

Nonparametric Bayes models for mixed-scale longitudinal surveys

T. KUNIHAMA ^{*} C. T. HALPERN [†] A. H. HERRING [‡]

June, 2016

Abstract

Modeling and computation for multivariate longitudinal surveys have proven challenging, particularly when data are not all continuous and Gaussian but contain discrete measurements. In many social science surveys, study participants are selected via complex survey designs such as stratified random sampling, leading to discrepancies between the sample and population, which are further compounded by missing data and loss to follow up. Survey weights are typically constructed to address these issues, but it is not clear how to include them in models. Motivated by data on sexual development, we propose a novel nonparametric approach for mixed-scale longitudinal data in surveys. In the proposed approach, the mixed-scale multivariate response is expressed through an underlying continuous variable with dynamic latent factors inducing time-varying associations. Bias from the survey design is adjusted for in posterior computation relying on a Markov chain Monte Carlo algorithm. The approach is assessed in simulation studies, and applied to the National Longitudinal Study of Adolescent to Adult Health.

Key words: Mixed-scale data; Mixture model; Multivariate longitudinal data; Nonparametric Bayes; Sample survey.

1 Introduction

The National Longitudinal Study of Adolescent to Adult Health (Add Health) has collected vast amounts of information about social, economic and health-related behaviours in adolescence over 20 years. One of the primary aims is to study trajectories of adolescent sexual development into adulthood. In particular, it is of interest to investigate associations among

^{*}Assistant Professor, Department of Economics, Nagoya University, Nagoya, Aichi, Japan. Email: kunihama@soec.nagoya-u.ac.jp

[†]Professor, Department of Maternal and Child Health, Gillings School of Global Public Health, and Carolina Population Center, The University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, U.S.A. Email: carolyn_halpern@unc.edu

[‡]Carol Remmer Angle Distinguished Professor of Children's Environmental Health, Department of Biostatistics, Gillings School of Global Public Health, and Carolina Population Center, The University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, U.S.A. Email: aherring@bios.unc.edu

different indicators of sexual orientation, including adolescent sexual behaviours, attraction and identity; changes in the associations over the transition to adulthood, and whether those changes may vary by demographic characteristics. Although Add Health data include all the sexual development indicators of interest, several characteristics of the data make it challenging to build a realistic statistical model.

First, the data consist of mixed-scale variables, e.g. partner counts or nominal sexual orientation identity. It is not straightforward to model mixed-scale multivariate data jointly, especially if they include both ordered variables and nominal variables. Second, individuals repeatedly answered a set of questions over time, leading to subject-specific time-dependence. Third, the Add Health study drew its sample via a complex survey design with stratified sampling of schools and oversampling of numerous groups. Hence, the data are not representative of the population until sampling weights are applied. Fourth, it is not efficient to build a statistical model separately for each subpopulation, because the number of respondents in certain groups can be relatively small, especially when considering weight adjustments that account for loss to follow-up. In order to make stable population-representative inferences and use small area estimation, it is essential to include information for all subpopulations. Fifth, missing values are common, and we observe both design-based and individual-specific missingness. For example, students were not asked as detailed questions on their sexuality in adolescence as in young adulthood by design. Therefore, we need to conduct statistical analysis of adolescent sexual development, taking fully into account all of these challenging characteristics of the Add Health data.

There is a rich literature on modeling of mixed-scale data. One approach is to apply generalized linear models for each outcome in which dependence among variables is induced through shared latent factors (Sammel et al. (1997); Dunson (2000); Moustaki and Knott (2000); Dunson and Herring (2005)). However, the robustness of the approaches based on generalised linear mixed models can be weak due to the dual role of the random effects structure in controlling the dependence and shape of marginal distributions. Another approach is to use underlying continuous variables, specified by a Gaussian model (Muthén (1984)) or a Dirichlet process mixture model (Kottas (2005); Canale and Dunson (2011); Kim and Ratchford (2013)). In this approach, discrete variables can be expressed by thresholding the

latent continuous variables. Avoiding specification of marginal distributions, Hoff (2007) proposed a semiparametric Gaussian copula model in which associations among mixed-scale variables are induced by correlations among the latent Gaussian variables. Murray et al. (2013) and Gruhl et al. (2013) extend this approach by incorporating factor structures but these copula models can incorporate only ordered variables. Also, Murray and Reiter (2015) propose a nonparametric Bayesian joint model for multiple imputation of missing values. McParland et al. (2014) develop a model-based clustering approach for mixed-scale data that combines item response theory models for ordered variables with factor analysis models for nominal variables.

Recently, a number of articles have addressed the analysis of multivariate longitudinal data (Bandyopadhyay et al. (2011), Verbeke et al. (2014)). Dunson (2003) proposed dynamic latent trait models in which autoregressive Gaussian latent factors incorporate subject-specific time-dependence and generalized linear models describe mixed-scale outcomes. It is routine to incorporate time effects into multivariate models through time-varying covariates, such as polynomial functions of age, with random coefficients (Gueorguieva and Sanacora (2006); Fieuws and Verbeke (2006); Luo and Wang (2014); Baghfalakia et al. (2014)). Liu et al. (2009) develop a joint model for longitudinal binary and continuous variables, consisting of a correlated probit model and a regression model with the Bartlett decomposition of a covariance matrix. Ghosh and Hanson (2010) propose a semiparametric approach with a mixture of Polya trees for random effect distributions. Das and Daniels (2014) develop a semiparametric model for bivariate sparse longitudinal data, applying a matrix stick-breaking process for a residual covariance matrix. However, it is not clear how to handle survey bias, missingness and small area estimation in these approaches. Therefore, none of these methods can capture the exact nature of longitudinal surveys in Add Health.

In the literature on survey data analysis, two major approaches are used: the design-based approach, which treats outcomes as fixed quantities and the model-based approach, which models outcomes and effectively predicts values for the non-sampled subjects in a population (Little (2004); Levy and Lemeshow (2008); Rao (2011)). Little (2004) and Gelman (2007) show difficulties with current methods in practice and highlight the importance

of including survey weights in model-based analyses. Zheng and Little (2003, 2005) propose a nonparametric method that flexibly models the outcome given inclusion probability using a penalized spline, and Chen et al. (2010) extend the model for binary variables. Si et al. (2015) propose a nonparametric Bayesian model, which jointly models an outcome and survey weights based on a Gaussian process regression. However, these approaches are developed mainly for inference of the population mean of a univariate response, and extensions to mixed-scale longitudinal data are not straightforward.

We propose a flexible nonparametric model for the analysis of Add Health sexual orientation development data. In the proposed approach, mixed-scale variables are expressed through transformation of latent continuous variables, for which a Dirichlet process mixture of Gaussian factor models is developed. For unordered categorical variables, we employ the concept of utilities in multinomial probit models, in which the nominal outcome is a manifestation of underlying continuous utility variables. The subject-specific dynamic variability can be captured by time-varying latent factors via Gaussian processes which can easily incorporate irregular time intervals for each respondent. Inferences valid for the source population are obtained by adjusting the mixture weights in the Dirichlet process mixtures using the survey weights. Because we build a joint model of the response variables and covariates, missing values can be easily imputed assuming missing at randomness. Also, from the proposed joint density, we construct a conditional density on covariates for small area inferences. For posterior computation, we develop an efficient Markov chain Monte Carlo (MCMC) algorithm in which we modify the Dirichlet process mixture with survey weights, taking into account uncertainty in the adjustment process.

Section 2 describes our data set from Add Health. Section 3 proposes a novel approach for mixed-scale longitudinal surveys. Section 4 develops an efficient MCMC algorithm. Section 5 assesses the performance of the proposed approach against competitors through simulation. Section 6 applies the proposed method to the adolescent sexual development questions of interest. Section 7 concludes the article.

2 National Longitudinal Study of Adolescent to Adult Health (Add Health) data

Add Health is a nation-wide longitudinal study of adolescents in grades 7 to 12 in the United States in the 1994-95 academic year (<http://www.cpc.unc.edu/projects/addhealth>). The cohort study has been following up respondents with four waves of in-home interviews conducted in 1994-95, 1996, 2001-02, and 2008-09. It collects data on respondents' social, behavioral, psychological and biological information, allowing researchers to study developmental trajectories across the life course of adolescence into adulthood.

The Add Health study design selects a stratified sample of 80 high schools and 52 middle schools from the United States with unequal probability of selection based on region, urbanicity, school type, ethnic mix, and size. Then, students in each school are stratified by grade and sex. In addition, the study drew supplemental samples, oversampling groups of particular interest based on ethnicity, genetic relatedness to siblings, adoption status, disability and parental education. Taking into account non-response as well as the sample selection, survey weights have been constructed for population-representative inferences via weighting adjustments. The full study design is described by Harris et al. (2009).

To understand adolescent sexual orientation development, our primary goal is studying associations among measures of sexual attraction, behavior, and identity in adolescence and trends in the transition to adulthood. We use three waves of surveys which were conducted when participants were in grades 7-12 in wave 1, young adults in wave 3 and adults in wave 4. Table 1 shows the list of variables in our data set. They are mixed-scale, and the sexual orientation identity is missing by design in wave 1. For longitudinal studies with these three waves, survey weights are constructed for $n = 12,288$ respondents. Figure 1 reports the histogram of the weights on a logarithmic scale.

Figure 2 plots the weighted sexual development variables and age in each wave. Although the total numbers of sex partners are categorized into five groups for the display, we treat them as count data in the real data analysis in Chapter 6. To obtain the weighted data, we simply resampled $n = 12,288$ respondents from the original data set with probability $\tilde{w}_i = w_i / \sum_{j=1}^n w_j$ where w_i is the survey weight for i th respondent. Figure 3 reports the plot of the weighted and unweighted covariates, indicating the gap induced by the sampling

	Type	Categories
Sexual development variables		
Attraction to opposite sex	binary	0. No, 1. Yes
Attraction to same sex	binary	0. No, 1. Yes
Total # of opposite sex partners	count	
Total # of same sex partners	count	
Sexual identity (missing at wave 1)	nominal	1. Heterosexual, 2. Mostly Heterosexual, 3. Mostly Homosexual, 4. Homosexual, 5. Bisexual
Covariates		
Gender	binary	0. Male, 1. Female
Race	nominal	1. Caucasian, 2. Hispanic, 3. African American, 4. Asian, 5. Other race
Parental education	ordered categorical	1. Less than high school, 2. High school diploma/GED, 3. Some college/voc education, 4. College graduate+

Table 1: List of sexual development variables and covariates.

design. For example, racial minorities are oversampled in the survey, leading to the relatively large decreased percentage of Caucasian in the raw data. Although the resampling method works for computing basic summary statistics, it may be inefficient in terms of discarding information of the unselected respondents in the original data set. Also, it is not clear how to incorporate uncertainty about the unsampled subjects in the population. We discuss an alternative adjustment method in Section 3.3, which takes fully into account the information of all respondents in our data set and reflects uncertainty about the unobserved subjects in the population.




Figure 1: Histogram of the logarithm of the survey weight.

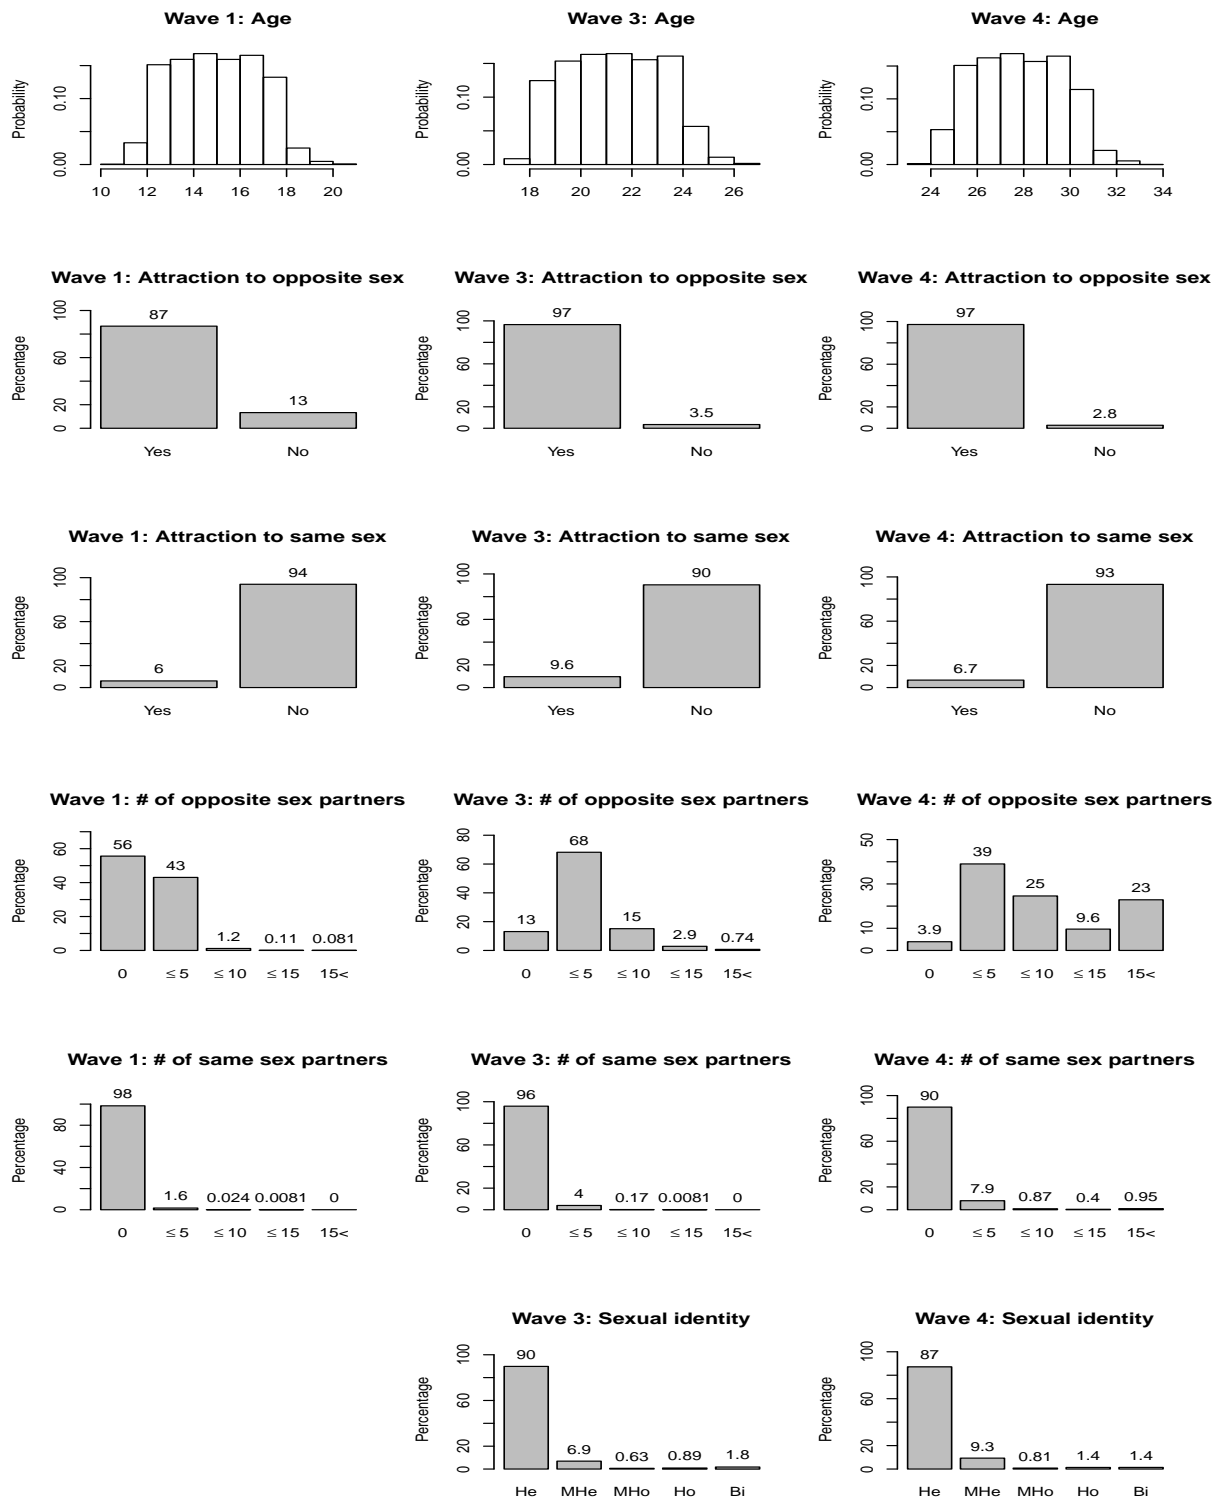


Figure 2: Histogram of age and response variables in wave 1 (left), 3 (middle) and 4 (right). In sex partners, ≤ 5 , ≤ 10 , ≤ 15 mean 1 to 5, 6 to 10, 11 to 15. In sexual identity, He, MHe, MHo, Ho, Bi represent heterosexual, mostly heterosexual, mostly homosexual, homosexual and bisexual.

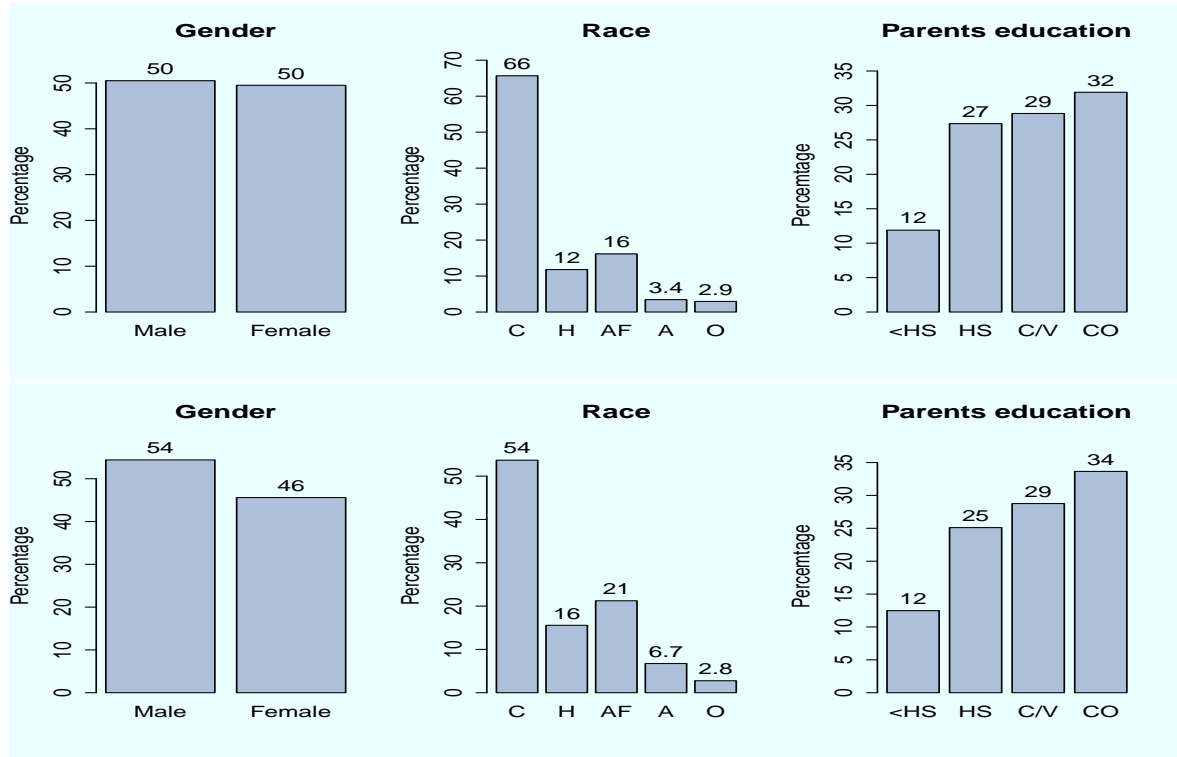


Figure 3: Histogram of the weighted (above) and unweighted (below) covariates. In race, C, H, AF, A, O represent Caucasian, Hispanic, African American, Asian, Other race. In parents education, <HS, HS, C/V, Co correspond to less than high school, high school, college/voc, college graduate.

3 Proposed modeling of mixed-scale longitudinal surveys

3.1 Modeling of mixed-scale data with latent continuous variables

Let $y = (y_1, \dots, y_p)'$ be a multivariate response with mixed-scale margins. We introduce a latent continuous variable $y^* = (y_1^*, \dots, y_p^*)' \in \mathbb{R}^p$ and express the response y by transforming the underlying variable y^* . For example, a binary variable $y_k \in \{0, 1\}$ can be induced by

$$\text{binary : } y_k = 1(y_k^* > 0),$$

where $1(\cdot)$ denotes an indicator function. A count variable $y_k \in \{0, \dots, \infty\}$ can be expressed as

$$\text{count : } y_k = \sum_{r=0}^{\infty} r \times 1(a_r < y_k^* \leq a_{r+1}),$$

where $-\infty = a_0 < a_1 < a_2 < \dots < \infty$. For a nominal variable with d categories, we translate it to d indicator variables y_{k_1}, \dots, y_{k_d} with $y_{k_l} \in \{0, 1\}$ for $l = 1, \dots, d$ and $\sum_{l=1}^d y_{k_l} = 1$. Then, we set

$$\text{nominal : } y_{k_l} = 1(y_{k_l}^* = \max_{1 \leq l' \leq d} y_{k_{l'}}^*), \quad l = 1, \dots, d. \quad (1)$$

In the equation (1), $y_{k_l}^*$ can be interpreted as utility for the l th category as in multinomial probit models (McCulloch and Rossi (1994); Imai and van Dyk (2005); Burgette and Nordheim (2012)). Because the order of $y_{k_1}^*, \dots, y_{k_d}^*$ is unchanged with respect to adding any constant and multiplying any positive scale by all of them, without loss of generality we may assume one of the utilities is zero.

3.2 Proposed nonparametric Bayes model

We apply the above framework to mixed-scale longitudinal data. Let $y_{ij} = (y_{ij1}, \dots, y_{ijp})'$ be the response variable for respondent i at age t_{ij} for $i = 1, \dots, n$ and $j = 1, \dots, n_i$ with $y_{ijk} \in \mathcal{Y}_k$ allowed to be any type of univariate random variable for $k = 1, \dots, p$. We

introduce $y_{ij}^* = (y_{ij1}^*, \dots, y_{ijp}^*)' \in \mathbb{R}^p$ as a latent continuous variable, which induces y_{ij} through a function g such that $y_{ij} = g(y_{ij}^*)$ as in Section 3.1. Also, $x_i = (x_{i1}, \dots, x_{iL})' \in \mathcal{X}$ denotes static covariates for i th respondent, from which target subpopulations may be constructed.

We propose a novel nonparametric Bayes model, in which time-varying factors flexibly capture a dynamic structure of the response by incorporating time effects and covariates into its mean and covariance. Let $\eta_{ij} = (\eta_{ij1}, \dots, \eta_{ijQ})'$ denote dynamic factors for respondent i at age t_{ij} with $Q < p$. We first consider modeling of the latent response variable given the covariates,

$$y_{ij}^* = Bx_i + \Omega\mu_{t_{ij}} + \Lambda\eta_{ij} + \epsilon_{ij}, \quad \epsilon_{ij} \sim N(0, \Sigma), \quad \Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_p^2), \quad (2)$$

where $\mu_t = (\mu_{t1}, \dots, \mu_{tQ_\mu})'$ is a dynamic effect for age t with $Q_\mu < p$, B is a $p \times L$ matrix of regression coefficients, Ω is a $p \times Q_\mu$ coefficient matrix for the time effect, and Λ is a $p \times Q$ factor loading matrix. In this model (2), we can capture age-specific effects for all individuals through μ_t and additional subject-specific time variability is induced by the dynamic factors η_{ij} . Relying on the covariance regression by Hoff and Niu (2012) and Fox and Dunson (2015), we assume that the dynamic factors can be expressed as

$$\eta_{ij} = Vx_i\eta_i^* + \xi_{t_{ij}}\tilde{\eta}_i, \quad \eta_i^* \sim N(0, 1), \quad \tilde{\eta}_i \sim N(0, I_{Q_\eta}), \quad (3)$$

where V is a $Q \times L$ coefficient matrix, $\xi_t = \{\xi_{tql}\}$ is a $Q \times Q_\eta$ matrix for $t = 1, \dots, T$ and η_i^* and $\tilde{\eta}_i$ are random effects. Integrating out the random effects, we obtain a covariance varying with covariates and age,

$$y_{ij}^* \sim N(Bx_i + \Omega\mu_{t_{ij}}, \Lambda Vx_i x_i' V' \Lambda' + \Lambda \xi_{t_{ij}} \xi_{t_{ij}}' \Lambda' + \Sigma). \quad (4)$$

Let $\mu_{\bullet q} = \{\mu_{tq}, t \in \mathcal{T}\}$ and $\xi_{\bullet ql} = \{\xi_{tql}, t \in \mathcal{T}\}$ be stochastic processes of the time effects

with time space \mathcal{T} , for which we apply Gaussian processes,

$$\mu_{\bullet q} \sim \text{GP}(0, c_\mu), \quad c_\mu(\mu_{tq}, \mu_{t'q}) = \text{Cov}(\mu_{tq}, \mu_{t'q}) = \exp(-\kappa_\mu |t - t'|^2), \quad (5)$$

$$\xi_{\bullet ql} \sim \text{GP}(0, c_\xi), \quad c_\xi(\xi_{tql}, \xi_{t'ql}) = \text{Cov}(\xi_{tql}, \xi_{t'ql}) = \exp(-\kappa_\xi |t - t'|^2), \quad (6)$$

where $\kappa_\mu > 0$, $\kappa_\xi > 0$ and $\mu_{\bullet q}$ and $\xi_{\bullet ql}$ are mutually independent with respect to q and l .

Next, relying on the conditional density of the latent variable (2)-(4), we develop a flexible joint model of the response and covariates. The approach of the response-covariate joint modeling has been widely studied in the literature (Müller et al. (1996); Müller and Quintana (2010); Hannah et al. (2011); Taddy and Kottas (2012); DeYoreo and Kottas (2015a,b)). In this framework, we can derive a marginal density $f(y)$ and a conditional density $f(y|x \in C)$ given $C \subset \mathcal{X}$ with covariate space \mathcal{X} from the joint model $f(y, x)$. Also, we can easily impute missing values from the conditional distribution given observed information assuming missing at randomness.

We construct a joint density of y_i^* and x_i with $y_i^* = \{y_{ij}^*, j = 1, \dots, n_i\}$, relying on Dirichlet process mixtures (Lo (1984); West et al. (1994); Escobar and West (1995)),

$$f(y_i^*, x_i) = \int f(y_i^* | x_i, \theta^y) f(x_i | \theta^x) dP(\theta), \quad P \sim \text{DP}(\alpha P_0), \quad (7)$$

$$= \sum_{h=1}^{\infty} \pi_h f(y_i^* | x_i, \theta_h^y) \prod_{l=1}^L f_l(x_{il} | \theta_h^x), \quad (8)$$

$$\pi_h = v_h \prod_{b < h} (1 - v_b), \quad v_h \sim \text{Beta}(1, \alpha), \quad (9)$$

where P denotes a prior distribution for $\theta = (\theta^y, \theta^x)$, $\alpha > 0$ is a precision parameter and P_0 is a base probability measure for the Dirichlet process, and $f(y_i^* | x_i, \theta_h^y)$ denotes the conditional density induced by (2)-(4). With respect to the density for the covariate $f_l(x_{il} | \theta_h^x)$ in (8), we assume a function depending on the type of the covariate, such as a multinomial distribution for a categorical variable. Equations (8) and (9) correspond to the stick-breaking representation of the Dirichlet process mixtures (Sethuraman (1994); Muliere and Tardella (1998)). The proposed model is the Dirichlet process mixture of Gaussian factor models. Finally, integrating out the latent variable over the space derived

from the function in Section 3.1, we obtain the joint density of the response and covariates,

$$f(y_i, x_i) = \int_{R(y_i)} f(y_i^*, x_i) dy_i^*, \quad (10)$$

where $R(y_i) = \{y_i^* \in \mathbb{R}^{p \times n_i} : y_{ij} = g(y_{ij}^*), j = 1, \dots, n_i\}$.

3.3 Population-representative inference on scale-free association

To adjust for the sample bias from complex survey designs, we incorporate survey weights into the proposed joint model. Let D be the finite population with N people, from which n subjects are sampled with inclusion probability ν_i for subject $i \in D$. Inverse of the inclusion probability ν_i^{-1} can be interpreted as the number of people the respondent i represents in the population (Horvitz and Thompson (1952)). Also, the survey weights can be defined as $w_i \propto \nu_i^{-1}$. For the finite approximation of the Dirichlet process mixtures with H components, Kuniyama et al. (2016) develop an adjustment method, which modifies the mixture weights $\{\pi_h\}$ in (8) using the survey weights. In the posterior computation, along with parameter update, we generate adjusted mixture weights $\tilde{\pi} = (\tilde{\pi}_1, \dots, \tilde{\pi}_H)'$ from

$$\tilde{\pi} \sim \text{Dirichlet} \left(a_1 + \sum_{i:s_i=1} \frac{w_i}{c}, \dots, a_H + \sum_{i:s_i=H} \frac{w_i}{c} \right), \quad (11)$$

where $c = \sum_{i=1}^n w_i/N$ and s_i is the cluster index variable for the i th respondent. With the prior distribution $\text{Dirichlet}(a_1, \dots, a_H)$, the posterior distribution (11) consists of the prior sample size, which represents uncertainty about the unsampled subjects in the population, and the enlarged observations based on the survey weights. As a default setting, the prior sample size is 1-2% of population size N . Compared to the posterior distribution $\text{Dirichlet}(a_1 + \sum_i 1(s_i = 1), \dots, a_H + \sum_i 1(s_i = H))$ in a standard Bayesian mixture model, the weight of subject i is changed from 1 to $w_i/c \approx \nu_i^{-1}$.

Let $\tilde{f}(y, x)$ be the adjusted joint density with the mixture weights $\tilde{\pi}$ from (11). Relying on the adjusted density \tilde{f} , we estimate associations of mixed-scale variables by generating samples from the posterior predictive distributions. As a scale-free measure of dependence between two variables, we employ Goodman and Kruskal's gamma (Goodman and Kruskal (1954); Goodman and Kruskal (1959); Goodman and Kruskal (1963); Goodman and Kruskal

(1972)),

$$\gamma = \frac{N_c - N_d}{N_c + N_d}, \quad (12)$$

where N_c is the number of concordant pairs, and N_d is the number of discordant pairs. A concordant pair can be defined as a pair of variables (X_1, Y_1) and (X_2, Y_2) such that $\text{sgn}(X_2 - X_1) = \text{sgn}(Y_2 - Y_1)$. On the other hand, a discordant pair means $\text{sgn}(X_2 - X_1) = -\text{sgn}(Y_2 - Y_1)$. The measure γ ranges from -1 (100% negative association) to 1 (100% positive association), and zero indicates the absence of association. Although other popular rank correlations such as Spearman's rho need an adjustment for the ties, which are neither concordant nor discordant pairs, Goodman and Kruskal's gamma is resistant to them.

To estimate associations for the subpopulation defined by $C \subset \mathcal{X}$, we generate \tilde{y}_r from $\tilde{f}(y | x \in C) = \tilde{f}(y, x \in C) / \tilde{f}(x \in C)$ with $r = 1, \dots, R$ where $\tilde{y}_r = \{\tilde{y}_{rt}, t \in \mathcal{T}\}$ and $\tilde{y}_{rt} = (\tilde{y}_{rt1}, \dots, \tilde{y}_{rtp})'$ is a mixed-scale response for time t . Then, we compute the concordant and discordant pairs for \tilde{y}_{rtj} and $\tilde{y}_{rtj'}$ with $j \neq j'$ and calculate the Goodman and Kruskal's gamma through (12) for each time point. For the overall population, we can estimate the gamma similarly by setting $C = \mathcal{X}$.

4 Posterior computation

Relying on the blocked Gibbs sampler by Ishwaran and James (2001), we develop an efficient MCMC algorithm for the proposed model. Letting $U_h = [B_h \ \Omega_h \ \Lambda_h]$ for the h th mixture component, we apply shrinkage priors for it with high density around zero to reduce effects of redundant elements while the tails are heavy enough to capture important signals. We assume $U_{hkl} \sim N(0, \delta_{kl}^2)$, $\delta_{kl}^2 \sim \text{IG}(0.5, 0.5)$ where IG denotes an inverse-gamma distribution for $k = 1, \dots, p$ and $l = 1, \dots, L^*$ with $L^* = L + Q_\mu + Q$. By integrating out δ_{kl}^2 , the prior corresponds to a Cauchy prior which is a commonly-used shrinkage prior with heavy tails. We also apply the same type of shrinkage prior to V_h , that is, $V_{hql} \sim N(0, \zeta_{ql}^2)$, $\zeta_{ql}^2 \sim \text{IG}(0.5, 0.5)$ for $q = 1, \dots, Q$ and $l = 1, \dots, L$. As a notation below, let $A_{k\bullet}$, $A_{\bullet l}$ and $a_{(-k)}$ be the k th row and l th column of a matrix A and a subvector of a vector a excluding k th component. Also, \hat{y}_{ij} denotes a standardized latent response, which is differently defined

at a different step. Then, we propose the following MCMC algorithm.

1. Update v_h from $\text{Beta}(1 + \sum_{i=1}^n 1(s_i = h), \alpha + \sum_{i=1}^n 1(s_i > h))$ for $h = 1, \dots, H - 1$.
2. Using the prior distribution $\text{Gamma}(a_\alpha, b_\alpha)$, update α from $\text{Gamma}(a_\alpha + H - 1, b_\alpha - \sum_{h=1}^{H-1} \log(1 - v_h))$.
3. Update s_i for $i = 1, \dots, n$ from

$$P(s_i = h \mid \dots) = \frac{\pi_h \prod_{j=1}^{n_i} f(y_{ij}^* \mid x_i, \eta_{ij}, \theta_h^y) \prod_{l=1}^L f_l(x_{il} \mid \theta_h^x)}{\sum_{m=1}^H \pi_m \prod_{j=1}^{n_i} f(y_{ij}^* \mid x_i, \eta_{ij}, \theta_m^y) \prod_{l=1}^L f_l(x_{il} \mid \theta_m^x)}.$$

4. Using the prior $\text{IG}(\tilde{a}_k, \tilde{b}_k)$, update σ_{hk}^2 for $k = 1, \dots, p$ and $h = 1, \dots, H$ from

$$\text{IG}\left(\frac{\sum_{i:s_i=h} n_i + \tilde{a}_k}{2}, \frac{\sum_{i:s_i=h} \sum_{j=1}^{n_i} \hat{y}_{ijk}^2 + \tilde{b}_k}{2}\right),$$

where $\hat{y}_{ijk} = y_{ijk}^* - B_{hk\bullet}x_i - \Omega_{hk\bullet}\mu_{t_{ij}} - \Lambda_{hk\bullet}\eta_{ij}$.

5. Update $U_{hk\bullet}$ for $k = 1, \dots, p$ and $h = 1, \dots, H$ from $N(\mu_U, \Sigma_U)$ where

$$\mu_U = \Sigma_U \left(\sigma_{hk}^{-2} \sum_{i:s_i=h} \sum_{j=1}^{n_i} z_{ij} y_{ijk}^* \right), \quad \Sigma_U = \left(\sigma_{hk}^{-2} \sum_{i:s_i=h} \sum_{j=1}^{n_i} z_{ij} z_{ij}' + \Sigma_0^{-1} \right)^{-1},$$

where $z_{ij} = (x_i', \mu_{t_{ij}}', \eta_{ij}')'$ and $\Sigma_0 = \text{diag}(\delta_{k1}^2, \dots, \delta_{kL}^2)$.

6. Update η_i^* from $N(\mu_{\eta^*}, \sigma_{\eta^*}^2)$ for $i = 1, \dots, n$ where

$$\mu_{\eta^*} = \sigma_{\eta^*}^2 x_i' V_{s_i}' \Lambda_{s_i}' \Sigma_{s_i}^{-1} \sum_{j=1}^{n_i} \hat{y}_{ij}, \quad \sigma_{\eta^*}^2 = (n_i x_i' V_{s_i}' \Lambda_{s_i}' \Sigma_{s_i}^{-1} \Lambda_{s_i} V_{s_i} x_i + 1)^{-1},$$

where $\hat{y}_{ij} = y_{ij}^* - B_{s_i}x_i - \Omega_{s_i}\mu_{t_{ij}} - \Lambda_{s_i}\xi_{t_{ij}}\tilde{\eta}_i$.

7. Update $\tilde{\eta}_i$ from $N(\mu_{\tilde{\eta}}, \Sigma_{\tilde{\eta}})$ for $i = 1, \dots, n$ where

$$\mu_{\tilde{\eta}} = \Sigma_{\tilde{\eta}} \left(\sum_{j=1}^{n_i} \xi_{t_{ij}}' \Lambda_{s_i}' \Sigma_{s_i}^{-1} \hat{y}_{ij} \right), \quad \Sigma_{\tilde{\eta}} = \left(\sum_{j=1}^{n_i} \xi_{t_{ij}}' \Lambda_{s_i}' \Sigma_{s_i}^{-1} \Lambda_{s_i} \xi_{t_{ij}} + I \right)^{-1},$$

with $\hat{y}_{ij} = y_{ij}^* - B_{s_i}x_i - \Omega_{s_i}\mu_{t_{ij}} - \Lambda_{s_i}V_{s_i}x_i\eta_i^*$.

8. Using the Griddy-Gibbs sampler by Ritter and Tanner (1992), update κ_μ from

$$P(\kappa_\mu = c_k \mid \dots) = \frac{\prod_{q=1}^{Q_\mu} f_N(\mu_{\bullet q} \mid 0, \Psi(c_k))}{\sum_{l=1}^G \prod_{q=1}^{Q_\mu} f_N(\mu_{\bullet q} \mid 0, \Psi(c_l))},$$

where c_1, \dots, c_G are grid points and $f_N(\cdot \mid 0, \Psi(c_k))$ denotes a normal density with mean 0 and covariance $\Psi(c_k)$, which is derived from the Gaussian process with the length-scale c_k .

9. Using the Griddy-Gibbs sampler, update κ_ξ from

$$P(\kappa_\xi = c_k \mid \dots) = \frac{\prod_{q=1}^Q \prod_{l=1}^{Q_\eta} f_N(\xi_{\bullet ql} \mid 0, \Psi(c_k))}{\sum_{j=1}^G \prod_{q=1}^Q \prod_{l=1}^{Q_\eta} f_N(\xi_{\bullet ql} \mid 0, \Psi(c_j))}.$$

10. Update $V_{h\bullet l}$ for $h = 1, \dots, H$ and $l = 1, \dots, L$ from $N(\mu_V, \Sigma_V)$ with

$$\mu_V = \Sigma_V \Lambda'_h \Sigma_h^{-1} \left(\sum_{i:s_i=h} x_{il} \eta_i^* \sum_{j=1}^{n_i} \hat{y}_{ij} \right), \quad \Sigma_V = \left(\sum_{i:s_i=h} n_i x_{il}^2 \eta_i^{*2} \Lambda'_h \Sigma_h^{-1} \Lambda_h + \Sigma_{0V}^{-1} \right)^{-1},$$

where $\hat{y}_{ij} = y_{ij}^* - B_h x_i - \Omega_h \mu_{t_{ij}} - \Lambda_h \xi_{t_{ij}} \tilde{\eta}_i - \Lambda_h V_{h\bullet(-l)} x_{i(-l)} \eta_i^*$ and $\Sigma_{0V} = \text{diag}(\zeta_{1L}^2, \dots, \zeta_{QL}^2)$

11. Update $\mu_{\bullet q} = (\mu_{1q}, \dots, \mu_{Tq})'$ for $q = 1, \dots, Q_\mu$ from $N(\mu_\mu, \Sigma_\mu)$ with $\mu_\mu = \Sigma_\mu b$, $\Sigma_\mu^{-1} = A + \Psi^{-1}(\kappa_\mu)$, $A = \text{diag}(a_1, \dots, a_T)$, $b = (b_1, \dots, b_T)'$ and

$$a_t = \sum_{(i,j):t_{ij}=t} \Omega'_{s_i \bullet q} \Sigma_{s_i}^{-1} \Omega_{s_i \bullet q}, \quad b_t = \sum_{(i,j):t_{ij}=t} \Omega'_{s_i \bullet q} \Sigma_{s_i}^{-1} \hat{y}_{ij},$$

where $\hat{y}_{ij} = y_{ij}^* - B_{s_i} x_i - \Lambda_{s_i} \eta_{ij} - \Omega_{s_i \bullet (-q)} \mu_{t_{ij}(-q)}$.

12. Update $\xi_{\bullet ql} = (\xi_{1ql}, \dots, \xi_{Tql})'$ for $q = 1, \dots, Q$ and $l = 1, \dots, Q_\eta$ from $N(\mu_\xi, \Sigma_\xi)$ with $\mu_\xi = \Sigma_\xi b^*$ and $\Sigma_\xi^{-1} = A^* + \Psi^{-1}(\kappa_\xi)$, $A^* = \text{diag}(a_1^*, \dots, a_T^*)$, $b^* = (b_1^*, \dots, b_T^*)'$ and

$$a_t^* = \sum_{(i,j):t_{ij}=t} \Lambda_{s_i \bullet q} \Sigma_{s_i}^{-1} \Lambda_{s_i \bullet q} \tilde{\eta}_{il}^2, \quad b_t^* = \sum_{(i,j):t_{ij}=t} \Lambda'_{s_i \bullet q} \Sigma_{s_i}^{-1} \hat{y}_{ij} \tilde{\eta}_{il},$$

where $\hat{y}_{ij} = y_{ij}^* - B_{s_i} x_i - \Omega_{s_i} \mu_{t_{ij}} - \Lambda_{s_i} V_{s_i} x_i \eta_i^* - c^*$ with $c^* = (c_1^*, \dots, c_p^*)'$ and $c_k^* = \sum_{(q^*, l^*) \neq (q, l)} \Lambda'_{s_i k q^*} \xi_{t q^* l^*} \tilde{\eta}_{il}^*$.

13. Update θ_{hl}^x in $\theta_h^x = (\theta_{h1}^x, \dots, \theta_{hL}^x)'$ for $h = 1, \dots, H$ and $l = 1, \dots, L$ from

$$f(\theta_{hl}^x | \dots) \propto \prod_{i:s_i=h} f_l(x_{il} | \theta_{hl}^x) \pi(\theta_{hl}^x),$$

where $\pi(\theta_{hl}^x)$ is the prior density.

14. Impute missing values x_{il} from

$$f(x_{il} | \dots) \propto \prod_{j=1}^{n_i} f(y_{ij}^* | x_i, \eta_{ij}, \theta_{s_i}^y) f_l(x_{il} | \theta_{s_i l}^x).$$

15. Impute missing values y_{ijk} by generating $y_{ijk} = g(y_{ijk}^*)$, $y_{ijk}^* \sim f(y_{ijk}^* | x_i, \eta_{ij}, \theta_{s_i}^y)$.

16. Update y_{ijk}^* for $i = 1, \dots, n$, $j = 1, \dots, n_i$ and $k = 1, \dots, p$ from

$$f(y_{ijk}^* | \dots) \propto 1\{y_{ij} = g(y_{ij}^*)\} f(y_{ijk}^* | x_i, \eta_{ij}, \theta_{s_i}^y).$$

17. Update δ_{kl}^2 from $\text{IG}(0.5(H+1), 0.5(\sum_{h=1}^H U_{hkl}^2 + 1))$ for $k = 1, \dots, p$ and $l = 1, \dots, L^*$.

18. Update ζ_{ql}^2 from $\text{IG}(0.5(H+1), 0.5(\sum_{h=1}^H V_{hql}^2 + 1))$ for $q = 1, \dots, Q$ and $l = 1, \dots, L$.

19. Generate the adjusted mixture weights $\tilde{\pi}$ and compute the Goodman and Kruskal's gamma as in Section 3.3.

5 Simulation study

We assess performance of the proposed method with simulation data by comparing to two existing approaches. We assume that the response consists of continuous, binary, count and three-category nominal variables, and the covariate is a binary variable. All respondents participate in the survey three times out of $T = 9$ time points, randomly taking $t_{ij} \in \{3j - 2, 3j - 1, 3j\}$ with equal probabilities for $j = 1, 2, 3$. We estimate associations in the mixed-scale response for groups with $x_i = 0$ and $x_i = 1$. For the priors in the proposed method, we use $\alpha \sim \text{Gamma}(0.25, 0.25)$, $\sigma_{hk}^2 \sim \text{IG}(2, \tilde{s}_k/200)$ where \tilde{s}_k is the sample variance of the k th response for the ordered variables and 1 for the nominal variable. Also, we set $H = 60$ and $Q = Q_\mu = Q_\eta = 4$, and put 25 grids on the $(0, 1]$ interval for κ_μ and κ_ξ .

For competitors, we consider two methods. The first one is based on covariance regression with a polynomial function of time (Hoff and Niu (2012)),

$$y_{ij} = \tilde{g}(y_{ij}^*), \quad y_{ij}^* = Az_{ij} + Bz_{ij}\eta_i + \varepsilon_{ij}, \quad \eta_i \sim N(0, 1), \quad \varepsilon_{ij} \sim N(0, \Psi),$$

where $z_{ij} = (1, x_i, t_{ij}, t_{ij}^2, t_{ij}^3)'$ and A and B are $p \times 5$ coefficient matrices. The function \tilde{g} transforms the latent variable y_{ij}^* with the identity map for a continuous variable, the probit function for a binary variable, discretization with unknown cut-points for a count variable and the multinomial probit function for a nominal variable. For the prior distributions, we assume $\Psi \sim \text{Inverse-Wishart}(\Psi_0^{-1}, \nu_0)$ with expectation $\Psi_0/(\nu_0 - p - 1)$ and $(A, B) | \Psi \sim \text{Matrix-Normal}(0, \Psi, V_0)$ where $\nu_0 = p + 2$, Ψ_0 is the sample covariance matrix of y_{ij} and V_0 is a block diagonal matrix where the two blocks are $n(Z'Z)^{-1}$ with $Z = (z_1, \dots, z_n)'$ and $z_i = (z_{i1}, z_{i2}, z_{i3})$. Also, we apply independent prior $N(0, 100)$ to thresholds for a count variable.

The second approach is a dynamic latent trait model (Dunson (2003)),

$$y_{ij} \sim f(y_{ij} | \zeta_{ij}), \quad \zeta_{ij} = A\tilde{x}_i + \lambda_{t_{ij}}\tau_{it_{ij}} + \phi_i, \quad \tau_{i\bullet} \sim \text{GP}(\beta\tilde{x}_i, c_\tau), \quad \phi_i \sim N(0, \Phi),$$

where $\tilde{x}_i = (1, x_i)'$, A is a $p \times 2$ matrix, λ_t is a factor loading vector at time t , $\tau_{i\bullet} = \{\tau_{it}, t \in \mathcal{T}\}$, $c_\tau(\tau_{it}, \tau_{it'}) = \exp(\kappa_\tau |t - t'|^2)$ and $\Phi = \text{diag}(\omega_1^2, \dots, \omega_p^2)$. For the density function f , we assume Gaussian regression with variance σ^2 for continuous, logistic regression for binary, Poisson regression for count and multinomial logistic regression for nominal variables. As for the priors, we assume $\sigma^2 \sim \text{IG}(2, \tilde{s}/2)$, \tilde{s} is the sample variance of the continuous response, $A \sim \text{Matrix-Normal}(0, I, 100I)$, $\lambda_t \sim N(0, 100I)$, $\omega_k^2 \sim \text{IG}(2, 5)$, $\beta \sim N(0, 100I)$ and discrete grid prior with 25 points on $(0, 1]$ for κ_τ . For all methods, we generated 10,000 samples after the initial 5,000 samples are discarded as a burn-in period and every 10th sample is saved. To compute the Goodman and Kruskal's gamma, we generated 4,000 sample from the posterior predictive distribution. We observed that the sample paths were stable and the sample autocorrelations dropped smoothly.

In case 1, we generated data with $n = 4,000$ subjects from

$$y_{ij} = g(y_{ij}^*), \quad y_{ij}^* = D\tilde{x}_i + F_{t_{ij}}\tilde{x}_i\eta_i + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim N(0, 0.1^2 I),$$

$$\eta_i \sim N(0, 1), \quad \tilde{x}_i = (1, x_i)', \quad x_i \sim \text{Bernoulli}(0.5),$$

where each element of D is independently generated from $N(0, 0.05^2)$. We assume $F_t = F_{t-1} + \varepsilon_t^*$, $\varepsilon_t^* \sim \text{Matrix-Normal}(0, I, 0.05^2 I)$ for $t = 4, 7$ and $F_t = F_{t-1}$ for the other time points with $F_0 \sim \text{Matrix-Normal}(0, I, 0.05^2 I)$. These matrices are given in the supplementary materials. The function g corresponds to the one in Section 3.1 with the identity map for the continuous variable and the zero utility of the last category for the nominal variable. Also, we use non-negative integers as cutting points for the count variable. Because it is not straightforward to analytically compute true values of the Goodman and Kruskal's gamma, we approximate them by generating 8,000 samples from the true data-generating function. Figure 4 reports boxplots of logarithm of mean absolute errors at each time point and averaged ones over time in case 1. Although the covariance regression model shows close values at some time points, the proposed method constantly produces the smallest errors over time in both groups $x = 0$ and $x = 1$. Also, the proposed method reports the lowest averaged errors over the period of time with relatively large differences.

Next, we simulated data with $n = 4,000$ subjects from the model in which time-dependence is induced by the covariance regression with a second-order polynomial function of time, and the sample is divided into two subgroups.

$$y_{ij} = g(y_{ij}^*), \quad y_{ij}^* = D_{s_i}z_{ij} + F_{s_i}z_{ij}\eta_i + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim N(0, \sigma_{s_i}^2 I), \quad \eta_i \sim N(0, 1),$$

$$z_{ij} = (1, x_i, \tilde{t}_{ij}, \tilde{t}_{ij}^2)', \quad x_i \sim \text{Bernoulli}(0.5), \quad s_i \sim \text{Bernoulli}(0.5),$$

where $\tilde{t} = t/T$ and $\sigma_0 = 0.1$, $\sigma_1 = 0.05$. Each component of D_s and F_s is independently generated from $N(0, 0.05^2)$ with $s = 0, 1$ and given in the supplementary materials. Figure 5 shows the estimation result of the logarithm of the mean absolute errors of associations for $x = 0$ and $x = 1$ in case 2. The covariance regression model reports smaller errors over time than the dynamic latent trait model, but the proposed method outperforms the competitors in both $x = 0$ and $x = 1$.

Finally, we assume data are collected in a stratified sampling design. We consider a finite population with $N = 1,000,000$ subjects consists of three subpopulations having $N_1 = 650,000$, $N_2 = 300,000$ and $N_3 = 50,000$ people. We sample $n_m = 1,500$ subjects from each stratum with $m = 1, 2, 3$ and construct survey weights by $w_i = N_{m_i}/n_{m_i}$ for i th subject where m_i represents the subpopulation to which i th subject belongs. Then, we assume the following data-generating function, expressing dynamic structures through the covariance regression as in case 2.

$$y_{ij} = g(y_{ij}^*), \quad y_{ij}^* = D_{m_i} z_{ij} + F_{m_i} z_{ij} \eta_i + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim N(0, 0.1^2 I), \quad \eta_i \sim N(0, 1),$$

$$z_{ij} = (1, x_i, \tilde{t}_{ij}, \tilde{t}_{ij}^2)', \quad x_i \sim \text{Bernoulli}(0.5),$$

where all component in D_m and F_m are independently generated from $N(0, 0.05^2)$ with $m = 1, 2, 3$ and given in the supplementary materials. In the proposed method, we set the prior sample size as 1% of population size N for the bias adjustment in Section 3.3. The estimation result of the log mean absolute errors is given in Figure 6 in case 3. In both cases with $x = 0$ and $x = 1$, the proposed method reports the smallest mean absolute errors at each time point and the lowest averaged value over time.

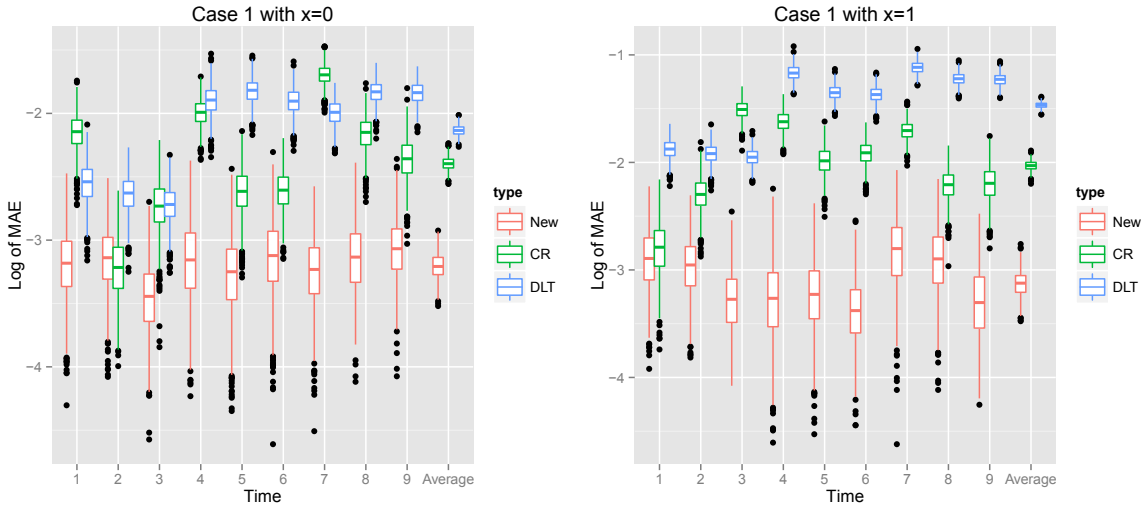


Figure 4: Boxplots of logarithm of mean absolute errors for the proposed model (New), covariance regression (CR) and dynamic latent trait model (DLT) in case 1 with $x = 0$ (left) and $x = 1$ (right). The x -axis shows time points and the average.

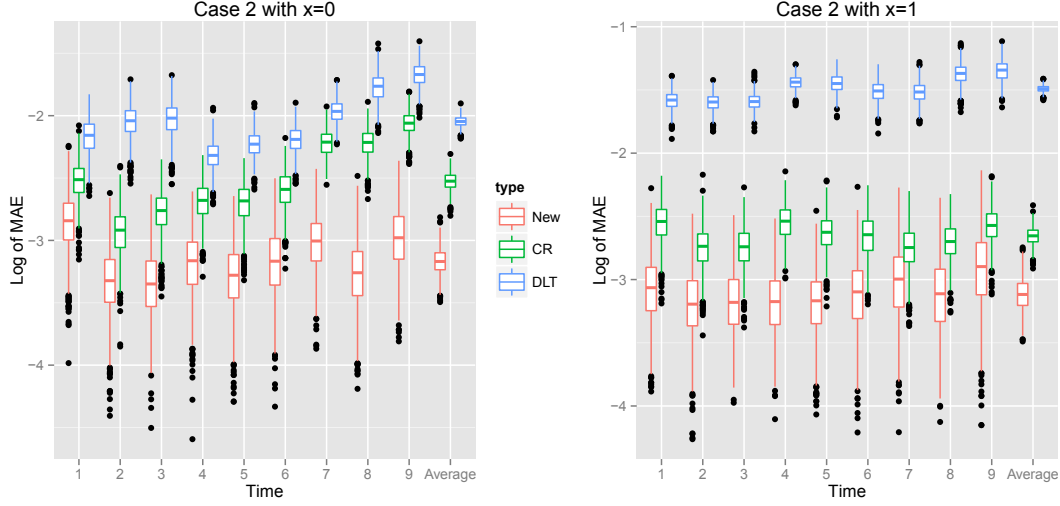


Figure 5: Boxplots of logarithm of mean absolute errors for the proposed model (New), covariance regression (CR) and dynamic latent trait model (DLT) in case 2 with $x = 0$ (left) and $x = 1$ (right). The x -axis shows time points and the average.

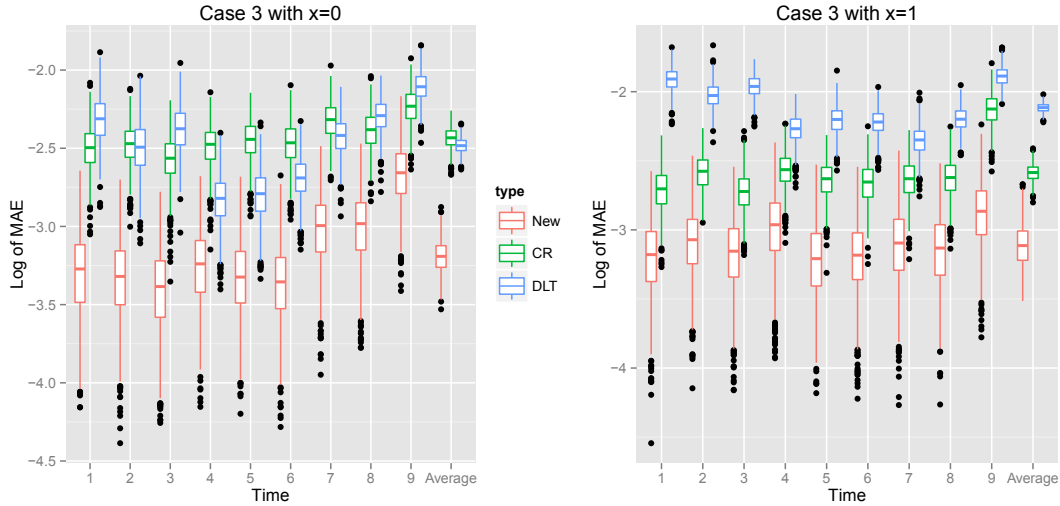


Figure 6: Boxplots of logarithm of mean absolute errors for the proposed model (New), covariance regression (CR) and dynamic latent trait model (DLT) in case 3 with $x = 0$ (left) and $x = 1$ (right). The x -axis shows time points and the average.

6 Analysis of adolescent sexual development data

This section applies the proposed method to the adolescent sexual development data. As for the function $f_l(x_{il} | \theta^x)$ for the covariates in (8), we assume multinomial distributions using Dirichlet priors with all concentration parameters 1. For the count variables, because we observe high right skewness in these data, we use log cut-points instead of nonnegative integers, which allow the Dirichlet process mixtures to efficiently approximate such distributions. Also, as the prior of the count variables, we assume $\sigma_{hk}^2 \sim \text{Inverse-Gamma}(2, \tilde{s}_k/200)$ where \tilde{s}_k is the sample variance of $\log(y_{ijk} + 0.5)$. For the other parameters, we use the same prior settings as in the simulation study in Section 5. In the sampling bias adjustment, we set the hyperparameter in the Dirichlet prior such that the prior sample size is equal to 1% of the population size $N = 14,677,347$. We estimated trajectories of associations among sexual development variables for the overall population and within subpopulations with respect to gender, race, parents' education and interactions of gender and race. We generate 10,000 samples after the initial 5,000 samples are discarded as a burn-in period, and every 10th sample is saved. At each MCMC iteration, we computed the Goodman and Kruskal's gamma by generating 2,500 subjects from the posterior predictive distribution. We observed that the sample paths were stable, and the sample autocorrelations dropped smoothly.

We study the association between attraction, behavior, and sexual orientation identity, the latter of which is included starting at age 18, because this variable was not included in wave 1. Also, we omitted the result for age 11 and 34 from the display below because of the small number of observations, which is 5 and 2 respectively. Figure 7 shows the posterior means and 95% credible intervals for the gamma statistic in the overall population for opposite sex attraction (Ao), same sex attraction (As), opposite sex partner count (#o), same sex partner count (#s), and sexual orientation identities of heterosexual (He), mostly heterosexual (MHe), mostly homosexual (MHo), homosexual (Ho), and bisexual (Bi). We observe various patterns of trajectories of the associations. Note that while there is no strong association between the attraction variables in early adolescence, with age a strongly negative association emerges and continues into adulthood. As one might expect, same sex attraction and same sex partner count have a strong positive association around 0.9,

and while the association between opposite sex attraction and partner count is positive, it is more modest, with estimated gamma values around 0.5. While the association is essentially null in the early teens, with age a strong negative association between opposite sex attraction and same sex partner count becomes apparent. However, the associations between the sexual orientation identity variables (not measured in the early to mid teens) and partner counts tend to be relatively constant over time.

While we saw few differences in the associations, or their changes over time, by race, we observed several strong differences by gender. Figures 8 and 9 report the comparison of the posterior means and 95% credible intervals for several pairs of sexual development variables among males and females. For example, while the association between counts of same and opposite sex partners is null or positive in the early teens, this association remains significantly positive among females into adulthood but becomes negative among males by the late teens. We observe a similar trend in the gender-specific associations between same sex attraction and opposite sex partner counts. This pattern is consistent with other findings that lesbian and bisexual-identified females are more likely than heterosexuals to experience both early vaginal intercourse and to engage in sexual risk behaviors such as more sexual partners and inconsistent contraception (Garofalo et al. (1998); Case et al. (2004); Charlton et al. (2011); McCabe et al. (2011); Herrick et al. (2013); Riskind et al. (2014); McCauley et al. (2014); McCauley et al. (2015)). We see opposite trends by gender between heterosexual sexual orientation identity and opposite sex partner counts, with this identity being associated with greater counts among men and lower counts among women from the late teens through early adulthood. In contrast, a mostly heterosexual identity has a null association with opposite sex partner counts for men, but a positive association among women. Among both men and women, homosexual identity is associated with lower opposite sex partner counts. Same sex partner counts are strongly associated with all identities except heterosexual among both men and women; the negative association between heterosexual identity and same sex partner count is stronger among men than among women.

Based on the developmental literature, we also examined whether changes in associations over time varied by gender within race but did not find strong evidence of this interaction in the association parameters. Similarly, we did not identify strong trends according to

parental education. These results are given in the supplementary materials.

7 Discussion

Using data from a longitudinal, naturally-representative sample of adolescents followed into adulthood, we were able to study how the associations between three components of sexual orientation: attraction, behavior, and sexual orientation identity, evolve with age. In order to study changes with time in the associations among these variables, measured longitudinally on a variety of categorical scales, we developed a flexible dynamic latent factor model allowing time-varying associations among the multivariate responses. This model also allowed the associations to depend on covariates of interest, such as gender. Population-representative inferences are obtained via an MCMC algorithm used for posterior computation. Previous related work either did not allow longitudinal data (Hoff and Niu (2012); Fox and Dunson (2015)), required variables to be ordinal (Hoff and Niu (2012)), or did not allow dependence of the association on covariates (Fox and Dunson (2015)).

References

- Baghfalaki, T., M. Ganjalina, and D. Berridge (2014). Joint modeling of multivariate longitudinal mixed measurements and time to event data using a Bayesian approach. *Journal of Applied Statistics* 41, 1934–1955.
- Bandyopadhyay, S., B. Ganguli, and A. Chatterjee (2011). A review of multivariate longitudinal data analysis. *Statistical Methods in Medical Research* 20, 299–330.
- Burgette, L. F. and E. V. Nordheim (2012). The trace restriction: An alternative identification strategy for the Bayesian multinomial probit model. *Journal of Business and Economic Statistics* 30, 404–410.
- Canale, A. and D. B. Dunson (2011). Bayesian kernel mixtures for counts. *Journal of the American Statistical Association* 106, 1528–1539.
- Case, P., S. B. Austin, D. J. Hunter, J. E. Manson, S. Malspeis, W. C. Willett, and

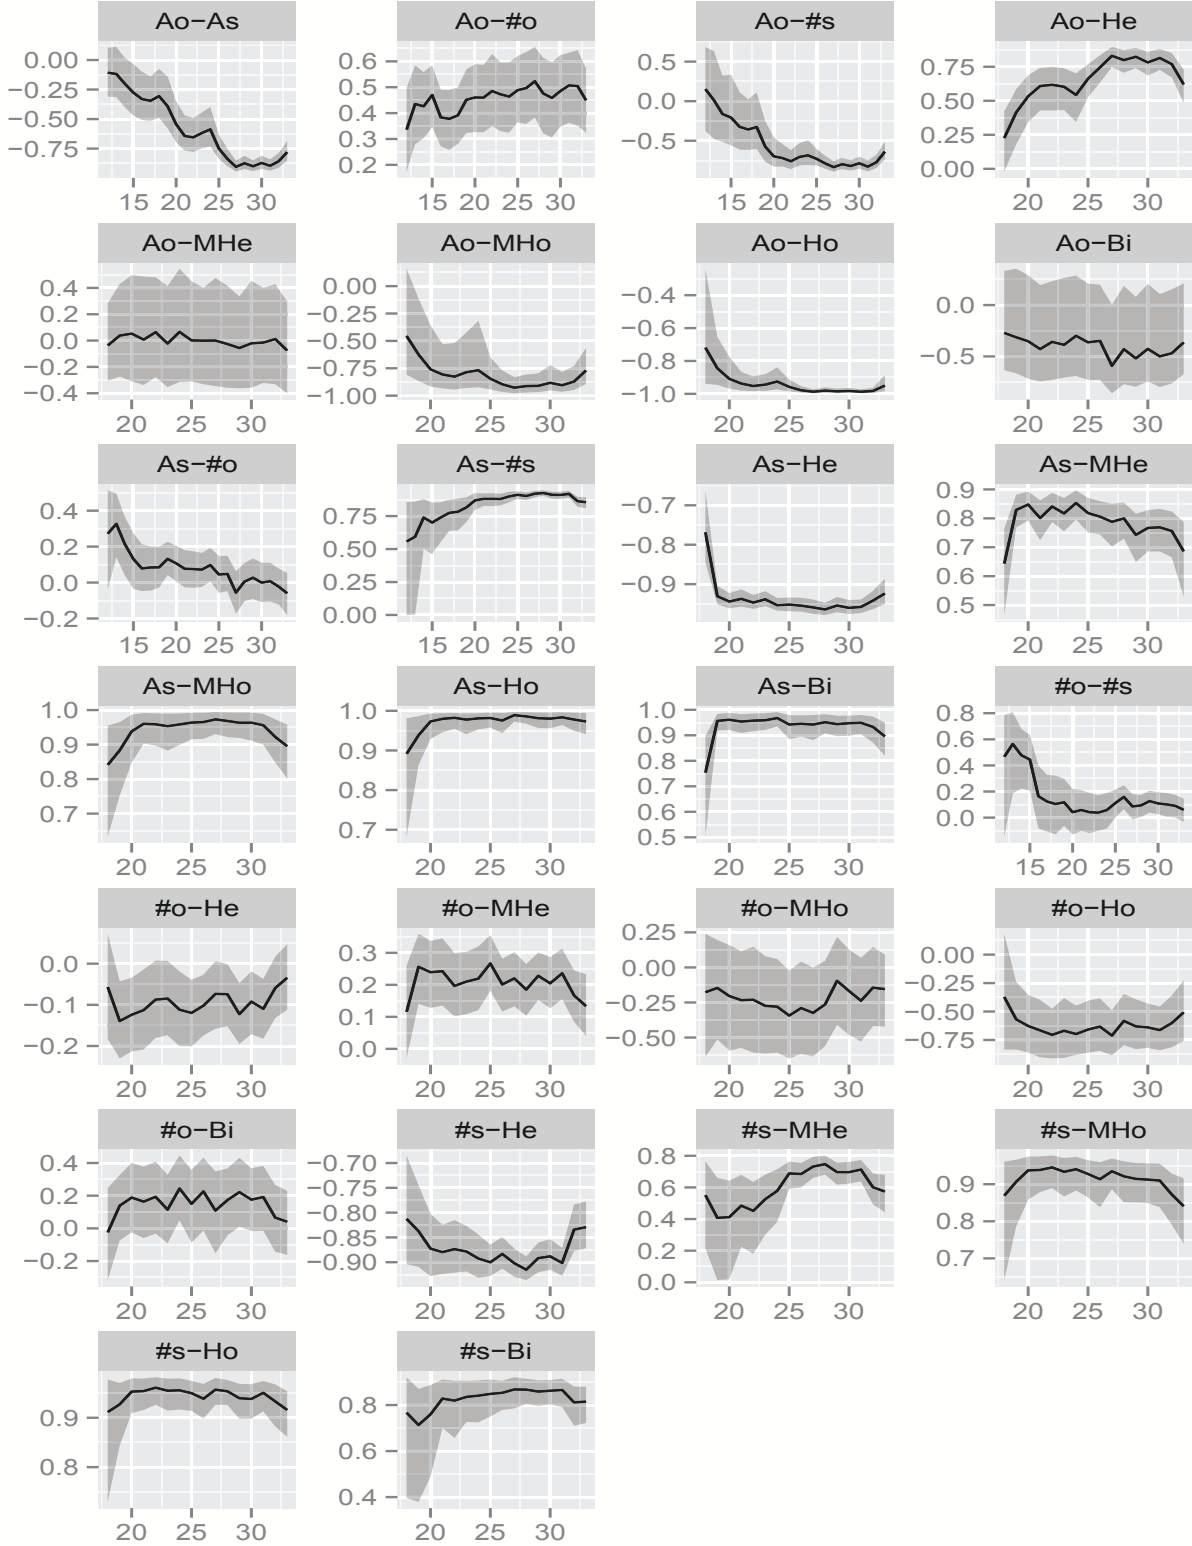


Figure 7: Posterior means and 95% credible intervals of associations for the overall population. y -axis shows association and x -axis age.

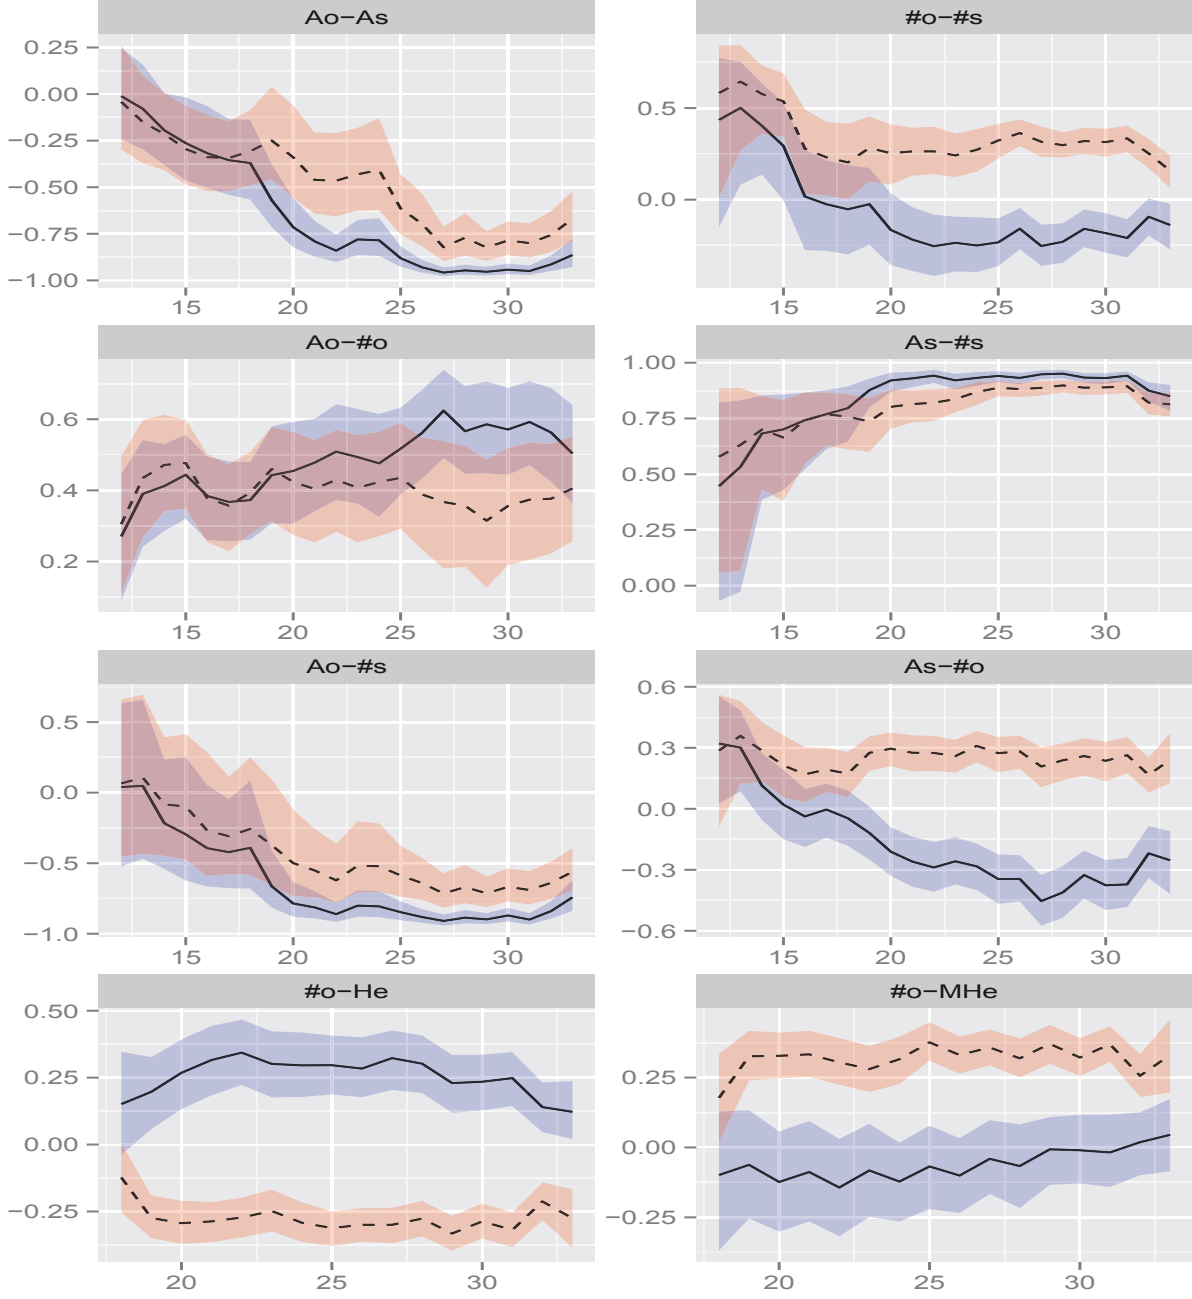


Figure 8: Comparison of posterior means and 95% credible intervals of associations for male (solid line with blue color) and female (dashed line with red color) [1]. y -axis shows association and x -axis age.

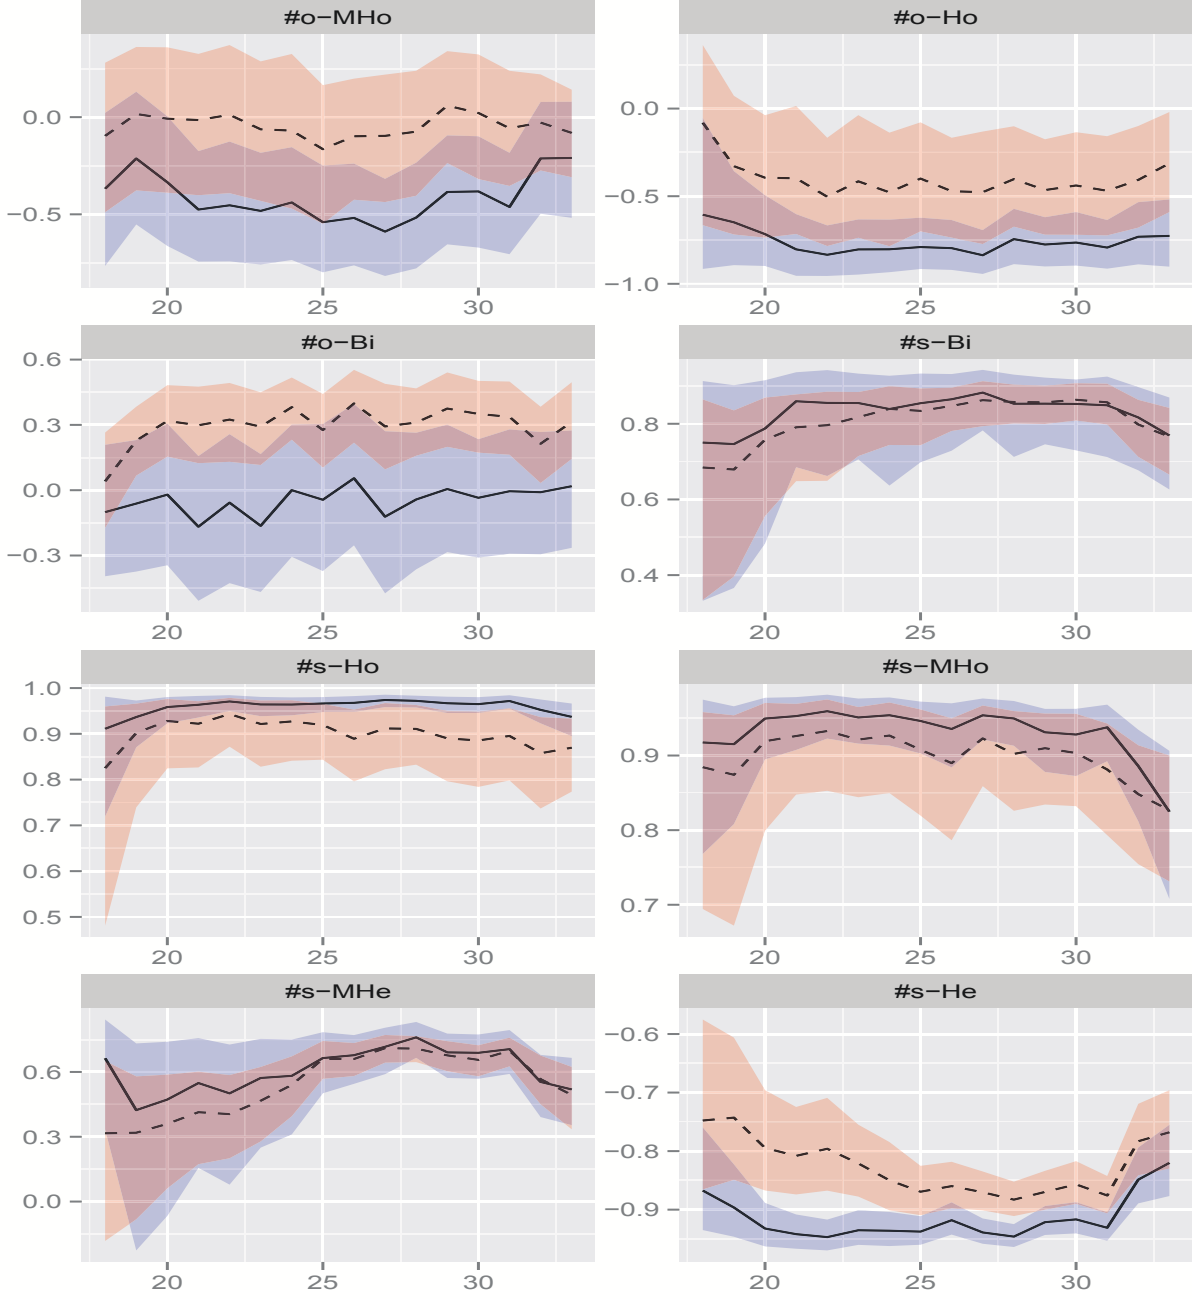


Figure 9: Comparison of posterior means and 95% credible intervals of associations for male (solid line with blue color) and female (dashed line with red color) [2]. y -axis shows association and x -axis age.

- D. Spiegelman (2004). Sexual orientation, health risk factors, and physical functioning in the nurses health study ii. *Journal of Women's Health* 13, 1033–1047.
- Charlton, B. M., H. L. Corliss, S. A. Missmer, A. L. Frazier, M. Rosario, J. A. Kahn, and S. B. Austin (2011). Reproductive health screening disparities and sexual orientation in a cohort study of U.S. adolescent and young adult females. *Journal of Adolescent Health* 49, 505–510.
- Chen, Q., M. R. Elliott, and R. J. A. Little (2010). Bayesian penalized spline model-based inference for finite population proportion in unequal probability sampling. *Survey Methodology* 36, 23–34.
- Das, K. and M. J. Daniels (2014). A semiparametric approach to simultaneous covariance estimation for bivariate sparse longitudinal data. *Biometrics* 70, 33–43.
- DeYoreo, M. and A. Kottas (2015a). Bayesian nonparametric modeling for multivariate ordinal regression. Technical report. <http://arxiv.org/abs/1408.1027>.
- DeYoreo, M. and A. Kottas (2015b). A fully nonparametric modeling approach to binary regression. *Biometrics* 10, 821–847.
- Dunson, D. B. (2000). Bayesian latent variable models for clustered mixed outcomes. *Journal of the Royal Statistical Society-B* 62, 355–366.
- Dunson, D. B. (2003). Dynamic latent trait models for multidimensional longitudinal data. *Journal of the American Statistical Association* 98, 555–563.
- Dunson, D. B. and A. H. Herring (2005). Bayesian latent variable models for mixed discrete outcomes. *Biostatistics* 1, 11–25.
- Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* 90, 577–588.
- Fieuws, S. and G. Verbeke (2006). Pairwise fitting of mixed models for the joint modeling of multivariate longitudinal profiles. *Biometrics* 62, 424–431.
- Fox, E. B. and D. B. Dunson (2015). Bayesian nonparametric covariance regression. *Journal of Machine Learning Research* 16, 2501–2542.

- Garofalo, R., R. C. Wolf, S. Kessel, J. Palfrey, and R. H. DuRant (1998). The association between health risk behaviors and sexual orientation among a school-based sample of adolescents. *Pediatrics* 101, 895–902.
- Gelman, A. (2007). Struggles with survey weighting and regression modeling. *Statistical Science* 22, 153–164.
- Ghosh, P. and T. Hanson (2010). A semiparametric Bayesian approach to multivariate longitudinal data. *Australian & New Zealand Journal of Statistics* 52, 275–288.
- Goodman, L. A. and W. H. Kruskal (1954). Measures of association for cross classifications. *Journal of the American Statistical Association* 49, 732–764.
- Goodman, L. A. and W. H. Kruskal (1959). Measures of association for cross classifications. ii: Further discussion and references. *Journal of the American Statistical Association* 54, 123–163.
- Goodman, L. A. and W. H. Kruskal (1963). Measures of association for cross classifications iii: Approximate sampling theory. *Journal of the American Statistical Association* 58, 310–364.
- Goodman, L. A. and W. H. Kruskal (1972). Measures of association for cross classifications, iv: Simplification of asymptotic variances. *Journal of the American Statistical Association* 67, 415–421.
- Gruhl, J., E. A. Erosheva, and P. K. Crane (2013). A semiparametric approach to mixed outcome latent variable models: Estimating the association between cognition and regional brain volumes. *Annals of Applied Statistics* 7, 2361–2383.
- Gueorguieva, R. V. and G. Sanacora (2006). Joint analysis of repeatedly observed continuous and ordinal measures of disease severity. *Statistics in Medicine* 25, 1307–1322.
- Hannah, L. A., D. M. Blei, and W. B. Powell (2011). Dirichlet process mixtures of generalized linear models. *Journal of Machine Learning Research* 12, 1923–1953.

- Harris, K. M., C. T. Halpern, E. Whitsel, J. Hussey, J. Tabor, P. Entzel, and J. R. Udry (2009). The national longitudinal study of adolescent health: Research design [www document]. URL: <http://www.cpc.unc.edu/projects/addhealth/design>.
- Herrick, A., L. Kuhns, S. Kinsky, A. Johnson, and R. Garofalo (2013). Demographic, psychosocial, and contextual factors associated with sexual risk behaviors among young sexual minority women. *Journal of the American Psychiatric Nurses Association* 19, 345–355.
- Hoff, P. D. (2007). Extending the rank likelihood for semiparametric copula estimation. *Annals of Applied Statistics* 1, 265–283.
- Hoff, P. D. and X. Niu (2012). A covariance regression model. *Statistica Sinica* 22, 729–753.
- Horvitz, D. G. and D. J. Thompson (1952). A generalization of sampling without replacement from a finite universe. *Journal of American Statistical Association* 47, 663–685.
- Imai, K. and D. van Dyk (2005). A Bayesian analysis of the multinomial probit model using marginal data augmentation. *Journal of Econometrics* 124, 311–334.
- Ishwaran, H. and L. F. James (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association* 96, 161–173.
- Kim, J. S. and B. T. Ratchford (2013). A Bayesian multivariate probit for ordinal data with semiparametric random-effects. *Computational Statistics & Data Analysis* 64, 192–208.
- Kottas, A., M. P. Q. F. (2005). Nonparametric Bayesian modeling for multivariate ordinal data. *Journal of Computational and Graphical Statistics* 14, 610–625.
- Kunihama, T., A. H. Herring, C. T. Halpern, and D. B. Dunson (2016). Nonparametric bayes modeling with sample survey weights. *Statistics & Probability Letters* 113, 41–48.
- Levy, P. S. and S. Lemeshow (2008). *Sampling of Populations: Methods and Applications*. Wiley.
- Little, R. J. A. (2004). To model or not to model? Competing modes of inference for finite population sampling. *Journal of the American Statistical Association* 99, 546–556.

- Liu, X., M. J. Daniels, and B. Marcus (2009). Joint models for the association of longitudinal binary and continuous processes with application to a smoking cessation trial. *Journal of the American Statistical Association* 104, 429–438.
- Lo, A. Y. (1984). On a class of Bayesian nonparametric estimates: I. Density Estimates. *Annals of Statistics* 12, 351–357.
- Luo, S. and J. Wang (2014). Bayesian hierarchical model for multiple repeated measures and survival data: an application to Parkinson’s disease. *Statistics in Medicine* 33, 4279–4291.
- McCabe, J., K. L. Brewster, and K. H. Tillman (2011). Patterns and correlates of same-sex sexual activity among U.S. teenagers and young adults. *Perspect Sex Reprod Health* 43, 142–150.
- McCauley, H. L., R. N. Dick, D. J. Tancredi, S. Goldstein, S. Blackburn, J. G. Silverman, E. Monasterio, L. James, and E. Miller (2014). Differences by sexual minority status in relationship abuse and sexual and reproductive health among adolescent females. *Journal of Adolescent Health* 55, 652–658.
- McCauley, H. L., J. G. Silverman, M. R. Decker, M. Agénor, S. Borrero, D. J. Tancredi, S. Zelazny, and E. Miller (2015). Sexual and reproductive health indicators and intimate partner violence victimization among female family planning clinic patients who have sex with women and men. *Journal of Women’s Health* 24, 621–628.
- McCulloch, R. and P. Rossi (1994). An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics* 64, 207–240.
- McParland, D., I. C. Gormley, T. H. McCormick, S. J. Clark, C. W. Kabudula, and M. A. Collinson (2014). Clustering south African households based on their asset status using latent variable models. *Annals of Applied Statistics* 8, 747–776.
- Moustaki, I. and M. Knott (2000). Generalized latent trait models. *Psychometrika* 65, 391–411.
- Muliere, P. and L. Tardella (1998). Approximating distributions of random functionals of Ferguson Dirichlet priors. *Canadian Journal of Statistics* 26, 283–297.

- Müller, P., A. Erkanli, and M. West (1996). Bayesian curve fitting using multivariate normal mixtures. *Biometrika* 83, 67–79.
- Müller, P. and F. A. Quintana (2010). Random partition models with regression on covariates. *Journal of Statistical Planning and Inference* 140, 2801–2808.
- Murray, J. S., D. B. Dunson, L. Carin, and J. E. Lucas (2013). Bayesian Gaussian copula factor models for mixed data. *Journal of the American Statistical Association* 108, 656–665.
- Murray, J. S. and J. P. Reiter (2015). Multiple imputation of missing categorical and continuous values via Bayesian mixture models with local dependence. *Journal of the American Statistical Association*. To appear.
- Muthén, B. (1984). A general structural equation model with dichotomous, ordered categorical and continuous latent variable indicators. *Psychometrika* 49, 115–132.
- Rao, J. N. K. (2011). Impact of frequentist and Bayesian methods on survey sampling practice: A selective appraisal. *Statistical Science* 26, 240–256.
- Riskind, R. G., S. L. Tornello, B. C. Younger, and C. J. Patterson (2014). Sexual identity, partner gender, and sexual health among adolescent girls in the United States. *American Journal of Public Health* 104, 1957–1963.
- Ritter, C. and M. A. Tanner (1992). Facilitating the Gibbs sampler: The Gibbs stopper and the Griddy Gibbs sampler. *Journal of American Statistical Association* 87, 861–868.
- Sammel, M. D., L. M. Ryan, and J. M. Legler (1997). Latent variable models for mixed discrete and continuous outcomes. *Journal of the Royal Statistical Society-B* 59, 667–678.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* 4, 639–650.
- Si, Y., N. Pillai, and A. Gelman (2015). Bayesian nonparametric weighted sampling inference. *Bayesian Analysis* 10, 605–625.
- Taddy, M. and A. Kottas (2012). Mixture modeling for marked Poisson processes. *Bayesian Analysis* 7, 335–362.

- Verbeke, G., S. Fieuws, G. Molenberghs, and M. Davidian (2014). The analysis of multivariate longitudinal data: A review. *Statistical Methods in Medical Research* 23, 42–59.
- West, M., P. Müller, and M. D. Escobar (1994). Hierarchical priors and mixture models, with application in regression and density estimation. *Aspects of uncertainty: A Tribute to DV Lindley*, 363–386.
- Zheng, H. and R. J. A. Little (2003). Penalized spline model-based estimation of finite population total from probability-proportional-to-size samples. *Journal of Official Statistics* 19, 99–117.
- Zheng, H. and R. J. A. Little (2005). Inference for the population total from probability-proportional-to-size samples based on predictions from a penalized spline nonparametric model. *Journal of Official Statistics* 21, 1–20.