

Bayesian Semiparametric Dynamic Frailty Models for Multiple Event Time Data

Michael L. Pennell^{1,2,*} and David B. Dunson²

¹Department of Biostatistics, University of North Carolina at Chapel Hill

²Biostatistics Branch, MD A3-03, National Institute of Environmental Health Sciences, P.O.
Box 12233, Research Triangle Park, NC 27709

**email:* pennell@niehs.nih.gov

August 22, 2005

SUMMARY

Many biomedical studies collect data on times of occurrence for a health event that can occur repeatedly, such as infection, hospitalization, recurrence of disease, or tumor onset. To analyze such data, it is necessary to account for within-subject dependency in the multiple event times. Motivated by data from studies of palpable tumors, this article proposes a dynamic frailty model and Bayesian semiparametric approach to inference. The widely used shared frailty proportional hazards model is generalized to allow subject-specific frailties to change dynamically with age while also accommodating non-proportional hazards. Parametric assumptions on the frailty distribution are avoided by using Dirichlet process priors for a shared frailty and for multiplicative innovations on this frailty. By centering the semiparametric model on a conditionally-conjugate dynamic gamma model, we facilitate posterior computation and lack of fit assessments of the parametric model. Our proposed method is demonstrated using data from a cancer chemoprevention study.

KEY WORDS: Breast cancer; Chemoprevention; Dirichlet process; Nonparametric Bayes; Palpable Tumors; Survival analysis; Tumor multiplicity data.

1. Introduction

Many biomedical studies are designed to assess covariate effects on the time to recurrence of health-related outcomes, such as infections, hospitalizations, or recurrences of disease. For example, data of this type are collected in chemoprevention and carcinogenicity studies measuring the rate of appearance of palpable tumors of the skin and breast of mice (Gail et al., 1980; Forbes and Sambuco, 1998; Dunson, 2000). In these experiments, a rich set of data are available for each mouse, including times of appearance of each lesion, total number of tumors, and time of death.

A number of methods have been proposed to analyze tumor multiplicity data including recent work by Dunson and Dinse (2000) and Sinha et al. (2002). These articles relied on frailty-type models (Vaupel et al., 1979; Clayton and Cuzick, 1985) to accommodate baseline heterogeneity in risk of developing tumors. In these models, random effects (or frailties) measure an animal's risk relative to that for other individuals in the population, accounting for covariates. Standard shared and multivariate (Sargent, 1998) frailty models treat the frailties as time-independent factors, and hence do not allow a subject's risk to evolve dynamically as they age. Such formulations may be overly-restrictive when applied to tumorigenicity data, since biological changes occurring with age result in complex and unanticipated trends in susceptibility to tumor development. A likely trend is that animals getting tumors relatively early may not be at higher risk later in life.

Recently, several authors have proposed more flexible, dynamic formulations. Yue and Chan (1997) and Yau and MacGilchrist (1998) introduced a proportional hazard model for inter-recurrence times in which a subject's frailty changes following each event, and Lam et al. (2002) developed a related approach for the proportional odds model. In tumorigenicity studies, a time-varying frailty structure may be more realistic since it is more natural to model individual-specific risk as changing with age instead of according to previous occurrences of

tumors. Relevant methods have been proposed by Henderson and Shimakura (2003), who developed a longitudinal Poisson regression model with gamma frailties which vary with time, and Paik et al. (1994), who proposed a proportional hazards frailty model with a time-specific random factor. Although promising, these methods involve complicated likelihoods and difficult computation, particularly when one considers generalizations (e.g., for joint modeling).

Bayesian approaches have several advantages, including ease of computation via MCMC, ability to incorporate prior information (e.g., from historical controls), and exact inferences on different aspects of the tumor response (time to first tumor, total tumor burden, etc). In the Bayesian literature, there has been limited consideration of dynamic frailty models and methods for multiple event time data in general. For recent Bayesian references on frailty models for multiple event time and multivariate survival data, refer to papers by Gustafson (1997), Sahu et al. (1997), Walker and Mallick (1997), Sargent (1998), Aslanidou et al. (1998), Sinha (1998), Chen et al. (2002), Dunson and Chen (2004), Sinha and Maiti (2004), as well as a review by Ibrahim, Chen, and Sinha (2001). Härkänen et al. (2003) proposed an innovative approach based on a model that allows subject-specific frailty trajectories to vary according to a latent class structure. In many settings, including animal tumorigenicity studies, it may be more natural to suppose that the age-specific risk trajectories vary according to a continuum, with each subject potentially having their own unique pattern.

Motivated by the tumor multiplicity application, we propose a Bayesian semiparametric dynamic frailty model. Our methodology generalizes the shared frailty model to allow time-varying frailties and regression coefficients. In addition, we use a multiplicative parameterization to introduce autocorrelation and smooth the time trajectories. To improve flexibility, we consider a non-parametric treatment of the frailty distribution using a Dirichlet process (DP) mixture (Antoniak, 1974). For references on related approaches for using DP mixtures in

Bayesian analyses, refer to West et al. (1994), Bush and MacEachern (1996), Mukhopadhyay and Gelfand (1997), Müller and Rosner (1997), Kleinman and Ibrahim (1998), and Dominici and Parmigiani (2001). In addition, an alternative nonparametric approach for recurrent event time data was proposed by Ishwaran and James (2004). In this paper, we use a DP mixture to allow uncertainty in the distributions of a shared frailty and multiplicative innovations on this frailty. By centering the semiparametric model on a conditionally-conjugate dynamic gamma model, we facilitate posterior computation and lack of fit assessments of the parametric model using predictive distributions.

Section 2 proposes the model and prior structure. Section 3 outlines an MCMC algorithm for posterior computation. Section 4 applies the method to data from a cancer chemoprevention study, and Section 5 discusses the results.

2. Dynamic Frailty Model

2.1 Model Specification and Frailty Structure

Consider a study measuring the times of occurrence of repeated events within n subjects. The rate of event occurrence for subject i ($i = 1, \dots, n$) at time t is denoted $\lambda_i(t)$. We partition the time axis into M finely-spaced intervals, T_1, T_2, \dots, T_M , where $T_j = (\tau_{j-1}, \tau_j]$, $0 = \tau_0 < \tau_1 < \tau_2 < \dots < \tau_M$, and τ_M is the maximum follow-up time in the study. The intervals are chosen to be sufficiently narrow so that it can be assumed that $\lambda_i(t) = \lambda_{ij}$ for all $t \in T_j$, $j = 1, \dots, M$.

Suppose that subject i is followed for t_i^* time units, where $t_i^* \in T_{M_i}$ and $M_i \leq M$. Under these specifications, let $\boldsymbol{\xi}_i = (\xi_{i1}, \dots, \xi_{iM_i})'$ denote a vector of time-varying frailties for subject i , and let $\mathbf{x}_{ij} = (x_{ij1}, \dots, x_{ijp})'$ be a vector of p predictors for subject i and time interval T_j . We focus on models having the following structure:

$$\lambda_{ij} = \xi_{ij} \lambda_{0j} g(\mathbf{x}_{ij}; \boldsymbol{\beta}_j), \tag{1}$$

where λ_{0j} is the baseline hazard for the j th interval, $g(\cdot)$ is a known link function mapping from $\mathfrak{R} \rightarrow \mathfrak{R}^+$, and β_j are interval-specific regression parameters. We further express the baseline hazard as $\lambda_{0j} = \widehat{\lambda}_{0j}\Delta_j$, where $\widehat{\lambda}_{0j}$ is an initial guess at the baseline hazard (e.g., estimated from historical control data or based on expert elicitation) and Δ_j is a multiplicative deviation from this guess.

Similar to what is done by Paik et al. (1994), the frailties are decomposed into time-independent and -dependent components:

$$\xi_{ij} = \phi_i \prod_{h=1}^j \phi_{ih}, \quad (2)$$

where ϕ_i is a subject-specific shared frailty and ϕ_{ih} is the multiplicative innovation over time interval h . This multiplicative structure provides a convenient framework for imposing autocorrelation amongst the frailties. To demonstrate this feature, consider a model in which $\phi_i \sim \mathcal{G}(\psi_1, \psi_1)$ independently of $\phi_{ih} \stackrel{iid}{\sim} \mathcal{G}(\psi_2, \psi_2)$, $h = 1, \dots, M_i$. Given these distributions, the correlation between frailties from intervals j and $j + d$ ($d > 0$) is

$$\text{corr}(\xi_{ij}, \xi_{i(j+d)}) = \sqrt{\frac{\psi_2^d \{(1 + \psi_1)(1 + \psi_2)^j - \psi_1 \psi_2^j\}}{(1 + \psi_1)(1 + \psi_2)^{j+d} - \psi_1 \psi_2^{j+d}}}. \quad (3)$$

Under the above model, ψ_1 provides an overall measure of between subject heterogeneity while ψ_2 can be thought of as both a smoothing parameter for the frailty trajectories and a measure of temporal heterogeneity within each subject.

Although we have observed that this dynamic gamma model performs well with some data, it can lead to spurious inferences when the actual distributions of the frailty components are distinctly non-gamma. For example, Walker and Mallick (1997) demonstrated that frailty distributions may differ across predictor-level resulting in multi-modal distributions when all the frailties are pooled together. With this in mind, we propose a Dirichlet process (DP) mixture model which is flexible to unanticipated trends in the frailty.

For subject i and interval T_j , let r_{ij} denote the time at risk and let N_{ij} denote the number of events experienced. Then, under expressions (1) and (2) with a DP specification for the distributions of ϕ_i and ϕ_{ij} , we have:

$$\begin{aligned}
N_{ij} &\stackrel{ind.}{\sim} \mathcal{P}\left(r_{ij}\phi_i\left(\prod_{h=1}^j\phi_{ih}\right)\lambda_{0j}g(\mathbf{x}_{ij},\boldsymbol{\beta}_j)\right) \\
\phi_i &\stackrel{ind.}{\sim} G_1 & \phi_{ij} &\stackrel{ind.}{\sim} G_2 \\
G_1 &\sim \mathcal{D}(\alpha_{01}G_{01}) & G_2 &\sim \mathcal{D}(\alpha_{02}G_{02}),
\end{aligned} \tag{4}$$

where $\mathcal{D}(\alpha_0G_0)$ denotes the Dirichlet process centered on G_0 with precision α_0 , and we assume G_{01} is $\mathcal{G}(\psi_1, \psi_1)$ and G_{02} is $\mathcal{G}(\psi_2, \psi_2)$. This structure is centered on the dynamic Gamma frailty model described above, but we allow the true frailty distribution to deviate from the parametric form. The amount of uncertainty in the gamma assumption for the two frailty components is controlled by the hyperparameters α_{01} and α_{02} , with small values of these parameters corresponding to little faith in the gamma forms.

2.2 Priors for Model Deviations and Regression Parameters

As with the frailties, the multiplicative deviations from the initial baseline hazard estimates are separated into time dependent and independent components:

$$\Delta_j = \nu_0 \prod_{h=1}^j \nu_h. \tag{5}$$

Assuming $\nu_0 \sim \mathcal{G}(\kappa, \kappa)$ and $\nu_j \stackrel{iid}{\sim} \mathcal{G}(\psi_3, \psi_3)$ for $j = 1, \dots, M$, we have a convenient structure for introducing autocorrelation amongst the model deviations

$$corr(\Delta_j, \Delta_{j+d}) = \sqrt{\frac{\psi_3^d \{(1 + \kappa)(1 + \psi_3)^j - \kappa\psi_3^j\}}{(1 + \kappa)(1 + \psi_3)^{j+d} - \kappa\psi_3^{j+d}}}, \tag{6}$$

where κ controls the degree of shrinkage of the posterior towards $\widehat{\boldsymbol{\lambda}}_0$, and ψ_3 measures smoothness in the deviations from the prior estimate. An appealing feature of this prior in the context

of the tumor multiplicity application is that the prior variance increases with time. In many applications, events are known to be rare early and one can obtain a good prior estimate early on, but at later times there is much more uncertainty.

Although many other correlated prior processes have been proposed for the baseline hazard (cf. Arjas and Gasbarra, 1994, Gamerman, 1991, and Gustafson et al., 2003), our proposed structure is practically appealing for complex models since it retains the conditional conjugacy properties of the Gamma process (Kalfleisch, 1978) without the need to introduce any additional latent variables (cf. Nieto-Barajas and Walker, 2002). These properties are highlighted in the following section on posterior computation.

A similar prior structure can be used to provide smooth trajectories for predictor effects. This technique is an alternative to dynamic covariate models based on random walks (cf. Gamerman, 1991 and Sargent, 1997) and can simplify computation in certain applications. We consider the typical exponential link function, $g(\mathbf{x}_{ij}; \boldsymbol{\beta}_j) = e^{\mathbf{x}'_{ij}\boldsymbol{\beta}_j}$. However, our priors are developed in terms of the following re-parameterization of $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jp})'$:

$$\gamma_{jk}^* = e^{\beta_{jk}} = \prod_{l=1}^j \gamma_{lk},$$

where $\gamma_{jk}^* = \lambda_i(t; x_{ijk} = c + 1) / \lambda_i(t; x_{ijk} = c)$ for $t \in T_j$ is the multiplicative change in hazard (i.e., hazard ratio) at time t attributable to a unit increase in the k th predictor, and $\gamma_{lk} = \gamma_{lk}^* / \gamma_{(l-1)k}^*$ is the multiplicative innovation in this hazard ratio experienced over time interval T_l . In the absence of prior information about the effect of the k th predictor, it is reasonable to assume $\gamma_{lk} \stackrel{iid}{\sim} \mathcal{G}(\psi_4, \psi_4)$ priors for the multiplicative increments on the hazard function. The prior is centered on no effect for a predictor with the degree of shrinkage and smoothness in the regression function controlled by the hyperparameter ψ_4 , as is clear from the following expression:

$$\text{corr}(\gamma_{jk}^*, \gamma_{(j+d)k}^*) = \sqrt{\frac{\psi_4^d \{(1 + \psi_4)^j - \psi_4^j\}}{(1 + \psi_4)^{j+d} - \psi_4^{j+d}}}. \quad (7)$$

In cases in which one has prior information about the regression function, the prior structure above could be used to model multiplicative deviations from these initial estimates. This approach is closely related to our priors for the baseline hazard. Another extension that may be useful in some applications, would be to allow the level of smoothing to vary across predictors. These generalizations are straightforward and we focus on the simple case for ease in exposition.

2.3 *Elicitation of Hyperpriors*

To reduce the sensitivity of analyses to subjectively chosen hyperparameters, we consider the use of hyperpriors in building our dynamic frailty model. A priori we assume that $\psi_1 \sim \mathcal{G}(a_1, b_1)$ and $\psi_2 \sim \mathcal{G}(a_2, b_2)$ to obtain some information from the data about the variance components of the frailties. Although a diffuse prior would be reasonable for ψ_1 , allowing the level of between subject heterogeneity to be determined by the data, we recommend using an informative prior for ψ_2 centered on values which reflect the expected amount of correlation between Poisson counts from adjacent time intervals.

One could additionally specify gamma hyperpriors for ψ_3 and ψ_4 , but preliminary results suggest that it can be difficult to specify a prior that ensures sufficient smoothing of the baseline hazard and predictor effects, since the likelihood tends to dominate the prior even for moderate sample sizes. Thus, to avoid over-fitting, we recommend choosing values of ψ_3 and ψ_4 which reflect both one's a priori intuition about the level of autocorrelation between time intervals and one's confidence in the initial guess of the baseline hazard. We also recommend performing a sensitivity analysis of this choice. Hyperpriors for α_{01} and α_{02} may also be elicited to determine the amount of deviation from the gamma frailty structure in the data, though we prefer to fix these parameters to avoid over-parameterization.

3. Posterior Computation

In applications with moderate to large sample sizes and lengthy followups, our model will result in a large number of latent variables and unknown parameters. Fortunately, due to the conditionally conjugate structure of the priors, computation can proceed fairly easily using a hybrid MCMC algorithm consisting of Gibbs and Metropolis-Hastings steps (Gelfand and Smith, 1990; Tierney, 1994). For simplicity, our methods assume that data are right censored and that censoring is non-informative. To modify our MCMC algorithm to accommodate interval-censored data, one can simply include data augmentation steps for sampling the exact event times for interval-censored observations. In addition, it is conceptually straightforward to allow dependent censoring by incorporating the time-varying frailty as a predictor in a model for the censoring time.

3.1 Full Conditional Posterior Distributions

Let $\mathbf{Y} = \{\mathbf{N}_1, \dots, \mathbf{N}_n; t_1^*, \dots, t_n^*; \mathbf{X}_1, \dots, \mathbf{X}_n\}$ denote the observed data, where $\mathbf{N}_i = (N_{i1}, \dots, N_{iM_i})'$ and $\mathbf{X}_i = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{iM_i})'$. Also, let $\Phi = \{\phi_i, \phi_{ij}, i = 1, \dots, n; j = 1 \dots, M_i\}$, $\nu = (\nu_0, \nu_1, \dots, \nu_M)'$, and $\Gamma = \{\gamma_{hk}, h = 1 \dots, M; k = 1, \dots, p\}$. Given the above notation, the likelihood is proportional to

$$L(\Phi, \nu, \Gamma | \mathbf{Y}) = \prod_{i=1}^n \prod_{j=1}^{M_i} \left(\phi_i \nu_0 \widehat{\lambda}_{0j} \prod_{h=1}^j \phi_{ih} \nu_h \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \right)^{N_{ij}} \exp \left\{ - r_{ij} \phi_i \nu_0 \widehat{\lambda}_{0j} \prod_{h=1}^j \phi_{ih} \nu_h \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \right\}. \quad (8)$$

Under the dynamic gamma model, the full conditional posterior distributions of the ϕ_i and ϕ_{ij} 's are also gamma due to the Poisson form of (8):

$$\pi(\phi_i | \cdot) = G_i = \mathcal{G} \left(\psi_1 + N_i^*, \psi_1 + m_i^* \right) \quad (9)$$

$$\pi(\phi_{is} | \cdot) = G_{is} = \mathcal{G} \left(\psi_2 + N_{is}^*, \psi_2 + m_{is}^* \right), \quad (10)$$

where the notation $a|\cdot$ denotes a given all other variables and $N_i^* = \sum_{j=1}^{M_i} N_{ij}$, $N_{is}^* = \sum_{j=s}^{M_i} N_{ij}$

$$\begin{aligned} m_i^* &= \nu_0 \sum_{j=1}^{M_i} r_{ij} \widehat{\lambda}_{0j} \prod_{h=1}^j \phi_{ih} \nu_h \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \\ m_{is}^* &= \phi_i \nu_0 \sum_{j=s}^{M_i} r_{ij} \widehat{\lambda}_{0j} \left(\prod_{h=1}^j \nu_h \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \right) \prod_{f \neq s}^j \phi_{if}, \end{aligned}$$

for $s = 1, \dots, M_i$. The full conditional posterior densities of the elements of $\boldsymbol{\nu}$ and Γ are

$$\pi(\nu_0|\cdot) = \mathcal{G} \left(\kappa + \sum_{i=1}^n N_i^*, \kappa + \sum_{i=1}^n \phi_i \sum_{j=1}^{M_i} r_{ij} \widehat{\lambda}_{0j} \prod_{h=1}^j \phi_{ih} \nu_h \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \right) \quad (11)$$

$$\pi(\nu_s|\cdot) = \mathcal{G} \left(\psi_3 + \sum_{i \in R_s} N_{is}^*, \psi_3 + \nu_0 \sum_{i \in R_s} \phi_i \sum_{j=s}^{M_i} r_{ij} \widehat{\lambda}_{0j} \left(\prod_{h=1}^j \phi_{ih} \prod_{k=1}^p \gamma_{hk}^{x_{ijk}} \right) \left(\prod_{f \neq s}^j \nu_f \right) \right) \quad (12)$$

$$\begin{aligned} \pi(\gamma_{sk}|\cdot) &\propto \exp \left\{ \left(\psi_4 + \sum_{i \in R_s} \sum_{j=s}^M N_{ij} x_{ijk} \right) \log(\gamma_{sk}) - \gamma_{sk} \psi_4 \right. \\ &\quad \left. - \nu_0 \sum_{i \in R_s} \sum_{j=s}^M r_{ij} \phi_i \widehat{\lambda}_{0j} \prod_{h=1}^j \phi_{ih} \nu_h \prod_{l=1}^p \gamma_{hl}^{x_{ijl}} \right\}, \end{aligned} \quad (13)$$

where $R_s = \{i : M_i \geq s\}$, for $s = 1, \dots, M$.

When the Dirichlet process is used to model the frailty terms, the full conditionals of ϕ_i and ϕ_{is} are not G_i and G_{is} . Using the Pólya urn representation of the Dirichlet process (Blackwell and MacQueen, 1973; MacEachern, 1994; West, 1990), one can show that the prior distribution of ϕ_i given $\boldsymbol{\phi}^{(i)} = (\phi_1, \dots, \phi_{i-1}, \phi_{i+1}, \dots, \phi_n)'$ is the mixture

$$\left(\frac{\alpha_{01}}{\alpha_{01} + n - 1} \right) G_{01} + \left(\frac{1}{\alpha_{01} + n - 1} \right) \sum_{l=1}^{h^{(i)}} n_l^{(i)} \delta_{\theta_l^{(i)}}, \quad (14)$$

where δ_θ denotes the degenerate distribution with all its mass at θ , and the prior for ϕ_{is} given $\boldsymbol{\phi}_s^{(i)} = \{\phi_{i's} : i' \in R_s, i' \neq i\}$ and $n_{(s)}$ total subjects in R_s is

$$\left(\frac{\alpha_{02}}{\alpha_{02} + n_{(s)} - 1} \right) G_{02} + \left(\frac{1}{\alpha_{02} + n_{(s)} - 1} \right) \sum_{l=1}^{h_s^{(i)}} n_{sl}^{(i)} \delta_{\theta_{sl}^{(i)}}, \quad (15)$$

where $\boldsymbol{\theta}^{(i)}$ and $\boldsymbol{\theta}_s^{(i)}$ denote the $h^{(i)}$ and $h_s^{(i)}$ unique values of $\boldsymbol{\phi}^{(i)}$ and $\boldsymbol{\phi}_s^{(i)}$, respectively, $n_l^{(i)}$ elements of $\boldsymbol{\phi}^{(i)}$ have value $\theta_l^{(i)}$, and $n_{sl}^{(i)}$ elements of $\boldsymbol{\phi}_s^{(i)}$ have value $\theta_{sl}^{(i)}$. After factoring in

the likelihood, the full conditional posterior of ϕ_i is

$$q_{i0}G_i + \sum_{l=1}^{h^{(i)}} q_{il}\delta_{\theta_l^{(i)}} \quad (16)$$

where

$$q_{il} = \begin{cases} \frac{c_1 \alpha_{01} C(\psi_1, \psi_1)}{C(\psi_1 + N_i^*, \psi_1 + m_i^*)} & l = 0 \\ c_1 n_l^{(i)} (\theta_l^{(i)})^{N_i^*} \exp\{-\theta_l^{(i)} m_i^*\} & l > 0, \end{cases}$$

$C(a, b) = b^a / \Gamma(a)$, and c_1 is a normalizing constant. Similarly, the full conditional posterior of ϕ_{is} ($s = 1, \dots, M_i$) is

$$q_{is0}G_{is} + \sum_{l=1}^{h_s^{(i)}} q_{isl}\delta_{\theta_{sl}^{(i)}} \quad (17)$$

where

$$q_{isl} = \begin{cases} \frac{c_2 \alpha_{02} C(\psi_2, \psi_2)}{C(\psi_2 + N_{is}^*, \psi_2 + m_{is}^*)} & l = 0 \\ c_2 n_{sl}^{(i)} (\theta_{sl}^{(i)})^{N_{is}^*} \exp\{-\theta_{sl}^{(i)} m_{is}^*\} & l > 0. \end{cases}$$

Thus, the full conditional distributions of ϕ_i and $(\phi_{i1}, \dots, \phi_{iM_i})'$ are mixtures of the gamma posteriors obtained under the dynamic gamma model and multinomial distributions with support on the unique values of each frailty component.

3.2 Updating Algorithm

In order to sample efficiently under the Dirichlet process mixture, we invoke a sampling scheme similar to that provided by MacEachern (1994) and West et al. (1994). Let there be h unique values in $(\phi_1, \dots, \phi_n)'$ and h_s unique values in $\{\phi_{is} : i \in R_s\}$, which we denote by $\boldsymbol{\theta} = (\theta_1, \dots, \theta_h)'$ and $\boldsymbol{\theta}_s = (\theta_{s1}, \dots, \theta_{sh_s})'$, respectively, for $s = 1, \dots, M$. We also define the discrete random variables \mathcal{S}_i and \mathcal{S}_{is} such that $\mathcal{S}_i = k$ if $\phi_i = \theta_k$ and $\mathcal{S}_{is} = l$ if $\phi_{is} = \theta_{sl}$, for $i \in R_s$. Our MCMC algorithm proceeds as follows:

Step 1. Sample ν_0 from (11), given the current values of $\boldsymbol{\Phi}$, ν_1, \dots, ν_M , and $\boldsymbol{\Gamma}$.

Step 2.1. Sample \mathcal{S}_i , for $i = 1, \dots, n$ from a multinomial distribution with $\Pr(\mathcal{S}_i = l) = q_{il}$, for $l = 0, 1, \dots, h^{(i)}$, with a new ϕ_i drawn from G_i if $\mathcal{S}_i = 0$.

Step 2.2. Given the updated values of h and $\mathcal{S}_1, \dots, \mathcal{S}_n$, generate a new $\boldsymbol{\theta}$ by sampling each θ_k from its full conditional posterior distribution, $\mathcal{G}(\psi_1 + \sum_{i:\mathcal{S}_i=k} N_i^*, \psi_1 + \sum_{i:\mathcal{S}_i=k} m_i^*)$, for $k = 1, \dots, h$. Assign the appropriate value of $\boldsymbol{\theta}^{(i)}$ to ϕ_i as indicated by \mathcal{S}_i .

For $s = 1, \dots, M$, perform the following steps:

Step 3.s(a). For $i \in R_s$, sample \mathcal{S}_{is} from the multinomial distribution with $\Pr(\mathcal{S}_{is} = l) = q_{isl}$, for $l = 0, 1, \dots, h_s^{(i)}$, with a new ϕ_{is} drawn from G_{is} if $\mathcal{S}_{is} = 0$.

Step 3.s(b). Update $\boldsymbol{\theta}_s$ by sampling each θ_{sl} from the full conditional posterior, $\mathcal{G}(\psi_2 + \sum_{i:\mathcal{S}_{is}=l} N_{is}^*, \psi_2 + \sum_{i:\mathcal{S}_{is}=l} m_{is}^*)$, for $l = 1, \dots, h_s$. Assign the appropriate value of $\boldsymbol{\theta}_s^{(i)}$ to ϕ_{is} .

Step 3.s(c-d). Sample ν_s from (12) and γ_{sk} from (13) for $k = 1, \dots, p$.

Step 4.1-4.2. Update ψ_1 and ψ_2 . Under the DP model, the full conditional posterior densities of ψ_1 and ψ_2 depend only on $\boldsymbol{\theta}$ and $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$,

$$\pi(\psi_1|\cdot) \propto \left(\frac{\psi_1^{\psi_1}}{\Gamma(\psi_1)} \right)^h \psi_1^{a_1-1} \exp \left\{ -\psi_1 \left(b_1 + \sum_{k=1}^h (\theta_k - \log \theta_k) \right) \right\} \quad (18)$$

$$\pi(\psi_2|\cdot) \propto \left(\frac{\psi_2^{\psi_2}}{\Gamma(\psi_2)} \right)^{\sum_{s=1}^M h_s} \psi_2^{a_2-1} \exp \left\{ -\psi_2 \left(b_2 + \sum_{s=1}^M \sum_{k=1}^{h_s} (\theta_{sk} - \log \theta_{sk}) \right) \right\}, \quad (19)$$

while under the dynamic gamma frailty model, the posteriors depend on each ϕ_i and ϕ_{ij} , respectively. Since (18) and (19) do not have closed forms, we recommend updating ψ_1 and ψ_2 using a Metropolis-Hastings random walk.

It is also fairly straightforward to sample from predictive distributions at each iteration of the Gibbs sampler. Given $\boldsymbol{\theta}$ and $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$, the frailty of a future subject, $\boldsymbol{\xi}_{n+1}$, may be

predicted by sampling from the distributions

$$\pi(\phi_{n+1}|\phi) = \left(\frac{\alpha_{01}}{\alpha_{01} + n}\right)G_{01} + \left(\frac{1}{\alpha_{01} + n}\right)\sum_{l=1}^h n_l\delta_{\theta_l} \quad (20)$$

$$\pi(\phi_{(n+1)s}|\phi_s) = \left(\frac{\alpha_{02}}{\alpha_{02} + n_s}\right)G_{02} + \left(\frac{1}{\alpha_{02} + n_s}\right)\sum_{l=1}^{h_s} n_{sl}\delta_{\theta_{sl}} \quad (21)$$

for $s = 1, \dots, M$. Count data for a future subject may then be simulated using the non-homogeneous Poisson process in (4).

Our proposed MCMC methodology performed well when applied to simulated data. Posterior estimates of the time-varying hazard ratios and future predictions of the hazard function tended to agree closely with the true values, even when the prior for the baseline hazard was poorly specified. As expected, results were somewhat sensitive to the prior when sample sizes were small. However, robustness improved with increasing sample sizes, likely reflecting the borrowing of information across time and hyperprior structure.

3.3 *Identifiability and Computational Issues*

As for most latent variable models, some identifiability restrictions are needed. In practice, ν_1 and ϕ_{i1} should be fixed at 1, since these parameters have the same contributions to likelihood as ν_0 and ϕ_i , $i = 1, \dots, n$. We also recommend fixing ν_0 since there may be a tendency for weak identifiability between ν_0 and the mean of G_1 , leading to slow mixing of the MCMC algorithm.

Another issue in implementation is the choice of knots for the piecewise constant hazard. The typical frequentist approach of choosing intervals at the unique failure times (e.g. Breslow, 1974) is inappropriate from a Bayesian perspective, since it involves using the data to choose priors. By choosing autocorrelated priors which borrow information across intervals, our approach allows for tightly spaced intervals. However, we recommend choosing knots so that intervals are at least as wide as the inter-exam times; data are not informative about changes

on a finer scale, and the computational burden increases with the number of intervals.

4. Chemoprevention Application

4.1 Data Analysis

We illustrate our approach using data from a study of the effect of canthaxanthin, a carotenoid found in fruits and vegetables, on chemically induced mammary carcinogenesis (Grubbs et al., 1991). This data set has been used as an example previously by both Kokoska et al. (1993) and Dunson and Dinse (2000). The study consisted of 119 Sprague-Dawley rats administered one of four diets beginning at 34 days of age: (1) 3390 mg/kg canthaxanthin, (2) 1130 mg/kg canthaxanthin, (3) 328 mg/kg retinyl acetate, or (4) vehicle control. There were 30 rats in treatment groups 1-3 and 29 in group 4. At age 55 days, each rat was administered 15 mg of a known carcinogen DMBA by gavage. Regular palpations began at 75 days of age and, in general, were performed twice a week until 235 days of age, or 180 days following administration of DMBA.

Our analysis focused on the tumor occurrence times, as measured from DMBA administration, of rats in treatment groups 1,2, and 4. We partitioned the study period into 24 time intervals, with each interval having a length of one week except for the first and last intervals (22 and 4 days respectively). When the animals died naturally or were sacrificed at the end of the study period, an extensive pathological examination was conducted to find tumors that may be undetectable by palpation. For this reason, we allowed for a multiplicative increase in the baseline hazard of tumor detection at the final examination by introducing an additional parameter, ω , into our model. Thus, the model we used has the following form:

$$\lambda_{ij} = \phi_i \nu_0 \hat{\lambda}_{0j} \omega^{I_{ij}} \prod_{h=1}^j \phi_{ih} \nu_h \gamma_{h1}^{x_{i1}} \gamma_{h2}^{x_{i2}}, \quad (22)$$

where $I_{ij} = 1$ if $j = M_i$ and 0 otherwise and $x_{ik} = 1$ if subject i is in treatment group k and

0 otherwise, $k = 1, 2$. Assigning a $\mathcal{G}(a_\omega, b_\omega)$ prior to ω results in the conditional posterior,

$$\pi(\omega|\cdot) = \mathcal{G}\left(a_\omega + \sum_{i=1}^n N_{iM_i}, b_\omega + \sum_{i=1}^n r_{iM_i} \phi_i \widehat{\lambda}_{0j} \nu_0 \prod_{h=1}^{M_i} \phi_{ih} \nu_h \gamma_{1h}^{x_{i1}} \gamma_{2h}^{x_{i2}}\right). \quad (23)$$

In this study of cancer initiation, the retinyl acetate group (3) can be considered as a second control group, since it is known that this treatment is only effective in decreasing promotion (Grubbs et al., 1991). Thus, these data were used to choose the prior for the baseline hazard. In particular, we fit a Poisson regression model with a cubic polynomial in time to the group 3 data and used the resulting fitted curve as $\widehat{\lambda}_0$. We chose $\mathcal{G}(5, 0.5)$ and $\mathcal{G}(25, 0.5)$ priors for ψ_1 and ψ_2 , respectively, to express belief in low between and within subject heterogeneity, but high levels of autocorrelation amongst the frailties. Since tumor incidence should be negligible in the first three weeks of the study, we assumed $\gamma_{11} = \gamma_{12} = 1$. In addition, to express modest confidence in the gamma frailty assumptions, we set $\alpha_{01} = \alpha_{02} = 5$. Finally, we let $\psi_3 = \psi_4 = 50$ to induce high levels of smoothing in both the baseline hazard and treatment effects, and set $a_\omega = 3$ and $b_\omega = 1$ to reflect our a priori belief that more tumors are observed in the exam following sacrifice.

Using our MCMC methodology, we ran a chain of 55,000 iterations with the first 5,000 discarded as a burn-in. To reduce autocorrelation, every 10th observation was saved to thin the chain. In addition, to speed up computation, we sampled each \mathcal{S}_i and \mathcal{S}_{is} conditional on the cluster memberships at the previous iteration.

Figure 1 provides the mean and pointwise 95% credible intervals for the hazard ratios comparing the canthaxanthin dose groups to control. Although the hazard ratios are initially near one, as time progresses they decrease toward a plateau between weeks 10-26. The posterior mean and 95% credible interval for the average hazard ratio over this interval is $\widehat{\gamma}_{[10,26]2}^* = 0.553$ (0.343, 0.863) in the low dose group and $\widehat{\gamma}_{[10,26]1}^* = 0.489$ (0.304, 0.758) in the high dose group. In addition, the estimated posterior probabilities of a chemopreventive effect

in the two groups for this interval are $\Pr(\gamma_{[10,26]2}^* < 1 \mid \mathbf{Y}) = 0.995$ and $\Pr(\gamma_{[10,26]1}^* < 1 \mid \mathbf{Y}) = 0.999$, with $\Pr(\gamma_{[10,26]2}^* > \gamma_{[10,26]1}^* \mid \mathbf{Y}) = 0.676$, suggesting highly significant, but similar, effects overall in the two dose groups.

[Figure 1 about here.]

As discussed by Kokoska et al. (1993) and Dunson and Dinse (2000), in the presence of a significant effect on tumor incidence, there is typically interest in assessing which aspects of the tumor response profile are most affected. In particular, tumor biologists wish to distinguish effects on multiplicity (total number of tumors) and latency (time to tumor onset). In our model, the effects of treatment group k on multiplicity may be evaluated by computing the posterior probability $P_{Mk} = \Pr(\Lambda_{kM}/\Lambda_{0M} < 1 \mid \mathbf{Y})$, where Λ_{kM} is the cumulative hazard at sacrifice for a typical animal (i.e. $\boldsymbol{\xi}_i = \mathbf{1}$) in treatment group k , $k = 1, 2$, and Λ_{0M} is the cumulative baseline hazard. Also, for $k = 1, 2$, a beneficial effect of treatment k on latency may be evaluated by computing $P_{Lk} = \Pr(\mu_0 < \mu_k \mid \mathbf{Y})$ where

$$\mu_k = \sum_{j=1}^M j \left(\frac{\lambda_{0j} \omega^{1(j=M)} \gamma_{jk}^* (\tau_j - \tau_{j-1})}{\sum_{h=1}^M \lambda_{0h} \omega^{1(h=M)} \gamma_{hk}^* (\tau_h - \tau_{h-1})} \right), \quad (24)$$

which is the expected interval of onset of a typical animal in treatment group k , and μ_0 is the expected interval of onset for a control animal. We found that both the low and high doses of canthaxanthin substantially decreased tumor burden, (P_{M1} and $P_{M2} > 0.99$), but there was no evidence of a beneficial effect on latency ($P_{L1} = 0.156, P_{L2} = 0.392$). These results are consistent with the findings of Dunson and Dinse (2000).

As seen in Figure 2, a substantial proportion of animals have frailties which evolve dynamically. In addition, two animals not depicted in the graphs have trajectories which increase toward values greater than four. These trends suggest that there are unmeasured factors which are affecting tumor incidence and a closer examination of the animals' genetic and physical traits should be made.

[Figure 2 about here.]

4.2 *Goodness of Fit and Sensitivity Analyses*

Goodness of fit can be assessed by comparing the predictive distribution of the weekly tumor counts with observed values in the data set. As seen in Figure 3, weekly predictions of tumor incidence prior to sacrifice agree with mortality adjusted means in the data. Assuming sacrifice at week 26, the posterior mean and 95% credible interval for the number of tumors discovered during the final exam are 0.520 (0.307, 0.808) for the vehicle control group, 0.285 (0.142, 0.509) for the low dose group, and 0.234 (0.111, 0.420) for the high dose group, which are in agreement with observed means (0.714, 0.167, and 0.222 tumors for the vehicle, low dose, and high dose groups, respectively).

[Figure 3 about here.]

We tested the sensitivity of our methodology to frailty assumptions by comparing the results using the priors from the previous section (denoted DPM1) against the fully parametric dynamic gamma model and the results obtained using an analysis with less confidence in the gamma assumption, expressed by letting $\alpha_{01} = \alpha_{02} = 2$ (DPM2). As seen in Figure 4, the predictive distributions of the frailties differ somewhat across each model. Most notably, low a priori confidence in the base model causes the distribution of ϕ_{n+1} to be more skewed and to have a fatter right tail. However, as seen in Table 1, parameter estimates and predictive probabilities are robust.

[Figure 4 about here.]

A second sensitivity analysis was performed for the smoothing parameters by repeating the analysis with $\psi_3 = \psi_4 = 25$ and $\alpha_{01} = \alpha_{02} = 5$ (denoted DPM3). Although Table 1

demonstrates that lower levels of smoothing decrease the estimated hazard ratios somewhat, conclusions regarding multiplicity and latency effects do not change. In addition, lowering ψ_3 and ψ_4 does not substantially improve goodness of fit.

[Table 1 about here.]

5. Discussion

In this paper, we have developed a flexible method for inference in multiple event time data. Our method provides a convenient framework for smoothing hazard functions and time-dependent frailty trajectories and covariate effects. An appealing feature of our approach is the incorporation of dynamic frailties with nonparametric distributions. Even in a data set comprised of fairly homogeneous animals, our model is capable of identifying age-dependent shifts in susceptibility. By comparing the individuals' frailty trajectories, one can identify unusual individuals for genotyping and further examination. However, even in applications where the frailty is of no interest to the researcher, it would still be necessary to account for these trends to improve the accuracy of future predictions.

By using Dirichlet process priors for a shared frailty and multiplicative innovations on the frailty, we have provided a less restrictive modeling framework than parametric alternatives. Unfortunately, these priors ensure that the posterior distributions of the frailties are almost surely discrete. This problem could be solved by using Pólya tree priors (Lavine, 1992, 1994) instead of the Dirichlet process, as was done by Walker and Mallick (1997) in the context of a shared frailty. However, a simpler alternative would be to model each ϕ_i and ϕ_{ij} using a Dirichlet process mixture. Although we did not consider such a methodology, this would be straightforward since it would just involve adding another level to the hierarchy.

Several extensions of our method would be interesting to pursue. For example, one could consider priors with a mixture structure to allow selection of the frailty form. In particular,

selection between no frailty, a shared frailty, and a dynamic frailty for each predictor. This approach would be a generalization of recent work by Dunson and Chen (2004). Another possible direction would be to apply our nonparametric framework to allow dynamic random effects in generalized linear models for longitudinal data, in which repeated outcomes can have any distribution in the exponential family, and for joint modeling of longitudinal and survival data.

ACKNOWLEDGEMENTS

The authors thank Dr. Clinton Grubbs, Univ. Alabama Birmingham, for providing the data and Bo Cai and Gregg Dinse for their helpful comments. This research was supported by the Intramural Research Program of the NIH, and NIEHS.

REFERENCES

- Antoniak, C.E. (1974). Mixtures of Dirichlet processes with applications to nonparametric problems. *Annals of Statistics* **2**, 1152-1174.
- Arjas, E. and Gasbarra, D. (1994). Nonparametric Bayesian inference from right censored survival data, using the Gibbs sampler. *Statistica Sinica* **4**, 505-524.
- Aslanidou, H., Dey, D.K., and Sinha, D. (1998). Bayesian analysis of multivariate survival data using Monte Carlo methods. *Canadian Journal of Statistics* **26**, 33-48.
- Blackwell, D. and MacQueen, J.B. (1973). Ferguson distributions via Pólya urn schemes. *Annals of Statistics* **1**, 353-355.
- Breslow, N.E. (1974). Covariance analysis of censored survival data. *Biometrics* **30**, 89-99.
- Bush, C.A. and MacEachern, S.N. (1996). A semi-parametric Bayesian model for randomized block designs. *Biometrika* **83**, 175-185.
- Chen, M.H., Ibrahim, J.G., and Sinha, D. (2002). Bayesian inference for multivariate survival data with a cure fraction. *Journal of Multivariate Analysis* **80**, 101-126.

- Clayton, D. and Cuzick, J. (1985). Multivariate generalizations of the proportional hazards model. *Journal of the Royal Statistical Society, Series A* **148**, 82-117.
- Dominici, F. and Parmigiani, G. (2001). Bayesian semiparametric analysis of developmental toxicology data. *Biometrics* **57**, 150-157.
- Dunson, D.B. (2000). Modeling of changes in tumor burden. *Journal of Agricultural, Biological and Environmental Statistics* **6**, 38-48.
- Dunson, D.B and Chen, Z. (2004). Selecting factors predictive of heterogeneity in multivariate event time data. *Biometrics* **60**, 352-358.
- Dunson, D.B. and Dinse, G.E. (2000). Distinguishing effects on tumor multiplicity and growth rate in chemoprevention experiments. *Biometrics* **56**, 1068-1075.
- Forbes, P.D. and Sambuco, C.P. (1998). Assays for photocarcinogenesis: relevance of animal models. *International Journal of Toxicology* **17**, 577-588.
- Gail, M.H., Santner, T.J., and Brown, C.C. (1980). An analysis of comparative carcinogenesis experiments based on multiple times to tumor. *Biometrics* **36**, 255-266.
- Gamerman, D. (1991). Dynamic Bayesian models for survival data. *Applied Statistics* **40**, 63-79.
- Gelfand, A.E. and Smith, A.F.M. (1990). Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**, 398-409.
- Grubbs, C.J., Eto, I., Juliana, M.M., and Whitaker, L.M. (1991). Effect of canthaxanthin on chemically induced mammary carcinogenesis. *Oncology* **48**, 239-245.
- Gustafson, P. (1997). Large hierarchical Bayesian analysis of multivariate survival data. *Biometrics* **53**, 230-242.
- Gustafson, P., Aeschliman, D., and Levy, A.R. (2003). A simple approach to fitting Bayesian survival models. *Lifetime Data Analysis* **9**, 5-19.
- Härkänen, T., Hausen, H., Virtanen, J.I., and Arjas, E. (2003). A non-parametric frailty model for temporally clustered multivariate failure times. *Scandinavian Journal of Statistics* **30**, 523-533.
- Henderson, R. and Shimakura, S. (2003). A serially correlated gamma frailty model for longitudinal count data. *Biometrika* **90**, 355-366.

- Ibrahim, J.G., Chen, M., and Sinha, D. (2001). Frailty models. In *Bayesian Survival Analysis*, 100-154. New York: Springer Verlag.
- Ishwaran, H. and James, L.F. (2004). Computational methods for multiplicative intensity models using weighted gamma processes: proportional hazards, marked point processes and panel count data. *Journal of the American Statistical Association* **99**, 175-190.
- Kalbfleisch, J.D. (1978). Nonparametric Bayesian analysis of survival time data. *Journal of the Royal Statistical Society, Series B* **40**, 214-221.
- Kleinman, K.P. and Ibrahim, J.G. (1998). A semiparametric Bayesian approach to the random effects model. *Biometrics* **54**, 921-938.
- Kokoska, S.M., Hardin, J.M., Grubbs, C.J., and Hsu, C. (1993). The statistical analysis of cancer inhibition/promotion experiments. *Anticancer Research* **13**, 1357-1364.
- Lam, K.F., Lee, Y.W., and Leung, T.L. (2002). Modeling multivariate survival data by a semiparametric random effects proportional odds model. *Biometrics* **58**, 316-323.
- Lavine, M. (1992). Some aspects of Pólya tree distributions for statistical modelling. *Annals of Statistics* **20**, 1161-1176.
- Lavine, M. (1994). More aspects of Pólya tree distributions for statistical modelling. *Annals of Statistics* **22**, 1222-1235.
- MacEachern, S.N. (1994). Estimating normal means with a conjugate style Dirichlet process prior. *Communications in Statistics - Simulation and Computation* **23**, 727-741.
- Mukhopadhyay, S. and Gelfand, A.E. (1997). Dirichlet process mixed generalized linear models. *Journal of the American Statistical Association* **92**, 633-639.
- Müller, P. and Rosner, G. (1997). A Bayesian population model with hierarchical mixture priors applied to blood count data. *Journal of the American Statistical Association* **92**, 1279-1292.
- Nieto-Barajas, L.E. and Walker, S.G. (2002). Markov beta and gamma processes for modelling hazard rates. *Scandinavian Journal of Statistics* **29**, 413-424.
- Paik, M.C., Tsai, W., and Ottman, R. (1994). Multivariate survival analysis using piecewise gamma frailty. *Biometrics* **50**, 975-988.

- Sahu, S.K., Dey, D.K., Aslanidou, H., and Sinha, D. (1997). A Weibull regression model with gamma frailties for multivariate survival data. *Lifetime Data Analysis* **3**, 123-137.
- Sargent, D.J. (1997). A flexible approach to time-varying coefficients in the Cox regression setting. *Lifetime Data Analysis* **3**, 13-25.
- Sargent, D.J. (1998). A general framework for random effects survival analysis in the Cox proportional hazards setting. *Biometrics* **54**, 1486-1497.
- Sinha, D. (1998). Posterior likelihood methods for multivariate survival data. *Biometrics* **54**, 1463-1474.
- Sinha, D., Ibrahim, J.G, and Chen, M.H. (2002). Models for survival data from cancer prevention studies. *Journal of the Royal Statistical Society, Series B* **64**, 467-477.
- Sinha, D. and Maiti, T. (2004). A Bayesian approach for the analysis of panel-count data with dependent termination. *Biometrics* **60**, 34-40.
- Tierney, L. (1994). Markov chains for exploring posterior distributions (with Discussions). *Annals of Statistics* **22**, 1701-1762.
- Vaupel, J.M., Manton, K.G., and Stallard, E. (1979). The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography* **16**, 439-454.
- Walker, S.G. and Mallick, B.K. (1997). Hierarchical generalized linear models and frailty models with Bayesian nonparametric mixing. *Journal of the Royal Statistical Society, Series B* **59**, 845-860.
- West, M., Müller, P., and Escobar, M.D. (1994). Hierarchical priors and mixture models with application in regression and density estimation. In *Aspects of Uncertainty: A Tribute to D.V. Lindley*, A. Smith and P. Freeman (eds), 363-386. New York: Wiley.
- Yau, K.K. and McGilchrist, C.A. (1998). ML and REML estimation in survival analysis with time dependent correlated frailty. *Statistics in Medicine* **17**, 1201-1213.
- Yue, H. and Chan, K.S. (1997). A dynamic frailty model for multivariate survival data. *Biometrics* **53**, 785-793.


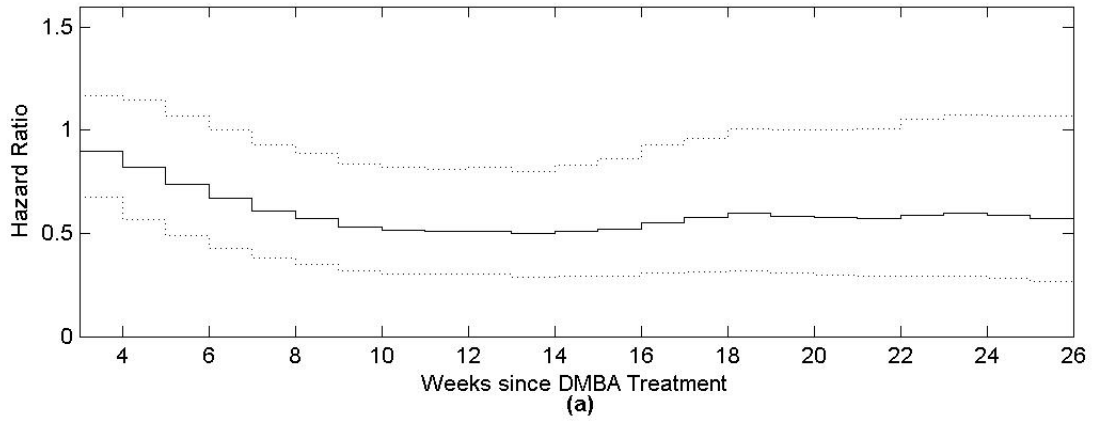


Figure 1. Posterior means (—) and 95% credible intervals (\cdots) for hazard ratios comparing mice fed (a) 1130 mg/kg and (b) 3390 mg/kg canthaxanthin to mice administered a vehicle control diet.

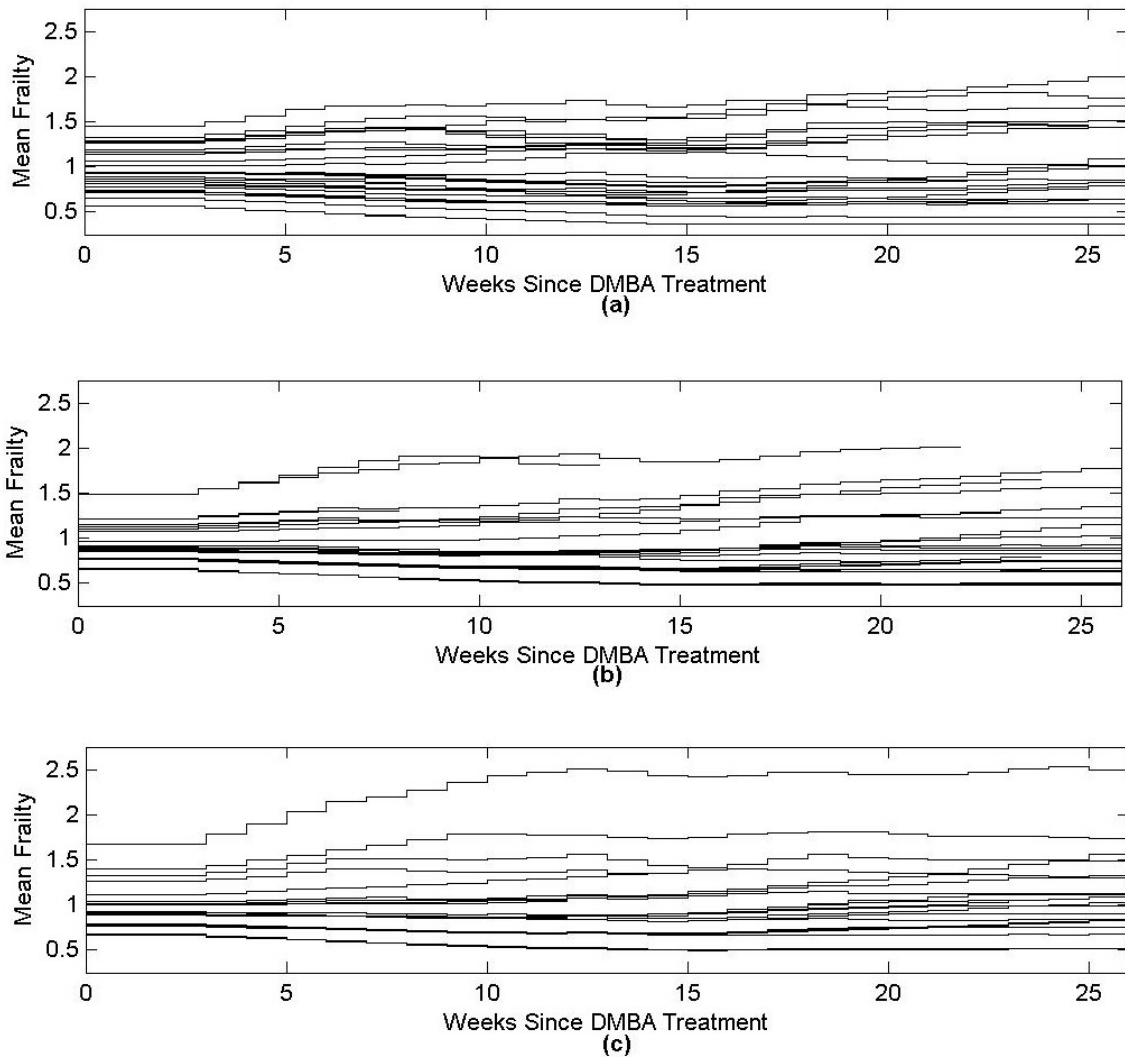


Figure 2. Posterior mean frailty trajectories for animals administered (a) vehicle, (b) 1130 mg/kg canthaxanthin, and (c) 3390 mg/kg canthaxanthin. Figure omits trajectories of animals 67 and 75 in the vehicle group due to their extreme values.

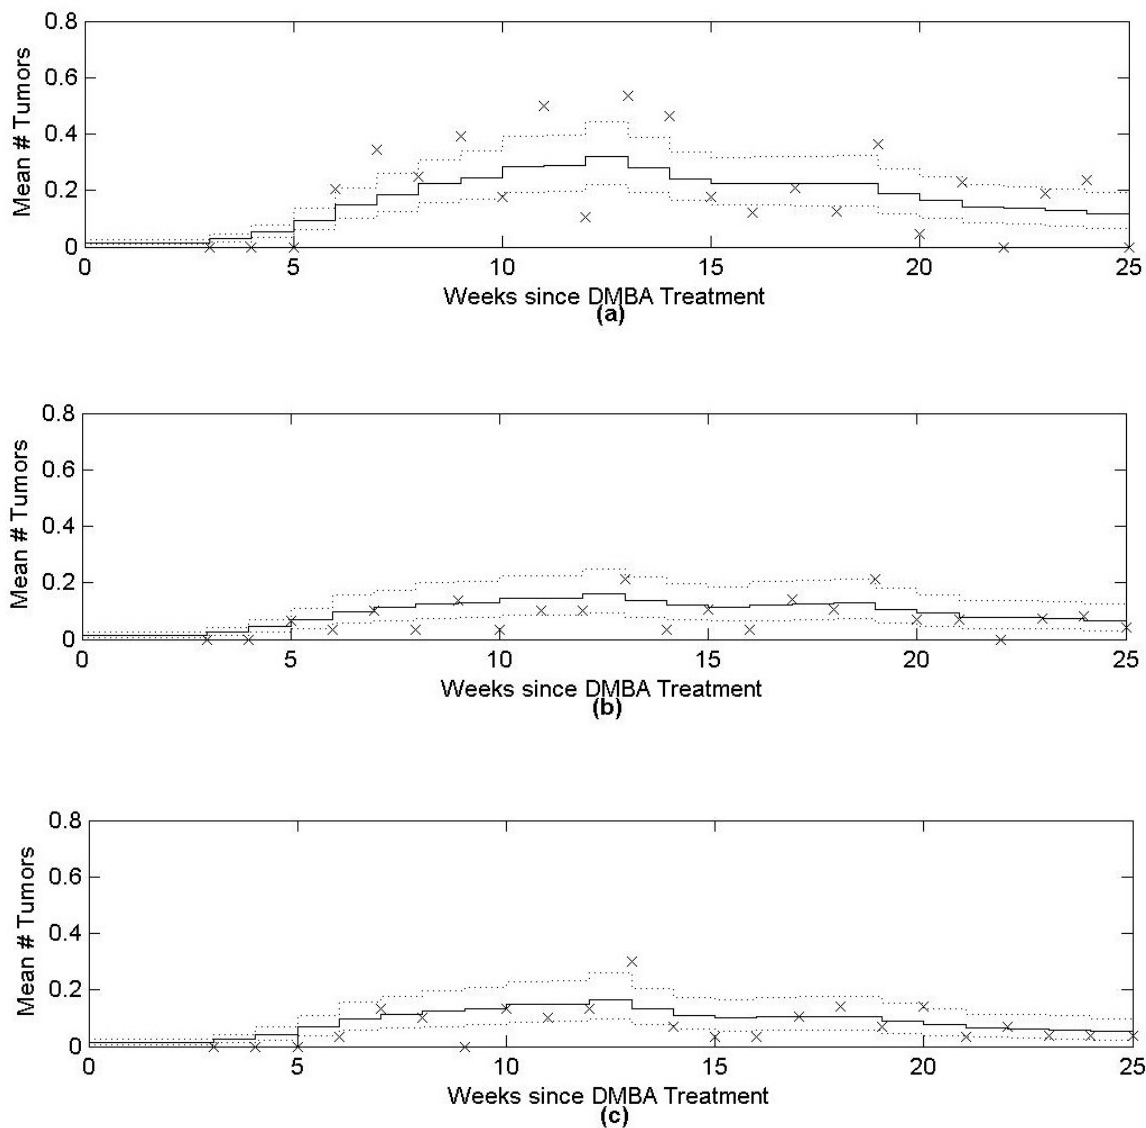


Figure 3. Observed (\times) and predicted weekly tumor incidence prior to sacrifice for mice administered **(a)** a vehicle, **(b)** 1130 mg/kg canthaxanthin, and **(c)** 3390 mg/kg canthaxanthin. The pointwise 95% credible intervals (\cdots) for the means ($—$) were calculated based on 1000 draws from the predictive distributions at each iteration of our MCMC algorithm.

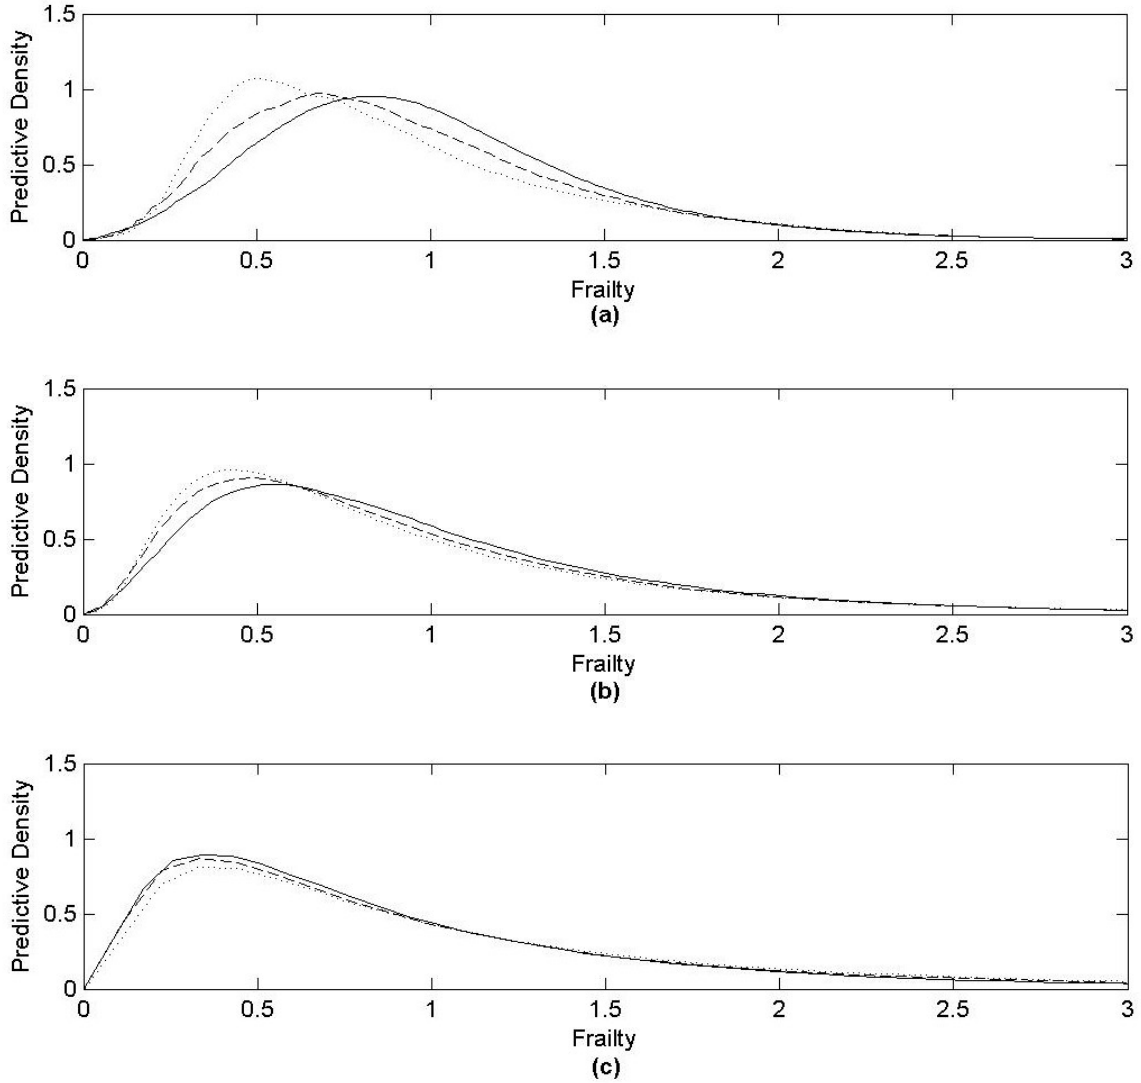


Figure 4. Comparison of predictive frailty distributions obtained using the dynamic gamma model (—), DPM1 (- - -), and DPM2 (\cdots). The densities of frailties from intervals **(a)** 1 (ϕ_{n+1}), **(b)** 11 ($\xi_{(n+1)11}$), and **(c)** 24 ($\xi_{(n+1)24}$) were approximated using a normal kernel smoother (width=0.05) applied to a posterior sample size of 500,000 (100 draws from the predictive distribution were taken at each iteration).

Table 1. Sensitivity analysis of parameter estimates (mean and 95% credible intervals) and posterior probabilities.

Parameter	Model			
	DPM1	DPM2	DPM3	Dynamic Gamma
$\gamma_{[10,26]1}^*$	0.489 (0.304, 0.758)	0.485 (0.304, 0.748)	0.420 (0.254, 0.669)	0.482 (0.305, 0.735)
$\gamma_{[10,26]2}^*$	0.553 (0.343, 0.863)	0.548 (0.341, 0.855)	0.480 (0.289, 0.758)	0.546 (0.345, 0.845)
ω	7.95 (5.34, 11.19)	7.87 (5.23, 11.26)	7.60 (5.01, 10.82)	8.01 (5.41, 11.26)
ψ_1	5.94 (2.03, 12.80)	5.99 (2.04, 12.97)	5.97 (2.03, 13.02)	5.26 (2.03, 12.36)
ψ_2	48.8 (32.4, 69.9)	48.0 (31.0, 68.9)	49.5 (32.2, 70.5)	50.7 (34.0, 70.8)
Probability				
P_{M1}	0.999	> 0.999	> 0.999	0.999
P_{M2}	0.998	0.999	> 0.999	0.998
P_{L1}	0.156	0.139	0.167	0.167
P_{L2}	0.392	0.360	0.399	0.394