Published in final edited form as:

Biometrics. 2014 March ; 70(1): 33–43. doi:10.1111/biom.12133.

A Semiparametric Approach to Simultaneous Covariance Estimation for Bivariate Sparse Longitudinal Data

Kiranmoy Das¹ and Michael J. Daniels²

¹Department of Statistics, Presidency University, Kolkata, India, 700073

²Department of Integrative Biology, Division of Statistics & Scientific Computation, University of Texas at Austin, Austin, TX, 78712

Summary

Estimation of the covariance structure for irregular sparse longitudinal data has been studied by many authors in recent years but typically using fully parametric specifications. In addition, when data are collected from several groups over time, it is known that assuming the same or completely different covariance matrices over groups can lead to loss of efficiency and/or bias. Nonparametric approaches have been proposed for estimating the covariance matrix for regular univariate longitudinal data by sharing information across the groups under study. For the irregular case, with longitudinal measurements that are bivariate or multivariate, modeling becomes more difficult. In this article, to model bivariate sparse longitudinal data from several groups, we propose a flexible covariance structure via a novel matrix stick-breaking process for the residual covariance structure and a Dirichlet process mixture of normals for the random effects. Simulation studies are performed to investigate the effectiveness of the proposed approach over more traditional approaches. We also analyze a subset of Framingham Heart Study data to examine how the blood pressure trajectories and covariance structures differ for the patients from different BMI groups (high, medium and low) at baseline.

Keywords

Covariance matrix; DIC; Dirichlet process mixture of normals; MCMC

1 Introduction

There is a rich literature on estimation of the covariance matrix for univariate longitudinal data. Most relevant to our work are a series of papers that use the modified Cholesky decomposition of the covariance structure. Pourahmadi (1999) first proposed (likelihood based) approaches to estimate the covariance structure for regular longitudinal data. Wu and Pourahmadi (2003) then proposed a nonparametric approach for estimating large longitudinal covariance matrices and Daniels and Pourahmadi (2002) proposed a Bayesian approach using (conditionally conjugate) priors which can shrink the underlying covariance matrices to some specific known structures. Pan and Mackenzie (2003) generalized these approaches to irregular sparse longitudinal data.

Although the literature for modeling the covariance structure in the univariate longitudinal case is rich, less work has been done for bivariate or multivariate settings. The challenge in

Supplementary Materials

Computational details and Tables and Figures referenced in Sections 4.1, 4.2, 5.3, 5.4 and 5.5 are available with this paper at the Biometrics website on Wiley Online Library. Code is available at <http://www.sbs.utexas.edu/mjdaniels/software>.

these situations is to handle the longitudinal dependence within and between the response features. Sy, Taylor and Cumberland (1997) proposed a parametric stochastic model for CD4 T-cells and beta-2 microglobulin in AIDS data based on integrated Ornstein-Uhlenbeck (IOU) process, which can handle unequally spaced observations. Theibaut et al. (2002) proposed linear mixed models for analyzing bivariate longitudinal data (using SAS). For regular bivariate longitudinal data, Sithole and Jones (2007) proposed a longitudinal model for detecting prescribing change in two drugs with correlated errors using a bivariate autoregressive process. Das et al. (2011, 2013b) extended Sithole and Jones' approach to irregular sparse bivariate longitudinal data. Bandyopadhyay et al. (2010) proposed linear mixed models that allow for skew normal independent distributions to model bivariate longitudinal (periodontal) responses in a Bayesian framework. Ghosh and Hanson (2010) proposed a semi-parametric Bayesian approach to model multivariate longitudinal data using a mixture of Polya trees prior distribution for subject-specific random effects.

In many situations, one has to deal with longitudinal measurements from multiple groups and it is often the case that features of the covariance structure are similar and/or the same across these groups and it is typically difficult to determine shared features 'manually' given positive definiteness constraints and the large space of possible models. Examples include treatments in a clinical trial and/or patterns in mixture models for longitudinal data with dropout. The matrix stick breaking process (Dunson et al., 2008) was recently adapted to covariance matrices by Gaskins and Daniels (2012). Gaskins and Daniels parameterized the covariance matrix for each group using the parameters of modified Cholesky decomposition which include the generalized autoregressive parameters (GARP) and the innovation variances (IV). We will use a similar approach for irregular longitudinal data building on work in Gaskins and Daniels (2012) and Pan and Mackenzie (2003). In particular, for each response feature, the generalized autoregressive parameters and innovation variances are first expressed as polynomial functions of time, which accommodates irregular longitudinal data (Pan and Mackenzie, 2003). For the coefficients of these polynomial functions, we develop matrix stick breaking process priors to allow information sharing across the parameters of different groups *and* to introduce sparsity (i.e. the order of the polynomial functions). The latter is a novel extension of the MSBP and will avoid the need for two-step procedures that first determine the polynomial order of the GARP and IV parameters via AIC or BIC and then fit the corresponding (Bayesian) models (Pourahmadi 1999, Pan and MacKenzie, 2003); such approaches underestimate uncertainty. Dependence between the response features at each time will be modeled using random effects whose distributions will be specified using Bayesian nonparametric methods (MacEachern and Mueller, 1998).

We use our approach to analyze data from the Framingham Heart Studies. These studies started in 1948 with 5,209 healthy men and women with an objective to study genetic effects on cardiovascular diseases. The dataset of interest here contain 977 participants whose blood pressures (systolic and diastolic) were measured at subject-specific time points. Given the known relationship between blood pressure and obesity (Wolk, Shamsuzzaman and Somers; 2003), we classify the subjects into three groups with respect to their baseline BMI; high ($BMI \geq 25$), medium ($18.5 < BMI < 25$) and low ($BMI < 18.5$), with sizes 245, 492 and 240 respectively. This categorization is based on a World Health Organization report (WHO, 1995). Our modeling approach and priors will allow us to find a parsimonious covariance structure that lies between a completely distinct structure for each group and the same structure for each group. The latter would be the most common approach.

The current article is organized as follows. In Section 2, we describe our modelling approach. We specify our proposed prior structure in Section 3. We describe the results of simulation studies to investigate the effectiveness of our proposed prior structure and

compare to simpler approaches in Section 4. Section 5 contains the analysis of the Framingham data. Concluding remarks and discussion are given in Section 6.

2 Model and Methods

Throughout this paper, we use the following notation. We consider M groups with the m -th

group consisting of n_m subjects for a total of $n = \sum_{m=1}^M n_m$ subjects; in our example, we have $M = 3$ groups defined based on BMI. We assume that a bivariate trait (here, systolic and diastolic blood pressure) is measured longitudinally for subject i at T_i subject-specific time points given by the vector $\mathbf{t}_i = (t_{i1}, \dots, t_{iT_i})$, $i = 1, \dots, n$. For the k -th response feature ($k=0,1$) from subject i who belongs to the m -th group, we denote the longitudinal vectors by $\mathbf{Y}_{ikm} = [Y_{ikm}(t_{i1}), \dots, Y_{ikm}(t_{iT_i})]$. We model the k -th response feature from subject i measured at time t_{ij} ($j = 1, \dots, T_i$) belonging to the m -th group as

$$Y_{ikm}(t_{ij}) = f_{km}(t_{ij}) + \mathbf{z}_{ijkm}^T \boldsymbol{\gamma}_{ikm} + e_{ijkm}, \quad (1)$$

where $\boldsymbol{\gamma}_{ikm}$ is a $(p+1) \times 1$ vector of subject-specific random effects with design vector $\mathbf{z}_{ijkm}$. The smooth function $\mathbf{f}_{km}$ is the general effect of time for k -th response in the m -th group which can be modeled using polynomials (Das et al. 2012) or penalized splines (Durban et al. 2005). In matrix notation, the above model can be represented as

$$\mathbf{Y}_{ikm} = \mathbf{f}_{km}(\mathbf{t}_i) + \mathbf{Z}_{ikm} \boldsymbol{\gamma}_{ikm} + \mathbf{e}_{ikm}, \quad (2)$$

where $\mathbf{Z}_{ikm}$ is the design matrix corresponding to the group-specific random effects $\boldsymbol{\gamma}_{ikm}$.

The residual vectors $\mathbf{e}_{im} = [\mathbf{e}_{i1m}^T, \mathbf{e}_{i0m}^T]^T$ are assumed to be normally distributed with mean 0 and covariance matrix $\boldsymbol{\Sigma}_{im}$. We assume independence between $\boldsymbol{\gamma}_{ikm}$ and $\mathbf{e}_{ikm}$ but consider dependence between the random effects $\boldsymbol{\gamma}_{i1m}$ and $\boldsymbol{\gamma}_{i0m}$ to explain the dependence between the bivariate response, $(\mathbf{Y}_{i1m}, \mathbf{Y}_{i0m})$ at each time point.

2.1 Modeling the Residual Covariance Structure

In the following development, for simplicity of notation and without loss of generality, we ignore the general effect of time $\mathbf{f}_{km}(\mathbf{t}_i)$ and suppress the subscript i . At time t , the residual for k -th response feature for a subject in group m from (2) is modeled as follows,

$$e_{tkm} = \sum_{t'=1}^{t-1} \varphi_{t,t',m}^{(k)} e_{t'km} + \sum_{t'=1}^{t-1} \psi_{t,t',m}^{(1-k)} e_{t'(1-k)m} + \varepsilon_{tkm}, \quad (3)$$

where $\varphi_{t,t',m}^{(k)}$ and $\psi_{t,t',m}^{(1-k)}$ are autoregressive coefficients and ε_{tkm} is the prediction error with mean 0 and variance σ_{tkm}^2 . Let $\mathbf{e}_m = [\mathbf{e}_{1m}^T, \mathbf{e}_{0m}^T]^T$ be the vector of residual errors of dimension $2T \times 1$ and $\boldsymbol{\varepsilon}_m = [\boldsymbol{\varepsilon}_{1m}^T, \boldsymbol{\varepsilon}_{0m}^T]^T$. Then equation (3) can be re-expressed as:

$$\mathbf{e}_m = \begin{bmatrix} \boldsymbol{\Phi}_{m1} & \boldsymbol{\Psi}_{m0} \\ \boldsymbol{\Psi}_{m1} & \boldsymbol{\Phi}_{m0} \end{bmatrix} \mathbf{e}_m + \boldsymbol{\varepsilon}_m, \quad (4)$$

where $\Phi_{\mathbf{m}\mathbf{k}}$ and $\Psi_{\mathbf{m}\mathbf{k}}$ are both $T \times T$ lower triangular matrices with 0's in the diagonal elements and $\varphi_{t,t',m}^{(k)}$ and $\psi_{t,t',m}^{(k)} (t > t')$ in the (t, t') -th position respectively. Hence equation (4) can be re-written as

$$\mathbf{L}_m \mathbf{e}_m = \boldsymbol{\varepsilon}_m, \quad (5)$$

where $\mathbf{L}_m = \begin{bmatrix} \mathbf{I} - \Phi_{\mathbf{m}1} & -\Psi_{\mathbf{m}0} \\ -\Psi_{\mathbf{m}1} & \mathbf{I} - \Phi_{\mathbf{m}0} \end{bmatrix}$. Here we assume that ε_{ikm} 's are uncorrelated

(Pourahmadi 1999) and obtain $\text{cov}(\boldsymbol{\varepsilon}_m) = \mathbf{E}_m = \begin{bmatrix} \mathbf{E}_m^{(1)} & 0 \\ 0 & \mathbf{E}_m^{(0)} \end{bmatrix}$, where $\mathbf{E}_m^{(1)}$ and $\mathbf{E}_m^{(0)}$ are $T \times T$ diagonal matrices with t -th diagonal element σ_{t1m}^2 and σ_{t0m}^2 , respectively. By taking the covariance on both sides of equation (5), we obtain

$$\text{cov}(\boldsymbol{\varepsilon}_m) = \mathbf{L}_m \text{cov}(\mathbf{e}_m) \mathbf{L}_m^T = \mathbf{L}_m \sum_m \mathbf{L}_m^T = \mathbf{E}_m. \quad (6)$$

For irregular longitudinal data, the subject-specific covariance matrices will be modeled using $\mathbf{E}_{im} = \mathbf{L}_{im} \sum_{im} \mathbf{L}_{im}^T$. The generalized autoregressive parameters ($\phi^{(k)}$'s and $\psi^{(k)}$'s) and the logarithm of the innovation variances ($\log \sigma_{t1m}^2$ and $\log \sigma_{t0m}^2$) are modeled as polynomial functions (Pan and Mackenzie, 2003) in lag $|t - t'|$ and time, t , respectively, as follows:

$$\log \sigma_{t1m}^2 = \lambda_{0m} + \lambda_{1m}t + \lambda_{2m}t^2 + \dots + \lambda_{gm}t^g, \quad (7)$$

$$\log \sigma_{t0m}^2 = \delta_{0m} + \delta_{1m}t + \delta_{2m}t^2 + \dots + \delta_{gm}t^g, \quad (8)$$

$$\varphi_{t,t',m}^{(1)} = \eta_{0m}^{(1)} + \eta_{1m}^{(1)}(t-t') + \eta_{2m}^{(1)}(t-t')^2 + \dots + \eta_{hm}^{(1)}(t-t')^h, \quad (t' = 1, 2, \dots, t-1) \quad (9)$$

$$\varphi_{t,t',m}^{(0)} = \eta_{0m}^{(0)} + \eta_{1m}^{(0)}(t-t') + \eta_{2m}^{(0)}(t-t')^2 + \dots + \eta_{hm}^{(0)}(t-t')^h, \quad (t' = 1, 2, \dots, t-1) \quad (10)$$

$$\psi_{t,t',m}^{(1)} = \theta_{0m}^{(1)} + \theta_{1m}^{(1)}(t-t') + \theta_{2m}^{(1)}(t-t')^2 + \dots + \theta_{hm}^{(1)}(t-t')^h, \quad (t' = 1, 2, \dots, t-1) \quad (11)$$

$$\psi_{t,t',m}^{(0)} = \theta_{0m}^{(0)} + \theta_{1m}^{(0)}(t-t') + \theta_{2m}^{(0)}(t-t')^2 + \dots + \theta_{hm}^{(0)}(t-t')^h, \quad (t' = 1, 2, \dots, t-1). \quad (12)$$

In (7)–(12), we have $2(g+2h+3)$ covariance parameters for each group. Even for moderate values of g , h and m , this results in quite a large number of parameters to estimate. As a result, we will modify the matrix stick breaking process (Dunson et al., 2008) for the covariance parameters λ_m , δ_m , $\eta_m^{(1)}$, $\eta_m^{(0)}$, $\theta_m^{(1)}$ and $\theta_m^{(0)}$ to effectively reduce the dimension by allowing equalities across groups and sparsity via reducing the order (g/h) of the polynomial functions.

3 Priors for the Covariance Structure

3.1 Priors for $\eta_m^{(1)}$, $\eta_m^{(0)}$, $\theta_m^{(1)}$ and $\theta_m^{(0)}$

The proposed prior for η_m is based on the matrix stick-breaking process (MSBP) introduced in Dunson et al. (2008). Information is borrowed across M related groups (based on baseline BMI in our example) by incorporating dependence in the prior distributions of $\eta_m^{(1)}$, which is an $(h+1)$ -dimension parameter vector for the m -th group. For $j' = 0, \dots, h$, assume

$$\eta_{j'm}^{(1)} \sim F_{j'm}^{\eta^{(1)}} = \sum_{q=1}^{N_{\eta^{(1)}}} \pi_{j'mq}^{\eta^{(1)}} \delta_{\varepsilon_{j'q}^{\eta^{(1)}}}(\cdot); \quad m=1, \dots, M; \quad j'=0, \dots, h;$$

$$\varepsilon_{j'q}^{\eta^{(1)}} \stackrel{\text{iid}}{\sim} F_{0j'}^{\eta^{(1)}},$$

where δ_x denotes a point mass at x . Define $\mathcal{E}^{\eta} = (\varepsilon_{j'q}^{\eta^{(1)}})$ as an $(h+1) \times N_{\eta^{(1)}}$ matrix of random atoms. Here the rows of $\mathcal{E}^{\eta}$ correspond to the parameters having base distribution $F_{0j'}^{\eta^{(1)}}$ and the columns correspond to the clusters. The weights, $\pi_{j'mq}^{\eta^{(1)}}$ are defined as

$$\pi_{j'mq}^{\eta^{(1)}} = V_{j'mq}^{\eta^{(1)}} \prod_{l < q} (1 - V_{j'ml}^{\eta^{(1)}}), \quad (13)$$

with $V_{j'mq}^{\eta^{(1)}} = U_{mq}^{\eta^{(1)}} W_{j'q}^{\eta^{(1)}}$ and $U_{mq}^{\eta^{(1)}} \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha_{\eta^{(1)}})$ and $W_{j'q}^{\eta^{(1)}} \stackrel{\text{iid}}{\sim} \text{Beta}(1, \beta_{\eta^{(1)}})$. These stick breaking weights $\pi_{j'mq}^{\eta^{(1)}}$ control the dependence among $F_{j'm}^{\eta^{(1)}}$ and are partitioned into two components; $U_{mq}^{\eta^{(1)}}$ and $W_{j'q}^{\eta^{(1)}}$ which allocate the coefficients from m -th group and j' -th

parameter component respectively to the q -th cluster. We take $V_{j'mN_{\eta^{(1)}}}^{\eta^{(1)}} = 1$; for all m and j'

to make $F_{j'm}^{\eta^{(1)}}$ a valid probability measure. The full matrix stick-breaking process corresponds to the limiting case, $N_{\eta^{(1)}} = \infty$. For finite $N_{\eta^{(1)}}$, it is called a truncated MSBP.

We consider a zero-inflated normal base distribution for $\eta_m^{(1)}$ as in Gaskins and Daniels (2012),

$$\varepsilon_{j'q}^{\eta^{(1)}} \sim (1 - B_{j'}^{\eta^{(1)}}) \delta_0(\cdot) + B_{j'}^{\eta^{(1)}} N(0, \sigma_{\eta^{(1)}}^2); \quad q=1, \dots, N_{\eta^{(1)}}; \quad j'=0, \dots, h;$$

$$B_{j'}^{\eta^{(1)}} = \prod_{l=0}^{j'} A_l^{\eta^{(1)}}; \quad A_l^{\eta^{(1)}} | \pi_{j'}^{\eta^{(1)}} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(\pi_{j'}^{\eta^{(1)}}); \quad (14)$$

$$\pi_{j'}^{\eta^{(1)}} \stackrel{\text{iid}}{\sim} \text{Uniform}(0, 1); \quad j'=0, \dots, h;$$

where $B_{j'}^{\eta^{(1)}}$ is a binary random variable (for each j') taking value 1 only when each of the independent Bernoulli random variables $\{A_l^{\eta^{(1)}} : l = 0, \dots, j'\}$ take value 1; (14) is a more complex specification than in Gaskins and Daniels (2012) who proposed a related prior for regular longitudinal data and assumed $B_{j'}^{\eta^{(1)}} = A_{j'}^{\eta^{(1)}}$. Here, for each j' , $B_{j'}^{\eta^{(1)}}$ is 1 only when all

the lower lag coefficients are non-zero. This specification implicitly allows the data to select the order of the polynomials in (9) in such a way that a non-zero higher order term can not appear with zero lower order terms and avoids the need for a two-step approach of selecting the order and then fitting the model (Pan and Mackenzie, 2003; Das et al., 2013a). The prior probability of order c (for (9)) is

$$Pr(Ord=c)=P(B_c^{\eta^{(1)}}=1, B_{c+1}^{\eta^{(1)}}=0)=(\prod_{i=0}^c \pi_i^{\eta^{(1)}})(1-\pi_{c+1}^{\eta^{(1)}}).$$

This prior also allows sharing information across the M related groups as in the original MSBP representation. An inverse gamma prior is specified for $\sigma_{\eta^{(1)}}^2$.

Various properties of the original matrix stick-breaking process specification have been discussed in detail in Dunson et al. (2008) and a formulation for multiple covariance matrices is given in Gaskins and Daniels (2012). The limiting cases ($N_{\eta^{(1)}} \rightarrow \infty$, $\alpha_{\eta^{(1)}} \rightarrow 0$ and $\beta_{\eta^{(1)}} \rightarrow \infty$, $\alpha_{\eta^{(1)}} \rightarrow 0$ and $\beta_{\eta^{(1)}} \rightarrow 0$) for the above specifications are very similar to the cases discussed in Dunson et al. (2008) and Gaskins and Daniels (2012). Properties specific to our formulation will be discussed in Section 3.3.

For $\eta_m^{(0)}$, $\theta_m^{(1)}$ and $\theta_m^{(0)}$, we assume the same prior specification.

3.2 Priors for λ_m and δ_m

Recall that λ_m and δ_m are the parameters to model the innovation variances. We consider the following matrix stick-breaking process priors for those parameters. For λ_m , we propose

$$\begin{aligned} \lambda_{mj'} &\sim F_{mj'}^{\lambda} = \sum_{q=1}^{N_{\lambda}} \pi_{mj'q}^{\lambda} \delta_{\varepsilon_{j'q}^{\lambda}}(\cdot); \quad m=1, \dots, M; \quad j'=0, \dots, g; \\ \varepsilon_{j'q}^{\lambda} &\sim (1-B_{j'}^{\lambda})\delta_0(\cdot) + B_{j'}^{\lambda}N(0, \sigma_{\lambda}^2); \quad q=1, \dots, N_{\lambda}; \quad j'=0, \dots, g; \\ B_{j'}^{\lambda} &= \prod_{l=0}^{j'} A_l^{\lambda}; \quad A_{j'}^{\lambda} | \pi_{j'}^{\lambda} \stackrel{\text{ind}}{\sim} \text{Bernoulli}(\pi_{j'}^{\lambda}); \quad \pi_{j'}^{\lambda} \stackrel{\text{iid}}{\sim} \text{Uniform}(0, 1); \quad j'=0, \dots, g; \\ \pi_{mj'q}^{\lambda} &= U_{mq}^{\lambda} W_{j'q}^{\lambda} \prod_{l < q} (1 - U_{ml}^{\lambda} W_{jl}^{\lambda}); \\ U_{mq}^{\lambda} &\stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha_{\lambda}), \quad W_{j'q}^{\lambda} \stackrel{\text{iid}}{\sim} \text{Beta}(1, \beta_{\lambda}); \quad q=1, \dots, N_{\lambda}-1. \end{aligned}$$

Similar to the prior for the autoregressive parameters, the above prior specification will choose the order of the polynomial functions from (7) and allow equality across groups. We consider an inverse gamma prior for σ_{λ}^2 . The same prior structure is proposed for δ_m .

3.3 Properties

Note that the novel features of the priors proposed here for $\eta_m^{(1)}$, $\eta_m^{(0)}$, $\theta_m^{(1)}$, $\theta_m^{(0)}$, λ_m and δ_m have the same form. As a result, for illustration, we focus only on the properties of the prior for $\eta_m^{(1)}$. We also concentrate on two properties unique to our specification. Details on other properties, including correlations between the probability measures for different groups and parameters, can be found in Dunson et al. (2008) and Gaskins and Daniels (2012).

Consider sets A in the Borel field of the real line. Denote

$G_{0j'}^* = (1 - p_{j'}^*)\delta_0(\cdot) + p_{j'}^*N(0, \sigma_{\eta(1)}^2)$, the probability measure for the mixture distribution of

the $\varepsilon_{lj'q}$'s, where $p_{j'}^* = P(B_{j'}^{\eta(1)} = 1) = \prod_{i=0}^{j'} \pi_i^{\eta(1)}$.

1. Behavior of $p_{j'}^*$: Since $\pi_i^{\eta(1)} \stackrel{\text{iid}}{\sim} \text{Uniform}(0, 1)$, $p_{j'}^*$ is decreasing in j' . As $j' \rightarrow \infty$, $p_{j'}^* \rightarrow 0$, and hence $P(B_{j'}^{\eta(1)} = 0) \rightarrow 1$. This property indicates that the higher order terms in the polynomial functions in (7){(12)} will be zero with high prior probability.
2. For $m = m'$; $Pr(\eta_{j'm}^{(1)} = \eta_{j'm'}^{(1)}) = (1 - p_{j'}^{*2}) + \frac{p_{j'}^{*2}}{(1 + \alpha_{\eta(1)})(2 + \beta_{\eta(1)}) - 1}$. Hence, as $j' \rightarrow \infty$, $Pr(\eta_{j'm}^{(1)} = \eta_{j'm'}^{(1)}) \rightarrow 1$. Intuitively this is sensible in our setting, since for larger values of j' , the parameters $\eta_{j'}^{(1)}$'s from different groups are zero with higher prior probabilities and this result corresponds to them being equal at value zero. Also as either $\alpha_{\eta(1)}$ or $\beta_{\eta(1)}$ tends to infinity, $Pr(\eta_{j'm}^{(1)} = \eta_{j'm'}^{(1)}) \rightarrow (1 - p_{j'}^{*2})$.

3.4 Priors for Random Effects

The modified MSBP priors model serial correlation within each longitudinal response. We now describe the modeling aspects for the correlation between the bivariate responses at each time. A (truncated) Dirichlet process mixture (DPM) prior is assumed for the distribution of the random effects (Ishwaran and James, 2001) to allow for correlation among the two response features Y_{i1m} and Y_{i0m} at each time. The correlation between these two response features (here, systolic and diastolic blood pressure) is captured by first considering a multivariate normal distribution for the random effects, $\gamma_{im} = [\gamma_{i1m}^T, \gamma_{i0m}^T]^T$. Note that since we have p (covariates with) random effects, the vectors γ_{i1m} and γ_{i0m} are of dimension $(p + 1) \times 1$. In our example, $p = 1$ and γ_{im} is the vector containing a random intercept and slope for subject i in group m . We assume that the random intercept and slope are independent and for the m -th group that the correlation of each between the two response features is ρ_{ism} (i.e., the correlation between γ_{is1m} and γ_{is0m}) for $s = 0, \dots, p$. The variance of the random effects is given by $\text{var}(\gamma_{iskm}) = \sigma_{iskm}^2$ for $k = 0, 1$ and $s = 0, \dots, p$. Thus, the DPM of normals model for each set of random effects, $\gamma_{ism} = (\gamma_{is0m}, \gamma_{is1m}) : s = 0, \dots, p$ is given by

$$\gamma_{ism} \sim \text{MVN}(\text{mean} = \mathbf{0}, \text{Cov} = \sum_i' (\rho_{ism}, \sigma_{is1m}^2, \sigma_{is0m}^2));$$

$$(\rho_{ism}, \sigma_{is1m}^2, \sigma_{is0m}^2) | G \sim G; G \sim \text{DP}(\alpha, G_0),$$

where α is the concentration parameter (unknown) and G_0 is the base distribution of the Dirichlet process. In the below, we assume the base distribution is the product of two inverse gamma distributions (for the variances) and a uniform distribution for the correlation.

By assuming a DP prior on the covariance parameters $(\rho_{ism}, \sigma_{is1m}^2, \sigma_{is0m}^2)$, we essentially classify them into several clusters. Define $C_1, C_2, \dots, C_n$ to be the classification variables induced by DP prior for n subjects under study and c denotes a generic value of C_i . The modeling approach can be shown to be based on the following hierarchical model (Escobar and West, 1995) for each $s = 0, \dots, p$,

$$\begin{aligned}
\gamma_{ism}|C_i=c &\sim N(\mathbf{0}, \sum'(\rho_{ismc}, \sigma_{is1mc}^2, \sigma_{is0mc}^2)), \\
C_i|\pi &\sim \sum_{j'=1}^{N_1} \pi_{j'} \delta_{j'}(\cdot), \\
\sigma_{iskmc}^2 &\sim IG(\nu_1, \nu_2), \text{ for } k=0, 1, \\
\rho_{ismc} &\sim \text{Uniform}(-1, 1), \alpha \sim \text{Gamma}(\kappa_1, \kappa_2),
\end{aligned}$$

where $IG(\nu_1, \nu_2)$ represents an inverse gamma distribution with $\text{mean} = \frac{\nu_2}{\nu_1 - 1}$ and the weights $\pi_{j'}$ are given by Sethuraman's (1994) stick breaking representation of the Dirichlet process given in (13). Note we could also allow dependence between the random effects for each response type. However to not further complicate the overall covariance model here, we assume the above block diagonal structure. In the analysis of the data example in Section 5, we compare to an unstructured covariance matrix.

3.5 Posterior Computations

Full details on the Markov chain Monte Carlo (MCMC) computational algorithm, including recommendations for truncating Dirichlet process priors which are necessary for the models for the random effects distribution and the residual covariance structure can be found in the supplementary materials.

4 Simulation Studies

We perform two simulation studies to investigate the operating characteristics of the proposed modelling approach. Although our approach of modelling simultaneous covariance matrices is designed for irregular sparse longitudinal data, for simplicity in the simulations we consider regular bivariate longitudinal data.

4.1 Simulation Study 1

We consider a sample of 75 subjects classified into 3 (related) groups ($M=3$) with sizes 20, 25 and 30 respectively. A bivariate response is measured for each subject at 5 different time points. We consider group specific linear functions as the general effect of time and simulate data using the following model:

$$Y_{ikm}(t) = f_{km}(t) + \gamma_{i0km} + \gamma_{i1km}t + e_{ikm}(t), \quad (15)$$

where $f_{km}(t) = a_{km} + b_{km}t$, $t = 1, \dots, 5$ and $m = 1, 2, 3$. We assume that the vector of random effects γ_{im} has a (mixture of normals) density with the exact form given in the supplementary materials. The parameters of the residual covariance matrices across the different groups are (somewhat) similar, but not the same. The values of the group-specific general effect of time and the group specific residual covariance matrices can also be found in the supplementary materials.

For the residual covariance parameters $(\eta^{(1)}, \eta^{(0)}, \theta^{(1)}, \theta^{(0)}, \lambda, \delta)$, we consider two alternative priors: priors as specified in Sections 3.1 and 3.2 with (i) zero-inflated normal base distributions and (ii) non-atomic base distribution (by removing the point mass at zero). For the random effects γ_{im} , we consider two different specifications; (i) a (multivariate) normal prior with mean 0 and unstructured covariance matrix Σ and (ii) Dirichlet process mixture of normals prior (as given in Section 3.4). We also fit models with the same residual covariance structure for all the groups and distinct residual covariance structure for each group. The different priors specifications are compared in terms of estimating Σ_m by

computing the risk associated with Bayes estimators under the following two common loss functions for covariance matrices (Yang and Berger, 1994)

$$L_1(\Sigma_m, \widehat{\Sigma}_{m1}) = \text{tr}(\Sigma_m^{-1} \widehat{\Sigma}_{m1}) - \log |\Sigma_m^{-1} \widehat{\Sigma}_{m1}| - P,$$

$$L_2(\Sigma_m, \widehat{\Sigma}_{m2}) = \text{tr}[(\Sigma_m^{-1} \widehat{\Sigma}_{m2} - I)^2],$$

where the matrix Σ_m is of order $P \times P$. Since we have a set of covariance matrices from different groups under study, we compute the weighted average of the losses for group-specific covariance matrices where weights are simply proportional to the groups' sample sizes (Gaskins and Daniels, 2012).

We generate 50 datasets. For each dataset, we run the MCMC algorithm for 55,000 iterations and discarded the first 5,000 iterations as burn-in. We thin the chains by keeping every 10th iteration. Model parameters are estimated as their respective posterior means. Specifying $N_\lambda = N_\delta = N_{\eta(1)} = N_{\eta(0)} = N_{\theta(1)} = N_{\theta(0)} = 20$ provided an adequate truncation approximation error (smaller than 0.01).

Table 1 shows the risks for estimating the set of Σ_m . The common Σ specifications with the normal random effects perform worst in terms of risk. Group-specific specifications performs better than the common Σ specifications. The estimated risks are higher for the normal random effects model. The proposed priors for the covariance parameters (as described in Sections 3.1 and 3.2) give lower estimated risks overall and lower than the non-atomic priors.

We also assess how well the proposed approach estimates the order of the polynomials in equations (7)–(12). For each simulated dataset, the posterior mode is taken as the optimal order. Figure S.1 in the supplementary materials shows the posterior probabilities for g and h . In over 70% of the datasets, the posterior mode corresponds to the true order and it is off by more than 1 in less than 10% of the simulated datasets.

Finally, we investigate the effect of the specification of the priors for the covariance structure on the estimated mean trajectories for each group. Table 2 shows the average width of the 95% credible intervals for a_{km} and b_{km} and the average mean effect parameters for group-1 when $k=1$. Corresponding bias and the estimated coverage probability of the credible intervals are also provided. These tables show that the proposed approach provides estimates with low biases and shorter credible intervals with comparable coverage probabilities. Similar results were observed for the other groups (results not shown).

4.2 Simulation Study 2

For the second simulation, we simulate data using equation (15) with the same values for the parameters for the general effect of time and the residual covariance matrices. For the random effects distribution, we assume a multivariate normal with an unstructured covariance matrix, Σ_{im} with a Wishart($V, 5$) prior; V is a 4×4 matrix with diagonal elements 2 and off-diagonal elements 1. We generate 50 datasets and compute risks using loss functions L_1 and L_2 . Table 3 shows the estimated risks for true model and the proposed DPM random effects model. We note that the true model provides slightly lower risk but the improvement is minimal. We also investigate the effect of the covariance structure on the group-specific estimated mean trajectories. Table S.1 (in the supplementary materials) shows estimated bias, average width of 95% CI and the respective coverage probabilities for a_{km} and b_{km} when $k=1$ for different models. Comparisons are similar to those in Simulation 1. Thus, our approach to modelling the covariance structure that assumes independence

between the random intercept and slope results in minimal increases in risk for estimating the covariance and mean structures. For the data example in Section 5, we do fit a DPM with an unstructured covariance matrix for comparison.

We also compare the performance of the proposed approach of modeling the covariance structure with a simpler and more traditional approach. For the simple approach, the

covariance matrix Σ is modelled as $\Sigma = \begin{bmatrix} \sigma_1^2 & \sigma_{10} \\ \sigma_{10} & \sigma_0^2 \end{bmatrix} \otimes \mathbf{Corr}$, where σ_1 and σ_0 denote respectively the standard deviation for two response features and σ_{10} denotes the covariance between them. The matrix $\mathbf{Corr}$ models the serial correlation structure using a first order autoregressive model; thus, the (j, j') -th element of this matrix is $\rho^{|t_j - t_{j'}|}$. We also consider group-specific covariance modeling where Σ_m is modeled by the group-specific parameters σ_{m1} , σ_{m0} , σ_{m10} and ρ_m , for $m = 1, 2, 3$. Estimated risks for these simpler models for the true random effects structure are given in Table 3. We note that the estimated risks are much higher for these simpler models. Also in Table S.1, we note that the performance of these simpler approaches of modeling the covariance matrices are worse in terms of bias and width of credible intervals compared to our proposed semiparametric approach.

5 Data Analysis

We analyze data from the Framingham Heart Study (FHS) using the proposed methodology. The longitudinal blood pressures (systolic and diastolic) were measured at subject-specific time points with the number of measurements varying from 3 to 19. We denote $k=1$ for systolic and $k=0$ for diastolic blood pressure.

5.1 Model and Priors

We fit the following model:

$$Y_{ikm}(t_{ij}) = f_{km}(t_{ij}) + \gamma_{i0km} + \gamma_{i1km}t_{ij} + e_{ijkm}, \quad (16)$$

where $f(\cdot)$ is the average effect of time and γ_{i0km} and γ_{i1km} are the subject-specific random intercept and random slope, respectively. The residual vectors $\mathbf{e}_{im} = [\mathbf{e}_{i1m}^T, \mathbf{e}_{i0m}^T]^T$ are assumed to follow $\text{MVN}(0, \Sigma_{im})$ as described in Section 2. We specify independent gamma(1,1) priors for α_x and β_x ; ($x \in \{\lambda, \delta, \eta^{(1)}, \eta^{(0)}, \theta^{(1)}, \theta^{(0)}\}$), inverse gamma(2, 2.5) priors for σ_λ^2 and σ_δ^2 and inverse gamma (1.5, 2.5) priors for $\sigma_{\eta^{(1)}}^2$, $\sigma_{\eta^{(0)}}^2$, $\sigma_{\theta^{(1)}}^2$ and $\sigma_{\theta^{(0)}}^2$. The distribution for the random effects γ_{im} is the same as described in Section 3.4 with $\nu_1=1.5$, $\nu_2=2$ and $\kappa_1=1$, $\kappa_2=1.5$.

We model $f(\cdot)$ by penalized splines of degree r with knots ($\tau_1, \dots, \tau_K$). The response-specific group means at different time points can be expressed as the following non-parametric regression model:

$$f_{km}(t_{ij}) = b_{km0} + b_{km1}t_{ij} + b_{km2}t_{ij}^2 + \dots + b_{kmr}t_{ij}^r + \sum_{s=1}^K c_{kms}(t_{ij} - \tau_s)_+^r, \quad (17)$$

where $(x)_+^r = x^r I(x > 0)$ and ($\tau_1 < \tau_2 < \dots < \tau_K$) is a fixed set of knots. The K knots are selected from the evenly spaced sample quantiles of time following Ruppert, Wand and Carroll (2003). In this analysis, we consider 5 knots ($K=5$) and third degree splines ($r=3$); this specification was chosen by picking the specification with the smallest DIC (defined in Section 5.2). Of course, we could have left these as unknown parameters in the model and specified priors for them.

5.2 Computational Details and Model Selection

Starting values for the MCMC algorithm are drawn from their respective prior distributions and we run chains with 110,000 iterations. The initial 10,000 burn-in iterations are discarded and the chains are thinned by keeping every 10th iteration. Convergence of the chains are assessed following Brooks and Gelman (1998) by considering 5 independent chains for each parameter. All the estimated potential scale reduction factors were found to be smaller than 1.1.

For the matrix stick-breaking priors, we assume $N_{\lambda} = N_{\delta}$ and $N_{\eta(1)} = N_{\eta(0)} = N_{\theta(1)} = N_{\theta(0)}$ and needed $N_{\lambda} = N_{\delta} = N_{\eta(1)} = N_{\eta(0)} = N_{\theta(1)} = N_{\theta(0)} = 25$ to keep the expected approximation error below our desired threshold of 0.01; for the DPM, we needed $N_1 = 30$.

Sensitivity analyses were performed to investigate the effects of different priors on the parameters estimates. As shown in the supplementary materials (Table S.2), for different gamma and/or inverse gamma priors for the hyperparameters, the estimates were very similar for $\eta^{(1)}$. Similar results were obtained for the other parameters (not shown).

We consider three different specifications for the covariance matrices of the residuals (common Σ , group-specific Σ_m , and the proposed prior) and three different specifications for the random effects (a normal distribution, a DPM of normals with unstructured covariance matrix, and a DPM of normal distributions with block diagonal covariance structure as described in Section 3.4). We compare the fit using the conditional deviance information criteria (DIC) proposed in Celeux et al. (2006). This DIC is based on the conditional likelihood $l(\mathbf{Y}|\boldsymbol{\gamma})$ and is given by $\text{DIC} = -4\mathbf{E}[\log l(\mathbf{Y}|\boldsymbol{\gamma})] + 2\log l(\mathbf{Y}|\hat{\boldsymbol{\gamma}})$, where $\boldsymbol{\gamma}$ denotes the vector of random effects and $\hat{\boldsymbol{\gamma}}$ is the corresponding estimate (posterior mean).

Table 4 provides DIC and the corresponding effective number of parameters (p_D) from different model combinations. The common Σ specification with normal random effects provides the largest DIC value while our proposed approach with block diagonal DPM random effects gives the smallest DIC.

The proposed modeling approach also estimates the posterior distribution of the optimal polynomial orders g and h (Figure 1); here we obtain $g=2$ and $h=3$ as the optimal order for the polynomial functions in (7)–(12). But note for inference, we integrate over the uncertainty in the order parameters.

5.3 Mean Structure Findings

Figures 2 and S.2 (in the supplementary materials) show the posterior mean of the curves for the systolic and diastolic blood pressure for the three groups with 95% pointwise credible intervals. Note that the curves for three BMI groups are quite different (minimal overlap in credible intervals). For the subjects with high BMI, diastolic blood pressure decreases until age 70 and then slightly increases over time although the systolic pressure increases until age 55 or so and then follows a decreasing trend. For the medium BMI group, the trends are slightly increasing but consistent for both response features. For the low BMI group, although the trend is increasing over time for the diastolic blood pressure, the systolic one shows an increasing trend until age 70 and then decreases.

5.4 Covariance Structure Findings

For the low BMI group, the correlation of the intercepts and slopes between two responses are estimated as 0.73 and 0.68 with the respective 95% CIs (0.53, 0.84) and (0.48, 0.81). For the medium BMI group, these correlations are 0.47 and 0.39 with respective CIs (0.23, 0.58) and (0.21, 0.53). For the high BMI group, the correlations are 0.71 and 0.76 with (0.49,

0.87) and (0.55, 0.89) as the respective CIs. We note that the estimated correlation values indicate high correlation between the systolic and diastolic BP for low and high BMI group. However, for the medium BMI group, the response features were only moderately correlated.

Figure 3 shows the grouping nature of the proposed priors for the parameter $\eta^{(1)}$. We show the posterior probabilities of $Pr(\eta_{j'm}^{(1)} = \eta_{j'm'}^{(1)})$ for each m, m', j' combinations. Larger boxes indicate higher grouping probabilities, with the boxes on the $y = x$ diagonal corresponding to probability one. We notice higher matching probabilities for the medium (group-2) and low BMI group (group-3). For other covariance parameters, we observe a similar pattern (as shown in supplementary materials, Figures S.3–S.7). Our priors provide an efficient data-dependent way to find such similarity in structure across groups.

5.5 Comparison with Simpler Models

To investigate the impact of the proposed covariance modeling on the estimated group-specific mean curves, we compare the results from the above analysis with the results from simpler, more standard, models. In (16), the subject-specific covariance matrix Σ_i is now

modelled using a Kronecker product structure, $\Sigma_i = \begin{bmatrix} \sigma_1^2 & \sigma_{10} \\ \sigma_{10} & \sigma_0^2 \end{bmatrix} \otimes \mathbf{Corr}_i$, similar to the specification in Simulation 2. The (j, j') -th element of the matrix $\mathbf{Corr}_i$ is $\rho^{t_{ij}-t_{ij'}}$. Penalized splines are used to model group-specific $f(\cdot)$ as in (16). For modeling the random effects γ_{im} , we consider a (multivariate) normal prior with mean 0 and an unstructured covariance matrix S . We also consider group-specific covariance modeling, Σ_{im} .

We compute the DIC for these simpler models. For common Σ , the DIC is 39,504; for group-specific covariance, the DIC is 38,317. We note that these DIC values are much higher than the DIC value for our proposed approach of semiparametric covariance modeling as given in Table 4.

The estimated group-specific mean curves for the simpler models are shown in Figures 4 and S.8 (in the supplementary materials). Solid curves are for common Σ while the dotted curves are for group-specific Σ 's. Note that these curves are quite different from Figures 2 and S.2. For example, in Figure 2, the systolic pressure for the high BMI group shows an increasing trend until age 55 and then decreases. This feature is not reflected in Figure 4, rather the solid curve in Figure 4 shows a decreasing trend throughout. Similarly, in Figure S.2, for low BMI group, the diastolic pressure trend is increasing over time. But in Figure S.8, although the dotted curve shows slightly increasing trend, the solid curve is more or less constant near age 65. This comparison shows the proposed approach provides different conclusions than the traditional approach of covariance modeling and provides improved fit as measured by DIC.

6 Discussion

In this article, we address the challenge of estimating the covariance structure from multiple groups when a bivariate trait is measured at subject-specific time points. In clinical trials or genetic experiments, this is not an uncommon scenario. Traditional approaches assume either the same covariance matrix for each group (homoscedasticity) or group-specific covariance matrices and often model the traits separately. Additionally, we address the issue of choosing the structure/order of the covariance matrix, avoiding the need for two-step procedures and appropriately accounting for the uncertainty. We assumed a block diagonal structure for the random effects covariance matrix in the DPM (which fit our data example best); however, a more general structure can easily be fit.

Here, we have considered the situations where both the response features are measured at each observed time point. For traits like systolic and diastolic blood pressure, this assumption is typically true, but it might not always be the case in other applications. If the missingness is ignorable, one can simply add a data augmentation step. For non-ignorable missingness, more complex approaches are needed (Daniels and Hogan, 2008). In addition, here we assume subjects in different groups are independent; this might not be the case in genetic experiments where there are often more than one participant from the same family.

A focus of biological research is to understand the biological system as a whole rather than studying different traits separately. Statistically, this requires a joint analysis of observed traits and consequent events like death or recovery. For example, higher blood pressure at younger age might result in cardiovascular disease at later age. But the expected time for the events (diseases) to occur can vary from one group of people to another (where subjects may be grouped on the basis of biomarkers). For such experiments, our proposed approach will be extremely useful for modeling observed bivariate (or multivariate) longitudinal traits.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Daniels was partially supported by NIH CA85295. The authors thank Prof. Rongling Wu, Penn State University, for sharing the data from the Framingham Heart Study.

References

- Bandyopadhyay D, Lachos VH, Abanto-Valle CA, Ghosh P. Linear Mixed Models for Skew-Normal/Independent bivariate responses with an application to Periodontal Disease. *Statistics in Medicine*. 2010; 29:2643–2655. [PubMed: 20740568]
- Brooks SP, Gelman A. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*. 1998; 7:434–455.
- Celeux G, Forbes F, Robert CP, Titterton DM. Deviance Information Criteria for Missing Data Models. *Bayesian Analysis*. 2006; 1(4):651–674.
- Daniels MJ, Pourahmadi M. Bayesian analysis of covariance matrices and dynamic models for longitudinal data. *Biometrika*. 2002; 89:553–566.
- Daniels, MJ.; Hogan, JW. *Strategies for Bayesian Modeling and Sensitivity Analysis*. Chapman and Hall (CRC Press); 2008. *Missing data in longitudinal studies*.
- Das K, Li J, Fu G, Wang Z, Wu R. Genome-wide association studies for bivariate sparse longitudinal data. *Human Heredity*. 2011; 72:110–120. [PubMed: 21996601]
- Das K, Li R, Huang Z, Gai J, Wu R. A Bayesian Framework for Functional Mapping through Joint Modeling of Longitudinal and Time-to-Event Data. *International Journal of Plant Genomics*. 2012.1155/2012/680634
- Das K, Li J, Fu G, Wang Z, Li R, Wu R. Dynamic semi-parametric Bayesian models for genetic mapping of complex traits with irregular longitudinal data. *Statistics in Medicine*. 2013a; 32(3): 509–523. [PubMed: 22903809]
- Das K, Li R, Sengupta S, Wu R. A Bayesian Semi-parametric Model for Bivariate Sparse Longitudinal Data. *Statistics in Medicine*. 2013b; 32(22):3899–3910. [PubMed: 23553747]
- Dunson DB, Xue Y, Carin L. The matrix stick-breaking process: Flexible Bayes meta-analysis. *Journal of American Statistical Association*. 2008; 103(481):317–327.
- Durban M, Harezlak J, Wand MP, Carroll RJ. Simple fitting of subject-specific curves for longitudinal data. *Statistics in Medicine*. 2005; 24:1153–1167. [PubMed: 15568201]
- Escobar M, West M. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*. 1995; 90:577–588.

- Gaskins JT, Daniels MJ. A nonparametric prior for simultaneous covariance estimation. *Biometrika*. 2012; 100:125–138.
- Ghosh P, Hanson TA. Semiparametric Bayesian approach to multivariate longitudinal data. *Australian and New Zealand Journal of Statistics*. 2010; 52:275–288. [PubMed: 21731424]
- Ishwaran H, James LF. Gibbs sampling methods for stick-breaking priors. *Journal of American Statistical Association*. 2001; 96(453):161–173.
- MacEachern SN, Mueller P. Estimating mixture of Dirichlet process models. *Journal of Computational and Graphical Statistics*. 1998; 6:223–238.
- Neal RM. Slice sampling. *The Annals of Statistics*. 2003; 31(3):705–767.
- Pan J, Mackenzie G. On modelling mean-covariance structures in longitudinal studies. *Biometrika*. 2003; 90:239–244.
- Pourahmadi M. Joint mean-covariance model with applications to longitudinal data: Unconstrained parameterisation. *Biometrika*. 1999; 86:677–690.
- Ruppert, D.; Wand, MP.; Carroll, RJ. *Semiparametric Regression*. Cambridge University Press; New York: 2003.
- Sethuraman J. A constructive definition of Dirichlet priors. *Statistica Sinica*. 1994; 4:639–650.
- Sithole JS, Jones PW. Bivariate Longitudinal Model for Detecting Prescribing Change in Two Drugs Simultaneously with Correlated Errors. *Journal of Applied Statistics*. 2007; 34:339–352.
- Sy JP, Taylor JMG, Cumberland WG. A Stochastic model for the analysis of bivariate longitudinal AIDS data. *Biometrics*. 1997; 53:542–555. [PubMed: 9192450]
- Thiebaut R, Jacqmin-Gadda H, Chene G, Leport C, Commenges D. Bivariate linear mixed models using SAS PROC MIXED. *Computer Methods and Programs in Biomedicine*. 2002; 69:249–256. [PubMed: 12204452]
- Wolk R, Shamsuzzaman ASM, Somers VK. Obesity, Sleep Apnea and Hypertension. *Hypertension*. 2003; 42:1067–1074. [PubMed: 14610096]
- Wu WB, Pourahmadi M. Nonparametric estimation of large covariance matrices for longitudinal data. *Biometrika*. 2003; 90:831–844.
- World Health Organization. Technical Report Series no 854. Geneva: World Health Organization; 1995. *Physical Status: the Use and Interpretation of Anthropometry*. WHO Global Database on body mass index 2006. Retrieved from <http://www.who.int/bmi/index>
- Yang R, Berger JO. Estimation of a covariance matrix using the reference prior. *The Annals of Statistics*. 1994; 22:1195–1211.

Figure 1.
Posterior probabilities for g and h in the FHS data

Figure 2.
Posterior means and 95% point wise credible intervals of the longitudinal mean trajectories for systolic blood pressure

Figure 3.

The posterior probabilities of matching for $\eta_0^{(1)}$, $\eta_1^{(1)}$, $\eta_2^{(1)}$ and $\eta_3^{(1)}$ respectively. The size of the boxes are proportional to the posterior probabilities of matching.

Figure 4. Posterior means of the longitudinal mean trajectories for systolic blood pressure for the simpler models

Risk estimates from different model specifications in the simulation study 1. L_1 and L_2 are the two loss functions introduced in Section 4.

Table 1

η_m	Priors			Risk	
	θ_m	λ_m	δ_m	L_1	L_2
Normal Random Effects					
proposed	proposed	proposed	proposed	1.04	1.35
proposed	proposed	non-atomic	non-atomic	2.38	3.88
non-atomic	non-atomic	proposed	proposed	2.63	3.74
non-atomic	non-atomic	non-atomic	non-atomic	3.55	4.98
Group-specific flat				4.92	6.79
Common Σ -flat				6.31	8.64
DPM Random Effects					
proposed	proposed	proposed	proposed	0.73	0.90
proposed	proposed	non-atomic	non-atomic	1.25	2.09
non-atomic	non-atomic	proposed	proposed	1.31	2.26
non-atomic	non-atomic	non-atomic	non-atomic	2.57	2.75
Group-specific flat				3.11	5.08
Common Σ -flat				4.47	6.83

Table 2

Effects of different covariance models on the average mean parameters for group-1 with different random effects for the response feature $k=1$ in the simulation study 1.

Covariance	a_1		b_1	
	Width of C.I.	Bias	Coverage prob.	Width of C.I.
Normal Random Effects				
Common	0.58	0.63	0.96	0.49
Group-specific	0.32	0.38	0.95	0.31
Proposed	0.21	0.17	0.95	0.18
DPM Random Effects				
Common	0.34	0.41	0.96	0.36
Group-specific	0.18	0.35	0.95	0.29
Proposed	0.10	0.16	0.95	0.14

Table 3

Risk estimates from different model specifications in the simulation study 2. L_1 and L_2 are the two loss functions introduced in Section 4.

η_m	Priors			Risk		
	θ_m	λ_m	δ_m	L_1	L_2	
MVN-UN Random Effects model						
proposed	proposed	proposed	proposed	0.69	0.87	
proposed	proposed	non-atomic	non-atomic	1.19	1.88	
non-atomic	non-atomic	proposed	proposed	1.28	2.17	
non-atomic	non-atomic	non-atomic	non-atomic	2.35	2.71	
	Group-specific flat			3.03	5.10	
	Common Σ -flat			4.19	6.74	
	(Simple) Group-specific flat			5.34	7.18	
	(Simple) Common Σ -flat			5.96	8.42	
DPM-BD Random Effects model						
proposed	proposed	proposed	proposed	0.78	0.95	
proposed	proposed	non-atomic	non-atomic	1.30	1.99	
non-atomic	non-atomic	proposed	proposed	1.36	2.38	
non-atomic	non-atomic	non-atomic	non-atomic	2.44	2.92	
	Group-specific flat			3.25	5.17	
	Common Σ -flat			4.34	7.08	

UN= Unstructured covariance matrix. BD= Block Diagonal covariance matrix

Table 4

DIC from different model specifications in the FHS data.

Covariance structure	Random effects	p_D	DIC
Common Σ -flat	MVN-UN	42.3	33,714
Group-specific flat	MVN-UN	51.7	34,901
Proposed structure	MVN-UN	46.8	30,182
Common Σ -flat	DPM-UN	48.2	34,418
Group-specific flat	DPM-UN	52.5	33,122
Proposed structure	DPM-UN	44.7	30,018
Common Σ -flat	DPM-BD	39.5	31,618
Group-specific flat	DPM-BD	45.2	34,013
Proposed structure	DPM-BD	41.3	29,816

p_D = Effective number of parameters in DIC. UN= Unstructured. BD= Block Diagonal