL-10 (w_6) classified in the Th1 group under the single membership model. The classification of IL-10 as primarily a Th2 cytokine in our partial membership model aligns with prior expectations of IL-10 classification, showing a possible, albeit unexpected, advantage of the partial membership model over the single membership model.

4.4   Seychelles Application for K = 3

We next fit our partial membership model to the Seychelles application with thirteen inflammatory markers and three latent groups. We assign IFN-γ as the sentinel cytokine in the first (i.e., "Th1") latent group, IL-4 as the sentinel cytokine in the second (i.e., "Th2") latent group, and MCP-1 as the sentinel cytokine in the third (i.e., "Chemokines") latent group. Lower maternal concentrations of MCP-1 may be associated with abnormal fetal growth, specifically intra-uterine growth restriction Briana et al. (2007). Select parameter estimates and posterior intervals for prior A (under $g = 5$ and $g = 10$) are presented in Table 6. Like the $K = 2$ model, results under prior B were similar to results under prior A, and are provided in Web Table 7 of the Supporting Information.

[Insert Table 6 about here]

TNF-α (w_2), IL-1β (w_3) and IL-2 (w_4) had majority membership in the Th1 group with IFN-γ (w_1), while IL-5 (w_5), IL-10 (w_6), and IL-6 (w_8) had majority membership in the Th2 group with IL-4 (w_7). TARC (w_{10}), sFLT-1 (w_{11}), VEGF-D (w_{12}) and CRP (w_{13}) had majority membership in the Chemokines group with MCP-1 (w_9). Neither latent

Table 6: Seychelles parameter estimates and 95% posterior intervals with $K = 3$ latent groups under prior A (IG(0.05, 0.01)) with $\rho^2 = 1$ for $g = 5$ and $g = 10$. β, λ, and γ are slopes for $(y \mid x, Z_y)$, $(w \mid x)$, and $(x \mid Z_x)$, respectively, and η^2 are variances for $(x \mid Z_x)$. For identifiability, λ_1, λ_7, and λ_9 are fixed at (1,0,0), (0,1,0), and (0,0,1) respectively. Partial membership assignments for each non-sentinel exposure are denoted with *.

	$g = 5$			$g = 10$		
	Mean	Median	95% Interval	Mean	Median	95% Interval
Coefficients for latent variables on outcome (birth weight)						
β_{Th1}	0.059	0.059	(-0.075, 0.194)	0.065	0.065	(-0.072, 0.204)
β_{Th2}	-0.110	-0.109	(-0.312, 0.088)	-0.117	-0.116	(-0.326, 0.089)
β_{Chem}	-0.289	-0.278	(-0.645, 0.004)	-0.304	-0.293	(-0.665, -0.006)
Partial membership estimates for observed exposure 2 (TNF-α)						
$\lambda_{2,1}$*	0.783	0.794	(0.442, 1.000)	0.823	0.869	(0.473, 1.000)
$\lambda_{2,2}$	0.115	0.000	(0.000, 0.462)	0.096	0.000	(0.000, 0.450)
$\lambda_{2,3}$	0.102	0.004	(0.000, 0.434)	0.081	0.001	(0.000, 0.421)
Partial membership estimates for observed exposure 3 (IL-1β)						
$\lambda_{3,1}$*	0.770	0.775	(0.421, 1.000)	0.811	0.852	(0.455, 1.000)
$\lambda_{3,2}$	0.127	0.000	(0.000, 0.479)	0.106	0.000	(0.000, 0.467)
$\lambda_{3,3}$	0.103	0.000	(0.000, 0.434)	0.083	0.000	(0.000, 0.420)
Partial membership estimates for observed exposure 4 (IL-2)						
$\lambda_{4,1}$*	0.760	0.758	(0.410,1.000)	0.801	0.832	(0.438, 1.000)
$\lambda_{4,2}$	0.132	0.008	(0.000, 0.493)	0.109	0.001	(0.000, 0.478)
$\lambda_{4,3}$	0.108	0.000	(0.000, 0.438)	0.090	0.000	(0.000, 0.428)
Partial membership estimates for observed exposure 5 (IL-5)						
$\lambda_{5,1}$	0.118	0.000	(0.000, 0.459)	0.096	0.000	(0.000, 0.444)
$\lambda_{5,2}$*	0.756	0.749	(0.406, 1.000)	0.798	0.825	(0.436, 1.000)
$\lambda_{5,3}$	0.126	0.000	(0.000, 0.475)	0.106	0.000	(0.000, 0.462)
Partial membership estimates for observed exposure 6 (IL-10)						
$\lambda_{6,1}$	0.133	0.006	(0.000, 0.497)	0.113	0.001	(0.000, 0.490)
$\lambda_{6,2}$*	0.768	0.776	(0.409, 1.000)	0.808	0.851	(0.436, 1.000)
$\lambda_{6,3}$	0.099	0.000	(0.000, 0.439)	0.079	0.000	(0.000, 0.422)
Partial membership estimates for observed exposure 8 (IL-6)						
$\lambda_{8,1}$	0.131	0.004	(0.000, 0.497)	0.111	0.000	(0.000, 0.488)
$\lambda_{8,2}$*	0.768	0.775	(0.403, 1.000)	0.808	0.851	(0.432, 1.000)
$\lambda_{8,3}$	0.101	0.000	(0.000, 0.443)	0.081	0.000	(0.000, 0.430)
Partial membership estimates for observed exposure 10 (TARC)						
$\lambda_{10,1}$	0.101	0.001	(0.000, 0.437)	0.081	0.000	(0.000, 0.428)
$\lambda_{10,2}$	0.098	0.000	(0.000, 0.437)	0.079	0.000	(0.000, 0.426)
$\lambda_{10,3}$*	0.802	0.826	(0.470,1.000)	0.840	0.898	(0.502, 1.000)
Partial membership estimates for observed exposure 11 (sFLT-1)						
$\lambda_{11,1}$	0.118	0.000	(0.000, 0.444)	0.100	0.000	(0.000, 0.435)
$\lambda_{11,2}$	0.116	0.007	(0.000, 0.449)	0.095	0.001	(0.000, 0.440)
$\lambda_{11,3}$*	0.766	0.757	(0.433, 1.000)	0.805	0.830	(0.462, 1.000)
Partial membership estimates for observed exposure 12 (VEGF-D)						
$\lambda_{12,1}$	0.110	0.000	(0.000, 0.438)	0.091	0.000	(0.000, 0.424)
$\lambda_{12,2}$	0.126	0.004	(0.000, 0.474)	0.105	0.000	(0.000, 0.464)
$\lambda_{12,3}$*	0.764	0.757	(0.425, 1.000)	0.804	0.832	(0.454, 1.000)
Partial membership estimates for observed exposure 13 (CRP)						
$\lambda_{13,1}$	0.119	0.000	(0.000, 0.461)	0.099	0.000	(0.000, 0.453)
$\lambda_{13,2}$	0.133	0.005	(0.000, 0.517)	0.112	0.001	(0.000, 0.501)
$\lambda_{13,3}$*	0.748	0.740	(0.383, 1.000)	0.789	0.820	(0.407, 1.000)
Coefficients for covariates on latent group 3 (chemokines)						
$\gamma_{Chem,1}$	-0.018	-0.017	(-0.064, 0.027)	-0.017	-0.017	(-0.063, 0.025)
$\gamma_{Chem,2}$	0.043	0.043	(0.006, 0.079)	0.042	0.042	(0.007, 0.078)
$\gamma_{Chem,3}$	0.031	0.031	(-0.008, 0.070)	0.031	0.031	(-0.007, 0.069)
$\gamma_{Chem,4}$	-0.018	-0.018	(-0.057, 0.021)	-0.019	-0.019	(-0.056, 0.019)
$\gamma_{Chem,5}$	0.051	0.051	(0.015, 0.089)	0.050	0.050	(0.015, 0.086)
$\gamma_{Chem,6}$	-0.029	-0.029	(-0.065, 0.006)	-0.029	-0.029	(-0.064, 0.005)
Variance estimates for latent group models						
η^2_{Th1}	0.378	0.375	(0.310, 0.465)	0.368	0.365	(0.303, 0.450)
η^2_{Th2}	0.240	0.238	(0.187, 0.304)	0.228	0.226	(0.179, 0.288)
η^2_{Chem}	0.107	0.105	(0.069, 0.153)	0.102	0.101	(0.066, 0.145)

Th1, latent Th2, nor latent Chemokines were associated with birth weight in the $K = 3$ model. Associations between the Z_y covariates and birth weight were similar to those in the $K = 2$ model, with child sex, gestational age at birth, maternal BMI, and maternal age all positively associated with birth weight (not shown). In the $K = 3$ model, Z_x covariate associations involving latent Th1 and latent Th2 were similar to those observed in the $K = 2$ model (also not shown). There were significant positive associations between the latent Chemokine group and both gestational age at time of blood draw ($\gamma_{Chem,2}$), and maternal Hg ($\gamma_{Chem,5}$).

4.4.1 Sensitivity Analyses

We performed two sensitivity analyses for the $K = 3$ model. First, we performed the same sensitivity analysis as Zavez et al. (2020) (i.e., selecting IL-1β as the sentinel Th1 cytokine instead of IFN-γ), and found that λ parameter estimates were similar under both models (see Web Table 8 of the Supporting Information for results). In a second sensitivity analysis, we changed the partial membership assignment (PMA) for IL-6 from the "Th2" latent group to the "Th1" latent group. While the IL-6 λ parameter estimates were altered by the change in partial membership, λ estimates for other exposures were not affected, and non-λ model parameter estimates were largely unchanged (see Web Table 9 of the Supporting Information for results).

4.5 Seychelles Application for K = 4

Lastly, we fit our partial membership model to the Seychelles application with the same thirteen inflammatory markers but with four latent groups. We assign IFN-γ as the sentinel cytokine in the first (i.e., "Th1") latent group, IL-4 as the sentinel cytokine in the second (i.e., "Th2") latent group, MCP as the sentinel cytokine in the third (i.e., "Chemokines") latent group, and VEGF-D as the sentinel cytokine in the fourth (i.e., "placental-fetal circulatory") latent group. Independently, VEGF-D was found to be a significant predictor of birth weight in the NC2 cohort (Yeates et al., 2020).

For the $K = 4$ model, traceplots and convergence diagnostics ($\hat{R}$, n_{eff}) indicated poor

convergence for this model, specifically the estimation of β_{x_4} (i.e., the "placental-fetal" slope in the model for birth weight) and η_4^2 (the variance of $x_4 | Z_x$). Increasing the model run time from 10,000 iterations per chain to 40,000 iterations per chain did improve $\hat{R}$ and n_{eff} metrics, but did not improve traceplots or alter estimates of β_{x_4} or η_4^2 (see Web Table 11 of the Supporting Information for details). These complications did not arise in simulated data with 13 observed exposures and three latent groups.

The extremely large estimate of β_{x_4} and extremely small estimate of η_4^2 may indicate overfitting with too many latent groups. Nonidentifiability and overfitting in finite mixture models is a known problem that may cause empty groups (Rousseau and Mengersen, 2011; Love et al., 2017). Since our model assumes each latent group has a sentinel marker with complete membership in that group, an empty group is impossible. However, in the $K = 4$ model, the slope for the fourth latent variable (β_{x_4}) was unrealistically large (≈ -2.57) and posterior variance of the fourth latent variable (x_4) was very close to zero (≈ 0.03). In contrast, the variances of latent x_1, x_2, and x_3 were approximately 0.41, 0.27, and 0.72 respectively, and their slopes were all less than 0.1 in absolute value. The extreme and unrealistic posterior means for β_{x_4} and η_4^2 may be an indication of overfitting in our model when fit with four latent groups.

4.6 Seychelles Application Model Comparison

For the Seychelles application, we computed the Watanabe-Akaike Information Criteria (WAIC; Watanabe and Opper (2010)) for the $K = 3$ and $K = 4$ models, using the formula recommended by Gelman et al. (2014). WAIC for the $K = 3$ model (51,469) was smaller than WAIC for the $K = 4$ model (59,111), suggesting that three latent groups provide a better fit for the Seychelles NC2 application compared to four latent groups.

5 Discussion

We proposed a novel Bayesian modeling approach that permits non-sentinel observed exposures to have membership across one or more latent groups. We compared the per-

formance of our model to two different SEM-based approaches, and applied our model to inflammatory marker data collected in the SCDS NC2 cohort. In simulations, our model estimated outcome model and latent group model parameters with little bias. Parameter intervals were tighter relative to those produced by the competing approaches, and coverage was close to the 95 percent level. Under reasonable partial membership assignments, our partial membership model estimated λ parameters with little bias. Under misspecified partial membership assignments, performance for outcome and latent group model parameters was preserved.

For the SCDS NC2 inflammatory marker application, we find little evidence of an association between child birth weight and any of the latent groups. Covariate associations were similar to those reported by Zavez et al. (2020), and were consistent with earlier Seychelles work (van Wijngaarden et al., 2014; McSorley et al., 2018; Yeates et al., 2020). The λ posterior point estimates suggest that each exposure has majority membership in one latent group. For the $K = 3$ model, one could interpret those exposures with majority membership in the first latent group as representing "Th1 expression," exposures with majority membership in the second latent group as representing "Th2 expression", and exposures with majority membership in the third latent group as representing "Chemokines expression."

Much of the information about these majority latent group memberships is likely coming from correlations between the observed cytokines, some of which are fairly large. However, the posterior intervals for the λ estimates are all quite large, and sometimes intervals for one exposure are actually overlapping between latent groups. For example, TNF-α appears to have majority membership in the first latent group ($\lambda_{2,1} = 0.783$), but the lower bound of the interval for $\lambda_{2,1}$ (0.442) is less than the upper bound of the interval for $\lambda_{2,2}$ (0.462). These non-negligible partial memberships may suggest that either the affected cytokines do not fully belong to one latent group, or that the latent groups are not fully distinct from each other. Consistently large λ posterior intervals were unexpected, and not observed in the simulated data. One possible explanation is that correlations between the NC2 inflammatory markers and child birth weight (range: -0.05 to 0.05)

were generally weaker than correlations between the simulated exposures and simulated outcome variable (range for one simulated dataset: -0.05 to 0.12). In addition, while within-group correlations for the simulated data were moderately positive (ρ's range: 0.22 to 0.50), many of the within-group correlations among the Seychelles non-Th1/Th2 exposures were weaker (ρ's range: -0.1 to 0.20). The one exception was MCP-1 and TARC, which were strongly correlated ($\rho = 0.65$). For the Xiao et al. (2014) Seychelles application, outcomes within the same domain were positively correlated, though the correlation strengths varied (ρ's range: 0.03 to 0.77).

In addition to considering the posterior intervals for the λ estimates, researchers may find it beneficial to compute the posterior probability that an observed exposure's partial membership in one latent group is greater than the observed exposure's partial membership in a different latent group. Since our partial membership model uses a Bayesian approach to obtain posterior draws for all λ parameters, this is straightforward to calculate for any set of λ's. For example, in the $K = 3$ Seychelles application, the posterior probability that TNF-α has majority membership in the first latent group (i.e., $Pr[\lambda_{2,1} > \lambda_{2,2}$ and $\lambda_{2,1} > \lambda_{2,3} \mid$ data]) was 99.2%.

Our $\lambda_{j,k}$ parameterization and corresponding prior specifications permit $\lambda_{j,k}$ values to be zero. While zero-valued $\lambda_{j,k}$ parameters were estimated accurately in simulations, in the NC2 application our model produced several small $\lambda_{j,k}$ parameter estimates for which the posterior median was non-zero. This also occurred, albeit less frequently, in the Seychelles application considered by Xiao et al. (2014) who modeled partial memberships of outcomes rather than exposures. Upon further investigation, we believe the differences in the number of non-zero medians may be partially attributable to differences between the simulated and Seychelles correlation structures, as discussed above.

Our proposed partial membership model has three general requirements. First, the number of latent groups must be specified. For the Seychelles application with thirteen observed exposures, we considered a model with three latent groups ($K = 3$), as well as a model with four latent groups ($K = 4$). While there was scientific rationale for both of these approaches, convergence diagnostics for the $K = 4$ model suggest evidence

of possible model misspecification in the form of overfitting (Rousseau and Mengersen, 2011; Love et al., 2017).

In simulated data with three latent groups ($K = 3$), our sensitivity analyses suggest that, when the number of groups is not known, it may be better to overfit ($K = 4$) rather than underfit ($K = 2$) when specifying K. The $K = 2$ model produced relatively large bias and poor coverage for β_{x_2}, which may alter the overall interpretation of the association between the second latent group and the model outcome. In comparison, the $K = 4$ model poorly estimated several variance parameters and convergence diagnostics for the empty group were poor.

To determine the appropriate K for a given application, one can first fit the partial membership model for two (or more) values of K and compare the candidate models using a model fit statistic, such as WAIC. To help determine the optimal model, we also recommend investigating convergence diagnostics (e.g., traceplots, $\hat{R}$, n_{eff}). For example, the Seychelles $K = 4$ traceplots and convergence metrics for the parameters related to the fourth latent group (i.e., β_{x_4}, η_4^2) were concerning (see Web Appendix G of the Supporting Information for details), as they were for the $K = 4$ simulation with three true latent groups. There were no convergence diagnostic concerns in the $K = 3$ model fit to simulated data with three true latent groups, or in the Seychelles $K = 3$ model. Overall, both the WAIC results and diagnostic checks indicate that the partial membership model with three latent groups provides a better fit than the partial membership model with four latent groups for this Seychelles NC2 application with 13 observed exposures.

Second, our partial membership model requires one observed exposure to be identified as the sentinel exposure for each latent group. While this was also a requirement of the single membership model (Zavez et al., 2020), single membership model estimates were not sensitive to choice of sentinel marker. Similarly, for our partial membership model, sensitivity analyses in the NC2 application showed similar parameter estimates under a different sentinel Th1 cytokine. While both sentinel marker arrangements considered for the NC2 application were scientifically motivated, our simulation study provides evidence

that parameter estimates can be sensitive to sentinel marker assignments, particularly when exposures with the weakest relationship with the outcome are selected as sentinel markers. Researchers should be aware of this limitation, and may want to consider an alternative method if sentinel marker assignments cannot be made confidently.

Lastly, our partial membership model requires one partial membership assignment (PMA) for each non-sentinel exposure. This is also the approach taken by Xiao et al. (2014) for non-sentinel outcomes and is reasonable when prior information about assigning the exposures to latent groups is available. While other partial membership approaches have been proposed, those methods often require further assumptions regarding the membership assignments for each exposure. For example, in a Bayesian SEM model, Muthén and Asparouhov (2012) assume both a primary membership and smaller partial memberships for each exposure by placing priors with mean 0 and small variance on cross-loadings for all partial memberships. However, the primary membership and all non-zero secondary partial memberships must be specified *a priori* for each exposure.

Our PMAs were based largely on earlier inflammatory marker work in the NC2 cohort (McSorley et al., 2018; Yeates et al., 2020) as well as additional work on Th1/Th2 expression (Stangou et al., 2016) and Th1/Th2 differentiation (Diehl and Rincón, 2002). While IL-6 is often classified as 'Th2" (Stangou et al., 2016), it is frequently referred to as pro-inflammatory due to its role in innate immune responses to pathogens (Sykes et al., 2012; Tanaka et al., 2014). Our simulation sensitivity analyses indicate that estimates of exposure model slope parameters (i.e. partial memberships, λ's) are generally not sensitive to PMA specifications. However, the NC2 IL-6 sensitivity results indicate a potential limitation of our model, in that λ parameter estimates for a given exposure may be strongly influenced by PMA in applications with weaker exposure correlation structures.

The partial membership model may be utilized when sufficient information is available to assign sentinel markers and establish PMAs. However, researchers should consider model sensitivity to misspecification when reporting and interpreting results. For some applications, it may be appropriate to first consider a model where each exposure belongs

to exactly one group. For example, grouped WQS regression (Wheeler et al., 2021) may be used when group memberships of each exposure are known *a priori*. Under grouped WQS regression, the direction and magnitude of association is permitted to vary across exposure groups. Multiple index models (McGee et al., 2021) is another alternative that estimates the relationship between groups of exposures and an outcome, while allowing a possibly non-linear exposure-response surface. Like grouped WQS, this method requires that the partitioning of exposures into groups is known. If the exposure group memberships are not known *a priori*, the single membership model of Zavez et al. (2020) can be considered. In that model, MCMC sampling is used to estimate the group membership for each non-sentinel exposure.

In future work, we intend to investigate the feasibility of a partial membership model that does not require partial membership assignments (PMAs). Xiao et al. (2014) briefly discuss a similar extension for multiple outcomes, which involves modifying the $z_{j,k}$ and $p_{j,k}$ priors. This may be advantageous for applications in which the PMAs are largely uncertain, or for applications such as the NC2 inflammatory markers, which show some sensitivity to PMA specification. Application of our partial membership model to data applications for which multiple measured exposures are believed to nest in (or across) two or more latent groups will be helpful for developing recommendations regarding applications that are best suited for the partial membership modeling approach.

References

Abi Berger. Science commentary: Th1 and Th2 responses: what are they? *Allergy*, 55: 2–10, 2000.

Jennifer F Bobb, Linda Valeri, Birgit Claus Henn, David C Christiani, Robert O Wright, Maitreyi Mazumdar, et al. Bayesian kernel machine regression for estimating the health effects of multi-pollutant mixtures. *Biostatistics*, 16(3):493–508, 2014.

Despina D Briana, Maria Boutsikou, Stavroula Baka, George Papadopoulos, Dimitrios

Gourgiotis, et al. Perinatal plasma monocyte chemotactic protein-1 concentrations in intrauterine growth restriction. *Mediators Inflamm*, 2007.

Caroline Carrico, Chris Gennings, David C Wheeler, and Pam Factor-Litvak. Characterization of weighted quantile sum regression for highly correlated data in a risk analysis setting. *J Agric Biol Environ Stat*, 20(1):100–120, 2015.

Raymond J Carroll, David Ruppert, Leonard A Stefanski, and Ciprian M Crainiceanu. *Measurement Error in Nonlinear Models: a Modern Perspective*. Chapman and Hall/CRC, 2006.

JM Cavaillon. Pro-versus anti-inflammatory cytokines: myth or reality. *Cell Mol Biol (Noisy-le-grand, France)*, 47(4):695, 2001.

Sean Diehl and Mercedes Rincón. The two faces of IL-6 on Th1/Th2 differentiation. *Mol Immunol*, 39(9):531–536, 2002.

Mei-Rong Du, Song-Cun Wang, and Da-Jin Li. The integrative roles of chemokines at the maternal–fetal interface in early pregnancy. *Cell Mol Immunol*, 11(5):438–448, 2014.

Andrew Gelman, Donald B Rubin, et al. Inference from iterative simulation using multiple sequences. *Stat Sci*, 7(4):457–472, 1992.

Andrew Gelman, Jessica Hwang, and Aki Vehtari. Understanding predictive information criteria for bayesian models. *Statistics and computing*, 24(6):997–1016, 2014.

Kamel Jedidi, Harsharanjeet S Jagpal, and Wayne S DeSarbo. Finite-mixture structural equation models for response-based segmentation and unobserved heterogeneity. *Mark Sci*, 16(1):39–59, 1997.

Heather H Keller. The SCREEN I (Seniors in the Community: Risk Evaluation for Eating and Nutrition) index adequately represents nutritional risk. *J Clin Epidemiol*, 59(8):836–841, 2006.

P. Kidd. Th1/Th2 balance: the hypothesis, its limitations, and implications for health and disease. *Altern Med Rev*, 8(3):223–46, 2003. ISSN 1089-5159 (Print) 1089-5159 (Linking). URL https://www.ncbi.nlm.nih.gov/pubmed/12946237.

Sik-Yum Lee and Xin-Yuan Song. Bayesian model selection for mixtures of structural equation models with an unknown number of components. *Br J Math Stat Psychol*, 56 (1):145–165, 2003.

Tanzy MT Love, Sally W Thurston, and Philip W Davidson. Finding vulnerable subpopulations in the Seychelles Child Development Study: effect modification with latent groups. *Stat Methods Med Res*, 26(2):809–822, 2017.

Glen McGee, Ander Wilson, Thomas F Webster, and Brent A Coull. Bayesian multiple index models for environmental mixtures. *Biometrics*, 2021.

Emeir M McSorley, Alison J Yeates, Maria S Mulhern, Edwin van Wijngaarden, Katherine Grzesik, Sally W Thurston, et al. Associations of maternal immune response with MeHg exposure at 28 weeks' gestation in the Seychelles Child Development Study. *Am J Reprod Immunol*, 80(5):e13046, 2018.

Toby J Mitchell and John J Beauchamp. Bayesian variable selection in linear regression. *J Amer Statist Assoc*, 83(404):1023–1032, 1988.

Bengt Muthén and Tihomir Asparouhov. Bayesian structural equation modeling: a more flexible representation of substantive theory. *Psychol Methods*, 17(3):313, 2012.

Martyn Plummer, Nicky Best, Kate Cowles, and Karen Vines. Coda: Convergence diagnosis and output analysis for mcmc. *R News*, 6(1):7–11, 2006. URL https://journal.r-project.org/archive/.

Marie Cecilie Paasche Roland, Camilla M Friis, Nanna Voldner, Kristin Godang, Jens Bollerslev, Guttorm Haugen, and Tore Henriksen. Fetal growth versus birthweight: the role of placenta versus other determinants. *PloS One*, 7(6):e39324, 2012.

Judith Rousseau and Kerrie Mengersen. Asymptotic behaviour of the posterior distribution in overfitted mixture models. *J R Stat Soc Ser B Methodol*, 73(5):689–710, 2011.

Brisa N Sánchez, Esben Budtz-Jørgensen, Louise M Ryan, and Howard Hu. Structural equation models: a review with applications to environmental epidemiology. *J Am Stat Assoc*, 100(472):1443–1455, 2005.

Jürgen Scheller, Athena Chalaris, Dirk Schmidt-Arras, and Stefan Rose-John. The pro- and anti-inflammatory properties of the cytokine interleukin-6. *Biochim Biophys Acta Mol Cell Res*, 1813(5):878–888, 2011.

M Stangou, C Bantis, M Skoularopoulou, L Korelidou, D Kouloukouriotou, M Scina, et al. Th1, Th2 and Treg/T17 cytokines in two types of proliferative glomerulonephritis. *Indian J Nephrol*, 26(3):159, 2016.

Matthew Stephens. Dealing with label switching in mixture models. *J R Stat Soc Series B Stat Methodol*, 62(4):795–809, 2000.

JJ Strain, Alison J Yeates, Edwin van Wijngaarden, Sally W Thurston, Maria S Mulhern, Emeir M McSorley, et al. Prenatal exposure to methyl mercury from fish consumption and polyunsaturated fatty acids: associations with child development at 20 mo of age in an observational study in the Republic of Seychelles. *Am J Clin Nutr*, 101(3):530–537, 2015.

L. Sykes, D. A. MacIntyre, X. J. Yap, T. G. Teoh, and P. R. Bennett. The Th1:Th2 dichotomy of pregnancy and preterm labour. *Mediators Inflamm*, 2012:967629, 2012. ISSN 1466-1861 (Electronic) 0962-9351 (Linking). doi:10.1155/2012/967629. URL https://www.ncbi.nlm.nih.gov/pubmed/22719180http://spiral.imperial.ac.uk/bitstream/10044/1/12876/.

Toshio Tanaka, Masashi Narazaki, and Tadamitsu Kishimoto. Il-6 in inflammation, immunity, and disease. *Cold Spring Harb Perspect Biol*, 6(10):a016295, 2014.

Edwin van Wijngaarden, Donald Harrington, Roni Kobrosly, Sally W Thurston, Todd O'Hara, Emeir M McSorley, et al. Prenatal exposure to methylmercury and LCPUFA in relation to birth weight. *Ann Epidemiol*, 24(4):273–278, 2014.

Sumio Watanabe and Manfred Opper. Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *J Mach Learn Res*, 11(12), 2010.

Jennifer Weuve, Brisa N Sánchez, Antonia M Calafat, Ted Schettler, Ronald A Green, Howard Hu, et al. Exposure to phthalates in neonatal intensive care unit infants: urinary concentrations of monoesters and oxidative metabolites. *Environ Health Perspect*, 114(9):1424–1431, 2006.

David C Wheeler, Salem Rustom, Matthew Carli, Todd P Whitehead, Mary H Ward, and Catherine Metayer. Assessment of grouped weighted quantile sum regression for modeling chemical mixtures and cancer risk. *Int J Environ Res Public Health*, 18(2): 504, 2021.

Luo Xiao, Sally W Thurston, David Ruppert, Tanzy MT Love, and Philip W Davidson. Bayesian models for multiple outcomes in domains with application to the Seychelles Child Development Study. *J Am Stat Assoc*, 109(505):1–10, 2014.

Alison J Yeates, Emeir M McSorley, Maria S Mulhern, Toni Spence, William Crowe, Katherine Grzesik, SW Thurston, GE Watson, GJ Myers, PW Davidson, et al. Associations between maternal inflammation during pregnancy and infant birth outcomes in the Seychelles Child Development Study. *J Reprod Immunol*, 137:102623, 2020.

Nengjun Yi, Varghese George, and David B Allison. Stochastic search variable selection for identifying multiple quantitative trait loci. *Genetics*, 164(3):1129–1138, 2003.

Alexis Zavez, Emeir M McSorley, Alison J Yeates, and Sally W Thurston. Modeling the effects of multiple exposures with unknown group memberships: a Bayesian latent variable approach. *J Appl Stat*, pages 1–27, 2020.

Supporting Information

Web Appendices and Web Tables, referenced in Sections 2, 3, 4, and 5 are available online.