

Mixtures of Varying Coefficient Models for Longitudinal Data with Discrete or Continuous Nonignorable Dropout

Joseph W. Hogan,^{1,*} Xihong Lin,² and Benjamin Herman¹

¹Center for Statistical Sciences, Department of Community Health,
Brown University, Box G-H, Providence, Rhode Island 02912, U.S.A.

²Department of Biostatistics, University of Michigan, Ann Arbor, Michigan 48109, U.S.A.

**email*: jwh@brown.edu

SUMMARY. The analysis of longitudinal repeated measures data is frequently complicated by missing data due to informative dropout. We describe a mixture model for joint distribution for longitudinal repeated measures, where the dropout distribution may be continuous and the dependence between response and dropout is semiparametric. Specifically, we assume that responses follow a varying coefficient random effects model conditional on dropout time, where the regression coefficients depend on dropout time through unspecified nonparametric functions that are estimated using step functions when dropout time is discrete (e.g., for panel data) and using smoothing splines when dropout time is continuous. Inference under the proposed semiparametric model is hence more robust than the parametric conditional linear model. The unconditional distribution of the repeated measures is a mixture over the dropout distribution. We show that estimation in the semiparametric varying coefficient mixture model can proceed by fitting a parametric mixed effects model and can be carried out on standard software platforms such as SAS. The model is used to analyze data from a recent AIDS clinical trial and its performance is evaluated using simulations.

KEY WORDS: Clinical trials; Equivalence trial; Linear mixed model; Missing data; Nonignorable dropout; Pattern-mixture model; Pediatric AIDS; Selection bias; Smoothing splines.

1. Introduction

1.1 Informative Dropout

Dropout and other types of missing data are common in long-term longitudinal studies; in many cases, dropout induces a missingness process that is *nonignorable* in the sense that missingness depends probabilistically on unobserved outcomes, even after conditioning on observable information. Most approaches to handling informative dropout in longitudinal data can be viewed as extensions of standard approaches such as multilevel modeling (Laird and Ware, 1982; Diggle, 1988; Breslow and Clayton, 1993) and marginal modeling (Liang and Zeger, 1986). Likelihood-based approaches include selection models (Wu and Carroll, 1988; Diggle and Kenward, 1994; Follman and Wu, 1995; Ten Have et al., 1998) and mixture models (Wu and Bailey, 1989; Little, 1993, 1994; Hogan and Laird, 1997a). Recent comprehensive surveys of parametric and likelihood-based approaches to handling dropout in longitudinal data can be found in Little (1995), Hogan and Laird (1997b), and Kenward and Molenberghs (1999). Moment-based methods also have been generalized to handle informative dropout under the selection modeling framework (see Robins, Rotnitzky, and Zhao, 1995; Rotnitzky, Robins, and Scharfstein, 1998; Scharfstein, Robins, and Rotnitzky, 1999).

In this article, we develop a general mixture modeling approach for continuous longitudinal repeated measures data where measurement times may be irregular across subjects

and where dropout might be at continuous times and potentially nonignorable. The conditional distribution of repeated measures given dropout follows a varying coefficient model (VCM) (Zhang et al., 1998) where regression coefficients such as intercepts and slopes depend on dropout through unspecified nonparametric functions. The shapes of the functions are estimated using step functions when dropout time is discrete (e.g., for panel data), and using natural cubic smoothing splines (Green and Silverman, 1994) when dropout is continuous. We show that estimation in the proposed varying coefficient mixture model can proceed by fitting an augmented parametric mixed effects model. The complete data distribution is a mixture of the VCMs over the dropout distribution, and the dropout distribution can be left completely unspecified.

This class of models can be viewed as an extension of pattern-mixture models (Little, 1993, 1994) and conditional linear models (CLM) (Wu and Bailey, 1989; Hogan and Laird, 1997a) described above. The former author mainly considered panel data, while the latter authors allowed for the dropout time to be continuous but assumed the regression coefficients to be parametric functions of the dropout time. Estimation can therefore be biased if the parametric functions are misspecified. The proposed varying coefficient mixture model relaxes the parametric assumption by providing a unified framework to allow for flexible dependence of the covariate effects on dropout patterns by assuming regression coefficients to be

nonparametric functions of dropout times. Hence estimation of the covariate effects is more robust to misspecification of the dependence between longitudinal responses and dropout.

1.2 Motivating Example

Protocol 128 of the ACTG was a randomized double-blind equivalency trial of high-dose (180 mg per square meter body surface area, six times daily) versus low-dose (90 mg) zidovudine (ZDV) for HIV-infected children (Brady et al., 1996). The study enrolled 424 children, randomized them to receive one of the two doses, and followed the children on a number of endpoints for up to 5 years. In this article, we are concerned with comparing longitudinal trajectory of CD4 cell counts. Children were scheduled for measurement of CD4 count every 12 weeks, but actual measurement times varied considerably. In addition, only about half of the participants completed 3 years of follow-up (113/216 [52%] on low dose, 93/208 [45%] on high dose).

A simple but reasonable approach to analyzing these data is to estimate treatment-group-specific CD4 trajectories using a linear random effects model (REM) (Laird and Ware, 1982). This model provides valid inference under ignorable dropout (Laird, 1988; Diggle and Kenward, 1994; Little, 1995). Basic exploratory analysis suggests that for the observed data, mean square root of CD4 is well described by a linear time trend in both treatment arms (Figure 1). Using the REM, estimated change from baseline to week 200 is -12.7 (SE 0.8) in the low-dose arm, -18.2 (SE 1.4) in the high-dose arm, for a difference of -5.5 (SE 1.6), favoring low dose.

To explore the potential for bias due to outcome-related dropout, we plotted estimated individual least-squares slopes versus follow-up time (Figure 2). A clear pattern is evident in both treatment arms, namely that lower slopes are associated with early dropout, which casts some doubt on the ignorable dropout assumption and hence on the validity of estimates from the linear REM, suggesting the need to utilize more elaborate models for addressing potential effects of informative dropout.

The remainder of our article is organized as follows. The model is described in Section 2, and estimation procedures are detailed in Section 3; this includes estimation for discrete and continuous dropout times, and for settings with censored dropout times. In Section 4, we apply the proposed model to the clinical trial described in previous sections. Section 5 presents a simulation study to evaluate the bias of the proposed method under departures from underlying assumptions. Summary and discussion follow in Section 6.

2. Mixtures of Varying Coefficient Models for Handling Informative Dropout

Suppose that the data consist of m subjects with the i th subject having n_i observations over time. For the i th subject, let \mathbf{Y}_i be an $n_i \times 1$ observed outcome vector, \mathbf{X}_i be an $n_i \times p$ covariate matrix associated with fixed effects, \mathbf{Z}_i be an $n_i \times q$ covariate matrix associated with random effects, and U_i be the dropout time. The complete data distribution for the response vector is the mixture obtained by integrating the joint distribution $f(\mathbf{y}, u)$ over u . Mixture model

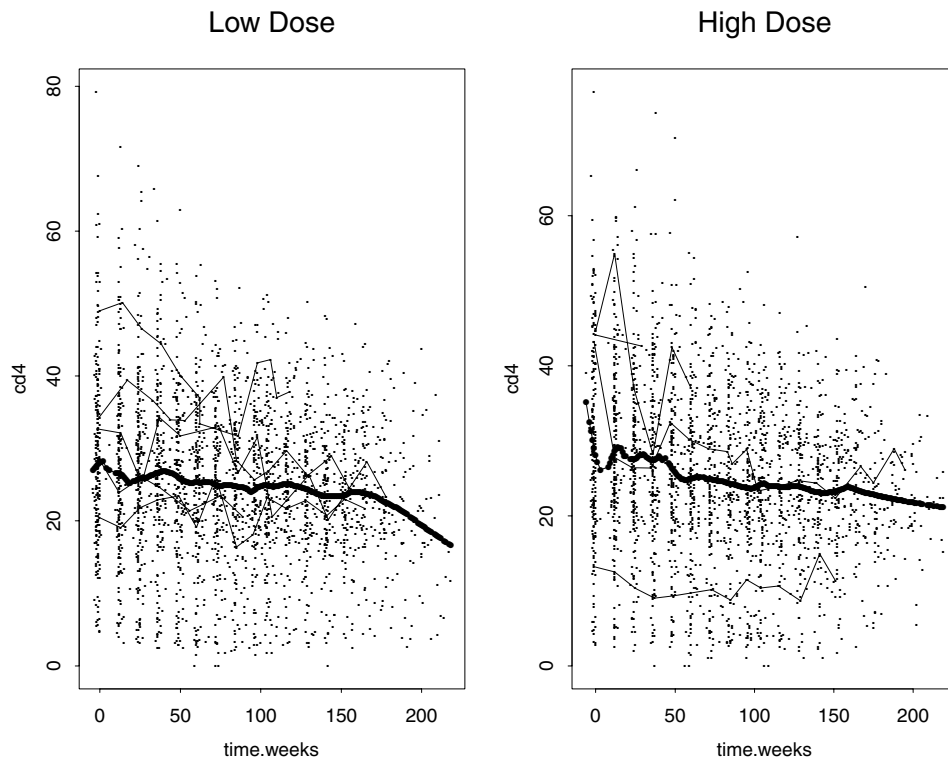


Figure 1. Observed (square root) CD4 counts versus time and stratified by dose, with lowest regression line fit to pooled sample and five individual profiles highlighted.




Figure 2. Individual-specific OLS slopes for square root of CD4 as a function of follow-up time, stratified by dose.

approaches, therefore, require specification of $f(\mathbf{y} | u)$ and $f(u)$. When dropout times are discrete, it is usual to leave $f(u)$ unspecified and estimate it nonparametrically; for continuous u , it is possible but not always desirable to use a parametric model such as log normal (Schluchter, 1992; DeGruttola and Tu, 1994). Our approach is to leave this marginal distribution unspecified.

To capture the dependence between \mathbf{Y} and U , we assume that repeated measurements \mathbf{Y}_i for those who drop out at u_i follow the varying coefficient REM

$$(\mathbf{Y}_i | U = u_i) = \mathbf{X}_i \boldsymbol{\beta}(u_i) + \mathbf{Z}_i \mathbf{b}_i + \boldsymbol{\epsilon}_i, \tag{1}$$

where $\boldsymbol{\beta}(u) = \{\beta_1(u), \dots, \beta_p(u)\}^T$ is a $p \times 1$ vector of unknown regression coefficient functions of the dropout time u , \mathbf{b}_i is a $q \times 1$ vector of random effects following $N\{\mathbf{0}, \mathbf{D}(\boldsymbol{\theta}, u_i)\}$, $\boldsymbol{\epsilon}_i$ is an $n_i \times 1$ vector of residuals following $N\{\mathbf{0}, \mathbf{R}_i(\boldsymbol{\theta}, u_i)\}$, and $\boldsymbol{\theta}$ is a $c \times 1$ vector of variance components. Note that we allow the covariance matrices \mathbf{D} and \mathbf{R} to depend on the dropout time u_i . To better understand model (1), we write $\mathbf{X}_i = (\mathbf{X}_{i1}, \dots, \mathbf{X}_{ip})$, where \mathbf{X}_{ij} is an $n_i \times 1$ vector of the values of the j th covariate measured over time for the i th subject. Equation (1) can be written as

$$(\mathbf{Y}_i | U = u_i) = \sum_{j=1}^p \mathbf{X}_{ij} \beta_j(u_i) + \mathbf{Z}_i \mathbf{b}_i + \boldsymbol{\epsilon}_i, \tag{2}$$

where $\beta_j(u)$ represents the j th covariate effect for those who drop out at time u .

In settings such as a panel design, where subjects are observed at prespecified finite time points (panels) and the underlying dropout times are discrete, the $\beta_j(u)$ are step functions (Little and Wang, 1996; Hogan and Laird, 1997a). If instead subjects are observed at different irregular time points and the underlying dropout times are continuous, we assume that each $\beta_j(u)$ is an unspecified smooth function. It follows that model (1) allows the covariate effects to vary with dropout times nonparametrically and therefore estimation of the covariate effects can be made more robust.

The varying coefficient specification (1) includes the pattern-mixture model (Little, 1993, 1994), random effects pattern-mixture model (Little, 1995; Hogan and Laird, 1997a), and CLM (Wu and Bailey, 1989; Schluchter, 1992) as special cases. For example, if U has a discrete distribution with finite support, then the $\beta_j(u)$ are step functions and (1) is a pattern-mixture model. For continuous dropout times, consider the case where $\mathbf{X}_i = (\mathbf{1}, \mathbf{T}_i)$ and $\mathbf{T}_i = (t_{i1}, \dots, t_{in_i})^T$, where t_{ik} is the k th follow-up time for the i th subject. If $\beta_1(u)$ and $\beta_2(u)$ are polynomial functions of u , model (1) reduces to a CLM (Wu and Bailey, 1989). If the $\beta_j(u)$ are constant in u , then the mixture distribution $f(\mathbf{y})$ has only one component and (1) reduces to a standard REM (Laird and Ware, 1982; Diggle, 1988).

Comparisons to selection models (e.g., Wu and Carroll, 1988; Diggle and Kenward, 1994) can be made by deriving the associated probability of dropout at time u as a function of \mathbf{y} . Let $\mathbf{u}^0 = (u_1^0, \dots, u_{r-1}^0)^\top$ denote the set of ordered, unique dropout times, and define an arbitrary time $u_r^0 > u_{r-1}^0$ for the completers. The selection function can be written in closed form as a logistic regression using baseline-category logits, where completers ($u = u_r^0$) define the baseline category. Let $h_s(\mathbf{y}_i) = \text{pr}(U = u_s^0 | \mathbf{y}_i) / \text{pr}(U = u_r^0 | \mathbf{y}_i)$, which characterizes the odds of dropping out at time u_s^0 relative to completing the study, given repeated measures \mathbf{y}_i . Taking logs,

$$\begin{aligned} \log h_s(\mathbf{y}_i) &= \log(\pi_r / \pi_s) + \log f(\mathbf{y}_i | U = u_s^0) \\ &\quad - \log f(\mathbf{y}_i | U = u_r^0), \end{aligned} \quad (3)$$

where $\pi_s = \text{Pr}(U_i = u_s^0)$, and

$$\begin{aligned} f(\mathbf{y}_i | U_i = u) &= |\mathbf{V}_i(\boldsymbol{\theta}, u)|^{-n_i/2} \\ &\quad \times \exp \left[-\{\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}(u)\}^\top \mathbf{V}_i(\boldsymbol{\theta}, u)^{-1} \right. \\ &\quad \left. \times \{\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}(u)\} / 2 \right], \end{aligned}$$

with $\mathbf{V}_i(\boldsymbol{\theta}, u) = \mathbf{Z}_i \mathbf{D}_i(\boldsymbol{\theta}, u) \mathbf{Z}_i^\top + \mathbf{R}_i(\boldsymbol{\theta}, u)$. In general, the selection model (3) is quadratic in \mathbf{y}_i . A more familiar version obtains when variance is constant across dropout times, i.e., when $\mathbf{V}_i(\boldsymbol{\theta}, u) = \mathbf{V}_i(\boldsymbol{\theta})$. Then (3) simplifies to

$$\begin{aligned} \log h_s(\mathbf{y}_i) &= w \{ \pi_s, \pi_r, \boldsymbol{\beta}(u_s^0), \boldsymbol{\beta}(u_r^0), \mathbf{V}_i(\boldsymbol{\theta}) \} \\ &\quad + \{ \boldsymbol{\beta}(u_s^0) - \boldsymbol{\beta}(u_r^0) \}^\top \mathbf{X}_i^\top \mathbf{V}_i(\boldsymbol{\theta})^{-1} \mathbf{y}_i, \end{aligned} \quad (4)$$

where $w(\cdot)$ is a function of terms that do not depend on \mathbf{y}_i . Under this formulation, the log relative probability of dropout at u_s^0 is linear in \mathbf{y}_i , with coefficients that depend on the difference $\boldsymbol{\beta}(u_s^0) - \boldsymbol{\beta}(u_r^0)$ increasing in magnitude as $\boldsymbol{\beta}(u)$ depends more strongly on u . For binary U , (4) is consistent with the selection model restrictions used by Little (1994) and Little and Wang (1996) to identify parameters in a two-component pattern-mixture model.

3. Estimation Procedures

In this section, we discuss estimation procedures when dropout times U_i are observed for all subjects. Under this circumstance, the likelihood of (\mathbf{Y}_i, U_i) for the i th subject can be partitioned as

$$L_i(\mathbf{Y}_i, U_i; \boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\pi}) = L_i(\mathbf{Y}_i | U_i; \boldsymbol{\beta}, \boldsymbol{\theta}) L_i(U_i; \boldsymbol{\pi}).$$

It follows that $(\boldsymbol{\beta}, \boldsymbol{\theta})$ can be estimated by maximizing the conditional likelihood $\prod_{i=1}^m L_i(\mathbf{Y}_i | U_i; \boldsymbol{\beta}, \boldsymbol{\theta})$, and $\boldsymbol{\pi}$ can be estimated by maximizing the marginal likelihood $\prod_{i=1}^m L_i(U_i; \boldsymbol{\pi})$. We consider estimation procedures first for the situation where subjects are observed at a common set of time points, and then where the set of observation times may be misaligned across subjects.

3.1 Estimation Procedures for Fixed and Common Observation Times

For a fixed design, such as a panel design, subjects are observed at n (often small) prespecified time points (t_1, \dots, t_n) and the number of possible dropout times is small. Thus, the nonparametric functions $\beta_j(u)$ are assumed to be step functions.

Let r be the number of observed distinct values of u , with $r \leq n$. As indicated in Section 2, let $\mathbf{u}^0 = (u_1^0, \dots, u_r^0)^\top$ be an $r \times 1$ vector of ordered distinct dropout times, where we assume completers take $u_r^0 = t_{n+1}$ for some value t_{n+1} . Hence $\beta_j(u)$ is fully determined by $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jr})^\top$, where β_{jk} represents the j th covariate effect for those dropping out at u_k^0 . Denote the incidence matrix by $\mathbf{N}_i = (N_{i1}, \dots, N_{ir})^\top$ where $N_{ik} = 1$ if $u_i = u_k^0$ and 0 otherwise. Some calculations show that model (2) can then be written as a linear mixed model

$$\mathbf{Y}_i = \sum_{j=1}^p \mathbf{X}_{ij} \mathbf{N}_i^\top \boldsymbol{\beta}_j + \mathbf{Z}_i \mathbf{b}_i + \boldsymbol{\epsilon}_i.$$

Equivalently, we have

$$\mathbf{Y} = \sum_{j=1}^p \tilde{\mathbf{X}}_j \boldsymbol{\beta}_j + \mathbf{Z} \mathbf{b} + \boldsymbol{\epsilon}, \quad (5)$$

where $\tilde{\mathbf{X}}_j = \mathbf{X}_j \otimes \mathbf{N}$, $\mathbf{X}_j = (\mathbf{X}_{1j}^\top, \dots, \mathbf{X}_{mj}^\top)^\top$, $\mathbf{N} = (\mathbf{N}_1^\top, \dots, \mathbf{N}_m^\top)^\top$, and \mathbf{Y} , \mathbf{Z} , \mathbf{b} , and $\boldsymbol{\epsilon}$ are defined similarly to \mathbf{N} . Here $\mathbf{A} \otimes \mathbf{B}$ denotes a direct product of matrices \mathbf{A} and \mathbf{B} .

It follows that estimation of $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_p^\top)^\top$ proceeds by solving the linear mixed model normal equation

$$(\tilde{\mathbf{X}}^\top \mathbf{V}^{-1} \tilde{\mathbf{X}}) \boldsymbol{\beta} = \tilde{\mathbf{X}}^\top \mathbf{V}^{-1} \mathbf{Y},$$

where $\tilde{\mathbf{X}} = (\tilde{\mathbf{X}}_1, \dots, \tilde{\mathbf{X}}_p)$, $\mathbf{V} = \text{diag}(\mathbf{V}_i)$, and $\mathbf{V}_i = \text{cov}(\mathbf{Y}_i) = \mathbf{Z}_i \mathbf{D} \mathbf{Z}_i^\top + \mathbf{R}_i$. The covariance matrix of the estimator of $\boldsymbol{\beta}$ is $(\tilde{\mathbf{X}}^\top \mathbf{V}^{-1} \tilde{\mathbf{X}})^{-1}$. Estimation of $\boldsymbol{\theta}$ can be obtained using the restricted maximum likelihood (REML) estimating equation

$$-\frac{1}{2} \text{tr} \left(\mathbf{P} \frac{\partial \mathbf{V}}{\partial \theta_j} \right) + \frac{1}{2} (\mathbf{Y} - \tilde{\mathbf{X}} \boldsymbol{\beta})^\top \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \theta_j} \mathbf{V}^{-1} (\mathbf{Y} - \tilde{\mathbf{X}} \boldsymbol{\beta}) = 0,$$

where $\mathbf{P} = \mathbf{V}^{-1} - \mathbf{V}^{-1} \tilde{\mathbf{X}} (\tilde{\mathbf{X}}^\top \mathbf{V}^{-1} \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}^\top \mathbf{V}^{-1}$. The (j, k) th component of the information matrix of the estimator of $\boldsymbol{\theta}$ is $(1/2) \text{tr} \{ \mathbf{P} (\partial \mathbf{V} / \partial \theta_j) \mathbf{P} (\partial \mathbf{V} / \partial \theta_k) \}$. If \mathbf{D} and \mathbf{R} depend on u , the identifiability of some components of $\boldsymbol{\theta}$ based on the observed data might require some constraints on $\boldsymbol{\theta}$ or a sensitivity analysis (Little, 1993; Daniels and Hogan, 2000). For example, consider the case with two time points ($n = 2$); if dropout is informative, there is no information in the observed data to fit a random intercept and slope model for those who drop out at time 2, and sensitivity analyses are needed.

The necessary and sufficient condition for the $\boldsymbol{\beta}_j$ to be identifiable from model (5) can be stated as follows:

Let $\mathcal{S}_k = \{i_{k1}, \dots, i_{km_k}\}$ denote the indexes of those m_k subjects who drop out at u_k^0 , and $\mathcal{X}_k = (\mathbf{X}_{i_{k1}}^\top, \dots, \mathbf{X}_{i_{km_k}}^\top)^\top$. If $\text{rank}(\mathcal{X}_k) = p$ for all $k = 1, \dots, r$, then the $\boldsymbol{\beta}_j (j = 1, \dots, p)$ are identifiable.

This condition states that when dropout times are discrete and finite, regression coefficients must be separately estimable for each dropout pattern. If some components of $\boldsymbol{\beta}_j$ are not identifiable for some dropout patterns, then the observed data do not have information about these components and either parameter constraints or sensitivity analyses are needed. For example, consider again the case with two time points ($n = 2$); if dropout is informative, there is no information in the observed data to estimate the mean of the outcome at time 2 for those who drop out at time 2, and sensitivity analysis

is hence needed. See Little and Wang (1996) and Daniels and Hogan (2000) for detailed discussion.

Note that the j th regression coefficient vector β_j measures the j th covariate effect conditional on dropout times. Primary interest is in the marginal j th covariate effect $\tilde{\beta}_j = E_U\{\beta_j(u)\} = \boldsymbol{\pi}^T \beta_j$, where $\boldsymbol{\pi} = (\pi_1, \dots, \pi_r)^T$ and $\pi_j = P(U = u_k^0)$ is the probability that a subject drops out at u_k^0 . It follows that one can estimate $\tilde{\beta}_j$ by $\hat{\boldsymbol{\pi}}^T \hat{\beta}_j$, where $\hat{\boldsymbol{\pi}}_j = m_j/m$, m_j is the number of subjects who drop out at u_k^0 , and $\hat{\beta}$ is the maximum likelihood estimator from fitting model (5).

3.2 Estimation for Random or Misaligned Measurement Times

For situations when the measurement times are random or misaligned across subjects, repeated measures of each subject are observed at different time points and the underlying dropout times are often continuous. We therefore assume that the regression coefficient functions $\beta_j(u)$ are twice-differentiable smooth functions and estimate them using cubic smoothing splines. A key feature of cubic smoothing spline estimation is that we can estimate all model components—including the nonparametric regression coefficient functions $\beta_j(u)$, variance components $\boldsymbol{\theta}$, and smoothing parameters—within an augmented parametric linear mixed model framework (Zhang et al., 1998).

Following the notation in Section 3.1, let r be the number of ordered distinct values of the U_i , $\mathbf{u}^0 = (u_1^0, \dots, u_r^0)^T$ be the ordered distinct values of the U_i , and $\boldsymbol{\beta}_j = \{\beta_j(u_1^0), \dots, \beta_j(u_r^0)\}^T$ be the values of $\beta_j(u)$ evaluated at \mathbf{u}^0 . Given the variance components $\boldsymbol{\theta}$, the conditional log likelihood of $\boldsymbol{\beta}_j$ given the u_i is

$$\begin{aligned} \ell(\mathbf{Y} | \mathbf{u}; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p) &= \sum_{i=1}^m -\frac{1}{2} \ln |\mathbf{V}_i| \\ &\quad - \frac{1}{2} \left\{ \mathbf{Y}_i - \sum_{j=1}^p \mathbf{X}_{ij} \beta_j(u_i) \right\}^T \\ &\quad \times \mathbf{V}_i^{-1} \left\{ \mathbf{Y}_i - \sum_{j=1}^p \mathbf{X}_{ij} \beta_j(u_i) \right\} \\ &= -\frac{1}{2} \ln |\mathbf{V}| - \frac{1}{2} \left(\mathbf{Y} - \sum_{j=1}^p \tilde{\mathbf{X}}_j \boldsymbol{\beta}_j \right)^T \\ &\quad \times \mathbf{V}^{-1} \left(\mathbf{Y} - \sum_{j=1}^p \tilde{\mathbf{X}}_j \boldsymbol{\beta}_j \right), \end{aligned}$$

where $\tilde{\mathbf{X}}_j$, \mathbf{V}_i , \mathbf{V} , and \mathbf{Y} were defined in Section 3.1.

Following O’Sullivan, Yandell, and Raynor (1986), one can show that the natural cubic smoothing spline estimators of the $\beta_j(u)$ maximize the following penalized conditional log likelihood

$$\begin{aligned} \ell(\mathbf{Y} | \mathbf{u}; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p) &- \frac{1}{2} \sum_{j=1}^p \lambda_j \int_{A_1}^{A_2} [\beta_j''(u)]^2 du \\ &= \ell(\mathbf{Y} | \mathbf{u}; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p) - \frac{1}{2} \sum_{j=1}^p \lambda_j \boldsymbol{\beta}_j^T \mathbf{K} \boldsymbol{\beta}_j, \end{aligned} \tag{6}$$

where the λ_j are smoothing parameters controlling the balance between goodness-of-fit and the smoothness of the estimated $\beta_j(u)$, A_1 and A_2 specify the range of u , and \mathbf{K} is the nonnegative definite natural cubic smoothing spline smoothing matrix constructed using \mathbf{u}_0 and defined in Green and Silverman (1994, equation [2.3]).

For fixed smoothing parameters $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)^T$ and the variance components $\boldsymbol{\theta}$, differentiation of (6) with respect to $(\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p)$ gives their estimating equations as

$$\begin{aligned} \begin{bmatrix} \tilde{\mathbf{X}}_1^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_1 + \lambda_1 \mathbf{K} & \tilde{\mathbf{X}}_1^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_2 & \cdots & \tilde{\mathbf{X}}_1^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_p \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{\mathbf{X}}_p^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_1 & \tilde{\mathbf{X}}_p^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_2 & \cdots & \tilde{\mathbf{X}}_p^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_p + \lambda_p \mathbf{K} \end{bmatrix} \\ \times \begin{bmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_p \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{X}}_1^T \mathbf{V}^{-1} \mathbf{Y} \\ \vdots \\ \tilde{\mathbf{X}}_p^T \mathbf{V}^{-1} \mathbf{Y} \end{bmatrix}. \end{aligned} \tag{7}$$

One can solve equation (7) using a backfitting algorithm as follows:

$$\hat{\boldsymbol{\beta}}_j = (\tilde{\mathbf{X}}_j^T \mathbf{V}^{-1} \tilde{\mathbf{X}}_j + \lambda_j \mathbf{K})^{-1} \tilde{\mathbf{X}}_j^T \mathbf{V}^{-1} \left(\mathbf{Y} - \sum_{k \neq j} \tilde{\mathbf{X}}_k \boldsymbol{\beta}_k \right)$$

for $j = 1, \dots, p$.

Following Zhang et al. (1998), we show that these cubic smoothing spline estimators of the β_j can be obtained by fitting an augmented linear mixed model. Specifically, $\boldsymbol{\beta}_j$ can be written via a one-to-one transformation as

$$\boldsymbol{\beta}_j = \mathbf{U} \boldsymbol{\gamma}_j + \mathbf{B} \mathbf{a}_j,$$

where $\mathbf{U} = (\mathbf{1}, \mathbf{u}^0)$, $\mathbf{B} = \mathbf{L}(\mathbf{L}^T \mathbf{L})^{-1}$, \mathbf{L} is an $r \times (r - 2)$ full rank matrix satisfying $\mathbf{K} = \mathbf{L} \mathbf{L}^T$ and $\mathbf{L}^T \mathbf{U} = \mathbf{0}$, $\boldsymbol{\gamma}_j$ is a 2×1 unknown vector, and \mathbf{a}_j is an $(r - 2) \times 1$ unknown vector. It can be shown that $\boldsymbol{\beta}_j^T \mathbf{K} \boldsymbol{\beta}_j = \mathbf{a}_j^T \mathbf{a}_j$. The penalized log likelihood (6) becomes

$$\ell(\mathbf{Y} | \mathbf{U}; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_p) - \frac{1}{2} \sum_{j=1}^p \lambda_j \mathbf{a}_j^T \mathbf{a}_j.$$

It follows that the nonparametric natural cubic spline estimators of the β_j can be obtained by fitting the parametric linear mixed model

$$\mathbf{Y} = \sum_{j=1}^p (\tilde{\mathbf{X}}_j \mathbf{U}) \boldsymbol{\gamma}_j + \sum_{j=1}^p (\tilde{\mathbf{X}}_j \mathbf{B}) \mathbf{a}_j + \mathbf{Z} \mathbf{b} + \boldsymbol{\epsilon}, \tag{8}$$

where $\boldsymbol{\gamma} = (\boldsymbol{\gamma}_1^T, \dots, \boldsymbol{\gamma}_p^T)^T$ is a vector of regression coefficients, $\mathbf{a} = (\mathbf{a}_1^T, \dots, \mathbf{a}_p^T)^T$ and \mathbf{b} are independent random effects with $\mathbf{a} \sim N(\mathbf{0}, \boldsymbol{\Lambda}(\boldsymbol{\tau}))$, $\mathbf{b} \sim N(\mathbf{0}, \text{diag}\{\mathbf{D}(\boldsymbol{\theta})\})$, $\boldsymbol{\Lambda}(\boldsymbol{\tau}) = \text{diag}(\tau_j \mathbf{I})$, $\tau_j = 1/\lambda_j$ and $\boldsymbol{\tau} = (\tau_1, \dots, \tau_p)^T$, and $\boldsymbol{\epsilon}_i \sim N(\mathbf{0}, \mathbf{R}_i(\boldsymbol{\theta}, u_i))$.

Estimation of $\boldsymbol{\gamma}$ and \mathbf{a} can proceed using the BLUP estimator by solving the normal equation

$$\begin{bmatrix} \mathbf{H}^T \mathbf{V}^{-1} \mathbf{H} & \mathbf{H}^T \mathbf{V}^{-1} \mathbf{G} \\ \mathbf{G}^T \mathbf{V}^{-1} \mathbf{H} & \mathbf{G}^T \mathbf{V}^{-1} \mathbf{G} + \boldsymbol{\Lambda}^{-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{\gamma} \\ \mathbf{a} \end{bmatrix} = \begin{bmatrix} \mathbf{H}^T \mathbf{V}^{-1} \mathbf{Y} \\ \mathbf{G}^T \mathbf{V}^{-1} \mathbf{Y} \end{bmatrix}, \tag{9}$$

where $\mathbf{H} = (\tilde{\mathbf{X}}_1 \mathbf{U}, \dots, \tilde{\mathbf{X}}_p \mathbf{U})$ and $\mathbf{G} = (\tilde{\mathbf{X}}_1 \mathbf{B}, \dots, \tilde{\mathbf{X}}_p \mathbf{B})$. Denoting $\hat{\boldsymbol{\gamma}}$ and $\hat{\mathbf{a}}$ as the solution of (9), the natural cubic

smoothing spline estimator of β_j is $\hat{\beta}_j = \mathbf{U}\hat{\gamma}_j + \mathbf{B}\hat{\mathbf{a}}_j$. One can show that the $\hat{\beta}_j$ from (9) are identical to those obtained from solving (7). The natural cubic spline estimators $\hat{\beta}_j$ are unique when \mathbf{H} is of full rank.

We have so far assumed that the smoothing parameters λ and the variance components θ are known when estimating the β_j . They are usually unknown in practice and need to be estimated from the data. Examination of the modified linear mixed model (8) suggests that τ behaves like variance components; therefore, following Zhang et al. (1998), we estimate the smoothing parameters τ and the variance components θ simultaneously using REML by treating τ as extra variance components in addition to θ in (8). The REML estimating equations for θ and τ are

$$\begin{aligned} & -\frac{1}{2}\text{tr}\left(\tilde{\mathbf{P}}\frac{\partial\mathbf{V}}{\partial\theta_k}\right) + \frac{1}{2}\left(\mathbf{Y} - \sum_{j=1}^p \mathbf{X}_j\beta_j\right)^{\text{T}} \\ & \quad \times \mathbf{V}^{-1}\frac{\partial\mathbf{V}}{\partial\theta_k}\mathbf{V}^{-1}\left(\mathbf{Y} - \sum_{j=1}^p \mathbf{X}_j\beta_j\right) = 0, \\ & -\frac{1}{2}\text{tr}\left(\tilde{\mathbf{P}}\tilde{\mathbf{X}}_j\mathbf{B}\mathbf{B}^{\text{T}}\tilde{\mathbf{X}}_j^{\text{T}}\right) + \frac{1}{2}\left(\mathbf{Y} - \sum_{j=1}^p \mathbf{X}_j\beta_j\right)^{\text{T}} \\ & \quad \times \mathbf{V}^{-1}\tilde{\mathbf{X}}_j\mathbf{B}\mathbf{B}^{\text{T}}\tilde{\mathbf{X}}_j^{\text{T}}\mathbf{V}^{-1}\left(\mathbf{Y} - \sum_{j=1}^p \mathbf{X}_j\beta_j\right) = 0, \end{aligned}$$

where $\tilde{\mathbf{P}} = \mathbf{V}^{-1} - \mathbf{V}^{-1}(\mathbf{H}, \mathbf{G})\mathbf{C}^{-1}(\mathbf{H}, \mathbf{G})^{\text{T}}\mathbf{V}^{-1}$, and \mathbf{C} is the coefficient matrix on the left-hand side of (9). Parameters from the varying coefficient mixture model (1) can therefore be obtained by fitting the parametric linear mixed model (8) by SAS Proc Mixed (Version 8.2). The smoothing matrix \mathbf{B} needs to be computed in advance.

3.3 Inference for Marginal Regression Coefficients

The marginal j th covariate effect $\tilde{\beta}_j$ is $\tilde{\beta}_j = E_U\{\beta_j(u)\} = \int \beta_j(u) dF(u)$, where $F(u)$ is the c.d.f. of u . Using the estimated cubic smoothing spline $\hat{\beta}_j(u)$ and the empirical c.d.f. $\hat{F}(u)$, one can estimate $\tilde{\beta}_j$ by $\int \hat{\beta}_j(u) d\hat{F}(u) = \hat{\pi}^{\text{T}}\hat{\beta}_j$, where $\hat{\pi} = \sum_{i=1}^m \mathbf{N}_i/m$. The delta method has been used for standard error estimation for mixture models where the support of U is very small relative to the number of subjects (Hogan and Laird, 1997a; Fitzmaurice and Laird, 2000). In our application with continuous dropout times, we found that the delta method performed poorly; the bootstrap was used instead, treating the subject as the basic resampling unit. Details are provided in the application.

4. Application to AIDS Clinical Trial

In this section, we analyze the ACTG data using several different mixture model formulations representing different assumptions about the missing data mechanism and provide detailed interpretation of the model parameters.

4.1 Variable Transformations and Candidate Models

As indicated in Section 1.2, observed CD4 data were transformed to the square root scale to reduce positive skewness. We fit three models (for computational tractability, each model is fit separately by dose). The first is a standard REM

with subject-specific intercepts and linear time trends. The REM, briefly summarized in Section 1.2, assumes that the complete data in each treatment arm follow the linear mixed model

$$Y_{il} = \beta_1 + \beta_2 t_{il}^* + b_{1i} + b_{2i} t_{il}^* + \epsilon_{il}, \quad (10)$$

where Y_{il} is square root of CD4 count at time t_{il} , $l = 1, \dots, n_i$, $\mathbf{b}_i = (b_{1i}, b_{2i}) \sim N(\mathbf{0}, \mathbf{D})$, and $\epsilon_{il} \sim N(0, \sigma^2)$ (i.e., $\mathbf{R}_i(\theta) = \sigma^2 \mathbf{I}_{n_i}$), with independence between ϵ_{il} and \mathbf{b}_i . The time axis is rescaled using $t_{il}^* = (t_{il} - \bar{U})/\text{range}(t_{il})$ so that the new time scale has range 1 and is centered at the sample mean of dropout times (the range is computed for the pooled sample from both dose arms). This is a standard REM, but can be viewed as a special case of (1) where $\beta_j(u)$ is constant in u for $j = 1, 2$.

In addition to the REM, we fit a CLM in which individual intercepts and slopes are linear functions of dropout time, and a VCM, where intercepts and slopes are unspecified smooth functions of dropout time. The CLM is precisely model (8) under the assumption that $\tau = \mathbf{0}$ (equivalently, $\mathbf{a}_j = \mathbf{0}$ for all j), and therefore it is just a specialized version of the VCM. Separately by treatment, the CLM elaborates (10) such that $\beta_1(u) = \gamma_1 + \gamma_2 u^*$ and $\beta_2(u) = \gamma_3 + \gamma_4 u^*$, where $u^* = (u - \bar{U})/\text{range}(t_{il})$. In our parameterization with the rescaled time axis, γ_1 is the mean of $(\text{CD4})^{1/2}$ at $u = \bar{U}$, γ_2 is the ‘‘main effect’’ of dropout, which is the slope of $\beta_1(u)$, γ_3 is mean change in $(\text{CD4})^{1/2}$ from baseline at $u = \bar{U}$, and γ_4 is the effect of interaction between dropout and change from baseline, representing the slope of $\beta_2(u)$ on u . These parameters are estimated in straightforward fashion by fitting with SAS PROC MIXED a standard linear mixed model $Y_{ij} = \mathbf{H}_{ij}^{\text{T}}\gamma + \mathbf{Z}_{ij}^{\text{T}}\mathbf{b}_i + \epsilon_{ij}$, where $\mathbf{H}_{ij} = (1, U_{ij}^*, t_{ij}^*, t_{ij}^* U_{ij}^*)^{\text{T}}$, and $\mathbf{Z}_{ij} = (1, t_{ij}^*)^{\text{T}}$.

Rescaling the time axis has some practical advantages in terms of model fitting and interpretation: (i) it makes the estimates more stable by increasing the variance of individual slopes away from zero; (ii) for the CLM, the intercept and slope main effect are already averaged over dropout time (Fitzmaurice, Laird, and Shneyer, 2001); (iii) for the REM and CLM, the slope parameter corresponds to average total change in square root CD4 from baseline to the longest follow-up time for the study; and (iv) for the CLM, the parameters γ_2 and γ_4 contrast mean intercept and slope, respectively, between those who drop out immediately (directly after $u = 0$) and those who complete the study protocol ($u = \max_i t_{il}$).

The VCM uses the same design matrix for γ as the CLM. Estimation of the \mathbf{a}_j is made by specifying the $r - 2$ columns of $\tilde{\mathbf{X}}_1\mathbf{B}$ and $\tilde{\mathbf{X}}_2\mathbf{B}$ as independent random effects with respective variances τ_1^2 and τ_2^2 .

4.2 Summary of Fitted Models

4.2.1 Parameter estimates. To get a crude understanding of whether MAR may be a valid assumption, we begin by summarizing regression coefficients for parameterizations of $f(\mathbf{y}|u)$ given by the REM and CLM; these appear in Table 1. For both models, the parameter γ_3 quantifies average change from baseline to the maximum follow-up time. Recall that in the CLM, $\beta_1(u) = \gamma_1 + \gamma_2(u^* - \bar{u}^*)$ and $\beta_2(u) = \gamma_3 + \gamma_4(u^* - \bar{u}^*)$, so that if we assume MAR is violated according to linear dependence between intercept and

Table 1

Parameter estimates from conditional part of joint model, assuming linear random effects structure (REM), and conditional linear model (CLM)

Model	Parameter ^a	AZT dose	
		Low (90 mg)	High (180 mg)
REM	γ_1	28.6 (0.8)	30.1 (0.9)
	γ_3	-12.7 (0.8)	-18.2 (1.4)
CLM	γ_1	28.9 (0.8)	30.3 (0.8)
	γ_2	14.6 (3.3)	13.0 (3.5)
	γ_3	-15.9 (0.9)	-21.4 (1.4)
	γ_4	26.1 (4.8)	33.7 (6.5)

^aSee Section 4 for definitions.

dropout and/or slope and dropout, the parameters γ_2 and γ_4 will quantify the degree to which MAR fails to hold. The parameter estimates from Model 1 suggest that intercepts $\beta_1(u)$ vary considerably by dropout time ($\hat{\gamma}_2 = 14.6$ and 13.0 [SE 3.3 and 3.5] for low and high dose, respectively); the same pattern is indicated for slopes $\beta_2(u)$, where $\hat{\gamma}_4 = 26.1$ and 33.7 (SE 4.8 and 6.5) for low and high dose. The parameters γ_2 and γ_4 represent the average difference in CD4 intercept (γ_2) and slope (γ_4) between those who dropped out immediately after enrolling and those who were followed for the maximum time (220 weeks). Hence the CLM indicates that dropouts have lower CD4 intercepts and slopes than completers, leading to selection bias for end-of-study comparisons due to missing data on less healthy participants. Because dropout time is centered at its mean in the CLM, $\hat{\gamma}_3$ is the expected change in square root CD4 from week 0 to 220 if a subject were followed for the whole study period (note, however, that the standard error is from the conditional distribution of \mathbf{Y} given U). A comparison of the estimates of γ under REM and CLM suggests that the effect of accounting for dropout via CLM is to correct the average change in $(CD4)^{1/2}$ downward.

In the VCM we allow both $\beta_1(u)$ and $\beta_2(u)$ to be completely unspecified for both treatment arms, and estimate them using cubic smoothing splines. Plots of the estimated functionals from the VCM, together with empirical Bayes estimates of individual CD4 intercepts and slopes, appear in Figure 3. In both treatment arms, CD4 intercept and slope are highly associated with dropout time: those who remain in the study for longer periods have higher values of both. Except for the CD4 intercept on the low-dose arm, the associations appear to be highly nonlinear.

4.2.2 Comparison of treatment effect inferences across models. Because the complete data likelihood factors over (β, θ) and π , and because our three models differ only in the specification of the conditional factor $f(\mathbf{y}|u; \beta, \theta)$, model selection can in principle be based on criteria for the conditional $(\mathbf{y}|u)$ model. However, formal model selection procedures comparing parametric and nonparametric mixed effects models are not currently well developed and are beyond the scope of this article; furthermore, there are potential difficulties associated with comparing likelihoods for models that are not properly nested in a traditional way, e.g., due to boundary-value problems.

Table 2 lists estimates and associated standard errors for the marginal regression coefficients $\hat{\beta}_j = \int \beta_j(u) dF(u)$, for $j = 1, 2$, estimated according to the procedure described in Section 3.3 (note that the marginal coefficients—and not the conditional coefficients reported in Table 1—are of direct scientific interest). Standard errors were calculated using bootstrap resampling based on 100 replicated datasets sampled with replacement. Quantile plots of the bootstrap parameter estimates showed no obvious departures from normality (not shown), so Z -statistics are used for inferences. Not surprisingly, standard errors associated with the VCM are increased relative to the CLM, reflecting uncertainty about the functional form of $\beta_j(u)$. For the low-dose arm, the increase is relatively modest. Comparing slopes in the low-dose arm, for example, $SE(\hat{\beta}_2) = 1.0$ for the CLM and 1.3 for the VCM. No appreciable difference in standard errors is seen in the estimated intercepts.

Table 2 also indicates the degree to which adjusting for selection bias affects the final inferences about treatment. Under the MAR assumption (REM), estimated mean difference in total change in $(CD4)^{1/2}$ is -5.5 , with Z -statistic $= -3.4$; adjustment under the CLM gives the same estimated effect, with Z -statistic $= -2.8$. Both lead to the conclusion that low dose is superior to high dose because the decline in CD4 is less steep. Under the VCM, the correction for possible selection biases on low dose changes the slope estimate from -12.7 (REM) to -17.1 ; the correction is less severe on high dose (-18.2 for REM, compared to -20.1 for VCM); the effect is to narrow the gap in treatment effect to -3.0 , with $Z = -1.1$, representing a change of 1.56 standard errors relative to the REM, and 1.25 standard errors relative to the CLM.

The VCM also provides an important substantive insight, namely that participants who drop out of low dose (the experimental dose in this trial) tend to have steeper decline in their CD4 counts, compared to those on high dose. The trial was designed to see whether the lower dose, known to be associated with fewer side effects in adults, would have efficacy equal to the high dose. The form of the $\beta_2(u)$ functions suggests that among the early dropouts, rate of change in CD4 for those on low dose is substantially less than for those on high dose. In an MAR analysis, early dropouts contribute less information to the estimate of population slope because they have fewer observed data points, leading to the potential selection bias seen in the REM.

5. Simulation Study

Our model gives the analyst considerable flexibility in specifying dependence between outcome and dropout in the context of a mixture model, and avoids biases that are possible if the functional form of this dependence is assumed to be known. The primary innovation of the VCM over CLM is that $\beta(u)$ can be left unspecified, but this generalization relies on the key assumption that $\beta(u)$ is a (vector of) smooth, twice differentiable functions of u . We designed a brief simulation study to investigate the performance of our model under violations of this assumption.

Each simulation uses datasets with $n = 50$ subjects having up to 15 unique dropout times, with $\beta(u)$ taking three different forms. We compare estimates of mean change from baseline from a standard REM, the CLM with components of

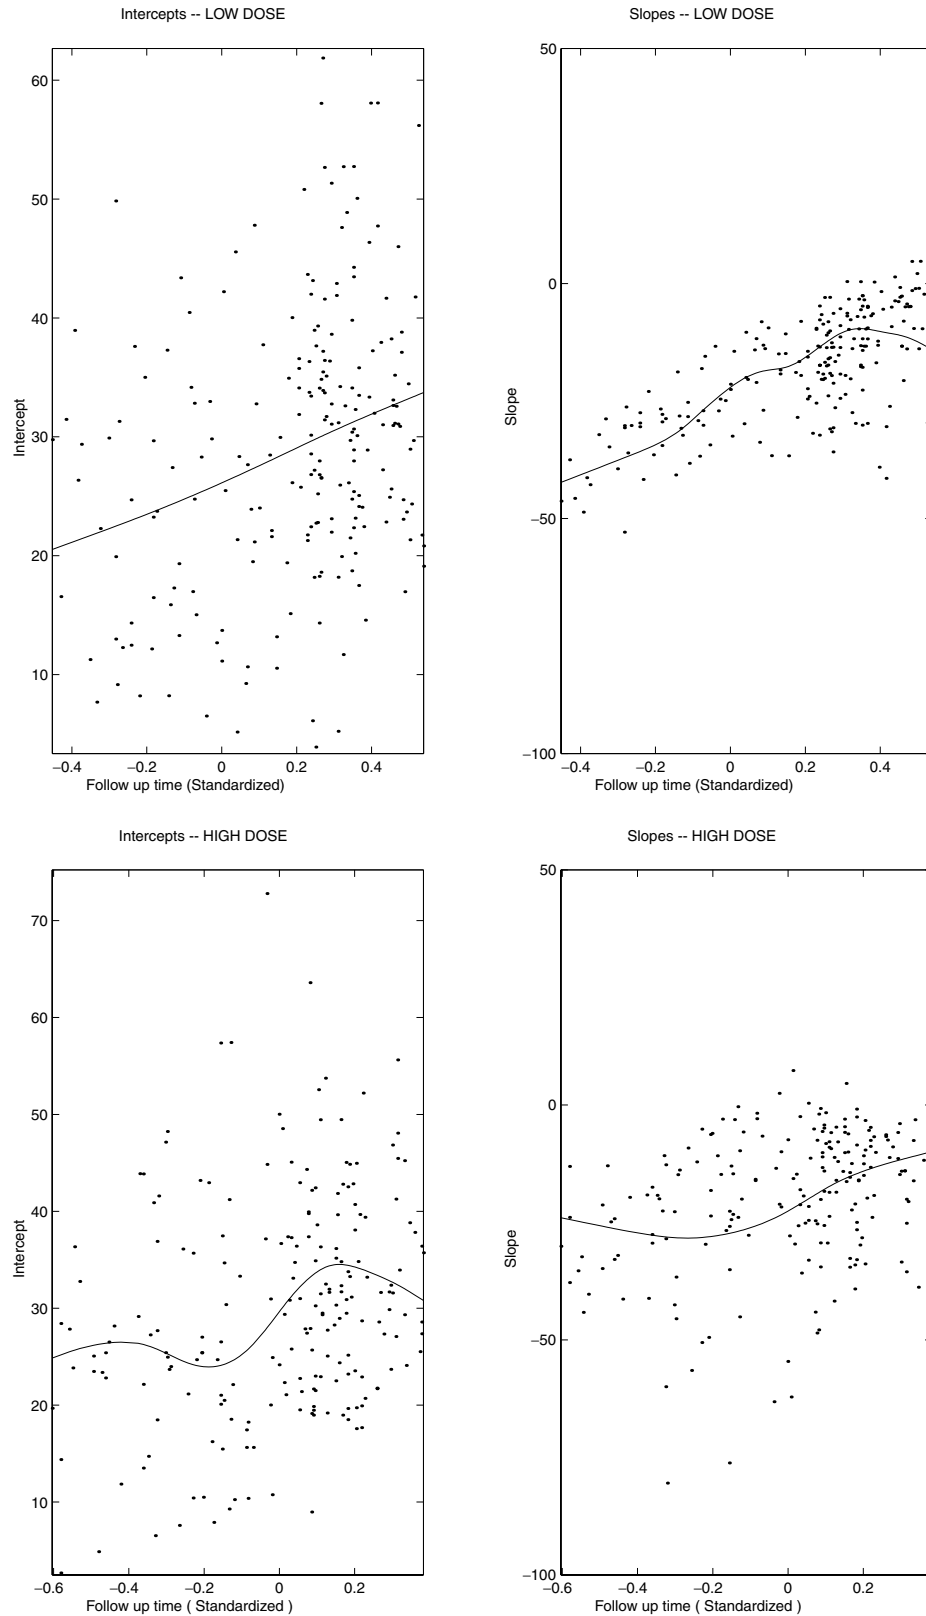


Figure 3. Estimated functions $\beta_1(u)$ and $\beta_2(u)$ for low- and high-dose ZDV arms, together with empirical Bayes estimates of individual intercepts and slopes (on square root of CD4 scale). Slopes correspond to change from baseline to week 200. Standardized follow-up times correspond to deviation from the average from the combined sample, and one unit represents 200 weeks.

Table 2

Estimated intercept and slope characterizing marginal mean of CD4 trajectory under three different specifications for conditional part of joint model, with standard errors estimated via bootstrap

Model	Parameter	Low dose	High dose	Difference (SE)	Z
REM	β_1	28.6 (0.8)	30.1 (0.9)		
	β_2	-12.7 (0.8)	-18.2 (1.4)	-5.5 (1.6)	-3.4
CLM	β_1	28.9 (0.8)	30.3 (0.9)		
	β_2	-15.9 (1.0)	-21.4 (1.8)	-5.5 (2.0)	-2.8
VCM	β_1	29.0 (0.7)	29.9 (0.9)		
	β_2	-17.1 (1.3)	-20.1 (2.4)	-3.0 (2.7)	-1.1

$\beta(u)$ assumed linear in u , and from the VCM with $\beta(u)$ left unspecified. Specifically, we assume

$$y_{il} = \beta_1(u_i) + \beta_2(u_i)t_{il} + b_{1i} + b_{2i}t_{il} + e_{il},$$

where $(b_{1i}, b_{2i})^T \sim N(\mathbf{0}, \mathbf{D})$, $e_{il} \sim N(0, \sigma^2)$. There are 15 time points $\{t_{il}\}$, equally spaced between 0 and 1. This simulation uses $d_{11} = 4$, $d_{22} = 0.1$, $d_{12} = -0.1$ (correlation ≈ -0.15), and $\sigma^2 = 1$, which implies that between-subject variation exceeds within-subject variation by a factor of about 4.

Dropout is generated from a beta mixture of binomial distributions as follows: $p \sim \text{Beta}(1.5, 1.5)$ (mean 0.5), $U^* \sim \text{Bin}(15, p)$, and dropout time $U = U^*/15 \in (0, 1)$. Finally, we assume $\beta_1(u) = 0$ and vary the functional form of $\beta_2(u)$; candidate functions are

- (i) $-\exp(\alpha u)$,
- (ii) $\exp(\alpha u) I(u < t^*) + \exp(\alpha t^*) I(u \geq t^*)$ (exponential with plateau effect for dropouts beyond t^*),
- (iii) $\alpha_1 I(u < t^*) + \alpha_2 I(u \geq t^*)$ (two-piece step function).

Case (i) actually meets the assumptions for the VCM, and is included for validating our simulation and estimation routines; case (ii) violates the smoothness assumption and case (iii) violates both smoothness and continuity assumptions.

For (i), at $\alpha = -4$, completers ($U = 1$) have mean change from baseline $\beta_2(1) = \exp(-4) \approx 0.02$ and early dropouts ($U = 0$) have mean change -1 , a difference of about 3 SD (because $d_{22} = 0.1$). Under (ii), we keep $\alpha = -4$ and invoke the plateau effect at $t^* = 2/3$, leading to a structure wherein those who complete 2/3 of the study or more have average change from baseline equal to $\exp(-4 \times 2/3) \approx 0.07$. For (iii), we keep $t^* = 2/3$ and set $\alpha_1 = 0$, $\alpha_2 = 1$.

Results are reported in Table 3. As expected, the VCM gives virtually unbiased estimation of the true slope for case (i), where $\beta_2(u)$ is both continuous and smooth, while both the REM and CLM show substantial upward bias. This comparison is not as trivial as it would appear, however, because exploratory plots of OLS slopes versus dropout time (e.g., Figure 2) do not always reveal an obvious functional form for $\beta(u)$, particularly in the early part of the time axis. One advantage to the VCM is its effectiveness in finding a signal from noisy data.

The VCM shows only very little bias for estimating the true slope for the continuous but not everywhere-differentiable

Table 3

Results from simulation to characterize bias. REM = linear random effects model; CLM = conditional linear model with $\beta_2(u)$ linear; VCM = varying coefficient model with $\beta_2(u)$ unspecified. Each estimate represents a sample average of estimated slopes over 100 replicated datasets, each having 100 subjects with up to 15 repeated measures. Standard errors for simulation-based estimated mean appear in parentheses.

Underlying model ^a	True slope (β_2)	Estimated slope		
		REM	CLM	VCM
(i) Continuous, smooth	-0.159 ^b	-0.073 (0.018)	-0.119 (0.035)	-0.160 (0.028)
(ii) Continuous, not smooth	-0.170 ^b	-0.062 (0.015)	-0.100 (0.027)	-0.166 (0.033)
(iii) Discontinuous	-0.587 ^b	-0.211 (0.018)	-0.622 (0.031)	-0.715 (0.038)

^aSee Section 5 for model descriptions.

^bComputed to nearest 0.001 via Monte Carlo simulation.

function from case (ii), but exhibits more bias than the CLM for the discontinuous function in case (iii). In all cases, however, the VCM outperforms the REM.

6. Discussion

6.1 Summary

We have proposed a mixture-modeling approach to analyzing longitudinal data with outcome-dependent dropout. Our model assumes that covariate effects depend on dropout time through unspecified functions $\{\beta_j(u)\}$, where u is dropout time. When dropout times are discrete, the $\beta_j(u)$ are step functions, and when dropout is continuous, $\beta_j(u)$ are assumed to be unspecified smooth functions of u . This formulation generalizes pattern-mixture models (Little and Baile, 1988, 1989; Hogan and Laird, 1997a; Albert and Follman, 2000) for continuous response data. Using an example from an AIDS clinical trial, we show that the model has the potential to adjust for selection biases induced by poor responders dropping out early.

The primary innovation in our approach is that the functional dependence between covariate effects and dropout time can be left unspecified. Our simulation study shows that when this relationship is misspecified, CLMs yield biased estimates, while varying coefficient mixture models still yield unbiased estimates and are more robust. In many applications, this is a decided advantage over the CLM model because the form of $\beta(u)$ rarely will be known or intuitive. Moreover, it is our experience that using polynomials leads to overfitting and/or extrapolations well outside the range of data, particularly when the polynomial has degree >2 . When u is continuous, our simulation also shows that inferences are unlikely to be sensitive to lack of smoothness in $\beta(u)$, but could be affected by discontinuities. In both cases, however, bias is substantially less than under an MAR analysis.

6.2 Strategies for Sensitivity Analysis

Another advantage to mixture modeling in general is that extrapolations of the missing data are transparent, and lend

themselves well both to substantive critique and to empirical sensitivity analysis (e.g., Rubin, 1977; Little and Wang, 1996; Daniels and Hogan, 2000; Rotnitzky et al., 2001). The nonidentifiable component of our model is the distribution of missing responses following dropout, $f(\mathbf{y}_{\text{mis}} | \mathbf{y}_{\text{obs}}, u)$.

In our data application, for example, we assume dropouts at time u have the same slope for $t > u$ as for $t \leq u$; on the surface this is a strong assumption but it is relatively easily modified and is a sensible starting point for sensitivity analysis. Specifically, the VCM in our data example takes the form

$$(Y_{ij} | U_i = u) = \beta_{1i}(u) + \beta_{2i}(u)t_{ij} + \epsilon_{ij}, \quad (11)$$

where $\beta_{1i}(u) = \beta_1(u) + b_{1i}$, $\beta_{2i}(u) = \beta_2(u) + b_{2i}$, $\mathbf{b}_i = (b_{1i}, b_{2i})^T \sim N(\mathbf{0}, \mathbf{D}(u))$, and $\epsilon_{ij} \sim N\{0, \sigma^2(u)\}$ and is independent of \mathbf{b}_i . Note that in our application, we assume $\mathbf{D}(u) = \mathbf{D}$ and $\sigma^2(u) = \sigma^2$, an assumption that can be relaxed by introducing models to characterize the variance components as a function of u (Daniels and Pourahmadi, 2002).

One approach to sensitivity analysis in model (11) is to assume a different slope on time for $t > u$, i.e., assume a continuous piecewise linear model with the change point u as

$$(Y_{ij} | U_i = u) = \beta_{1i}(u) + \beta_{2i}(u)t_{ij} + \delta_i(u)(t_{ij} - u)_+ + \epsilon_{ij}, \quad (12)$$

where $a_+ = a$ if $a > 0$ and 0 otherwise, $\delta_i(u) = \delta(u) + d_i$, and $\mathbf{b}_i^* = (b_{1i}, b_{2i}, d_i)^T \sim N(\mathbf{0}, \mathbf{D}^*)$ is a 3×3 matrix for random effects. Model (12) assumes the slope changes from $\beta_{2i}(u)$ for $t \leq u$ to $\beta_{2i}(u) + \delta_i(u)$ for $t > u$. The nonidentifiable sensitivity parameters are therefore $\delta(u)$ and the variance components comprising the third row (column) of \mathbf{D}^* . Observed data provide no information about these parameters; one approach is to fix them at various values and recompute quantities of interest (such as expected change in CD4 from beginning to end of the study) and examine their sensitivity with respect to $\delta(u)$ and the unidentifiable parameters of \mathbf{D}^* . Details of the sensitivity analysis will be investigated in future work.

ACKNOWLEDGEMENTS

Work on this project was funded by grants R01-AI-50505 and P30-AI-42853 (Hogan) and CA76404 (Lin) from the U.S. National Institutes of Health. The authors are grateful to Rusty Tchernis for assistance with computing related to the simulation and implementation of the bootstrap, and to Jason Roy and two anonymous reviewers for helpful comments.

RÉSUMÉ

L'analyse d'observations longitudinales répétées est souvent compliquée par l'absence de données consécutives à des sorties d'étude informatives. Nous décrivons un modèle par mélange pour la distribution conjointe de mesures longitudinales répétées où la distribution des sorties peut être continue et où la relation entre réponse et sortie est semi-paramétrique. Plus précisément, nous supposons que les réponses suivent un modèle mixte conditionnellement au temps de sortie, dont les coefficients dépendent de la sortie par l'intermédiaire de fonctions non paramétriques estimées à l'aide de fonctions de saut, quand les dates de sortie sont discrètes (études de panel), et à l'aide de fonctions de splines quand les sorties sont continues. L'inférence à partir de ce modèle est ainsi plus robuste

qu'à partir du modèle linéaire conditionnel paramétrique. La distribution non conditionnelle des mesures répétées est un mélange de distributions conditionnelles. Nous montrons que l'on peut estimer le modèle par mélange en ajustant un modèle paramétrique à effets mixtes à l'aide d'un logiciel standard comme SAS. Le modèle est appliqué aux données d'un essai thérapeutique récent contre le SIDA et ses performances sont évaluées par simulations.

REFERENCES

- Albert, P. S. and Follmann, D. (2000). Modeling repeated count data subject to informative dropout. *Biometrics* **56**, 667–677.
- Brady, M. T., McGrath, N., Brouwers, P., et al., for the Pediatric AIDS clinical trial. (1996). Randomized study of the tolerance and efficacy of high- versus low-dose zidovudine in human immunodeficiency virus-infected children with mild to moderate symptoms (ACTG 128). *Journal of Infectious Disease* **173**, 1097–1106.
- Breslow, N. E. and Clayton, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* **88**, 9–25.
- Daniels, M. J. and Hogan, J. W. (2000). Reparameterizing the pattern mixture model for sensitivity analyses under informative dropout. *Biometrics* **56**, 1241–1248.
- Daniels, M. and Pourahmadi, M. (2002). Bayesian analysis of covariance matrices and dynamic models for longitudinal data. *Biometrika* **89**, 553–566.
- DeGruttola, V. and Tu, X. M. (1994). Modeling the progression of CD4-lymphocyte count and its relationship to survival time. *Biometrics* **50**, 1003–1014.
- Diggle, P. J. (1988). An approach to the analysis of repeated measurements. *Biometrics* **44**, 959–971.
- Diggle, P. and Kenward, M. G. (1994). Informative drop-out in longitudinal data analysis. *Applied Statistics* **43**, 49–73.
- Fitzmaurice, G. M. and Laird, N. M. (2000). Generalized linear mixture models for handling nonignorable dropouts in longitudinal studies. *Biostatistics* **1**, 141–156.
- Fitzmaurice, G. M., Laird, N. M., and Shneyer, L. (2001). An alternative parameterization of the general linear mixture model for longitudinal data with non-ignorable drop-outs. *Statistics in Medicine* **20**, 1009–1021.
- Follman, D. and Wu, M. C. (1995). An approximate generalized linear model with random effects for informative missing data. *Biometrics* **51**, 151–168.
- Green, P. J. and Silverman, B. W. (1994). *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. London: Chapman & Hall.
- Hogan, J. W. and Laird, N. M. (1997a). Mixture models for the joint distribution of repeated measures and event times. *Statistics in Medicine* **16**, 239–257.
- Hogan, J. W. and Laird, N. M. (1997b). Model-based approaches to analysing incomplete longitudinal and failure time data. *Statistics in Medicine* **16**, 259–272.
- Kenward, M. G. and Molenberghs, G. (1999). Parametric models for incomplete continuous and categorical longitudinal data. *Statistical Methods in Medical Research* **8**, 51–83.

- Laird, N. M. (1988). Missing data in longitudinal studies. *Statistics in Medicine* **7**, 305–315.
- Laird, N. M. and Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics* **38**, 963–974.
- Liang, K. Y. and Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika* **73**, 13–22.
- Little, R. J. A. (1993). Pattern-mixture models for multivariate incomplete data. *Journal of the American Statistical Association* **88**, 125–134.
- Little, R. J. A. (1994). A class of pattern-mixture models for normal incomplete data. *Biometrika* **81**, 471–483.
- Little, R. J. A. (1995). Modeling the drop-out mechanism in repeated-measures studies. *Journal of the American Statistical Association* **90**, 1112–1121.
- Little, R. J. A. and Wang, Y. (1996). Pattern-mixture models for multivariate incomplete data with covariates. *Biometrics* **52**, 98–111.
- O’Sullivan, F., Yandell, B. S., and Raynor, W. J., Jr. (1986). Automatic smoothing of regression functions in generalized linear models. *Journal of the American Statistical Association* **81**, 96–103.
- Robins, J., Rotnitzky, A., and Zhao, L. P. (1995). Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association* **90**, 106–121.
- Rotnitzky, A., Robins, J. M., and Scharfstein, D. O. (1998). Semiparametric regression for repeated outcomes with non-ignorable non-response. *Journal of the American Statistical Association* **93**, 1321–1339.
- Rotnitzky, A., Scharfstein, D., Su, T. L., and Robins, J. (2001). Methods for conducting sensitivity analysis of trials with potentially nonignorable competing causes of censoring. *Biometrics* **57**, 103–113.
- Rubin, D. B. (1977). Formalizing subjective notions about the effect of nonrespondents in sample surveys. *Journal of the American Statistical Association* **72**, 538–543.
- Scharfstein, D., Robins, J., and Rotnitzky, A. (1999). Adjusting for nonignorable nonresponse using semiparametric nonresponse models with time dependent covariates (with discussion). *Journal of the American Statistical Association* **94**, 1096–1146.
- Schluchter, M. D. (1992). Methods for the analysis of informatively censored longitudinal data. *Statistics in Medicine* **11**, 1861–1870.
- Ten Have, T. R., Kunselman, A. R., Pulkstenis, E. P., and Landis, J. R. (1998). Mixed effects logistic regression models for longitudinal binary response data with informative drop-out. *Biometrics* **54**, 367–383.
- Wu, M. C. and Bailey, K. (1988). Analysing changes in the presence of informative right censoring caused by death and withdrawal. *Statistics in Medicine* **7**, 337–346.
- Wu, M. C. and Bailey, K. (1989). Estimation and comparison of changes in the presence of informative right censoring: Conditional linear model (corr: V46 p. 889). *Biometrics* **45**, 939–955.
- Wu, M. C. and Carroll, R. J. (1988). Estimation and comparison of changes in the presence of informative right censoring by modeling the censoring process. *Biometrics* **44**, 175–188.
- Zhang, D., Lin, X., Raz, J., and Sowers, M. (1998). Semiparametric stochastic mixed models for longitudinal data. *Journal of the American Statistical Association* **93**, 710–719.

Received May 2003. Revised May 2004.

Accepted May 2004.