Published in final edited form as:

Comput Stat Data Anal. 2008 August 15; 52(12): 5100–5120. doi:10.1016/j.csda.2008.05.008.

A Flexible Approach to Bayesian Multiple Curve Fitting

Carsten H. Botts and

Department of Mathematics and Statistics, Williams College, Williamstown, Massachusetts 01267, USA, cbotts@williams.edu

Michael J. Daniels

Department of Epidemiology and Biostatistics, and Department of Statistics, University of Florida, Gainesville, Florida 32611, USA, mdaniels@stat.ufl.edu

Summary

We model sparse functional data from multiple subjects with a mixed-effects regression spline. In this model, the expected values for any subject (conditioned on the random effects) can be written as the sum of a population curve and a subject-specific deviate from this population curve. The population curve and the subject-specific deviates are both modeled as free-knot b-splines with k and k' knots located at $\mathbf{t}_k$ and $\mathbf{t}_{k'}$, respectively. To identify the number and location of the “free” knots, we sample from the posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ using reversible jump MCMC methods. Sampling from this posterior distribution is complicated, however, by the flexibility we allow for the model’s covariance structure. No restrictions (other than positive definiteness) are placed on the covariance parameters ψ and σ^2 and, as a result, no analytical form for the likelihood $p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ exists. In this paper, we consider two approximations to $p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ and then sample from the corresponding approximations to $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$. We also sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ which has a likelihood that is available in closed form. While sampling from this larger posterior is less efficient, the resulting marginal distribution of knots is exact and allows us to evaluate the accuracy of each approximation. We then consider a real data set and explore the difference between $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ and the more accurate approximation to $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$.

Keywords

B-splines; Laplace approximation; Reversible jump MCMC; Unit-information prior

1 Introduction

Social and physical scientists are often interested in how certain variables depend on one another. They often assume that a functional relationship exists between these two variables, and they run experiments or collect samples to learn about this relationship. They are aware, however, that the data does not follow a deterministic equation, but follows the general stochastic equation

$$Y = f(x) + \varepsilon, \quad (1)$$

where Y is a dependent variable, x is an independent or controlled variable, and ε is random error with mean 0. One of the simplest models relating a response, Y , to a single predictor, x , is the p^{th} -order polynomial regression model

$$Y(x) = \beta_0 + \sum_{j=1}^p \beta_j x^j + \varepsilon.$$

While this model is appealing and can be very useful, it may be inadequate since it assumes that only one polynomial describes the average relationship between Y and x . A more complex model would allow several polynomials to do this. An example of such a model is a piecewise polynomial regression spline. A piecewise polynomial regression spline with k knots divides the domain of $f(x)$ into $k + 1$ mutually exclusive regions, and to each region corresponds a unique polynomial describing the average relationship between Y and x . With different polynomials describing different parts of the function, $f(x)$ is not restricted to have the same smoothness throughout its domain. A piecewise polynomial regression spline of order p with k knots at $\mathbf{t}_k = (t_1, t_2, \dots, t_k)$ is written as

$$f(x) = \beta_0 + \beta_1 x + \sum_{i=1}^k (x - t_i)_+^p \beta_{1+i} + \sum_{j=2}^p \beta_{k+j} x^j, \quad (2)$$

where $\beta = (\beta_0, \beta_1, \dots, \beta_{p+k})$ is a fixed set of parameters and $(z)_+^p = z^p \mathbb{I}(z \geq 0)$.

A natural polynomial regression spline takes a form nearly identical to the piecewise regression spline seen above. It is different, however, in that it restricts the function to be linear at the boundaries. This sets the last term on the right-hand side of (2) to 0. A natural polynomial regression spline of order p thus takes the form

$$f(x) = \beta_0 + \beta_1 x + \sum_{i=1}^k (x - t_i)_+^p \beta_{1+i}.$$

If $f(x)$ were modeled as such, (1) could then be re-written as

$$Y(x) = \mathbf{c}(x, \mathbf{t}_k) \beta + \varepsilon,$$

where $\mathbf{c}(x, \mathbf{t}_k) = (1, x, (x - t_1)_+^p, \dots, (x - t_k)_+^p)$ and $\beta = (\beta_0, \beta_1, \beta_2, \dots, \beta_{1+k})$, and for m observations collected at $(x_1, x_2, \dots, x_m)$, it would be written as

$$Y(\mathbf{x}) = \begin{pmatrix} Y(x_1) \\ Y(x_2) \\ \vdots \\ Y(x_m) \end{pmatrix} = \begin{bmatrix} \mathbf{c}(x_1, \mathbf{t}_k) \\ \mathbf{c}(x_2, \mathbf{t}_k) \\ \vdots \\ \mathbf{c}(x_m, \mathbf{t}_k) \end{bmatrix} \beta + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_m \end{pmatrix}. \quad (3)$$

The natural polynomial regression spline written in Equation (3) illustrates the concept behind splines, but the design matrix composed of the row vectors $\mathbf{c}(x_1, \mathbf{t}_k), \dots, \mathbf{c}(x_m, \mathbf{t}_k)$ is unstable and thus rarely used. B-splines, which are computationally more stable, are preferred (Hastie and Tibshirani, 1990; Zhou and Shen, 2001). The functional form of a b-spline is more complex than that of the natural polynomial regression spline, but is readily available (deBoor, 1978). We write it as

$$f(x) = \mathbf{b}(x, \mathbf{t}_k) \alpha,$$

where $\mathbf{b}(x, \mathbf{t}_k)$ is the design vector associated with a b-spline evaluated at x with k knots at $\mathbf{t}_k = (t_1, t_2, \dots, t_k)$, and α is a fixed set of parameters. For m observations collected at $(x_1, \dots, x_m)$, the model would be written as

$$Y(\mathbf{x}) = \begin{pmatrix} Y(x_1) \\ Y(x_2) \\ \vdots \\ Y(x_m) \end{pmatrix} = \begin{bmatrix} \mathbf{b}(x_1, \mathbf{t}_k) \\ \mathbf{b}(x_2, \mathbf{t}_k) \\ \vdots \\ \mathbf{b}(x_m, \mathbf{t}_k) \end{bmatrix} \alpha + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_m \end{pmatrix}.$$

With this model, estimating $f(x)$ becomes a problem of estimating the number of knots, k , the locations of those knots, $\mathbf{t}_k$, and α .

A variety of methods have been developed in estimating k and $\mathbf{t}_k$. Halpern (1973) approached this problem using Bayesian methods. Allowing knots to only be placed at the design points in the experiment, he considered all of the subsets of the design points. Halpern assigned prior probabilities to all of these subsets, and calculated the corresponding posterior probabilities of these subsets. His estimator of the function was based on these posterior probabilities. Denison, Mallick, and Smith (1998) placed priors on the number of knots, k , and their locations, $\mathbf{t}_k$. With these priors, they calculated a joint posterior distribution which included the variables k and $\mathbf{t}_k$ and then sampled from this posterior distribution using reversible jump MCMC methods. They too restricted the knots to be located only at the design points of the experiment. DiMatteo, Genovese, and Kass (2001) proposed a method similar to that of Denison, et al. They did not restrict the knots to be located only at the design points of the experiment and they penalized models with unnecessarily large values of k by averaging over the fixed effect coefficients, α .

Friedman (1991) and Stone et al. (1997) try to solve the problem of knot selection using backward and/or forward knot selection. This process continues until the “best” model has been identified. Zhou and Shen (2001) used an alternative method in identifying the locations of the knots. They constructed an algorithm which favored adding knots in locations where several knots had already been added. Lindstrom (1999) used similar methods when selecting knot locations.

Splines have also been used to model curves that vary within and between subjects (Shi et al., 1996, Rice and Wu, 2001, Behseta et al., 2005, Crainiceanu et al., 2005). Brumback and Rice (1998) use splines to model nested functions which vary between subjects and between groups of subjects. Functional models which only account for variation between and within subjects take the general form

$$\begin{aligned} Y_i(x_j) &= f_i(x_j) + \varepsilon_{ij} \\ &= f(x_j) + G_i(x_j) + \varepsilon_{ij}, \end{aligned} \tag{4}$$

where $Y_i(x_j)$ is the observation of the i^{th} individual at x_j , $f(x)$ can be thought of as a population curve, and $G_i(x)$ is a random curve specific to subject i . While these functions can be modeled in a variety of ways, we model these two functions as free-knot b-splines with k and k' knots located at $\mathbf{t}_k$ and $\mathbf{t}_{k'}$, respectively. Setting $f(x) = \mathbf{b}(x; \mathbf{t}_k) \alpha$ and $G_i(x) = \mathbf{b}(x; \mathbf{t}_{k'}) \gamma_i$, where γ_i are independent random vectors such that $\gamma_i \sim N(0, \Sigma_\gamma)$, equation (4) becomes

$$Y_i(x_j) = \mathbf{b}(x_j; \mathbf{t}_k) \alpha + \mathbf{b}(x_j; \mathbf{t}_{k'}) \gamma_i + \varepsilon_{ij}. \quad (5)$$

Bigelow and Dunson (2007) consider a variation of this model and identify the location of the knots by sampling from the full posterior distribution. They lessen the computational burden of this sampling procedure by setting $\mathbf{t}_{k'} = \mathbf{t}_k$ and restricting Σ_γ to be diagonal. In this paper, we consider a more flexible model and explore the computational issues associated with it. To be more specific, we eliminate the restrictions imposed by Bigelow and Dunson and then try to identify the number and location of the knots by sampling from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$. This posterior can not be sampled from exactly, however, because eliminating the diagonal restriction on Σ_γ makes the likelihood $p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ intractable. We thus sample from two approximations to this posterior; each approximation corresponds to a different approximation to the likelihood $p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'})$. To assess the accuracy of these approximations, we compare the posterior distribution of knots in each case to the marginal distribution of knots when sampling from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$, where ψ is a set of covariance parameters and σ^2 is the within-subject variability. The marginal distribution of knots when sampling from this higher-dimensional posterior is exact as its corresponding likelihood is available in closed form. Sampling from this larger posterior, however, is less efficient because of the additional parameters that need to be sampled.

In Section 2, we discuss the linear model given in (5) in more detail. In Section 3, the likelihood of interest, $p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'})$, is discussed as are the two methods we use to approximate this likelihood. Section 4 shows the algorithm used to sample from the resulting approximate posterior distributions, and Section 5 shows the algorithm we use to sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$. Section 6 theoretically compares the approximations, and in Section 7, we describe a simulation study performed to explore the differences in these approximations. In Section 8, we apply the preferred approximation to a real data set.

2 The model and priors

Let $Y_i(x_j)$ be the observed value of the i^{th} curve ($i = 1, \dots, n$) at x_j . We specify the following mixed model for functional data,

$$Y_i(x_j) = \mathbf{b}(x_j, \mathbf{t}_k) \alpha + \mathbf{b}(x_j, \mathbf{t}_{k'}) \gamma_i + \varepsilon_{ij}, \quad (6)$$

where $\mathbf{b}(x_j, \mathbf{t}_q)$ is the design vector of a b-spline with q knots at $\mathbf{t}_q$, k is the number of fixed effect knots at locations $\mathbf{t}_k$, k' is the number of random effect knots at $\mathbf{t}_{k'}$, α is a fixed set of parameters with $\dim(\alpha) = k+2$, γ_i are independent and identically distributed random vectors such that $\gamma_i \sim N(0, \Sigma_\gamma)$ and $\dim(\gamma_i) = k'+2$, ε_{ij} are independent random variables with $\varepsilon_{ij} \sim N(0, \sigma^2)$ for all pairs (i, j) , and the random vectors γ_i are independent of all random variables ε_{ij} . The appealing feature of this model is its flexibility; recall that no restriction is being placed on Σ_γ (other than $|\Sigma_\gamma| > 0$), and that the random and fixed effect knots are not set equal to one another. We do assume, however, that all individuals have the same random effect knots, and that $k' \leq k$. When the number of observations per subject is small, such assumptions are needed, as they limit the number of complicated models proposed in the Markov chain we consider.

We place prior distributions on $k, k', \mathbf{t}_k, \mathbf{t}_{k'}, \alpha, \Sigma_\gamma$, and σ^2 . Poisson priors are given to k and k' . Additionally, since there is no reason to favor knots at any particular location on the domain of f and G , flat priors are placed on both $\mathbf{t}_k$ and $\mathbf{t}_{k'}$. These priors are written as

$$k \sim \text{Poi}(\mu_k), \quad p(\mathbf{t}_k | k) \propto \mathbb{I}(a < t_{k(1)} < t_{k(2)} < \dots < t_{k(k)} < b), \\ k' | k \sim \text{Poi}(\mu_{k'}) \mathbb{I}(k' \leq k), \quad p(\mathbf{t}_{k'} | k') \propto \mathbb{I}(a < t_{k'(1)} < t_{k'(2)} < \dots < t_{k'(k')} < b),$$

where $t_{q(j)}$ is the j^{th} smallest knot in the vector $\mathbf{t}_q$, and (a, b) is the domain of f and G .

For Σ_γ and σ^2 , we place prior distributions on transformed values of these parameters. The transformations we consider create a set of unconstrained parameters which are preferable to work with since their sampling distributions can be well approximated with the normal density. For Σ_γ , we applied the modified Cholesky parameterization proposed by Pourhamadi (1999). Pourhamadi decomposes the inverse of the covariance matrix Σ_γ as $\Sigma_\gamma^{-1} = TDT'$, where

$$T = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ -\phi_{2,1} & 1 & 0 & \dots & 0 \\ -\phi_{3,1} & -\phi_{3,2} & 1 & \dots & 0 \\ \vdots & & & \ddots & \vdots \\ -\phi_{p,1} & -\phi_{p,2} & \dots & -\phi_{p,p-1} & 1 \end{bmatrix}, \text{ and } D = \begin{bmatrix} \sigma_1^2 & 0 & \dots & 0 \\ 0 & \sigma_2^2 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & 0 & \sigma_p^2 \end{bmatrix}.$$

These parameters can be calculated directly from the elements of the covariance matrix using the functions $g(\cdot; \cdot)$ and $h(\cdot; \cdot)$, where $\sigma_j^2 = g(\Sigma_\gamma; j) = \Sigma_\gamma[j, j] - \sigma_j' \Sigma_{\gamma,j}^{-1} \sigma_j$ for $2 \leq j \leq p$ ($\sigma_1^2 = \Sigma_\gamma[1, 1]$), $\phi_j = (\phi_{j,1}, \phi_{j,2}, \dots, \phi_{j,j-1})' = h(\Sigma_\gamma; j) = \Sigma_{\gamma,j}^{-1} \sigma_j$, $\Sigma_{\gamma,j}$ is the minor matrix within Σ_γ composed of its first $j-1$ columns and rows, and σ_j is the vector within Σ_γ composed of the first $j-1$ elements of the j^{th} column. To ease our notational burden, we will refer to these modified Cholesky parameters collectively as ψ , where

$\psi = (\phi_{2,1}, \log(\sigma_1^2), \phi_{3,1}, \phi_{3,2}, \log(\sigma_2^2), \dots, \log(\sigma_p^2))$, and we will replace the symbol Σ_γ with ψ (or, to illustrate that Σ_γ is a function of ψ , with $\Sigma_\gamma(\psi)$). All of the elements in the vector ψ are unconstrained, and a multivariate normal unit-information prior is assigned to them. The mean of the unit-information prior on ψ is the maximum likelihood estimate of the parameters, $\hat{\psi}$, and the variance is inversely proportional to approximately one unit of information on ψ . The prior can be written as

$$\psi | \mathbf{t}_k, \mathbf{t}_{k'} \sim MVN\left(\hat{\psi}, -\left[\frac{1}{n - (k+2)} I_\psi\right]^{-1}\right),$$

where

$$I_\psi = \frac{\partial^2 \log \{p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \psi, \sigma^2)\}}{\partial \psi' \partial \psi} \Big|_{\alpha = \hat{\alpha}, \psi = \hat{\psi}, \sigma^2 = \hat{\sigma}^2},$$

$p(y | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \psi, \sigma^2)$ is the distribution of all the observed data in the study, n is the number of subjects in the study, and $\hat{\alpha}$, $\hat{\psi}$, and $\hat{\sigma}$ are the maximum likelihood estimates of α , ψ , and σ , respectively. The exact formula for I_ψ is given in the Appendix.

For σ^2 , a unit-information prior was placed on $\log(\sigma^2)$. The prior can be written as

$$\log(\sigma^2)|\mathbf{t}_k, \mathbf{t}_{k'} \sim N \left(\log(\widehat{\sigma}^2), -\left(\frac{1}{n} \sum_{i=1}^n \frac{1}{m_i - \frac{k+2}{n}} I_{\log(\sigma^2),i} \right)^{-1} \right),$$

where

$$I_{\log(\sigma^2),i} = \frac{\partial^2 \log \{p(y_i|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \psi, \sigma^2)\}}{\partial (\log(\sigma^2))^2} \bigg|_{\alpha=\widehat{\alpha}, \psi=\widehat{\psi}, \sigma^2=\widehat{\sigma}^2},$$

and m_i is the number of observations for subject i . The exact formula for $I_{\log(\sigma^2)}$, i is given in the Appendix.

The prior placed on α is another unit-information prior and can be written as

$$\alpha|\mathbf{t}_k, \mathbf{t}_{k'} \sim MVN \left(\widehat{\alpha}, -\left(\frac{1}{n} \sum_{i=1}^n \frac{1}{m_i} I_{\alpha,i} \right)^{-1} \right),$$

where

$$I_{\alpha,i} = \frac{\partial^2 \log \{p(y_i|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \psi, \sigma^2)\}}{\partial \alpha' \partial \alpha} \bigg|_{\alpha=\widehat{\alpha}, \psi=\widehat{\psi}, \sigma^2=\widehat{\sigma}^2}.$$

The exact formula for $I_{\alpha,i}$ is given in the Appendix.

Kass and Wasserman (1994) and Pauler (1998) defend the use of unit-information priors, especially in the context of model selection. They show that when using such priors, the resulting Bayes factors can be well approximated with the Bayesian information criterion. Although centering unit-information priors at the maximum likelihood estimates has been suggested in Bayesian adaptive regression splines (Paciorek, 2006), the final results presented in this paper were not sensitive to the means selected for these three prior distributions.

Estimating the fixed and random effect functions now becomes a problem in estimating k , $\mathbf{t}_k$, k' , and $\mathbf{t}_{k'}$. We do this by sampling from the posterior distribution $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$. While this procedure is intuitively appealing, a problem arises since the likelihood $p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ can not be calculated analytically. We consider two approximations to this likelihood in the next section.

3 Approximating $p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$

The likelihood of interest can be expressed as

$$\begin{aligned}
p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}) &= \int_{\sigma^2} \int_{\psi} \int_{\alpha} \prod_{i=1}^n \left[\int_{\gamma_i} p(y_i|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \gamma_i, \sigma^2) p(\gamma_i|\psi) d\gamma_i \right] \times p(\alpha|\mathbf{t}_k, \mathbf{t}_{k'}) p(\psi|\mathbf{t}_k, \mathbf{t}_{k'}) p(\sigma^2|\mathbf{t}_k, \mathbf{t}_{k'}) d\alpha d\psi d\sigma^2 \\
&= \int_{\sigma^2} \int_{\psi} p(y_i|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2) p(\psi|\mathbf{t}_k, \mathbf{t}_{k'}) p(\sigma^2|\mathbf{t}_k, \mathbf{t}_{k'}) d\psi d\sigma^2,
\end{aligned}
\tag{7}$$

where

$$\begin{aligned}
p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2) &= |\mathbf{A}^{-1}|^{\frac{1}{2}} \prod_i (2\pi)^{-\frac{m_i}{2}} |\mathbf{B}_{R_i} \sum_{\gamma} (\psi) \mathbf{B}'_{R_i} + \sigma^2 \mathbf{I}_{m_i}|^{-\frac{1}{2}} |\sum_{\alpha}|^{-\frac{1}{2}} \times \exp \left(-\frac{1}{2} \left(\sum_{i=1}^n y'_i (\mathbf{B}_{R_i} \sum_{\gamma} (\psi) \mathbf{B}'_{R_i} + \sigma^2 \mathbf{I}_{m_i})^{-1} y_i + \widehat{\alpha} \sum_{\alpha}^{-1} \widehat{\alpha} - \mathbf{d}' \right) \right) \\
\mathbf{d} &= \mathbf{A}^{-1} \sum_{i=1}^n \mathbf{B}'_{F_i} (\mathbf{B}_{R_i} \sum_{\gamma} (\psi) \mathbf{B}'_{R_i} + \sigma^2 \mathbf{I}_{m_i})^{-1} y_i + \sum_{\alpha}^{-1} \widehat{\alpha} \\
\mathbf{A} &= \sum_{i=1}^n \mathbf{B}'_{F_i} (\mathbf{B}_{R_i} \sum_{\gamma} (\psi) \mathbf{B}'_{R_i} + \sigma^2 \mathbf{I}_{m_i})^{-1} \mathbf{B}_{F_i} + \sum_{\alpha}^{-1}, \\
\mathbf{B}_{F_i} &= (\mathbf{b}(x_{1_i}; \mathbf{t}_k)', \mathbf{b}(x_{2_i}; \mathbf{t}_k)', \dots, \mathbf{b}(x_{m_i}; \mathbf{t}_k'))', \\
\mathbf{B}_{R_i} &= (\mathbf{b}(x_{1_i}; \mathbf{t}_{k'})', \mathbf{b}(x_{2_i}; \mathbf{t}_{k'})', \dots, \mathbf{b}(x_{m_i}; \mathbf{t}_{k'}'))',
\end{aligned}
\tag{8}$$

$\mathbf{x}_i = (x_{1_i}, x_{2_i}, \dots, x_{m_i})$ is the vector containing the x -values corresponding to subject i , n is the total number of subjects (curves) in the study, and m_i is the total number of observations for subject i .

The integral in (7) can not be calculated analytically. We thus explore two different approximations to $p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$. The first simply plugs in the maximum likelihood estimate of ψ and σ^2 . The second approximates (7) using a Laplace approximation.

The first approximation that we consider estimates $p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ with

$$p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \widehat{\psi}(\mathbf{t}_k, \mathbf{t}_{k'}), \widehat{\sigma}^2(\mathbf{t}_k, \mathbf{t}_{k'})),$$

where $\widehat{\psi}(\mathbf{t}_k, \mathbf{t}_{k'})$ and $\widehat{\sigma}^2(\mathbf{t}_k, \mathbf{t}_{k'})$ are the maximum likelihood estimators of ψ and σ^2 corresponding to the model with fixed effect knots at $\mathbf{t}_k$ and random effect knots at $\mathbf{t}_{k'}$. The dependence of these maximum likelihood estimators on $(\mathbf{t}_k, \mathbf{t}_{k'})$ will be suppressed, and the estimators will now simply be denoted as $\widehat{\psi}$ and $\widehat{\sigma}^2$. We refer to this likelihood approximation as the “plugged-in” approximation and denote it as $\hat{p}_{\text{Plugged In}}(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$. Although this approximation ignores the penalty from increasing the dimension of ψ , it is computationally inexpensive and is similar to the methods employed by Taplin (1993), Draper (1995), Raftery et al. (1996), and Dominici et al. (2002). They approximate marginal likelihoods by plugging in the maximum likelihood estimates of the nuisance parameters that can not be integrated out.

An alternative, but computationally more complex approach, is to integrate out ψ and σ^2 using an approximation method. We use a Laplace approximation to estimate $p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'})$. The approximate likelihood is given by

$$\widehat{p}_{\text{Laplace}}(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}) = p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \widehat{\psi}, \widehat{\sigma}^2) (2\pi)^{\frac{\dim(\widehat{\psi})+1}{2}} \left| \widehat{\text{Cov}}(\widehat{\psi}, \log(\widehat{\sigma}^2)) \right|^{\frac{1}{2}},$$

where

$$\widehat{\text{Cov}}(\widehat{\psi}, \log(\widehat{\sigma}^2)) = - \left(\frac{\partial^2 \log \{ p(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha, \psi, \sigma^2) p(\alpha|\mathbf{t}_k) p(\psi|\mathbf{t}_k, \mathbf{t}_{k'}) p(\sigma^2|\mathbf{t}_k, \mathbf{t}_{k'}) \}}{\partial(\psi, \log(\sigma^2))' \partial(\psi, \log(\sigma^2))} \right)^{-1} \Big|_{\alpha=\widehat{\alpha}, \psi=\widehat{\psi}, \sigma^2=\widehat{\sigma}^2},$$

and $\dim(\psi) = (k' + 3)(k' + 2)/2$. We chose the unconstrained parameterization for Σ_γ , ψ , and σ^2 , $\log(\sigma^2)$, to make this Laplace approximation more accurate. The optimal

parameterization is an area of future research. Calculating $\widehat{\text{Cov}}(\widehat{\psi}, \log(\widehat{\sigma}^2))$ makes the computational expense of the Laplace approximation greater than that of the “plugged-in” approximation; the matrix of second derivatives takes time to compute and can be numerically unstable when $\dim(\psi)$ is large and the total number of subjects, n , is small. We still expect, however, that the Laplace approximation will be more accurate than the “plugged-in” approximation. A major question of interest addressed in this paper is whether the Laplace’s gain in accuracy is worth its computational cost.

4 Sampling from $\widehat{p}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$

The two likelihood approximations induce the following approximations to the posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$. Let

$$\widehat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y) \propto \widehat{p}_{\text{Plugged In}}(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}) p(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$$

and

$$\widehat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y) \propto \widehat{p}_{\text{Laplace}}(y|k, \mathbf{t}_k, k', \mathbf{t}_{k'}) p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}).$$

To sample from these posterior approximations to $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ (which we denote as $\widehat{p}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$), we use reversible jump MCMC methods. Reversible jump MCMC methods (Green, 1995; DiMatteo et al., 2001) can be used when sampling from a distribution of a random variable θ and $\dim(\theta)$. In this particular case, the dimensions of the vectors $\mathbf{t}_k$ and $\mathbf{t}_{k'}$ vary with the values of k and k' .

Rather than sampling the fixed and random effect knots together in one MCMC iteration (which results in very low acceptance rates), we sample from the posterior $\widehat{p}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ using a Gibbs Sampling type algorithm. In the algorithm, one value of the pair $(k, \mathbf{t}_k)$ is sampled from $\widehat{p}(k, \mathbf{t}_k|k', \mathbf{t}_{k'}, y)$ using the Metropolis-Hastings algorithm. Conditioned on this sampled pair of $(k, \mathbf{t}_k)$, a pair $(k', \mathbf{t}_{k'})$ is sampled from $\widehat{p}(k', \mathbf{t}_{k'}|k, \mathbf{t}_k, y)$ using identical Metropolis-Hastings methods. This algorithm is given below, and will be referred to as Algorithm I.

The jump probabilities and details of this algorithm (including derivation of the acceptance weights) are given in the Appendix.

5 Sampling from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$

An alternative way of sampling from the posterior distribution of $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ is to sample from the joint posterior distribution which contains both the covariance parameters, $(\psi$ and $\sigma^2)$, and $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$. This posterior is written as $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$. Although sampling from this posterior is expected to be less efficient than sampling from the approximations to $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ (ψ and σ^2 have to be sampled in addition to $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$), the posterior sample will be exact since $p(y/k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2)$ can be calculated (see (8)). The resulting posterior sample of knots will also allow us to evaluate the accuracy of the two approximations proposed in Section 4. The accuracy of the approximations will be judged by comparing the distribution of $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ when sampling from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ and $\hat{p}_{\text{Plugged-In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ to the marginal distribution of $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ when sampling from the “true” posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$. The algorithm used to sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$ is given below, and will be referred to as Algorithm II.

The jump probabilities and details of this algorithm are given in the Appendix.

Note that Algorithm II is nearly identical to Algorithm I with the exception of Step 3. In Step 3 of Algorithm II, new values of ψ and σ^2 are proposed in addition to new values of k' and $\mathbf{t}_{k'}$. To be more specific, the candidates of ψ and σ^2 (conditioned on the candidates of k' and $\mathbf{t}_{k'}$) are proposed from a multivariate normal distribution with covariance equaling $\widehat{\text{Cov}}(\widehat{\psi}, \log(\widehat{\sigma}^2))$. Calculating this estimated covariance matrix in Step 3 makes this algorithm's computation time per iteration comparable to that of Algorithm I when $\hat{p}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y) = \hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$. Although the computational cost of these two algorithms are similar, we expect Algorithm II to mix slower and have a lower acceptance rate, as it proposes higher dimensional moves in Step 3.

6 A Theoretical Comparison of the Approximations and Algorithms

To sample from the posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$, three options have been proposed: (1) sample from $\hat{p}_{\text{Plg}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ using Algorithm I, (2) sample from $\hat{p}_{\text{LP}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ using Algorithm I, and (3) exactly sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$ using Algorithm II. We refer to options (1), (2), and (3) as PLG, LP, and EX, respectively. In this section, we theoretically justify why LP is preferred to PLG and EX.

While LP will take longer to implement than PLG (as calculating $\hat{p}_{\text{LP}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ will take longer than calculating $\hat{p}_{\text{Plg}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$), it is expected that PLG will result in unnecessarily large values of k' . We anticipate this positive bias in the number of random effect knots since the plugged-in approximation to $p(y/k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ does not average over the random effect covariance parameters, ψ ; it simply plugs in the maximum likelihood estimate of ψ . The tendency of PLG to estimate more random effect knots than LP becomes clear after close examination of the ratio

$$\frac{\widehat{p}(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}, \quad (9)$$

calculated in step 3(a) of Algorithm I. In PLG, this ratio is calculated as

$\frac{\widehat{p}_{\text{Plg}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{Plg}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}$, and in LP, this ratio is calculated as $\frac{\widehat{p}_{\text{LP}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{LP}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}$. In general, we expect

$$\frac{\widehat{p}_{\text{LP}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{LP}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)} \leq \frac{\widehat{p}_{\text{Plg}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{Plg}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)} \quad (10)$$

when $k'_{\text{new}} > k'_{\text{old}}$. If this inequality holds, then it does follow that PLG will over estimate the number of random effect knots. In Theorem 1, it is proven that when $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}})$ is the true set of knots, the inequality given in (10) asymptotically holds.

Theorem 1. *Let the true values of k , k' and $\mathbf{t}_{k'}$ be k_{new} , $\mathbf{t}_{k_{\text{new}}}$, k'_{old} , and $\mathbf{t}'_{k'_{\text{old}}}$, respectively. If the four conditions below are met, then as n (the total number of subjects) goes to ∞ ,*

$$\frac{\widehat{p}_{\text{LP}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{LP}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)} \leq \frac{\widehat{p}_{\text{Plg}}(k'_{\text{new}}, \mathbf{t}'_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{Plg}}(k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}.$$

1. $\mathbf{t}'_{k'_{\text{old}}} \subset \mathbf{t}'_{k'_{\text{new}}}$ with $k'_{\text{new}} > k'_{\text{old}}$.
2. m_i , the number of observations recorded for individual i , is random and the random variables $m_1, m_2, \dots, m_n$ are i.i.d. for all n .
3. $\sum_{i=1}^n (m_i - (k'_{\text{new}} + 2)) > k_{\text{new}} + 2$
4. $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ are i.i.d. vectors, where $\mathbf{x}_j = (\mathbf{x}_{j1}, \mathbf{x}_{j2}, \dots, \mathbf{x}_{jm_j})$ are the m_j points in time at which subject j is observed.

The proof of this theorem is given in the Appendix.

LP is not only preferred to PLG, but also to EX. The Markov chain resulting from LP converges to the stationary distribution more quickly than that corresponding to EX. This is the case because, in LP, transitions are only made in the space of the fixed and random effect knots. In EX, transitions are made in the space of the fixed knots, the random effect knots, and the covariance parameters. The acceptance probability of the chain corresponding to LP is thus expected to be larger than that corresponding to EX.

While the arguments given in this section suggest that LP is preferred to PLG and EX, it is still of interest to see how these options vary in practice. Of particular interest is to see the extent of the inequality given in (10). If the bias in the number of random effect knots fitted by PLG is only marginal, it might be the preferred option (since it is the easiest approximation to compute). To compare these three options, we conducted a simulation study. This is the subject of Section 7

7 Simulation studies

The accuracies of the posterior distributions constructed using the two likelihood approximations were determined using simulated data sets. Five data sets were simulated

from two models at two different sample sizes. Each model had the general form given in (5), but different values of k , $\mathbf{t}_k$, k' , $\mathbf{t}_{k'}$ and α . The specific values of these parameters are given in Table (1).

The population curve and five randomly sampled subject specific curves corresponding to each model are given in Figure 1. The simulation experiments were performed at two different values of n ($n = 25$ and $n = 40$, where n is the number of subjects). The subjects in all simulations each had 20 observations, 5 of which were randomly selected ($m_i = 5$ for all i). With such a small sample size per subject, it was not feasible to smooth observations separately for each individual as has been done in previous work (Behseta et al., 2005).

For each simulated data set, we sampled from $\hat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ and $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ using Algorithm I, and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ using Algorithm II. We refer to the posterior sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ as the “truth” or the “true” posterior sample. The number of iterations for each chain of $\hat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ and $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ was at least 10,000, and for $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ the chains were twice as long. All chains were run using R 2.5.0, and the average time it took to complete 1000 iterations at $n = 25$ was 18 minutes for $\hat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$, and 54 minutes for $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$. At $n = 40$ the average time it took to complete 1000 iterations was 30 minutes for $\hat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$, and 78 minutes for $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$. Convergence to a stationary distribution did not seem to be an issue with any of the chains run. The distributions of k and k' corresponding to each approximation and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ are shown in Figures (2) – (5). The captions of each figure also give the average acceptance rates for the plugged-in, Laplace, and “true” chains. The acceptance rate of the Laplace chains are, in each case, higher than those for the “true” or plugged-in chains. Figure 6 gives the posterior distributions of $\mathbf{t}_k$ and $\mathbf{t}_{k'}$ for data simulated from model 1 at $n = 25$. The accuracy of the Laplace and plugged-in as shown in Figure 6 is similar to the accuracy seen in data simulated from model 1 at $n = 25$ and from data simulated from model 2 at $n = 25$ and $n = 40$ (not shown).

The marginal distributions of k' show that the Laplace method is more accurate than the plugged-in method. In each case, the marginal distributions of k' for the Laplace method essentially matches the “true” marginal distributions of k' . The plugged-in estimator greatly overestimates the true number of random effect knots (see Figures 2 – 5). This is to be expected since both ψ and σ^2 were not averaged over (averaging over the parameters would penalize vectors of ψ with large dimensions). The marginal distributions of $\mathbf{t}_{k'}$ in Figure 6 also show that the Laplace method is more accurate than the plugged-in method. The posterior distributions of $\mathbf{t}_{k'}$ corresponding to $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$ are essentially the same. The posterior distribution of $\mathbf{t}_{k'}$ corresponding to $\hat{p}_{\text{Plugged-In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ is far from the true posterior distribution of $\mathbf{t}_{k'}$; the plugged-in method fits far too many random effect knots and thus can not correctly identify their true locations. Thus, in terms of accuracy, the Laplace approximation is definitely preferred to the plugged-in approximation.

The simulation results also suggest that sampling from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$ is preferred to sampling from the exact posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2 | y)$. The distribution of knots corresponding to both posteriors are very similar and the time per iteration is essentially the same. However, sampling from the exact posterior requires that the nuisance parameters ψ and σ^2 be sampled jointly with k' and $\mathbf{t}_{k'}$ in Step 3 of Algorithm II. The acceptance rate of Algorithm II is thus less than that of Algorithm I when sampling from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'} | y)$. The lower acceptance rate of Algorithm II also affected the mixing of the chain. Figure 7 gives the trace plots of $\mathbf{t}_{k'}$ for the Laplace and true chains (for iterations 2000–3000) on one (representative) simulated data set. From the plots in Figure 7, it is clear that Algorithm I

(applied to the Laplace approximation) explores the posterior distribution of $\mathbf{t}_{k'}$ more efficiently than Algorithm II does.

8 Application

To gain further insight into how the Laplace approximation compares to the “true” posterior, we considered a real data set and sampled from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$. The data set measures the protein content of milk collected from cows who were on a barley diet (Diggle, 1990). Twenty-five cows were in the experiment ($n = 25$), and milk was collected weekly from each cow for 19 weeks. Five observations from each cow were randomly selected and the reversible jump MCMC algorithms were run on the resulting data set. A plot of the data for all cows is given in Figure 8. Just as with the simulated data sets, 10,000 iterations of Algorithm I were used to sample from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ and 20,000 iterations of Algorithm II were used to sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$. The posterior distributions of k and k' corresponding to the $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ and $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$ are given in Figure 9. The posterior distributions of $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ corresponding to $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ once again matches the marginal distribution of $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$ when exactly sampling from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$.

In addition to examining the two posterior distributions of knots, we also examine the two corresponding posterior distributions of one randomly selected cow's curve. Recall that the curve of one cow is the sum of the fixed curve and that cow's random effect curve. To sample one cow's curve from the chain corresponding to $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$, values of ψ and σ^2 were sampled using an independence sampler. Conditioned on these sampled values, a fixed effect curve was sampled, and conditioned on this sampled fixed effect curve, a random effect curve was sampled. The algorithm used to do this requires a slight addition to Algorithm I. This addition is noted below: This addition to Algorithm I slows it down only marginally.

To sample random effects from the chain corresponding to $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$, steps (2) and (3) given above were just added to Algorithm II. In other words, a fixed and random effect curve were drawn for every random effects covariance matrix drawn. Figure 10 shows the pointwise median curve for the selected cow drawn for each chain, along with pointwise 95% credible intervals. The estimated cow-specific curves and credible intervals for both methods are similar.

9 Discussion

We have proposed a flexible approach to model multiple curves using free-knot splines. To identify the number and location of fixed and random effect knots in this flexible mixed-effects regression spline, sampling from $\hat{p}_{\text{Laplace}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ is highly preferred to sampling from $\hat{p}_{\text{Plugged In}}(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$ and marginally preferred to sampling from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$. The posterior sample corresponding to the Laplace approximation is more accurate than that of the plugged-in approximation. The disparity between the Laplace approximation and the plugged-in approximation is most pronounced in the different number of random effect knots fit to the simulated data sets. The Laplace method clearly seems to penalize models with unnecessarily large values of k' more appropriately than the plugged-in method. In addition, as measured by our simulations, the posterior based on the Laplace approximation is much more efficient to sample from (and just as accurate) as the exact posterior sample from $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi, \sigma^2|y)$.

We are currently investigating other approaches to the knot selection problem. Rather than sampling from the posterior $p(k, \mathbf{t}_k, k', \mathbf{t}_{k'}|y)$, we set $\mathbf{t}_k = \mathbf{t}_{k'}$ and then try to reduce the number of principal component curves associated with the random effects by sampling from

the posterior $p(k, \mathbf{t}_k, r/y)$ where r = the number of random effect principal components retained (Botts, 2005). This is similar in spirit to the methods employed by Shi et al. (1996), and James et al. (2000).

Behseta et al. (2005) propose an alternative hierarchical approach that, at the first level, fits curves separately for each subject, allowing each subject to have her own set of knots. However, their methods are only appropriate for large within-subject sample sizes.

References

- deBoor, C. A Practical Guide to Splines. New York City: Springer-Verlag; 1978.
- Behseta S, Kass RE, Wallstrom GL. Hierarchical models for assessing variability among functions. *Biometrika* 2005;92:419–434.
- Bigelow JL, Dunson DB. Bayesian adaptive regression splines for hierarchical data. *Biometrics*. 2007 in press.
- Botts, C. Ph.D. Thesis. Iowa State University; 2005. Bayesian Methods in Single and Multiple Curve Fitting.
- Brumback BA, Rice J. Smoothing spline models for the analysis of nested and crossed samples of curves. *Journal of the American Statistical Association* 1998;93:961–976.
- Crainiceanu C, Ruppert D, Wand MP. Bayesian analysis for penalized spline regression using WinBUGS. *Journal of Statistical Software* 2005;14 Article 14.
- Demidenko, E. Mixed Models: Theory and Applications. Hoboken, New Jersey: John Wiley & Sons, Inc.; 2004.
- Densin DGT, Mallick BK, Smith AFM. Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Series B* 1998;60:333–350.
- Diggle, PJ. Time Series: A Biostatistical Introduction. Oxford: Oxford University Press; 1990.
- DiMatteo I, Genovese CR, Kass RE. Bayesian curve-fitting with free-knot splines. *Biometrika* 2001;88:1055–1071.
- Dominici F, Daniels M, Zeger SL, Samet JM. Air pollution and mortality: Estimating regional and national dose-response relationships. *Journal of the American Statistical Association* 2002;97:100–111.
- Friedman JH. Multivariate adaptive regression splines. *The Annals of Statistics* 1991;19:1–141.
- Gelfand AE, Dey DK. Bayesian model choice: Asymptotics and exact calculations. *Journal of the Royal Statistical Society, Series B* 1994;56:501–514.
- Green PJ. Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination. *Biometrika* 1995;82:711–732.
- Halpern EF. Bayesian spline regression when the number of knots is unknown. *Journal of the Royal Statistical Society, Series B* 1973;35:347–360.
- Hastie, TJ.; Tibshirani, RJ. Generalized Additive Models. Boca Raton, Florida: Chapman & Hall; 1990.
- Hooper PM. Iterative weighted least squares estimation in heteroscedastic linear models. *Journal of the American Statistical Association* 1993;88:179–184.
- James GM, Hastie TJ, Sugar CA. Principal component models for sparse functional data. *Biometrika* 2000;87:587–602.
- Kass RE, Wasserman L. A reference Bayesian test for nested hypotheses and its relationship to the Schwarz Criterion. *Journal of the American Statistical Association* 1994;90:928–934.
- Lindstrom MJ. Penalized estimation of free-knot splines. *Journal of Computational and Graphical Statistics* 1999;8:333–352.
- Paciorek CJ. Misinformation in the conjugate prior for the linear model with implications for free-knot spline modeling. *Bayesian Analysis* 2006;1:375–383. [PubMed: 18185854]
- Pauler DK. The Schwarz criterion and related methods for normal linear models. *Biometrika* 1998;85:13–27.

- Pourhamadi M. Joint mean-covariance covariance models with applications to longitudinal data, I: Unconstrained parameterization. *Biometrika* 1999;86:677–690.
- Raftery, AE.; Madigan, D.; Volinsky, CT. Accounting for model uncertainty in survival analysis improves predictive performance (with discussion). In: Bernardo, J.; Berger, J.; Dawid, A.; Smith, A., editors. *Bayesian Statistics 5*. Oxford University Press; 1996. p. 323–349.
- Rice J, Wu C. Nonparametric mixed effects models for unequally sampled noisy curves. *Biometrics* 2001;57:253–259. [PubMed: 11252607]
- Shi M, Weiss RE, Taylor JMG. An analysis of pediatric CD4 counts for AIDS patients using flexible random curves. *Applied Statistics* 1996;45:151–163.
- Stone CJ, Hansen M, Kooperberg C, Truong YK. Polynomial splines and their tensor products in extended linear modeling. *The Annals of Statistics* 1997;25:1371–1470.
- Taplin RH. Robust likelihood calculation for time series. *Journal of the Royal Statistics Society, Series B* 1993;55:829–836.
- Zhou S, Shen X. Spatially adaptive regression splines and accurate knot selection schemes. *Journal of the American Statistical Association* 2001;96:247–259.

Appendix

Exact formulas for I_ψ , I_α , and $I_{\log(\sigma^2)}$

The $(j, k)^{th}$ element in the matrix I_ψ , can be calculated as

$$\begin{aligned}
 [I_\psi]_{j,k} &= -\frac{1}{2} \sum_{i=1}^n \text{tr} \left\{ -\mathbf{L}_i^{-1} \mathbf{G}_{j,i} (\mathbf{L}_i^{-1})' \mathbf{G}_{k,i} + \mathbf{L}_i^{-1} \mathbf{H}_{(j,k),i} \right\} + \frac{1}{2} \sum_{i=1}^n (y_i - \mathbf{G}_{F_i} \hat{\alpha})' \left(-\mathbf{L}_i^{-1} (\mathbf{G}_{j,i} \mathbf{L}_i^{-1} \mathbf{G}_{k,i}' - \mathbf{H}_{(j,k),i} + \mathbf{G}_{k,i} \mathbf{L}_i^{-1} \mathbf{G}_{j,i}') (\mathbf{L}_i^{-1})' \right) (y_i - \mathbf{G}_{F_i} \hat{\alpha}) \\
 \mathbf{L}_i &= \mathbf{B}_{R_i} \Sigma_\gamma(\hat{\psi}) \mathbf{B}_{R_i}' + \hat{\sigma}^2 \mathbf{I}, \\
 \mathbf{G}_{j,i} &= \mathbf{B}_{R_i} \mathbf{W}_j(\hat{\psi}) \mathbf{B}_{R_i}', \\
 \mathbf{W}_j(\hat{\psi}) &= \frac{\partial(TDT')}{\partial \psi_j} \bigg|_{\psi=\hat{\psi}}, \text{ and} \\
 \mathbf{H}_{(j,k),i} &= \mathbf{B}_{R_i} \left(\frac{\partial^2(TDT')}{\partial \psi_j \partial \psi_k} \bigg|_{\psi=\hat{\psi}} \right) \mathbf{B}_{R_i}' \\
 I_{\log(\sigma^2),i} &= -\frac{1}{2} \text{tr} \left\{ -(\hat{\sigma}^2)^2 \mathbf{L}_i^{-1} (\mathbf{L}_i^{-1})' + \sigma^2 \mathbf{L}_i^{-1} \right\} + \frac{1}{2} (y_i - \mathbf{B}_{F_i} \hat{\alpha})' \left\{ -\mathbf{L}_i^{-1} \left((\hat{\sigma}^2)^2 \mathbf{L}_i^{-1} - \hat{\sigma}^2 \mathbf{I} + (\hat{\sigma}^2)^2 \mathbf{L}_i^{-1} \right) (\mathbf{L}_i^{-1})' \right\} (y_i - \mathbf{B}_{F_i} \hat{\alpha}) \\
 I_{\alpha,i} &= -\mathbf{B}_{F_i}' \left(\mathbf{B}_{R_i} \Sigma_\gamma(\hat{\psi}) \mathbf{B}_{R_i}' + \hat{\sigma}^2 \mathbf{I}_{m_i} \right)^{-1} \mathbf{B}_{F_i}
 \end{aligned}$$

RJMCMC details for Algorithm I

This section gives the details on Algorithm I. The details show how a move is made from a set of old knots to a new set of knots. Note that the constants which appear in this algorithm (.4 and .2) were selected to optimize its mixing and acceptance probability. The old set of

knots will be denoted as $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}, k'_{\text{old}}, \mathbf{t}'_{k'_{\text{old}}})$.

The first step of the algorithm involves a step in the space of the fixed effect knots, Ω^k . Note that Ω^k is the sum of spaces $\Omega_1, \Omega_2, \dots$ where $\Omega_j = [a, b]^j$. Therefore Ω^k can be written as $\Omega^k = \bigoplus_{j=1}^{\infty} \Omega_j$. The first decision to be made is whether to give birth, relocate, or kill a fixed knot. These probabilities are labeled as b_F , r_F , and d_F , respectively, and are calculated as

$$b_F(k_{\text{old}}) = .4 \times \min \left(1, \frac{p_k(k_{\text{old}}+1)}{p_k(k_{\text{old}})} \right), \quad d_F(k_{\text{old}}) = \begin{cases} 0 & k_{\text{old}} = k'_{\text{old}} \\ .4 \times \min \left(1, \frac{p_k(k_{\text{old}}-1)}{p_k(k_{\text{old}})} \right) & \text{o.w.} \end{cases},$$

and $r_F(k_{old}) = 1 - b_F(k_{old}) - d_F(k_{old})$, where $p_k(k_{old})$ is the prior distribution assigned to k evaluated at k_{old} .

If birth is chosen, a fixed effect knot, $t_k^{selected}$, is selected at random and a new knot is given birth to. The new knot, t_k^{new} , is sampled from the distribution $t_k^{new} \sim TN_a^b(t_k^{selected}, .2)$ where TN_a^b is the normal distribution truncated at a and b with mean at $t_k^{selected}$ and variance $.2$. If death is selected, then one of the fixed effect knots is uniformly selected and killed. When relocation is chosen, a fixed effect knot, $t_k^{relocated}$, is uniformly selected and then relocated to a position $t_k^{new loc}$ where $t_k^{new loc} \sim TN_a^b(t_k^{relocated}, .2)$. The jump probabilities corresponding to each move is given in Table 2. Note: $h(y/x) = \phi(y; x, .2) (\Phi(b; x, .2) - \Phi(a; x, .2))^{-1}$ where $\phi(\cdot; \mu, \tau^2)$ is the normal density with mean μ and variance τ^2 , and $\Phi(\cdot; \mu, \tau^2)$ is the cumulative distribution function of a normal density with mean μ and variance τ^2 .

A move must then be made with the random effect knots. This is the second step in the algorithm. As with the fixed effect knots, one of three moves can be made in Ω^k , the space of random effect knots: birth, death, or relocation. The probabilities associated with each are denoted as b_R , d_R , and r_R , respectively. They are calculated as

$$b_R(k'_{old}) = \begin{cases} 0 & k'_{old} = k_{old} \\ .4 \times \min\left(1, \frac{p_{k'}(k'_{old}+1)}{p_{k'}(k'_{old})}\right) & o.w. \end{cases},$$

$$d_R(k'_{old}) = \begin{cases} 0 & k'_{old} = 1 \\ .4 \times \min\left(1, \frac{p_{k'}(k'_{old}-1)}{p_{k'}(k'_{old})}\right) & o.w. \end{cases},$$

and $r_R(k'_{old}) = 1 - b_R(k'_{old}) - d_R(k'_{old})$, where $p_{k'}(k'_{old})$ is the prior distribution of k' evaluated at k'_{old} . If death is selected, then one of the random effect knots is uniformly selected and then killed. If birth is selected, then one of the random effect knots is uniformly selected (call this $t_{k'}^{selected}$), and a knot is added to the set of random effect knots (call this knot $t_{k'}^{new}$). Note that $t_{k'}^{new} \sim TN_a^b(t_{k'}^{selected}, .2)$. If relocation is selected, then one of the random effect knots is uniformly selected, $t_{k'}^{relocated}$, and moved to another location (call this $t_{k'}^{new loc}$). Again, $t_{k'}^{new loc} \sim TN_a^b(t_{k'}^{relocated}, .2)$. The jump probabilities associated with these moves are given in Table 3.

Derivation of Metropolis-Hastings Acceptance Probabilities to Satisfy Detailed Balance of Algorithm I

We show that the acceptance probabilities given in Section 4 are such that detailed balance is satisfied. In this case, it needs to be shown that

$$P_{Trans} \left(k_{old}, t_{k_{old}} k'_{old}, t_{k'_{old}} \rightarrow k_{new}, t_{k_{new}}, t_{k'_{new}} \right) P \left(k_{old}, t_{k_{old}}, k'_{old}, t_{k'_{old}} | y \right) \\ = P_{Trans} \left(k_{new}, t_{k_{new}}, k'_{new}, t_{k'_{new}} \rightarrow k_{old}, t_{k_{old}}, k'_{old}, t_{k'_{old}} \right) P \left(k_{new}, t_{k_{new}}, k'_{new}, t_{k'_{new}} | y \right)$$

(Note: throughout the rest of this proof, $t_{k_{new}}$ will be denoted as t_{k_n} , k_{old} will be denoted as k_o , and so on). Assume that both a fixed and random knot are being added. In this case,

$$\begin{aligned}
& P_{\text{Trans}}(k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o} \rightarrow k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n}) p(k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o} | y) \\
&= b_F(k_o) \frac{1}{k_o} \sum_{x \in \mathbf{t}_{k_o}} h(\mathbf{t}_k^{\text{new}} | x) \times \min \left(1, \frac{P_{\text{Jump}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) p(k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o}, y)}{P_{\text{Jump}}(k_o, \mathbf{t}_{k_o} \rightarrow k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o}) p(k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}, y)} \right) \\
&\quad \times b_R(k'_o) \frac{1}{k'_o} \sum_{x \in \mathbf{t}_{k'_o}} h(\mathbf{t}_{k'}^{\text{new}} | x) \times \min \left(1, \frac{P_{\text{Trans}}(k_o, \mathbf{t}_{k_o} \rightarrow k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o})}{P_{\text{Jump}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) p(k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n}, y)} \right) \\
&\quad \times p(k'_o, \mathbf{t}_{k'_o} \rightarrow k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n}) \\
&\quad \times p(k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o} | y).
\end{aligned}$$

Assume that the ratios within each min argument are < 1 (if this is not the case, the reciprocal of the ratios will be less than 1, and the proof can be done in the other direction), and recall that

1. $b_F(k_o) \frac{1}{k_o} \sum_{x \in \mathbf{t}_{k_o}} h(\mathbf{t}_k^{\text{new}} | x) = P_{\text{Jump}}(k_o, \mathbf{t}_{k_o} \rightarrow k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o})$
2. $b_R(k'_o) \frac{1}{k'_o} \sum_{x \in \mathbf{t}_{k'_o}} h(\mathbf{t}_{k'}^{\text{new}} | x) = P_{\text{Jump}}(k'_o, \mathbf{t}_{k'_o} \rightarrow k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n})$
3. $p(k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o}, y) / p(k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}, y) = \frac{p(y | k_n, \mathbf{t}_{k_n}, k'_o, \mathbf{t}_{k'_o}) p(k_n) p(\mathbf{t}_{k_n} | k_n)}{p(y | k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o}) p(k_o) p(\mathbf{t}_{k_o} | k_o)}$
4. $p(k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n}, y) / p(k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}, y) = \frac{p(y | k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n}) p(k'_n) p(\mathbf{t}_{k'_n} | k'_n)}{p(y | k'_o, \mathbf{t}_{k'_o}, k_n, \mathbf{t}_{k_n}) p(k'_o) p(\mathbf{t}_{k'_o} | k'_o)}$

With these four observations it is clear that the above expression reduces to

$$\begin{aligned}
& P_{\text{Jump}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) \frac{p(y | k_n, \mathbf{t}_{k_n}, k'_o, \mathbf{t}_{k'_o}) p(k_n)}{p(y | k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o}) p(k_o)} \frac{p(\mathbf{t}_{k_n} | k_n)}{p(\mathbf{t}_{k_o} | k_o)} P_{\text{Jump}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) \\
& \times \frac{p(y | k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n}) p(k'_n)}{p(y | k_n, \mathbf{t}_{k_n}, k'_o, \mathbf{t}_{k'_o}) p(k'_o)} \frac{p(\mathbf{t}_{k'_n} | k'_n)}{p(\mathbf{t}_{k'_o} | k'_o)} \frac{p(k_o) p(k'_o) p(\mathbf{t}_{k_o} | k_o) p(\mathbf{t}_{k'_o} | k'_o)}{p(y)} \\
& = P_{\text{Jump}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) P_{\text{Jump}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) \\
& \times \frac{p(y | k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n}) p(k_n)}{p(y | k_n, \mathbf{t}_{k_n}, k'_o, \mathbf{t}_{k'_o}) p(k'_o)} \frac{p(k_n)}{p(k'_n)} \frac{p(\mathbf{t}_{k_n} | k_n) p(\mathbf{t}_{k'_n} | k'_n)}{p(y)} \\
& = P_{\text{Jump}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) \\
& \quad = d_F(k_n)^{\frac{1}{k_n+1}} \\
& \quad \times \min \left(1, \frac{P_{\text{Jump}}(k_o, \mathbf{t}_{k_o} \rightarrow k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o}) p(k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}, y)}{P_{\text{Jump}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) p(k_n, \mathbf{t}_{k_n} | k'_o, \mathbf{t}_{k'_o}, y)} \right) \\
& \quad = 1 \\
& \quad \times P_{\text{Jump}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) \\
& \quad = d_R(k'_n)^{\frac{1}{k'_n+1}} \\
& \quad \times \min \left(1, \frac{P_{\text{Jump}}(k'_o, \mathbf{t}_{k'_o} \rightarrow k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n}) p(k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}, y)}{P_{\text{Jump}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) p(k'_n, \mathbf{t}_{k'_n} | k_n, \mathbf{t}_{k_n}, y)} \right) \times p(k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n} | y) \\
& \quad = 1 \\
& = P_{\text{Trans}}(k_n, \mathbf{t}_{k_n} \rightarrow k_o, \mathbf{t}_{k_o} | k'_o, \mathbf{t}_{k'_o}) P_{\text{Trans}}(k'_n, \mathbf{t}_{k'_n} \rightarrow k'_o, \mathbf{t}_{k'_o} | k_n, \mathbf{t}_{k_n}) p(k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n} | y) \\
& = P_{\text{Trans}}(k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n} \rightarrow k_o, \mathbf{t}_{k_o}, k'_o, \mathbf{t}_{k'_o}) p(k_n, \mathbf{t}_{k_n}, k'_n, \mathbf{t}_{k'_n} | y).
\end{aligned}$$

This argument shows detailed balance holds for the given acceptance ratios when a fixed and random effect knot are added. Similar arguments can be made to prove detailed balance when a fixed and/or random effect knot is deleted and/or relocated.

RJMCMC details for Algorithm II

As mentioned in the main text of this paper, Algorithm II is only a slight variation of Algorithm I. In Algorithm II, candidate values of $(k, \mathbf{t}_k)$ are sampled just as they are in Algorithm I. Conditioned on each sampled pair of knots, candidate values of $k', \mathbf{t}_{k'}, \psi$ and σ^2 are sampled. These candidate values are denoted ψ_{new} and σ_{new}^2 and they are sampled from the multivariate normal distribution

$$\begin{pmatrix} \psi_{\text{new}} \\ \log(\sigma_{\text{new}}^2) \end{pmatrix} \sim \text{MVN} \left(\begin{pmatrix} \hat{\psi} \\ \log(\hat{\sigma}^2) \end{pmatrix}, \widehat{\text{Cov}}(\hat{\psi}, \log(\hat{\sigma}^2)) \right).$$

Note that $\hat{\psi}$ and $\hat{\sigma}^2$ are the maximum likelihood estimates of ψ and σ^2 corresponding to the model with fixed effect knots at $\mathbf{t}_{k_{\text{old}}}$ and random effect knots at $\mathbf{t}_{k'_{\text{new}}}$.

The jump probabilities corresponding to step 2 of Algorithm II are the same as those in step 2 of Algorithm I. The jump probabilities corresponding to step 3 of Algorithm II are given in Table 4. Note: $\text{MVN}_3(z, w)$ is the multivariate density, with mean $(\hat{\psi}, \log(\hat{\sigma}^2))$ and covariance $\widehat{\text{Cov}}(\hat{\psi}, \log(\hat{\sigma}^2))$ evaluated at the vector (z, w) . In this case, $\hat{\psi}$ and $\hat{\sigma}^2$ are the maximum likelihood estimates calculated under the model with fixed effect knots at $\mathbf{t}_{k_{\text{old}}}$ and random effect knots at $\mathbf{t}_{k'_{\text{new}}}$.

It can be shown, using methods similar to those given before, that these jump probabilities (along with the acceptance probabilities of Algorithm II given in Section 5) guarantee detailed balance.

Proof of Theorem 1

We first prove that the maximum likelihood estimators of α , Σ_γ , and σ^2 are consistent with respect to n , the number of subjects in the study. To be more specific, we prove that $\hat{\alpha}$, $\hat{\psi}$ and $\hat{\sigma}^2$ converge to their true values as n goes to infinity. Proving the consistency of these estimators is not straight-forward since the observed data for all n subjects, $y_1, y_2, \dots, y_n$, are not i.i.d. (recall that they are independent yet not identically distributed since

$$y_i \sim MVN(\mathbf{B}_{F_i}\alpha, \mathbf{B}_{R_i}\Sigma_\gamma(\psi)\mathbf{B}_{R_i}' + \sigma^2\mathbf{I}_{m_i}).$$

Hooper (1993) and Demidenko (2004) show, however, that if one considers the design matrices of the i^{th} individual ($\mathbf{B}_{F_i}, \mathbf{B}_{R_i}$) and the number of observations for the i^{th} individual (m_i) as random, independent of ε_i and γ_i , and part of the observed data, then the i.i.d. assumption of the data is satisfied. For the problem presented in this paper, such considerations make sense. The design matrices corresponding to individual i , $\mathbf{B}_{F_i}$ and $\mathbf{B}_{R_i}$, can be thought of as random since the points in time at which the m_i observations occur, $\mathbf{x}_i$, vary. If the number of observations recorded for all individuals, $m_1, m_2, \dots, m_n$, are i.i.d. (as we assume in Condition (2)), and the times at which these observations occur, $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, are i.i.d. (as we assume in Condition (4)), then the sets $\{y_i, \mathbf{B}_{F_i}, \mathbf{B}_{R_i}, m_i\}_{i=1}^n$ are i.i.d. It follows that the maximum likelihood estimators of α , ψ , and σ^2 converge to their true values as n , the number of subjects, approaches ∞ .

We now arrive at the result in Theorem 1 by examining the behavior of log

$$\left(\frac{\widehat{p}_{LP}(k'_{old}, \mathbf{t}'_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y)}{\widehat{p}_{LP}(k'_{new}, \mathbf{t}'_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y)} \right) \text{ as } n \rightarrow \infty. \text{ Recall that } k'_{old} < k'_{new} \text{ and } \mathbf{t}'_{k'_{old}} \subset \mathbf{t}'_{k'_{new}}, \text{ making the}$$

model with knots at $(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}})$ nested within the model with knots at $(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}})$.

$$\begin{aligned} & \log \left(\frac{\widehat{p}_{LP}(k'_{old}, \mathbf{t}'_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y)}{\widehat{p}_{LP}(k'_{new}, \mathbf{t}'_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y)} \right) \\ &= \log \left(\frac{\widehat{p}_{LP}(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}'_{k'_{old}}) p(k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}'_{k'_{old}})}{\widehat{p}_{LP}(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}'_{k'_{new}}) p(k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}'_{k'_{new}})} \right) \\ &= \log \left(\frac{\widehat{p}(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}'_{k'_{old}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}}), \widehat{\sigma}^2(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}}))}{\widehat{p}(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}'_{k'_{new}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}}), \widehat{\sigma}^2(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}}))} \right) \\ & \quad + \log \left(\frac{p(\widehat{\psi} | \mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}}) p(\widehat{\sigma}^2 | \mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}})}{p(\widehat{\psi} | \mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}}) p(\widehat{\sigma}^2 | \mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}})} \frac{|\text{Cov}(\widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{old}})))|^{1/2}}{|\text{Cov}(\widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{new}}, \mathbf{t}'_{k'_{new}})))|^{1/2}} \right)^{\frac{d_1 - d_2}{2}} \\ & \quad + \log \left(\frac{p(k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}'_{k'_{old}})}{p(k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}'_{k'_{new}})} \right), \end{aligned}$$

$$\begin{aligned}
d_1 &= \dim \left(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}) \right) \\
&= \frac{(k'_{\text{old}} + 2)(k'_{\text{new}} + 3)}{2}, \\
d_2 &= \dim \left(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}) \right) \\
&= \frac{(k'_{\text{new}} + 2)(k'_{\text{new}} + 3)}{2}, \\
d_2 &> d_1 \text{ and } \widehat{\text{Cov}} \left(\widehat{\psi}, \log(\widehat{\sigma}^2) \right)
\end{aligned}$$

$$\text{where} \quad = -\frac{1}{n} H \left(\widehat{\psi}, \log(\widehat{\sigma}^2) \right)^{-1}, \text{ with}$$

$$H \left(\widehat{\psi}, \log(\widehat{\sigma}^2) \right) = \left(\frac{\partial^2 \log \{ p(y|k, \mathbf{t}_k, k', \mathbf{t}_k, \alpha, \psi, \sigma^2) p(\alpha|\mathbf{t}_k) p(\psi|\mathbf{t}_k, \mathbf{t}_{k'}) p(\sigma^2|\mathbf{t}_k, \mathbf{t}_{k'}) / n \}}{\partial(\psi, \log(\sigma^2))' \partial(\psi, \log(\sigma^2))} \right)^{-1} \Big|_{\alpha=\widehat{\alpha}, \psi=\widehat{\psi}, \sigma^2=\widehat{\sigma}^2}.$$

From this, we get that

$$\begin{aligned}
&\log \left(\frac{\widehat{p}_{\text{LP}}(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{\widehat{p}_{\text{LP}}(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}}, k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)} \right) \\
&= \log \left(\frac{\widehat{p}(y|k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, \widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}), \widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}))}{\widehat{p}(y|k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}}, \widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}), \widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}))} \right) \\
&\quad + \log \left(\frac{p(\widehat{\psi}|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}) p(\widehat{\sigma}^2|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}})}{p(\widehat{\psi}|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}) p(\widehat{\sigma}^2|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}})} \frac{\left| -H(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}})) \right|}{\left| -H(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}})) \right|} \right)^{\frac{1}{2}} \left(\frac{n}{2\pi} \right)^{\frac{d_2 - d_1}{2}} \right) \\
&\quad + \log \left(\frac{p(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})}{p(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}})} \right) \\
&= \log \left(\frac{\widehat{p}(y|k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, \widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}), \widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}))}{\widehat{p}(y|k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}}, \widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}), \widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}))} \right) + K_n + \frac{d_2 - d_1}{2} \log(n) \\
&\quad + \log \left(\frac{p(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})}{p(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}})} \right),
\end{aligned}$$

where

$$K_n = \log \left(\frac{p(\widehat{\psi}|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}) p(\widehat{\sigma}^2|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}})}{p(\widehat{\psi}|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}) p(\widehat{\sigma}^2|\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}})} \frac{\left| -H(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{old}}})) \right|}{\left| -H(\widehat{\psi}(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}}), \log(\widehat{\sigma}^2(\mathbf{t}_{k_{\text{new}}}, \mathbf{t}_{k'_{\text{new}}})) \right|} \right)^{\frac{1}{2}} (2\pi)^{\frac{d_1 - d_2}{2}} \right).$$

Gelfand and Dey (1994) show that if the models are nested (which is guaranteed by Condition (1)) and if the maximum likelihood estimators exist, then $K_n = O(1)$. Note that existence of the maximum likelihood estimators follows from Condition (3). Demidenko (2004) shows that if $\sum_{i=1}^n (m_i - (k'_{\text{new}} + 2)) > k_{\text{new}} + 2$, then the maximum likelihood estimators

of α , ψ , and σ^2 exist. Since K_n is bounded, for sufficiently large values of n , K_n becomes negligible compared to $\frac{d_2 - d_1}{2} \log(n)$, and we can say that

$$\begin{aligned}
 & \log \left(\frac{\widehat{p}_{LP} \left(k'_{old}, \mathbf{t}_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)}{\widehat{p}_{LP} \left(k'_{new}, \mathbf{t}_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)} \right) \\
 & \approx \log \left(\frac{\widehat{p} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{old}}), \widehat{\sigma}^2 \left(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{old}} \right) \right)}{\widehat{p} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{new}}), \widehat{\sigma}^2 \left(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{new}} \right) \right)} \right) + \frac{d_2 - d_1}{2} \log(n) \\
 & \quad + \log \left(\frac{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}} \right)}{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}} \right)} \right) \\
 & \geq \log \left(\frac{\widehat{p} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{old}}), \widehat{\sigma}^2 \left(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{old}} \right) \right)}{\widehat{p} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}}, \widehat{\psi}(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{new}}), \widehat{\sigma}^2 \left(\mathbf{t}_{k_{new}}, \mathbf{t}_{k'_{new}} \right) \right)} \right) + \log \left(\frac{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}} \right)}{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}} \right)} \right) \\
 & = \log \left(\frac{\widehat{p}_{PLg} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}} \right)}{\widehat{p}_{PLg} \left(y | k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}} \right)} \frac{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{old}, \mathbf{t}_{k'_{old}} \right)}{p \left(k_{new}, \mathbf{t}_{k_{new}}, k'_{new}, \mathbf{t}_{k'_{new}} \right)} \right) \\
 & = \log \left(\frac{\widehat{p}_{PLg} \left(k'_{old}, \mathbf{t}_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)}{\widehat{p}_{PLg} \left(k'_{new}, \mathbf{t}_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)} \right),
 \end{aligned}$$

and it follows that

$$\frac{\widehat{p}_{LP} \left(k'_{new}, \mathbf{t}_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)}{\widehat{p}_{LP} \left(k'_{old}, \mathbf{t}_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)} \leq \frac{\widehat{p}_{PLg} \left(k'_{new}, \mathbf{t}_{k'_{new}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)}{\widehat{p}_{PLg} \left(k'_{old}, \mathbf{t}_{k'_{old}} | k_{new}, \mathbf{t}_{k_{new}}, y \right)}.$$

Model 1

Model 2

Figure 1.

Plots of Models 1–2. (—): $f(x) + G_i(x)$, (—): $f(x)$. The location of the fixed effect knots are denoted with a • and the location of the random effect by a Δ .

Figure 2.
 Marginal Distributions of k and k' for Model 1 at $n = 25$. Average acceptance rates: Laplace - 72%, Plugged-In - 44%, Truth - 42%.

Figure 3.
Marginal Distributions of k and k' for Model 2 at $n = 25$. Average acceptance rates: Laplace - 65%, Plugged-In - 39%, Truth - 44%.

Figure 4.
Marginal Distributions of k and k' for Model 1 at $n = 40$. Average acceptance rates: Laplace - 68%, Plugged-In - 46%, Truth - 39%

Figure 5.
Marginal Distributions of k and k' for Model 2 at $n = 40$. Average acceptance rates: Laplace - 59%, Plugged-In - 51%, Truth - 38%.

Figure 6.
Marginal Distributions of t_k and $t_{k'}$ for Model 1 at $n = 25$.

Figure 7.

Trace plots of $\mathbf{t}_{k'}$ for the Laplace and true chains for one simulated data set (for iterations 2000 – 3000). Note: The figures trace the positions of all elements in $\mathbf{t}_{k'}$, not just the first. Thus, when the trace splits into two lines, a random effect knot has been added.

Figure 8.
Protein content of milk measured in 25 cows on Barley diet

Figure 9.
Marginal distributions of k and k'

Figure 10.
Median and 95% pointwise credible intervals for cow-specific curve. (—): Laplace.
(.....): Truth.

Table 1

True Models Simulated From

	Model 1	Model 2
$\mathbf{t}_k$	(.35, 2.3, 2.4, 3)	(.2, .58, 1, 1.7, 2.8)
$\mathbf{t}_{k'}$	(2.4)	(.8, 2.58)
α	(5, -5.5, 1.4, -6.2, 1, 8)	(5, 10, 4, 6, 14.2, 10, 7)
Σ_γ	$5 \times \mathbf{I}_{3 \times 3}$	$5 \times \mathbf{I}_{4 \times 4}$
σ^2	.5	.5

Table 2

Jump Probabilities in Ω^k

Move in Ω^k	$P_{\text{Jump}}\left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \rightarrow k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \mid k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}\right)$
Birth	$b_F(k_{\text{old}}) \frac{1}{k_{\text{old}}} \sum_{x \in t_{k_{\text{old}}}} h(t_k^{\text{new}} \mid x)$
Relocation	$r_F(k_{\text{old}}) \frac{1}{k_{\text{old}}} h(t_k^{\text{new loc}} \mid t_k^{\text{relocated}})$
Death	$d_F(k_{\text{old}}) \frac{1}{k_{\text{old}}}$

Table 3

Jump Probabilities in $\Omega^{k'}$

Move in $\Omega^{k'}$	$P_{\text{Jump}}\left(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} \rightarrow k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} \mid k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}\right)$
Birth	$b_R(k'_{\text{old}}) \frac{1}{k'_{\text{old}}} \sum_{x \in \mathcal{I}_{k'_{\text{old}}}} h(t_{k'}^{\text{new}} \mid x)$
Relocation	$r_R(k'_{\text{old}}) \frac{1}{k'_{\text{old}}} h(t_{k'}^{\text{new loc}} \mid t_{k'}^{\text{selected}})$
Death	$d_R(k'_{\text{old}}) \frac{1}{k'_{\text{old}}}$

Table 4

Jump Probabilities For Step 3 in Algorithm II

	$P_{\text{Jump}} \left(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \rightarrow k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \right)$
Birth	$\left(b_F(k'_{\text{old}}) \frac{1}{k'_{\text{old}}} \sum_{x \in \mathcal{f}_{k'_{\text{old}}}} h(t_{k'}^{\text{new}} x) \right) \times MVN_3(\psi_{\text{new}}, \log(\sigma_{\text{new}}^2))$
Relocation	$\left(r_F(k_{\text{old}}) \frac{1}{k'_{\text{old}}} h(t_{k'}^{\text{new loc}} t_{k'}^{\text{relocated}}) \right) \times MVN_3(\psi_{\text{new}}, \log(\sigma_{\text{new}}^2))$
Death	$d_F(k'_{\text{old}}) \frac{1}{k'_{\text{old}}}$

Algorithm I

Repeat the Following Steps:

- 1 Call the current set of knots $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})$
- 2 Sample from $\hat{p}(k, \mathbf{t}_k | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, y)$
 - a. Propose a move for the fixed effect knots. Accept with probability

$$p_{\text{Fixed}} = \min \left(1, \frac{P_{\text{Jump}}(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \rightarrow k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}) \hat{p}(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, y)}{P_{\text{Jump}}(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \rightarrow k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}) \hat{p}(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}}, y)} \right),$$

where $P_{\text{Jump}}(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \rightarrow k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} | k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})$ is the probability of jumping from $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}})$ to $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}})$ conditioned on the current value of $(k', \mathbf{t}_{k'})$.

- b. If not accepted, change $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}})$ to the older values $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}})$.
- 3 Sample from $\hat{p}(k', \mathbf{t}_{k'} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)$

- a. Propose a move for the random effect knots. Accept with probability

$$p_{\text{Random}} = \min \left(1, \frac{P_{\text{Jump}}(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} \rightarrow k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}) \hat{p}(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)}{P_{\text{Jump}}(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} \rightarrow k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}) \hat{p}(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, y)} \right),$$

where $P_{\text{Jump}}(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}} \rightarrow k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}} | k_{\text{new}}, \mathbf{t}_{k_{\text{new}}})$ is the probability of jumping from $(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})$ to $(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}})$ conditioned on the current value of $(k, \mathbf{t}_k)$.

- b. If not accepted, change the value of $(k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}})$ to the older pair $(k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})$.
 - 4 Identify the new set of knots $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}_{k'_{\text{new}}})$ as $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}, k'_{\text{old}}, \mathbf{t}_{k'_{\text{old}}})$.
-

Algorithm II

Repeat the Following Steps:

- 1 Call the current set of knots and the current value of ψ and $\sigma^2 \left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}, k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)$
- 2 Sample from $p \left(k, \mathbf{t}_k, \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2, \mathcal{Y} \right)$
 - a. Propose a move for the fixed effect knots. Accept with probability

$$p_{\text{Fixed}} = \min \left(1, \frac{P_{\text{Jump}} \left(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \rightarrow k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right) p \left(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)}{P_{\text{Jump}} \left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \rightarrow k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right) p \left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)} \right)$$

where $P_{\text{Jump}} \left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}} \rightarrow k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \mid k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)$ is the probability of jumping from $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}})$ to $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}})$ conditioned on the current value of $(k', \mathbf{t}_k)$.

- b. If not accepted, change $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}})$ to the older values $(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}})$.
- 3 Sample from $p \left(k', \mathbf{t}_k, \psi, \sigma^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, \mathcal{Y} \right)$
 - a. Propose a move for the random effect knots, ψ , and σ^2 . Accept with probability

$$p_{\text{Random}}, \text{ where } p_{\text{Random}} = \min \left(1, \frac{P_{\text{Jump}} \left(k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 \rightarrow k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \right) p \left(k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \right)}{P_{\text{Jump}} \left(k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \rightarrow k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \right) p \left(k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \right)} \right)$$

and $P_{\text{Jump}} \left(k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \rightarrow k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 \mid k_{\text{new}}, \mathbf{t}_{k_{\text{new}}} \right)$ is the probability of jumping from $\left(k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)$ to $\left(k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2 \right)$ conditioned on the current value of $(k, \mathbf{t}_k)$.

- b. If not accepted, change the value of $(k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2)$ to the older pair $\left(k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)$.
- 4 Identify the new set $(k_{\text{new}}, \mathbf{t}_{k_{\text{new}}}, k'_{\text{new}}, \mathbf{t}'_{k_{\text{new}}}, \psi_{\text{new}}, \sigma_{\text{new}}^2)$ as $\left(k_{\text{old}}, \mathbf{t}_{k_{\text{old}}}, k'_{\text{old}}, \mathbf{t}'_{k_{\text{old}}}, \psi_{\text{old}}, \sigma_{\text{old}}^2 \right)$.

For each sampled set of knots, $(k, \mathbf{t}_k, k', \mathbf{t}_{k'})$, do the following:

- 1 For $t = 1 : L$ (we set $L = 100$)
 - a. Sample the modified Cholesky parameters ψ and $\log(\sigma^2)$ from

$$\begin{pmatrix} \psi^t \\ \log(\sigma^2, t) \end{pmatrix} \sim MVN \left(\begin{pmatrix} \hat{\psi} \\ \log(\hat{\sigma}^2) \end{pmatrix}, -I(\hat{\psi}, \log(\hat{\sigma}^2))^{-1} \right),$$

where $(\hat{\psi}, \hat{\sigma}^2)$ are the maximum likelihood estimates of the Pourhamadi parameters and σ^2 corresponding to the current set of knots, and $I(\hat{\psi}, \log(\hat{\sigma}^2))$ is the Hessian matrix of the log-likelihood function evaluated at these estimates.

- b. Accept this draw with probability

$$\min \left(1, \frac{\frac{p(\mathcal{Y} | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi^t, \sigma^2, t) p(\psi^t | k', \mathbf{t}_{k'}) p(\sigma^2, t)}{\xi(\psi^t, \log(\sigma^2, t))}}{\frac{p(\mathcal{Y} | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \psi^{t-1}, \sigma^2, t-1) p(\psi^{t-1} | k', \mathbf{t}_{k'}) p(\sigma^2, t-1)}{\xi(\psi^{t-1}, \log(\sigma^2, t-1))}} \right)$$

where $p(\psi^t | k', \mathbf{t}_{k'})$ is the prior distribution of ψ evaluated at ψ^t , $p(\sigma^2, t)$ is the prior distribution of σ^2 evaluated at σ^2, t , and $\xi(\psi^t, \log(\sigma^2, t))$ is the multivariate gamma function evaluated at $(\psi^t, \log(\sigma^2, t))$.

- c. If the move is not accepted, let $\psi^t = \psi^{t-1}$, $\log(\sigma^2, t) = \log(\sigma^2, t-1)$
- 2 Draw a fixed effect, α^* , from the conditional posterior distribution

$$a | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \Sigma_{\gamma}^t \sigma^2, t \text{ where } a | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \Sigma_{\gamma}^t \sigma^2, t \sim MVN \left((\mathbf{F}' \mathbf{C}^{-1} \mathbf{F})^{-1} \mathbf{F}' \mathbf{C}^{-1} (\mathbf{Y}, \hat{a})', (\mathbf{F}' \mathbf{C}^{-1} \mathbf{F})^{-1} \right), \mathbf{F}$$

$\mathbf{C}$ is a block-diagonal matrix composed of the n matrices $\left\{ (\mathbf{B}_{R_i} \Sigma_{\gamma}^t \mathbf{B}_{R_i}' + \sigma^2, t \mathbf{I}_{m_i}) \right\}_{i=1}^n$ and the matrix $\left(\frac{1}{n} \sum_{i=1}^n \frac{1}{m_i} I_{a,i}(\hat{a}) \right)^{-1}$, $\mathbf{B}_{F_i}$, and $\mathbf{B}_{R_i}$ are defined as in the text. $\hat{a}$ is the maximum likelihood estimate of α .

- - 3 Draw a random effect, γ^* , from the conditional posterior distribution

$$\gamma | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha^*, \Sigma_{\gamma}^t \sigma^2, t \text{ where } \gamma | k, \mathbf{t}_k, k', \mathbf{t}_{k'}, \alpha^*, \Sigma_{\gamma}^t \sigma^2, t \sim MVN(\lambda, \Lambda), \lambda = \Sigma_{\gamma}^t \mathbf{B}_{R_i}' (\sigma^2, t \mathbf{I}_{m_i})^{-1} (\mathbf{Y}_i - \mathbf{B}_{F_i} \alpha^*)$$

$$\Lambda = \Sigma_{\gamma}^t \mathbf{B}_{R_i}' \left\{ (\sigma^2, t \mathbf{I}_{m_i})^{-1} - (\sigma^2, t \mathbf{I}_{m_i})^{-1} \mathbf{B}_{F_i} \left(\frac{1}{\sigma^2, t} \sum_{i=1}^n \mathbf{B}_{F_i}' \mathbf{B}_{F_i} \right)^{-1} \mathbf{B}_{F_i} (\sigma^2, t \mathbf{I}_{m_i})^{-1} \right\} \mathbf{B}_{R_i} \Sigma_{\gamma}^t$$