

A Functional Random Effects Model for
Flexible Assessment of Susceptibility in
Longitudinal Designs

Brent A. Coull*

*Harvard School of Public Health, bcoull@hsph.harvard.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/harvardbiostat/paper90>

Copyright ©2008 by the author.

A Functional Random Effects Model for Flexible Assessment of Susceptibility in Longitudinal Designs

Brent A Coull

Department of Biostatistics
Harvard School of Public Health
655 Huntington Avenue
Boston, MA 02115

bcoull@hsph.harvard.com
phone: 617.432.2376
fax: 617.432.5619

SUMMARY

In many biomedical investigations, a primary goal is the identification of subjects that are susceptible to a given exposure or treatment of interest. We focus on methods for addressing this question in longitudinal studies when interest focuses on relating susceptibility to a subject's baseline or mean outcome level. In this context, we propose a functional random effects model that relaxes simplistic assumptions in existing mixed models and yields an estimate of the functional form of this relationship. We propose a penalized spline formulation for the nonparametric function that represents this relationship, and implement a fully Bayesian approach to model fitting. We investigate the frequentist performance of our method via simulation, and apply the model to data on the effects of particulate matter on coronary blood flow from an animal toxicology study. The general principles introduced here apply more broadly to settings in which interest focuses on the relationship between baseline and change over time.

keywords: latent variable, particulate matter, Laird-Ware model, random intercept - random slope model; semiparametric regression, penalized spline.

1 Introduction

A key advantage of longitudinal study designs is their ability to identify subject-specific factors that influence an individual's susceptibility to a time-varying exposure or treatment. In some settings, scientific interest focuses on whether or not pre-existing disease, as manifested by either baseline level or the mean level of the outcome over the course of the study, is one such factor. For example, Dubowsky et al. (2006) investigated whether within-subject associations between c-reactive protein (CRP) levels and particulate matter (PM) exposure in a cohort of elderly patients differed between subjects with high and low levels of CRP, and Zanobetti et al. (2008) asked a similar question when investigating associations between PM and T-wave alternans, an endpoint reflecting cardiac electrical instability, in a cohort of patients recovering from a coronary event. Similarly, Bartoli et al. (2008) assessed whether associations between cardiac blood flow and PM exposure in an animal model of coronary artery disease was more severe in myocardial tissue exhibiting a low level of flow under control conditions, compared with myocardial tissue exhibiting normal flow under control conditions.

Mixed effect regression models represent an attractive approach to addressing this question due to the subject-specific interpretation of the random effects. In studies focusing on susceptibility, the random intercepts and slopes model is a natural starting point. This class of models specifies susceptibility as being randomly distributed across the population of interest, with the model most often assuming a bivariate normal distribution for the subject-specific intercepts and slopes and parameterizing the association between baseline and susceptibility as the correlation parameter in the variance-covariance matrix.

The bivariate normal assumption in a standard random intercepts and slopes model implies that the conditional model for the slopes given the intercepts is a simple linear model. In cases in which interest focuses on estimating the functional form of this relationship, one does not wish to make this assumption *a priori*, as the shrinkage towards linearity induced by the linear mixed model may obscure the true relationship. Applied work generalizing this assumption is based on relatively simple approaches, such as classifying subjects according to percentiles of the distribution of

the average response level and fitting models that include interactions between subject class and exposure (Dubowsky et al. 2006; Zanobetti et al. 2008; Bartoli et al. 2008). However, this approach assumes a step function for this relationship between random intercepts and slopes, which may also be overly restrictive in some situations.

More modern approaches relax the multivariate normal assumption for the random effects, yielding inferences robust to this choice. For instance, one can use a copula model (Nelsen 1999) to specify the marginal distributions of the random effects without specifying a linear form for the intercept-slope association. Another approach that avoids the multivariate normal assumption is the conditional score method of Li, Zhang and Davidian (2004). These approaches, however, do not *directly estimate* the functional form of the intercept-slope relationship, making them less useful when this relationship is of primary scientific interest. Our interest lies in methods that yield such an estimate but do not assume a functional form for this relationship *a priori*.

In this article we propose a functional random effects model to directly address this scientific question. The model generalizes the standard random intercepts and slopes model by replacing the linear model for the slopes given the intercepts with an unspecified, smooth function that is nonparametrically estimated from the data. The proposed model generalizes existing additive mixed model formulations by allowing one to estimate a smooth function of a latent variable, specifically random intercepts, rather than observed covariates. Because we formulate unknown regression functions on unobserved, latent variables, an immediate question that arises is whether all model parameters are identifiable. In Section 3 we show that the model is identifiable, and propose a Bayesian approach that samples from the joint posterior distribution of the model parameters via Markov chain Monte Carlo sampling.

Our proposal is related to that of Berry, Carroll, and Ruppert (2002), who considered regression spline models for measurement error problems. These authors considered a univariate outcome with a single covariate, whereas the measurement error in our problem arises out of a longitudinal setting in which the latent model parameters are covariates in a second stage nonparametric regression model. In this sense this work extends that of Wang, Wang, and Wang (2000) and Li, Zhang and Davidian (2004) to the nonparametric regression setting, using the popular mixed model formulation of penalized splines (Brumback, Ruppert, and Wand 1999).

The remainder of the paper is organized as follows. Section 2 describes data from an environmental toxicology study that partially motivated this work and generated the data considered in this article. Section 3 reviews the random intercepts and slopes linear mixed model (LMM) of Laird and Ware (1982), presents the functional random effects extension, considers model identifiability, and presents a Bayesian approach to model fitting. Section 4 presents a simulation study that compares the performance of the proposed approach to *ad-hoc* approaches based on existing models. Section 5 applies the methods to analyze the coronary blood flow data, and Section 6 presents further discussion.

2 Health Effects of PM on Coronary Ischemia

We develop these methods in the context of analyzing the effects of particulate matter (PM) exposure on coronary blood flow in an animal toxicology study. Associations between increased deaths due to ischemic heart disease and elevated levels of ambient PM are well-established. Recent laboratory studies seek to pinpoint the biologic mechanisms underlying these effects. In a recent experiment conducted at the Harvard School of Public Health, investigators (Bartoli et al. 2008) used fluorescent microspheres to repeatedly measure regional myocardial blood flow (hereafter blood flow) in a conscious canine model of myocardial ischemia via coronary occlusion. In layman's terms, investigators implanted laboratory animals with arterial balloons, which could then be inflated to occlude that artery and initiate an ischemic episode – that is, restrict blood flow into the heart. First assessed at baseline (no occlusion), animals were then exposed under occlusion conditions to filtered air and concentrated ambient particles (CAPs) in a crossover protocol (five hr/day exposures, 3 filtered air, 3 CAPs exposures). During each exposure occasion, investigators injected colored microspheres into an animal's bloodstream, with these spheres eventually traveling to the heart. Each exposure corresponded to a different microsphere color. After the entire study, multiple pieces of myocardium tissue from each animal were assessed for blood flow under baseline, filtered air and CAPs conditions by assessing the number of spheres with color belonging to that particular exposure, normalized by a reference blood flow rate, present in each tissue piece. A primary question of interest was

whether or not there was a difference in blood flow between filtered air and CAPs exposures overall. Bartoli et al. (2008) reported a strong overall negative association between blood flow and exposure, suggesting ischemia as a potentially important biologic mechanism of the cardiovascular effects of particulate matter (PM).

A second major question of interest, and the one we focus on here, is whether exposure affected blood flow inside and outside the ischemic zone differentially. More specifically, define the ischemic zone as the region of the myocardium affected by the occluded artery, independent of exposure. Investigators are interested in knowing whether or not the effect of exposure on blood flow in ischemic tissue was significantly greater than that in the tissue outside of the ischemic zone or whether the maximal exposure effect occurs on the edges of this zone. The former possibility would imply that exposure could make the effects of an adverse cardiac event more severe, whereas the latter possibility would imply that exposure makes the affected area larger, but not necessarily more severe. Thus, in this setting, interest focuses on the susceptibility of a given piece of myocardium tissue as a function of the degree of ischemia induced in that piece by occlusion. Such a determination would yield further insight into the exact mechanism of PM effects in a potentially susceptible subpopulation, that of patients with coronary artery disease.

One complication in this analysis is that an objective determination of whether or not a tissue piece is in the ischemic zone is not available. Instead, one must make this determination by comparing blood flow under baseline (non-occlusion) and control (i.e. under filtered air exposure) occlusion conditions. Bartoli et al. (2008) estimated the interaction between CAPs exposure and whether or not a tissue piece was in the ischemic zone, where a piece was classified as being in the ischemic zone if the blood flow under filtered air conditions was less than blood flow at baseline. This classification was motivated by the consideration that blood flow is likely to increase outside of the ischemic zone since the damaged tissue will cause the blood that does reach the heart to disperse to the unaffected zones of the heart. This assumption assumes two levels of exposure susceptibility, one inside the ischemic zone and another outside this affected area. One question that immediately arises is whether or not this categorization is sufficient, or whether it obscures patterns of susceptibility in different types of tissue. Therefore, we seek to estimate the form of relationship more

flexibly.

3 Functional Random Effects Model

Let Y_{it} and E_{it} denote the outcome and exposure of interest, respectively, for subject i , $i = 1, \dots, n$ at time $t = 1, \dots, T_i$. Ignoring other covariates for now, the Laird-Ware linear mixed model (Laird and Ware 1982) with random intercepts and slopes is

$$Y_{it} = b_{0i} + b_{1i}E_{it} + \varepsilon_{it} \quad (1)$$

where

$$\mathbf{b}_i = (b_{0i}, b_{1i})^\top \stackrel{iid}{\sim} N \left(\begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix}, \begin{bmatrix} \sigma_0^2 & \sigma_0\sigma_1\rho \\ \sigma_0\sigma_1\rho & \sigma_1^2 \end{bmatrix} \right) \quad (2)$$

and $\varepsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2)$, with $\mathbf{b}_i \perp \varepsilon_{it}$. The bivariate normal assumption (2) for the random effects implies that the random slopes and random intercepts satisfy a linear model, such that model (1) can be re-expressed as

$$Y_{it} = b_{0i} + (\gamma_0 + \gamma_1 b_{0i} + \eta_i) E_{it} + \varepsilon_{it},$$

where $\gamma_1 = (\sigma_1\rho)/\sigma_0$ and $\eta_i \stackrel{iid}{\sim} N(0, \sigma_{1.0}^2)$ with $\sigma_{1.0}^2 = \sigma_1^2(1 - \rho^2)$. Just as the empirical distribution of the best linear unbiased predictions (BLUPs) of these random effects will be shrunk towards normality (Carlin and Louis 2000), the observed relationship between these BLUPs will be shrunk towards linearity, with the amount of shrinkage depending on the influence of this distributional assumption relative to the data. Thus, this shrinkage will be relatively strong in situations where the estimates obtained from the subject-specific data are relatively imprecise, which occurs when there is a relatively small number of repeated measures, the exposure variance is not too large, the heterogeneity in susceptibility is relatively large, or any combination of the above.

We seek to relax the linearity assumption between subject-specific intercepts and slopes in the Laird and Ware (1982) model by specifying a nonparametric term for

this functional form. We propose the functional random effects (FRE) model taking form (1), but with

$$b_{0i} \stackrel{iid}{\sim} N(\beta_0, \sigma_0^2), \quad b_{1i}|b_{0i} \stackrel{iid}{\sim} N(f(b_{0i}), \sigma_{1.0}^2), \quad \text{and} \quad \varepsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2). \quad (3)$$

Here, $f(b)$ is an unspecified smooth function that is estimated from the data. Letting η_i denote the residual error of b_{1i} around its conditional mean $f(b_{0i})$, we assume $b_{0i} \perp \eta_i \perp \varepsilon_{it}$.

We propose a low-rank thin-plate spline representation for $f(\cdot)$,

$$f(b) = \gamma_0 + \gamma_1 b + \sum_{k=1}^K \tilde{u}_k |b - \kappa_k|^3,$$

where $(\gamma_0, \gamma_1, \tilde{u}_1, \dots, \tilde{u}_K)^\top$ is the vector of regression coefficients and $\kappa_1 < \kappa_2 < \dots < \kappa_K$ are fixed knots (Fan, Leslie, and Wand 2007; Crainiceanu, Ruppert, and Wand 2005).

To avoid over-fitting, we penalize the magnitude of the $\{\tilde{u}_k\}$. Let $\mathbf{b}_0 = (b_{01}, \dots, b_{0n})^\top$ and $\mathbf{b}_1 = (b_{11}, \dots, b_{1n})^\top$. Further, let $\mathbf{B}_{n \times 2}$ be the matrix with i^{th} row $B_i = (1 \quad b_{0i})$, $\tilde{\mathbf{Z}}$ be an $n \times K$ matrix with i^{th} row $[|b_{0i} - \kappa_1|^3 \dots |b_{0i} - \kappa_K|^3]$, $\mathbf{\Omega}$ be a $K \times K$ matrix with $(l, k)^{\text{th}}$ element equal to $|\kappa_l - \kappa_k|^3$, $\boldsymbol{\gamma} = (\gamma_0, \gamma_1)^\top$, and $\tilde{\mathbf{u}} = (\tilde{u}_1, \dots, \tilde{u}_K)^\top$. The thin-plate spline penalty is induced by assuming a multivariate normal distribution for $\tilde{\mathbf{u}}$, so that we can write

$$\mathbf{f}(\mathbf{b}_0) = \mathbf{B}\boldsymbol{\gamma} + \tilde{\mathbf{Z}}\tilde{\mathbf{u}}$$

where $\tilde{\mathbf{u}} \sim N\left[\mathbf{0}, \sigma_u^2 \mathbf{\Omega}^{-1/2} (\mathbf{\Omega}^{-1/2})^\top\right]$. Under the reparameterization $\mathbf{u} = \mathbf{\Omega}^{1/2} \tilde{\mathbf{u}}$ and $\mathbf{Z} = \tilde{\mathbf{Z}} \mathbf{\Omega}^{-1/2}$, the model for \mathbf{b}_1 is equivalent to

$$\mathbf{b}_1 = \mathbf{B}\boldsymbol{\gamma} + \mathbf{Z}\mathbf{u} + \boldsymbol{\eta}, \quad \text{cov} \begin{pmatrix} \mathbf{u} \\ \boldsymbol{\eta} \end{pmatrix} = \begin{pmatrix} \sigma_u^2 \mathbf{I}_K & \mathbf{0} \\ \mathbf{0} & \sigma_{1.0}^2 \mathbf{I}_n \end{pmatrix}.$$

While this basis is slightly more complex to compute than the simpler truncated polynomial basis that is popular in penalized spline formulations of smooth functions

(Berry et al. 2002), it is more stable numerically. Our experience suggests that this advantage leads to superior mixing of the Markov Chain Monte Carlo algorithm proposed in Section 2.2.

An advantage of the Laird-Ware model is interpretability of β_0 and β_1 , which represent the mean baseline and exposure slope, respectively, in the study population (i.e. μ_0 and μ_1). In the FRE model, β_0 again represents the mean baseline level, μ_0 . However, the mean exposure slope is $\mu_1 = \int f(b)\phi(b)db$, where $\phi(b)$ is the standard normal density. While this latter quantity is not available in closed form for general f , the proposed Monte Carlo algorithm proposed in Section 3.2 for simulating from the joint posterior of the model parameters yields an estimate of this mean susceptibility value as a by-product.

In standard applications of penalized spline models, the nonparametric term $f(\cdot)$ is a function of an observed covariate, and so knots are typically chosen to be spread throughout the range of the covariate (Ruppert, Wand, and Carroll 2003). An issue that arises in our functional slope formulation is that the argument of $f(\cdot)$, \mathbf{b}_0 , is unobserved. To address this issue, we first fit the Laird-Ware random intercepts and slopes model to the data, obtain the best linear unbiased predictions (BLUPs; Robinson 1991) of the random intercepts from this preliminary fit, and use the approach of Ruppert et al. (2003) treating these predictions as fixed. This approach is similar to the strategy advocated by Berry et al. (2002), who place knots within the estimated range of a latent variable. We investigate the effectiveness of this approximation in a simulation study described in Section 4.

Model (1) and (3) can be easily extended to accommodate systematic subject-specific factors that are associated with the baseline level of the outcome. The general form of the FRE model is

$$Y_{it} = b_{0i} + b_{1i}E_{it} + \mathbf{W}_{it}\boldsymbol{\alpha} + \varepsilon_{it}, \quad (4)$$

where $b_{0i} \sim N(\mathbf{X}_i\boldsymbol{\beta}, \sigma_0^2)$, $b_{1i} | b_{0i} \sim N(f(b_{0i}), \sigma_{1.0}^2)$, with \mathbf{X}_i a $p \times 1$ vector containing between-subject covariates that remain constant throughout the duration of the study and \mathbf{W}_{it} a $q \times 1$ vector of time-varying covariates. This formulation assumes that between-subject factors relate to a subject's overall outcome level, which is in turn

associated with that subject's susceptibility to exposure. This formulation yields a natural partitioning of covariates into between- and within-subject covariates, as recommended by Neuhaus and McCulloch (2006) and others.

3.1 Model Identifiability

An important question concerns identifiability of model (4). This issue amounts to a problem of identifiability in a nonparametric regression measurement error setting (Carroll et al. 2004). Consider fixed effects estimates $\{\widehat{b}_{0i}\}$ and $\{\widehat{b}_{1i}\}$ of the subject-specific intercepts and slopes, respectively. Consider estimation of the parameters in the model $\widehat{b}_{1i} = f(b_{0i}) + w_i$. Instead of having in hand the covariate values $\{b_{0i}\}$, we have versions $\{\widehat{b}_{0i}\}$ measured with error. Specifically, we have $\widehat{b}_{0i} = b_{0i} + v_i$, where v_i is estimation error. Berry et al. (2002) considered this measurement error framework when one has independent replicates of the covariate measured with error (b_{0i} in this case), and noted that the model was identifiable in this case. Carroll et al. (2004) showed that the model is identifiable if, instead of independent replicates, an instrumental variable for the true values of the covariates are available. That is, the model is identifiable if there exists a variable S such that $S_i = \alpha_0 + \alpha_1 b_{0i} + \nu_i$, where ν_i is independent of (b_{0i}, v_i, w_i) , and $\text{Cov}[b_{0i}, f(b_{0i})] \neq 0$.

In the functional random effects model, we do not have either replicates or an instrument for the random intercepts, but the model is identifiable due to the form of the variance of v_i . That is, in the functional slope setting, the error structure takes the form

$$\text{Var}(v_i) = \sigma_\varepsilon^2 (D_i^\top D_i)^{-1},$$

where D_i is the $n_i \times 2$ covariate matrix from the subject-specific regression $Y_{it} = \beta_{0i} + \beta_{1i} X_{it} + \varepsilon_i$. Because σ_ε^2 is estimable, the model is identifiable. See also Wang et al. (2000) and Wall and Amemiya (2000) for discussions of this important difference between this longitudinal setting, in which the error component is identified, and the general measurement error setting.

3.2 Estimation

We take a fully Bayesian approach to model fitting, using Monte Carlo Markov chain (MCMC) methods (Gelfand and Smith 1990) to sample from the joint posterior distribution of all model parameters. Specifically, let $[A]$ and $[A|B]$ represent densities and conditional densities, respectively, and let $N = \sum_{i=1}^n T_i$ denote the total number of observations taken on the n subjects in the study. Define \mathbf{Y} as an $N \times 1$ vector containing all responses, \mathbf{X} as an $n \times p$ matrix with rows \mathbf{X}_i , \mathbf{W} as an $N \times q$ matrix with rows \mathbf{W}_{it} , and \mathbf{E} as an $N \times 1$ vector of all exposures. Further, for notational convenience, let \mathbf{Z}_0 be an $N \times n$ matrix whose i^{th} column contains 1's in those rows corresponding to subject i and 0's otherwise. Let \mathbf{Z}_1 be an $N \times n$ matrix defined similarly, except the i^{th} column contains $\mathbf{E}_i = (E_{i1}, \dots, E_{iT_i})^\top$ in those rows corresponding to subject i . Let $\mathbf{C}_b = [\mathbf{B}|\mathbf{Z}]$, with the b subscript denoting the dependence of this matrix on \mathbf{b}_0 , \mathbf{C}_{bi} denote the i^{th} row of \mathbf{C}_b , and let $\boldsymbol{\omega} = (\boldsymbol{\gamma}, \mathbf{u})^\top$. Consider the general model (4), and recall $\boldsymbol{\theta} = (\mathbf{b}_0, \mathbf{b}_1, \boldsymbol{\omega}, \boldsymbol{\beta}, \boldsymbol{\alpha}, \sigma_0^2, \sigma_{1.0}^2, \sigma_u^2, \sigma_\varepsilon^2)^\top$. The joint posterior density of the parameters given the data \mathbf{y} , \mathbf{X} , \mathbf{E} , and \mathbf{W} is

$$\begin{aligned} [\boldsymbol{\theta}|\mathbf{y}, \mathbf{X}, \mathbf{E}, \mathbf{W}] &\propto [\mathbf{y}|\mathbf{X}, \mathbf{E}, \mathbf{W}, \mathbf{b}_0, \mathbf{b}_1, \sigma_\varepsilon^2] [\mathbf{b}_1|\mathbf{b}_0, \boldsymbol{\omega}, \sigma_{1.0}^2] [\mathbf{b}_0|\sigma_0^2] [\mathbf{u}|\sigma_u^2] \times \\ &[\boldsymbol{\beta}] [\boldsymbol{\gamma}] [\boldsymbol{\alpha}] [\sigma_0^2] [\sigma_u^2] [\sigma_{1.0}^2] [\sigma_\varepsilon^2]. \end{aligned}$$

To implement our MCMC approach, we use flat, non-informative priors for the elements of $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\alpha}$, and inverse gamma priors $\sigma_\varepsilon^2 \sim IG(A_\varepsilon, B_\varepsilon)$, $\sigma_0^2 \sim IG(A_0, B_0)$, $\sigma_u^2 \sim IG(A_u, B_u)$, and $\sigma_{1.0}^2 \sim IG(A_{1.0}, B_{1.0})$ for the variance components. Here, IG denotes the inverse-gamma distribution as defined by Casella and Berger (1992):

$$f(x|A, B) = \frac{1}{\Gamma(A)B^A x^{A+1}} \exp\left(-\frac{1}{Bx}\right) I_{(0, \infty)}(x).$$

In the simulations and data analysis, we fixed the hyperparameters at $A_\varepsilon = A_0 = A_{1.0} = A_u = .1$ and $B_\varepsilon = B_0 = B_{1.0} = B_u = 10$. These choices represent relatively vague priors, and simulations presented in the next section show that they yield an algorithm that is able to improve upon the LMM fit for a variety of nonlinear functional forms for the intercept-slope association. In our simulation study, we evaluated the sensitivity of results to these choices, re-running all simulation scenarios using the variance component priors proposed by Berry et al. (2002) for a general nonparametric regression problem with measurement error.

Under model (4), the joint posterior of $\boldsymbol{\theta}$ given the data is proportional to

$$\begin{aligned} & \exp \left\{ -\frac{1}{2\sigma_\varepsilon^2} \|\mathbf{Y} - (\mathbf{Z}_0\mathbf{b}_0 + \mathbf{Z}_1\mathbf{b}_1 + \mathbf{W}\boldsymbol{\alpha})\| - \frac{1}{2\sigma_{1.0}^2} \|\mathbf{b}_1 - \mathbf{C}_b\boldsymbol{\omega}\| \right\} \\ & \times \exp \left\{ -\frac{1}{2\sigma_0^2} \|\mathbf{b}_0 - \mathbf{X}\boldsymbol{\beta}\| - \frac{1}{2\sigma_u^2} \|\mathbf{u}\| - \frac{1}{B_\varepsilon\sigma_\varepsilon^2} - \frac{1}{B_u\sigma_u^2} - \frac{1}{B_0\sigma_0^2} - \frac{1}{B_{1.0}\sigma_{1.0}^2} \right\} \\ & \times \sigma_\varepsilon^{-2(N/2+A_\varepsilon+1)} \sigma_u^{-2(K/2+A_u+1)} \times \sigma_0^{-2(n/2+A_0+1)} \sigma_{1.0}^{-2(n/2+A_{1.0}+1)}. \end{aligned} \quad (5)$$

It is possible to set up a Gibbs sampler to sample from (5) using the software package WinBUGS (Spiegelhalter, Thomas, and Best 2000). This package offers general MCMC sampling strategies that tend to work well in a wide variety of problems. The major advantage of this implementation is that it is extremely easy to program. Experience, however, shows that these standard strategies tend to mix slowly for our FRE formulation and that one can significantly improve upon this performance using a Metropolis-Hastings-within-Gibbs algorithm tailored specifically for this problem. Letting $\boldsymbol{\theta}_{-0} = \boldsymbol{\theta} \setminus \mathbf{b}_0$, $\lambda_u = \sigma_{1.0}^2/\sigma_u^2$, $\lambda_1 = \sigma_\varepsilon^2/\sigma_{1.0}^2$, and $\mathbf{D} = \text{Diag}[0, 0, \mathbf{1}_K]$, the full conditional distributions for the parameters are

1. $\boldsymbol{\beta} | \mathbf{b}_0, \sigma_0^2, \mathbf{X} \sim N \left[(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{b}_0, \sigma_0^2 (\mathbf{X}^\top \mathbf{X})^{-1} \right]$
2. $\boldsymbol{\alpha} | \mathbf{b}_0, \mathbf{b}_1, \mathbf{y}, \mathbf{E}, \mathbf{W}, \sigma_\varepsilon^2 \sim N \left[(\mathbf{W}^\top \mathbf{W})^{-1} \mathbf{W}^\top [\mathbf{y} - \mathbf{Z}_0\mathbf{b}_0 - \mathbf{Z}_1\mathbf{b}_1], \sigma_\varepsilon^2 (\mathbf{W}^\top \mathbf{W})^{-1} \right]$

3. $\boldsymbol{\omega}|\mathbf{b}_0, \mathbf{b}_1, \sigma_{1.0}^2, \sigma_u^2 \sim N \left\{ \left[\mathbf{C}_b^\top \mathbf{C}_b + \lambda_u \mathbf{D} \right]^{-1} \mathbf{C}_b^\top \mathbf{b}_1, \sigma_{1.0}^2 \left[\mathbf{C}_b^\top \mathbf{C}_b + \lambda_u \mathbf{D} \right]^{-1} \right\}$
4. Sample from $\mathbf{b}_0|\boldsymbol{\theta}_{-0}, \mathbf{y}, \mathbf{E}, \mathbf{X}, \mathbf{W}$ via Metropolis-Hastings; see below.
5. $\mathbf{b}_1|\mathbf{b}_0, \boldsymbol{\omega}, \boldsymbol{\alpha}, \sigma_\varepsilon^2, \sigma_{1.0}^2, \mathbf{y}, \mathbf{E}, \mathbf{W} \sim N \left[\left(\mathbf{Z}_1^\top \mathbf{Z}_1 + \lambda_1 \mathbf{I}_n \right)^{-1} \left[\mathbf{Z}_1^\top (\mathbf{y} - \mathbf{Z}_0 \mathbf{b}_0 + \mathbf{W} \boldsymbol{\alpha}) + \lambda_1 \mathbf{C}_b \boldsymbol{\omega} \right], \right. \\ \left. \sigma_\varepsilon^2 \left(\mathbf{Z}_1^\top \mathbf{Z}_1 + \lambda_1 \mathbf{I}_n \right)^{-1} \right]$
6. $\sigma_\varepsilon^2|\mathbf{b}_0, \mathbf{b}_1, \boldsymbol{\alpha}, \mathbf{y}, \mathbf{E}, \mathbf{W} \sim IG \left(A_\varepsilon + N/2, \right. \\ \left. \{1/B_\varepsilon + \|\mathbf{y} - (\mathbf{Z}_0 \mathbf{b}_0 + \mathbf{Z}_1 \mathbf{b}_1 + \mathbf{W} \boldsymbol{\alpha})\|\}^{-1} \right)$
7. $\sigma_0^2|\boldsymbol{\beta}, \mathbf{b}_0, \mathbf{X} \sim IG \left(A_0 + n/2, \{1/B_0 + (1/2) \|\mathbf{b}_0 - \mathbf{X} \boldsymbol{\beta}\|\}^{-1} \right)$
8. $\sigma_{1.0}^2|\mathbf{b}_0, \mathbf{b}_1, \boldsymbol{\omega} \sim IG \left(A_{1.0} + n/2, \{1/B_{1.0} + (1/2) \|\mathbf{b}_1 - \mathbf{C}_b \boldsymbol{\omega}\|\}^{-1} \right)$
9. $\sigma_u^2|\boldsymbol{\omega} \sim IG \left(A_u + K/2, \{1/B_u + (1/2) \|\mathbf{u}\|\}^{-1} \right)$

Because the full conditionals for the b_{0i} 's do not have closed form, we use a random walk Metropolis-Hastings (Chib and Greenberg 1993) step to sample these random intercepts. We generate a candidate value from a normal distribution with mean equal to the current value of b_{0i} and variance equal to a scaling factor τ times the variance estimate of \hat{b}_{0i} from the analogous LMM. Because it is typically more efficient to overestimate the posterior variance of b_{0i} when generating candidate values, in our simulations and data analysis we set $\tau = 4$. We then use the standard Metropolis-Hastings rejection step to decide whether to accept the candidate value in the current iteration of the chain.

The resulting Gibbs sampler samples from each of these full conditional distributions in the order presented. For both the simulations and data analysis, we initialized each chain using the parameter estimates from the analogous LMM fit. We used 1000 iterations as burn-in, and kept 500 iterations for posterior sample processing. We assessed the mixing of the proposed Metropolis-within-Gibbs algorithm for a random

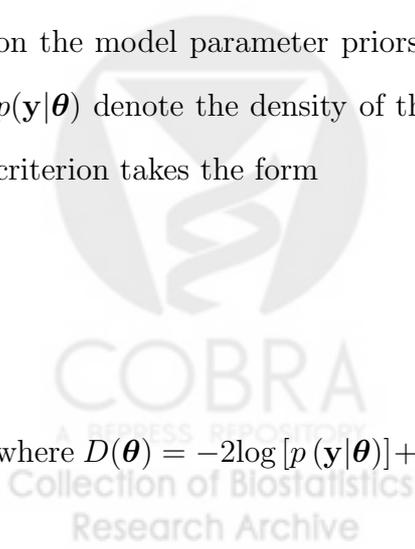
sample of simulated data sets as well as for the blood flow data, and observed proper mixing in all cases. As a second check, we re-ran a sample of the simulated datasets using twice the number of Gibbs samples and observed no change in parameter estimates.

We note that one might also consider integrating out the random slopes to reduce the dimensionality of the model, which could potentially improve the mixing properties of the resulting Markov chain. We investigated this option, but this alternative approach did not significantly improve upon the original Gibbs sampler outlined above, primarily because the mixing properties of this original chain perform quite well without any modification. R functions for implementing this sampler are available from the author.

The FRE model is a natural extension of the Laird-Ware LMM in that it reduces to the LMM in the special case when f is linear. In the penalized spline formulation, this corresponds to $\sigma_u^2 = 0$. Thus, one can evaluate the evidence of a nonlinear association, and hence lack of fit of the LMM, by evaluating the evidence that $\sigma_u^2 \neq 0$. In the Bayesian framework we outline in Section 3.2, we use the *Deviance Information Criterion* (DIC; Spiegelhalter et al. 2002) to compare models. Like other information criterion, the DIC compares likelihoods from two competing models after adjusting for the “effective” number of parameters in the model, with this number depending on the model parameter priors. Let $\boldsymbol{\theta} = (\mathbf{b}_0, \mathbf{b}_1, \mathbf{u}, \boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\alpha}, \sigma_0^2, \sigma_{1.0}^2, \sigma_u^2, \sigma_\varepsilon^2)^\top$, and let $p(\mathbf{y}|\boldsymbol{\theta})$ denote the density of the observed data \mathbf{y} given the model parameters. This criterion takes the form

$$DIC = D(\bar{\boldsymbol{\theta}}) + 2 * p_D,$$

where $D(\boldsymbol{\theta}) = -2\log [p(\mathbf{y}|\boldsymbol{\theta})] + 2\log [p(\mathbf{y})]$ is the usual deviance of the model. That is,



$p(\mathbf{y}) = p(\mathbf{y}|\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbf{y})$, where $\boldsymbol{\mu}(\boldsymbol{\theta})$ denotes the mean of \mathbf{y} given all model parameters. The quantities $\bar{\boldsymbol{\theta}}$ and $\overline{D(\boldsymbol{\theta})}$ denote the posterior means of the model parameters and deviance, respectively, and $p_D = \overline{D(\boldsymbol{\theta})} - D(\bar{\boldsymbol{\theta}})$. All necessary quantities can be calculated from the posterior samples generated by the MCMC algorithm described above.

4 Simulation Study

We performed simulations to compare the ability of the proposed FRE model to accurately estimate the functional relationship between subject-specific intercepts and slopes, as well as the loss of efficiency of our approach when the random effects are truly multivariate normal. Because recently developed nonparametric approaches do not provide direct estimates of this relationship, we compared the performance of the proposed approach to two simpler ad-hoc approaches to this problem that yield smoothed estimates of this relationship, and are hence competitors to the proposed approach. The first competitor, which we denote GAM-LMM, is a two-stage approach that fits a linear mixed model with random intercepts and slopes in a first stage, and fits a penalized regression spline model to the resulting BLUPs with smoothing parameter selected via generalized cross-validation (GCV). The second competitor, which we denote GAM-FIXED, is similar to the first, except the subject-specific effects in the first stage are treated as fixed effects, which avoids smoothing the resulting estimates towards a bivariate normal prior for the random effects.

For each scenario, we generated 100 simulated data sets, with $T_i = T = 5$, $E_{it} \sim N(0, 0.6)$, $b_{0i} \sim N(0, 0.8)$, $\eta_i \sim N(0, 0.1)$, and $\epsilon_{it} \sim N(0, 0.3)$. We chose four functional forms for f that represent realistic functional relationships between overall level and susceptibility in a variety of biomedical contexts (Figure 1): (a)

Symmetric: $f(b) = \phi(b/0.6)$, where ϕ is a normal pdf, (b) Multi-modal $f(b) = 2.5\phi((\text{abs}(x/3) - 0.35)/0.15)$, (c) Sinusoidal: $f(b) = \sin(b\pi/1.8)$, and (d) Linear: $f(b) = 0.5b$. Scenario (a) represents a setting in which individuals with moderate levels have high susceptibility, whereas individuals with extreme levels have low susceptibility. Scenario (b) represents a setting with high susceptibility among individuals with moderately high or low levels, and low susceptibility among individuals with extreme or average levels. Scenario (c) implies different biologic mechanisms for individuals with moderately high and low levels, with exposure among subjects with low levels driving the response lower, exposure among subjects with high levels driving the response higher, and no susceptibility among individuals with either extreme or average levels.

For each method, we computed the mean squared bias (Bias^2) and mean squared error (MSE) of the f estimators over a grid defined by 100 evenly-spaced points in the interval from $[-2\sigma_0, 2\sigma_0]$, or approximately 95% of the range of b_{0i} . These results, shown in Table 1, suggest that for the scenarios in which the true functional form between subject-specific level and susceptibility is nonlinear, the proposed FRE significantly outperforms the two simpler approaches, both in terms of MSE and bias. Depending on the true f , relative to the GAM-LMM approach the FRE offers an improvement in MSE of between 10%-35% for $n = 100$ and 50%-60% for $n = 200$, and even more for squared bias. In general, for scenarios (a)-(c), the parametric distributional assumption in the first stage of the GAM-LMM approach results in the largest bias among the three approaches, but a variance smaller than its fixed counterpart and hence smaller MSE. Figure 2 illustrates a typical result for a single randomly selected dataset generated under scenario (a). The procedure based on the LMM BLUPs shrinks the estimate towards linearity, the GAM based on the fixed effect estimates does somewhat better, and the functional slopes formulation improves

upon this latter estimate by correcting for the measurement error associated with the estimated intercepts. Upon re-running the simulations using the variance component priors proposed by Berry et al. (2002) for a general nonparametric regression problem with measurement error, we observed the same ordering among the three methods for all scenarios.

In addition to assessing the f estimators from the different models, we also compared the ability of each model to predict the susceptibility of the subjects in the study, or \mathbf{b}_1 . For each simulated dataset and for each approach, we calculated the average conditional (on the random effects) mean squared error (CMSE) across subjects (Table 2). The fixed effects approach performed the worst, producing the largest CMSE for all scenarios. This is to be expected, as it does not borrow strength across subjects, but rather estimates a subject's slope using data from that subject only. Perhaps a bit more surprising, the LMM BLUPs outperformed the FRE predictions with respect to CMSE in all scenarios. However, the linearity assumption places a parametric structure on the marginal distribution of the slopes, and this results in more precise predictions of \mathbf{b}_1 . Once one relaxes this linearity assumption, one essentially takes a nonparametric approach to estimating this marginal distribution, and this flexibility results in much larger uncertainty associated with these predictions with no reduction in bias.

Taken together, these simulation results suggest that the FRE model is preferable to existing subject-specific models when scientific interest focuses on the functional relationship between level and susceptibility and this true association is nonlinear, but is not necessarily preferable to a more structured parametric model when interest focuses on prediction of an individual's susceptibility, even when the simpler model does not hold. Although there is a large body of work on mixed models with nonparametric assumptions for the distribution of the random effects, to our knowl-

edge this bias-variance tradeoff for the random effect predictions has not previously been explored.

5 Data Analysis

Here we analyze the cardiac blood flow data first considered by Bartoli et al. (2008) and described in Section 2. The costs of these microsphere experiments are prohibitive, so the methodology was applied to only a very few animals. In this analysis we use data from $n = 109$ tissue pieces from two laboratory animals, considering the longitudinal responses defined by the piece-specific values

$$y_{it} = (\text{Occlusion Blood Flow} - \text{Baseline Blood Flow})_{it}, \quad i = 1, \dots, 109, \quad t = 1, \dots, 6.$$

In preliminary analyses, we inspected scatterplots of the fixed-estimates of the intercepts and slopes from the standard model

$$Y_{it} = b_{0i} + b_{1i}CAPS_{it} + \varepsilon_{it},$$

constructed for each animal separately. Observing little difference in the form of this relationship in the two animals, we fit a model that allowed the mean of the intercept to differ by animal but for the intercept-slope relationship to be constant across animal. That is, we fit the FRE model

$$Y_{it} = b_{0i} + [f(b_{0i}) + \eta_i]CAPS_{it} + \varepsilon_{it}, \quad (6)$$

where $b_{0i} \sim N(\beta_0 + \beta_1 \text{Animal}_i, \sigma_0^2)$. Here, Animal_i is an animal indicator that accounts for the fact that multiple tissue pieces are taken from the same animal, $CAPS_{it}$

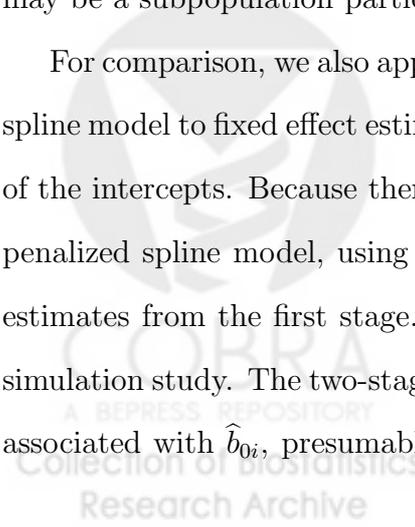
is a binary indicator of CAPs (vs. filtered air) exposure during occlusion, and

$$\eta_i \stackrel{iid}{\sim} N(0, \sigma_{1,0}^2), \quad \varepsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2), \quad b_{0i} \perp \eta_i \perp \varepsilon_{it}.$$

In this formulation, b_{0i} represents the difference between the average occlusion blood flow under filtered air conditions and baseline blood flow, which as described in the Introduction is thought to reflect the degree of ischemia for a given piece. In this framework, the analysis presented in Bartoli et al. (2008) corresponds to estimation of f under the assumption that this function is a step function, with the step defined by dichotomizing the sufficient statistic for b_{0i} , $s_i = (1/|\mathcal{F}_i|) \sum_{t \in \mathcal{F}_i} y_{it}$ at 0, where \mathcal{F}_i denotes the set of outcomes y_{it} corresponding to filtered air exposure for subject i . Model (6) relaxes this assumption and estimates f nonparametrically.

Panel (a) of Figure 3 shows the posterior mean of f , and the corresponding 95% credible interval for this estimate, from the fit of model (6). This estimate suggests that the effects of exposure occur in those pieces that are most ischemic under control conditions, with pieces that are not severely ischemic showing little effect of exposure. This result suggests that pollution exposure may not only make the myocardial region affected by coronary artery disease larger, but may also further restrict flow in the already affected region. This finding suggests that subjects with existing heart disease may be a subpopulation particularly susceptible to PM exposure.

For comparison, we also applied the *ad-hoc* two-stage approach that fits a penalized spline model to fixed effect estimates of the slopes as a function of fixed effect estimates of the intercepts. Because there were missing data for some pieces, we fit a weighted penalized spline model, using as weights the inverse variance estimates of the slope estimates from the first stage. The results illustrate those patterns observed in the simulation study. The two-stage approach does not account for the measurement error associated with \hat{b}_{0i} , presumably leading to biased estimate of f . More importantly,



in terms of the conclusions of the study, the FRE model estimates that the effect of particle exposure on those pieces with blood flow increasing during occlusion is null, whereas the naive two-stage approach estimates an increase in blood flow associated with exposure, which is less biologically plausible.

Finally, we compared the estimates of μ_1 , the mean of the random slopes, from the functional random effects and the linear mixed models. The estimates were not appreciably different, with the standard LMM yielding $\hat{\mu}_1 = 0.03$ (S.E. = 0.02) and the functional random effects model yielding $\hat{\mu}_1 = 0.02$ (posterior S.D. = 0.04). This suggests that on average, for these two animals, there was no exposure effect averaged over all tissue pieces.

We compared the DIC value from the FRE model to that from the simpler random intercepts and slopes LMM that results in the special case when $\sigma_u^2 = 0$. These values were $DIC_{FRE} = 117.3$ and $DIC_{LMM} = 142.3$, respectively. In view of the fact that Spiegelhalter et al. (2002) suggested that a difference in DIC of more than three represents a true difference between two model fits, this observed difference suggests the proposed functional random effects model fits significantly better than the standard Laird-Ware model for these data.

6 Discussion

Although there are many methods for relaxing the assumption of normality for the random effects in a linear mixed model (e.g. Magder and Zeger 1996; Zhang and Davidian 2001; Agresti et al. 2004), there are no good methods available when scientific interest focuses on *direct* estimation of a potentially nonlinear association between baseline and susceptibility. We argue that simple approaches, such as dichotomizing the average response and existing mixed model formulations, can place somewhat

strong assumptions on this relationship that can obscure the true relationship. Recently proposed nonparametric methods do not yield a direct estimate of this relationship. The proposed FRE model can avoid these issues since it does not shrink this relationship towards linearity and properly accounts for all sources of uncertainty in the estimation process.

Although the application we considered involves a binary exposure, the general principles introduced here apply more broadly to settings in which interest focuses on the relationship between baseline and change over time, or between baseline and the effect of some other time-varying continuous covariate. However, in this more general setting, additional subtleties arise. For instance, one should take care in centering the time-varying covariate in this case (Edland 2000). Moreover, incorrectly assuming a linear form for the effect of time can spuriously induce correlation between intercepts and slopes (Rabinowitz and Shea 1997). Thus, in the general case, it is important to diagnose any lack of fit of the underlying linear mixed effects model before extending it to include functional slopes. The diagnostics proposed by Rabinowitz and Shea (1997) are likely to be useful for this purpose. Moreover, because the application considered here was a randomized toxicology study, there were no additional time-varying confounders to consider. In the general case, it would be important to check whether such effects should be modeled as fixed effects constant across subjects or subject-specific random effects (Kinney and Dunson 2007).

There are several potentially interesting directions for future research. A natural direction is the extension of the proposed FRE framework to the generalized case for non-normal responses. Interesting questions include how well can one estimate a true relationship between level and susceptibility when one has binary outcomes, which contain much less information than continuous endpoints. Model fitting is also likely to be more challenging, as the development of MCMC algorithms with good mixing

properties in the generalized additive framework is delicate (Zhao et al. 2006). It would also be interesting to investigate to what degree this formulation can serve as a type of nonparametric treatment for the distributional assumptions for the random effects, and whether this yields more robust fixed effect estimates in the generalized case.

7 Acknowledgements

This research was supported by NIEHS grants ES012044, ES012972, and ES009825, and EPA grant R827353. The author thanks John Godleski for use of the microsphere data, Chris Paciorek and Eric Tchetgen for helpful discussions, and Louise Ryan for thoughtful comments on a previous draft of the manuscript.

8 References

- Agresti, A., Caffo, B., and Ohman-Strickland, P. (2004). Examples in which misspecification of a random effects distribution reduces efficiency, and possible remedies *Computational Statistics & Data Analysis* **47**, 639–653.
- Bartoli, C. R., Akiyama, I., Wellenius, G. A., Coull, B. A., Diaz, E. A., Lawrence, J., Okabe, K., Verrier, R. L., and Godleski, J. J. (2008). Acute inhalation of particulate air pollution causes inappropriate coronary vasoconstriction during ongoing myocardial ischemia. *Environmental Health Perspectives*, in press.
- Berry, S. M., Carroll, R. J., and Ruppert, D. (2002). Bayesian smoothing and regression splines for measurement error problems. *Journal of the American Statistical Association* **97**, 160–169.
- Brumback, B. A., Ruppert, D., and Wand, M. P. (1999). Comment to “Variable selection and function estimation in additive nonparametric regression using a data-based prior”. *Journal of the American Statistical Association* **94**, 794–797.
- Carlin, B. P. and Louis, R. A. (2000). *Bayes and Empirical Bayes Methods for Data Analysis*. 2nd ed. Boca Raton: Chapman and Hall/CRC Press.

- Carroll R. J., Ruppert D., Crainiceanu, C. M., Tosteson, T. D., Karagas, M. R. (2004). Nonlinear and nonparametric regression and instrumental variables. *Journal of the American Statistical Association* **99**, 736–750.
- Casella, G. and Berger, R. L. *Statistical Inference*. Pacific Grove, CA: Wadsworth & Brooks/Cole.
- Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *The American Statistician* **49**, 327–335.
- Crainiceanu, C. M., Ruppert, D. and Wand, M. P. (2005). Bayesian analysis for penalized spline regression using WinBUGS. *Journal of Statistical Software* **14**, 1548–766.
- Dubowsky, S. D., Suh, H., Schwartz, J., Coull, B. A., Gold, D. R. (2006). Diabetes, obesity, and hypertension may enhance associations between air pollution and markers of systemic inflammation. *Environmental Health Perspectives* **114**, 992–998.
- Edland, S. D. (2000). Blomqvist revisited: how and when to test the relationship between level and longitudinal rate of change. *Statistics in Medicine*, **19**, 1441–1452.
- Fan, Y., Leslie, D. S. and Wand, M. P. (2007). Generalised linear mixed model analysis via sequential Monte Carlo sampling. Submitted for publication. Available at <http://www.uow.edu.au/~mwand/papers.html>.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**, 398–409.
- Kinney, S. K. and Dunson, D. B. (2007). Fixed and random effects selection in linear and logistic models. *Biometrics* **63**, 690–698.
- Laird, N. M. and Ware, J. H. (1982). Random effects models for longitudinal data. *Biometrics* **38**, 963–974.
- Li, E., Zhang, D., and Davidian, M. (2004). Conditional estimation for generalized linear models when covariates are subject-specific parameters in a mixed model for longitudinal measurements. *Biometrics* **60**, 1–7.
- Magder, L. S. and Zeger, S. L. (1996). A smooth nonparametric estimate of a mixing distribution using mixtures of Gaussians. *Journal of the American Statistical Association* **91**, 1141–1151.
- Nelsen, R. (1999). *An Introduction to Copulas*. New York: Springer-Verlag.

- Neuhaus, J. M. and McCulloch, C. E. (2006). Separating between- and within-cluster covariate effects by using conditional and partitioning methods. *Journal of the Royal Statistical Society, Series B* **68** 859–872.
- Rabinowitz, D., and Shea, S. (1997). Random effects analysis of children’s blood pressure data. *Statistical Science* **12**, 185–194.
- Robinson, G. K. (1991). That BLUP is a good thing: the estimation of random effects. *Statistical Science* **6**, 15–51.
- Ruppert, D., Wand, M. P. and Carroll, R. J. (2003). *Semiparametric Regression*. Cambridge, UK: Cambridge University Press.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B* **64**, 583–639.
- Spiegelhalter, D., Thomas, A., and Best, N. (2000). *WinBUGS Version 1.3. User’s Manual*. MRC Biostatistics Unit. Institute of Public Health, Cambridge. <http://www.mrc-bsu.cam.ac.uk/bugs>.
- Wall, M. M. and Amemiya, Y. (2000). Estimation for polynomial structural equation models. *Journal of the American Statistical Association* **95**, 929–940.
- Wang, C. Y., Wang, N. and Wang, S. (2000). Regression analysis when covariates are regression parameters of a random effects model for observed longitudinal measurements. *Biometrics* **56**, 487–495.
- Zanobetti, A., Speizer, F. E., Schwartz, J., Stone, P., Coull, B. A., Suh, H., Verrier, R., Nearing, B., Gold, D. R. (2008). T-wave alternans, air pollution and traffic in high-risk subjects. Submitted for publication.
- Zhang, D. and Davidian, M. (2001). Linear mixed models with flexible distributions of random effects for longitudinal data. *Biometrics* **57**, 795–802.
- Zhao, Y., Staudenmayer, J., Coull, B. A., Wand, M.P. (2006). Towards general design Bayesian generalized linear mixed models. *Statistical Science* **21**, 35–51.

Table 1: Simulated MSE and Bias² for f estimators ($\times 10^{-2}$)

Function	Method	$n = 100$		$n = 200$	
		MSE	Bias ²	MSE	Bias ²
M-Modal	FRE	6.25	1.12	2.86	0.86
	GAM-LMM	6.93	5.31	5.67	4.88
	GAM-Fixed	7.11	3.24	4.24	2.41
Norm	FRE	3.91	0.14	1.76	0.06
	GAM-LMM	4.66	3.49	3.52	2.94
	GAM-Fixed	3.97	1.11	2.03	0.68
Sine	FRE	5.52	0.61	2.47	0.25
	GAM-LMM	8.46	5.64	6.33	4.87
	GAM-Fixed	6.71	2.87	3.74	1.97
Linear	FRE	4.64	0.03	2.00	0.07
	GAM-LMM	0.94	0.00	0.50	0.00
	GAM-Fixed	2.31	0.21	1.35	0.32

Lowest value per scenario in bold.

Simulation SE ≈ 0.02



Table 2: Simulated MSE and Bias² for \mathbf{b}_1 predictions

Function	Method	$n = 100$		$n = 200$	
		MSE	Bias ²	MSE	Bias ²
M-Modal	FRE	0.56	0.22	0.52	0.20
	GAM-LMM	0.32	0.22	0.29	0.19
	GAM-Fixed	0.81	0.22	0.82	0.20
Norm	FRE	0.45	0.14	0.43	0.15
	GAM-LMM	0.20	0.14	0.21	0.14
	GAM-Fixed	0.67	0.14	0.71	0.15
Sine	FRE	1.36	0.64	1.29	0.60
	GAM-LMM	1.09	0.63	1.03	0.60
	GAM-Fixed	1.62	0.64	1.61	0.60
Linear	FRE	0.83	0.37	0.75	0.31
	GAM-LMM	0.60	0.37	0.52	0.31
	GAM-Fixed	1.05	0.37	1.02	0.31

Lowest MSE value per scenario in bold.

Simulation SE ≈ 0.01



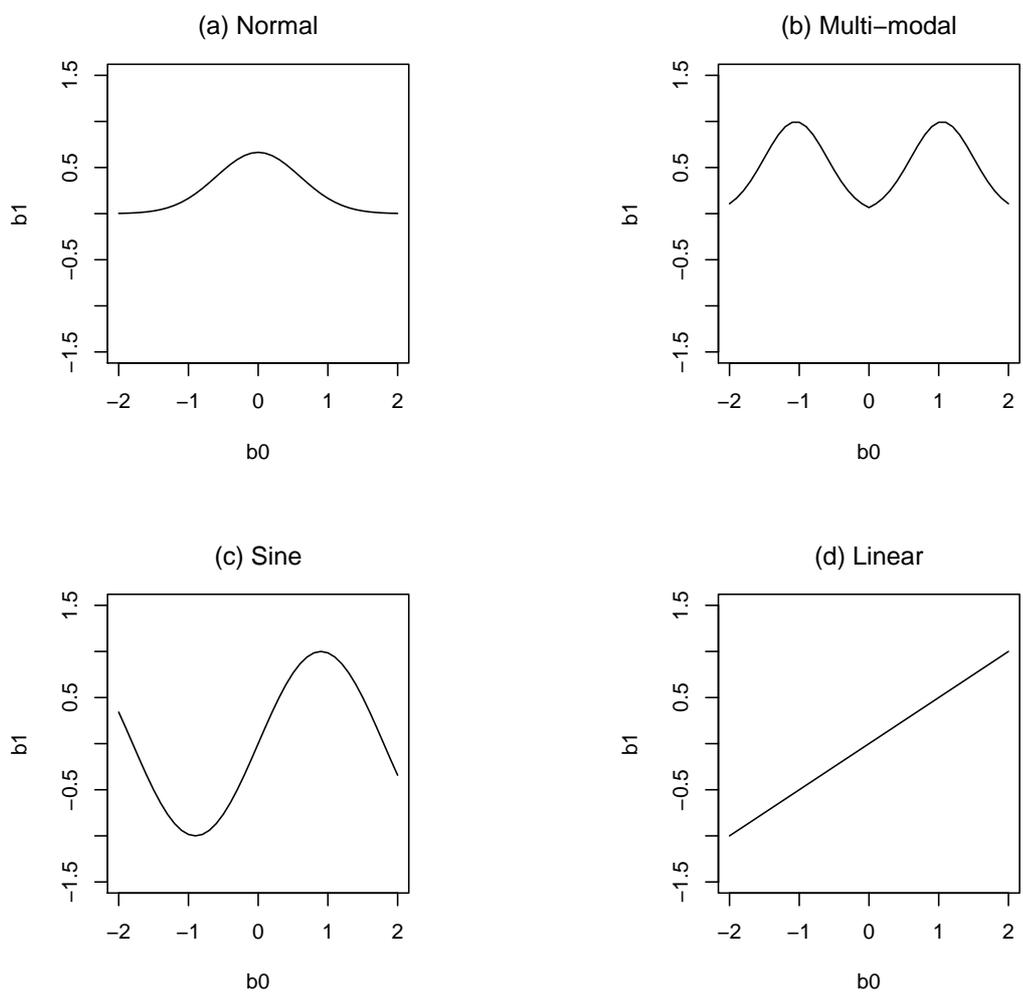


Figure 1: Simulation scenarios for the intercept-slope association.



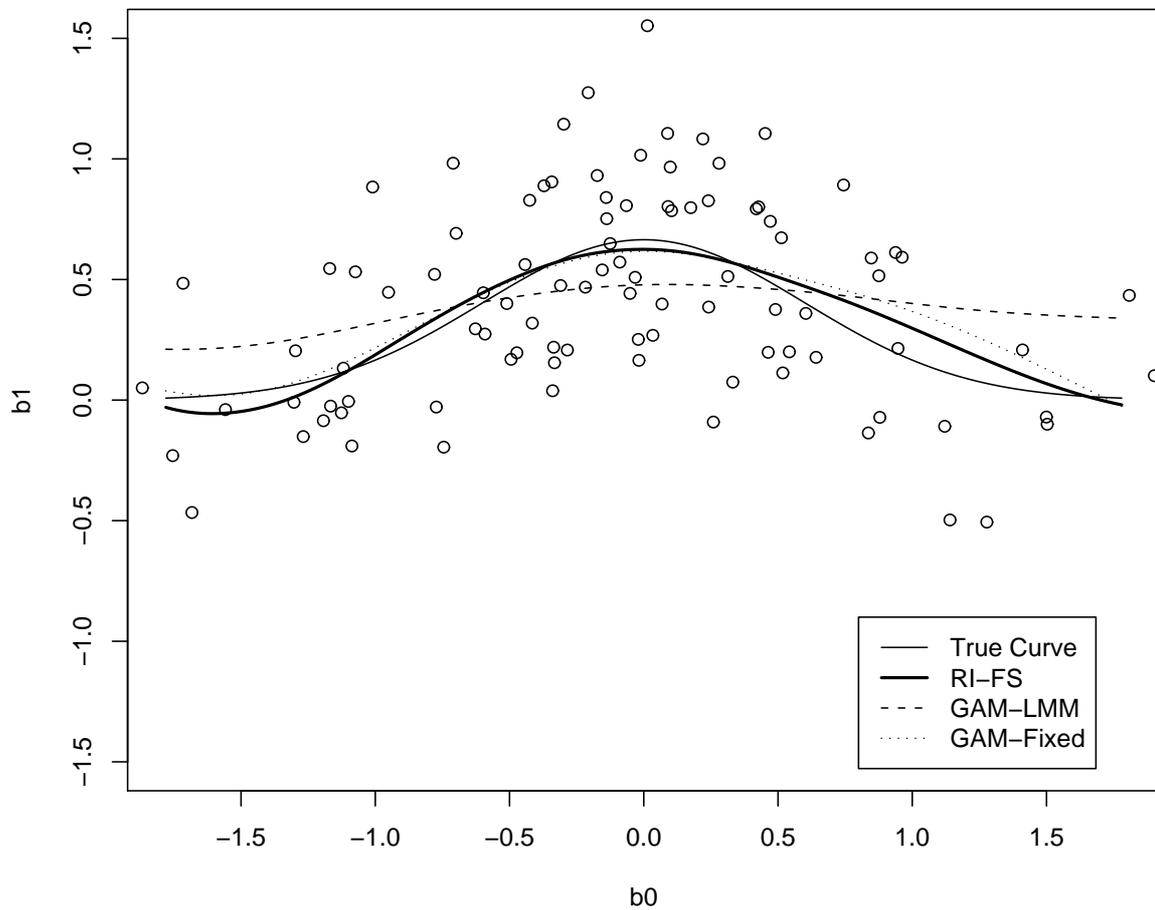


Figure 2: Typical simulation result for scenario (a). Points represent 100 random slopes generated for this particular data set.



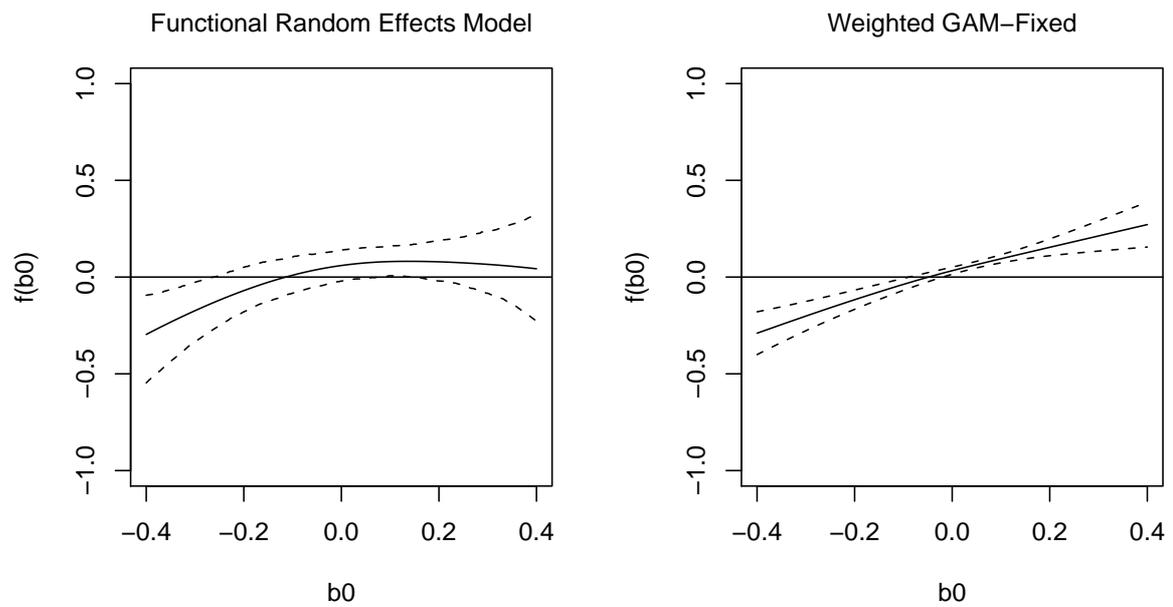


Figure 3: Results of blood flow analysis. Left panel shows FRE estimate. Right panel shows estimate from weighted GAM-Fixed analysis.

