

Combining Dynamic Predictions from Joint Models for Longitudinal and Time-to-Event Data using Bayesian Model Averaging

Dimitris Rizopoulos^{1,*}, Laura A. Hatfield², Bradley P. Carlin³ and Johanna J.M. Takkenberg⁴

¹ Department of Biostatistics, Erasmus University Medical Center

² Department of Health Care Policy, Harvard Medical School

³ Division of Biostatistics, School of Public Health, University of Minnesota

⁴ Department of Cardiothoracic Surgery, Erasmus University Medical Center

Abstract

The joint modeling of longitudinal and time-to-event data is an active area of statistics research that has received a lot of attention in the recent years. More recently, a new and attractive application of this type of models has been to obtain individualized predictions of survival probabilities and/or of future longitudinal responses. The advantageous feature of these predictions is that they are dynamically updated as extra longitudinal responses are collected for the subjects of interest, providing real time risk assessment using all recorded information. The aim of this paper is two-fold. First, to highlight the importance of modeling the association structure between the longitudinal and event time responses that can greatly influence the derived predictions, and second, to illustrate how we can improve the accuracy of the derived predictions by suitably combining joint models with different association structures. The second goal is achieved using Bayesian model averaging, which, in this setting, has the very intriguing feature that the model weights are not fixed but they are rather subject- and time-dependent, implying that at different follow-up times predictions for the same subject may be based on different models.

Keywords: Prognostic Modeling; Risk Prediction; Random Effects; Time-Dependent Covariates.

*Correspondance at: Department of Biostatistics, Erasmus University Medical Center, PO Box 2040, 3000 CA Rotterdam, the Netherlands. E-mail address: d.rizopoulos@erasmusmc.nl.

1 Introduction

In recent years it has been recognized that personalized medicine forms the future of medical care. This has increased interest in development of prognostic models for many different types of diseases. Examples are numerous and include, prognostic models for various types of cancer, such as breast and prostate cancer; risk scores for cardiovascular diseases, such as the Framingham score; and prognostic models applied in AIDS research to assess the risk of HIV infected patients. However, even though there is a wealth of patient data available, the majority of prognostic models in the literature provide risk predictions using only a small portion of the recorded information. This is especially true for patient outcomes that are repeatedly measured in time where typically only the last measurement is utilized. A clear advantage of such simple models is that they can be applied in everyday clinical practice. However, an important limitation is that valuable information is discarded, which if appropriately used, could offer a better insight into the dynamics of the disease's progression. In particular, an inherent characteristic of many medical conditions, such as those described above, is their dynamic nature. That is, the rate of progression is not only different from patient to patient but also dynamically changes in time for the same patient. Hence, it is medically relevant to investigate whether repeated measurements of a biomarker can provide a better understanding of disease progression and a better prediction of the risk for the event of interest than a single biomarker measurement (e.g., at baseline or the last available). It is evident that markers with this capability would become a valuable tool in everyday medical practice, because they would provide physicians with a better understanding of disease progression for a particular patient, and allow them to make more informed decisions. This is also the aim of our motivating case study on patients who underwent aortic valve allograft implantation, described in detail in Section 2. In particular, even though human tissue valves have some advantages, they are also more susceptible to tissue degeneration and require re-intervention. Hence, cardiologists wish to use an accurate prognostic tool that will inform them about the future prospect of a patient with a human tissue valve in order to adjust medical care and postpone re-operation and/or death.

Motivated by the Aortic Valve case study, this paper aims to provide a flexible modeling framework for producing risk predictions that utilize all available patient information. We will explicitly model the longitudinal history of each subject by basing our developments on the framework of joint models for longitudinal and time-to-event data (Faucett and Thomas 1996; Wulfsohn and Tsiatis 1997; Henderson et al. 2000; Tsiatis and Davidian 2004; Guo and Carlin 2004; Rizopoulos 2012). In particular, an attractive use of joint models is to derive dynamic predictions for either the survival or longitudinal outcomes (Yu et al. 2008; Proust-Lima and Taylor 2009; Rizopoulos 2011). The advantageous features of these predictions are that they are individualized, are updated as extra longitudinal information is recorded for each subject, and are based on all past values of the longitudinal outcome. Our novel contributions are two-fold. First, the subjects under consideration often exhibit complex longitudinal trajectories with nonlinearities and plateaus. In these settings, it is relevant to consider which characteristics of a subject’s trajectory best predict the event of interest. To this end, we will investigate how predictions are affected by assuming different types of association structure between the longitudinal and event time processes. We go beyond the standard formulation of joint models (Henderson et al. 2000), and we postulate functional forms that allow the rate of increase/decrease of the longitudinal outcome or a suitable summary of the whole longitudinal trajectory to determine the risk for an event. The consideration of competing association structures to describe the link between the two processes raises the issue of model uncertainty, which is typically ignored in prognostic modeling. Motivated by this, our second contribution is to derive predictions based not on a single model but on a collection of models simultaneously, combining them using Bayesian model averaging. As we will show later this approach bases predictions are based on the available data of a subject and the model weights are both individual- and time-dependent.

The rest of the paper is organized as follows. Section 2 gives a background on aortic allograft implantation and describes the Aortic Valve dataset that motivates this research. Section 3 briefly introduces the joint modeling framework, presents estimation under a Bayesian approach and shows how dynamic individualized predictions can be derived under

a joint model. Section 4 introduces several formulations of the association structure between the longitudinal and survival processes. Section 5 presents the Bayesian model averaging methodology to combine predictions, and Section 6 illustrates the use of this technique in the Aortic Valve dataset. Section 7 refers to the results of a simulation study, and Section 8 concludes the paper.

2 Background on the Aortic Valve Dataset

The motivation for this research comes from a study, conducted by the Department of Cardio-Thoracic Surgery of the Erasmus Medical Center in the Netherlands that includes 286 patients who received a human tissue valve in the aortic position in the hospital from 1987 until 2008 (Bekkers et al. 2011). Aortic allograft implantation has been widely used for a variety of aortic valve or aortic root diseases. Initial reports on the use of either fresh or cryopreserved allografts date from the early years of heart valve surgery. Major advantages ascribed to allografts are the excellent hemodynamic characteristics as a valve substitute; the low rate of thrombo-embolic complications, and, therefore, absence of the need for anticoagulant treatment; and the resistance to endocarditis. A major disadvantage of using human tissue valves, however is the susceptibility to (tissue) degeneration and need for re-interventions. The durability of a cryopreserved aortic allograft is age-dependent, leading to a high lifetime risk of re-operation, especially for young patients. Re-operations on the aortic root are complex, with substantial operative risks, and mortality rates in the range 4–12%.

In our study, a total of 77 (26.9%) patients received a sub-coronary implantation (SI) and the remaining 209 patients a root replacement (RR). These patients were followed prospectively over time with annual telephone interviews and biennial standardized echocardiographic assessment of valve function until July 8, 2010. Echo examinations were scheduled at 6 months and 1 year postoperatively, and biennially thereafter, and at each examination, echocardiographic measurements of aortic gradient (mmHg) were taken. By the end of

follow-up, 1241 aortic gradient measurements were recorded with an average of 5 measurements per patient (s.d. 2.3 measurements), 59 (20.6%) patients had died, and 73 (25.5%) patients required a re-operation on the allograft. Following the discussion in Section 1, our aim here is to use the existing data to construct a prognostic tool that will provide accurate risk predictions for future patients from the same population, utilizing their baseline information, namely age, gender and the type of operation they underwent, and their recorded aortic gradient levels. The composite event re-operation or death was observed for 125 (43.7%) patients, and the corresponding Kaplan-Meier estimator for the two intervention groups is shown in Figure 1. We can observe minimal differences in the re-operation-free survival rates

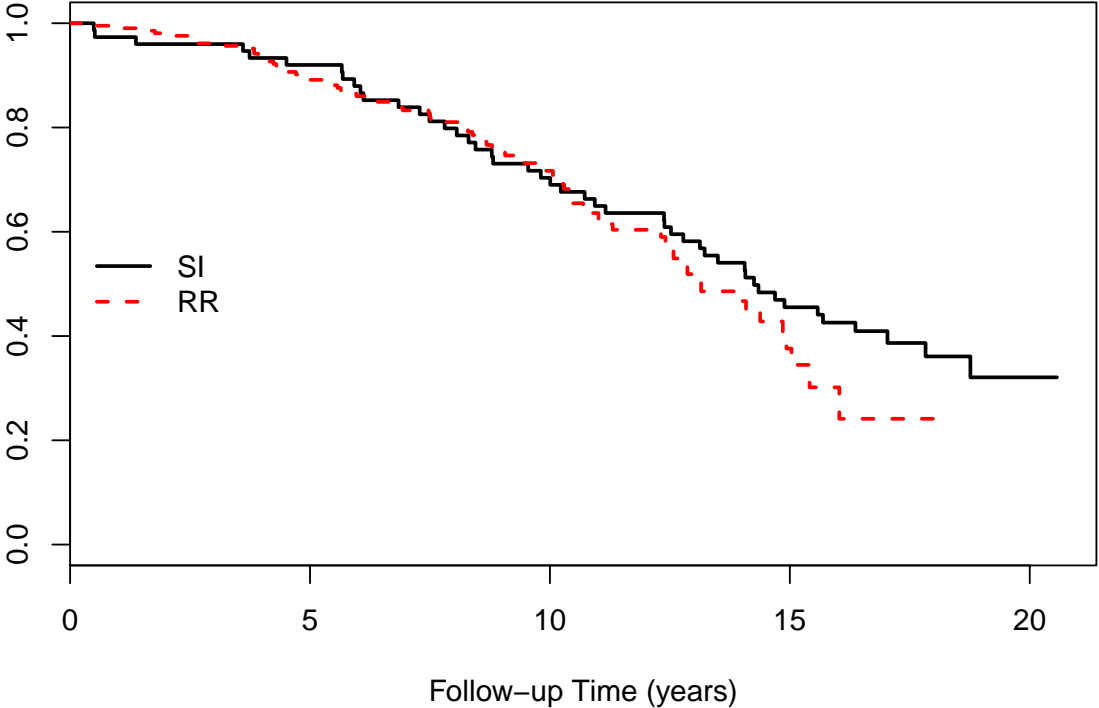


Figure 1: Kaplan-Meier estimates of the survival functions for re-operation-free survival for the sub-coronary implantation (SI) and root replacement (RR) groups.

between sub-coronary implantation and root replacement, with only a slight advantage of sub-coronary implantation towards the end of the follow-up. For the longitudinal outcome,

Figure 2 depicts the subject-specific longitudinal profiles for the two intervention groups. Because aortic gradient exhibits right skewness, we use the square root transform of aortic

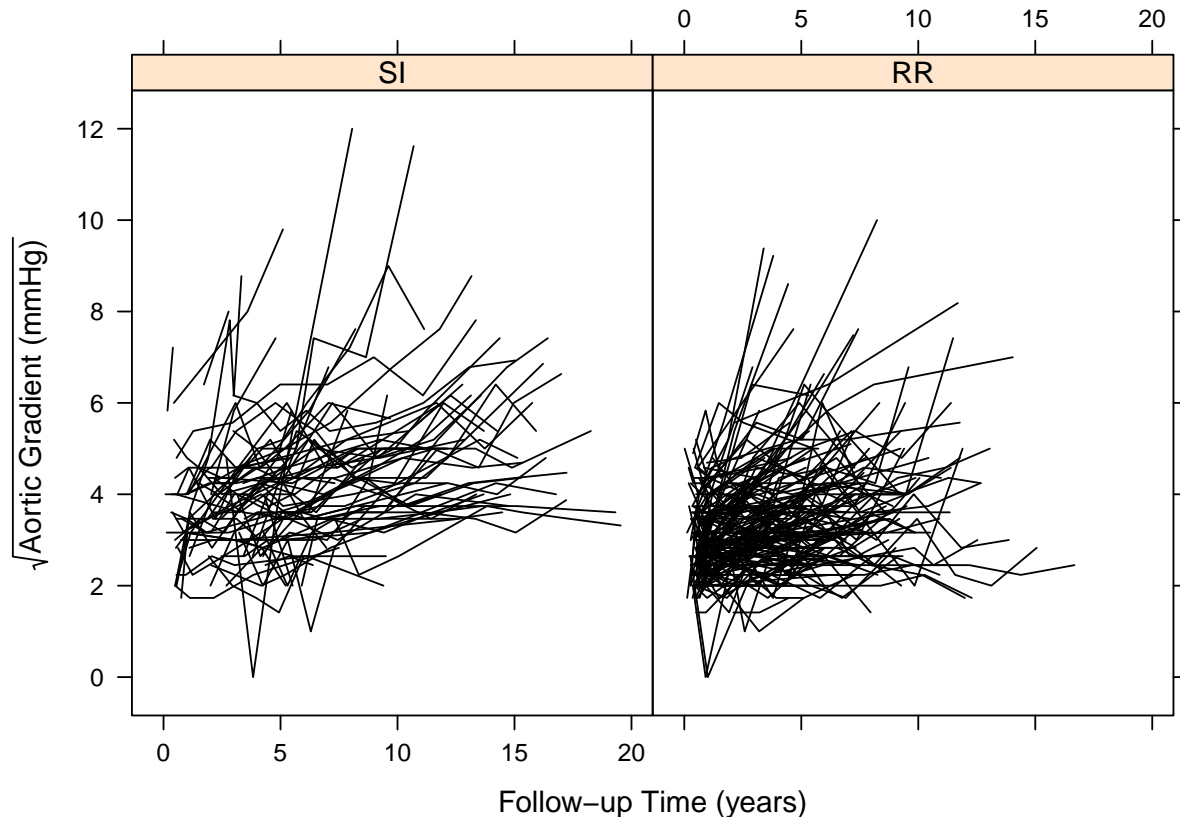


Figure 2: Subject-specific profiles for the square root aortic gradient separately for the sub-coronary implantation (SI) and root replacement (RR) groups.

gradient. We observe considerable variability in the shapes of these trajectories, but there are no systematic differences apparent between the two groups.

3 Joint Model Specification and Predictions

3.1 Submodels

Let T_i^* denote the true event time for the i -th subject ($i = 1, \dots, n$), C_i the censoring time, and $T_i = \min(T_i^*, C_i)$ the corresponding observed event time. In addition, we let $\delta_i = I(T_i^* \leq C_i)$ denote the event indicator, with $I(\cdot)$ being the indicator function that takes

the value 1 when $T_i^* \leq C_i$, and 0 otherwise. For the longitudinal process, we let \mathbf{y}_i denote the $n_i \times 1$ longitudinal response vector for the i -th subject, with element y_{il} denoting the value of the longitudinal outcome taken at time point t_{il} , $l = 1, \dots, n_i$. Our aim is to postulate a suitable joint model that associates these two processes.

Focusing on normally distributed longitudinal outcomes, we use a linear mixed-effects model to describe the subject-specific longitudinal trajectories. Namely, we have

$$\begin{aligned} y_i(t) &= m_i(t) + \varepsilon_i(t) = \mathbf{x}_i^\top(t)\boldsymbol{\beta} + \mathbf{z}_i^\top(t)\mathbf{b}_i + \varepsilon_i(t), \\ \mathbf{b}_i &\sim \mathcal{N}(\mathbf{0}, \mathbf{D}), \quad \varepsilon_i(t) \sim \mathcal{N}(0, \sigma^2), \end{aligned} \tag{1}$$

where $y_i(t)$ denotes the value of the longitudinal outcome at any particular time point t , $\mathbf{x}_i(t)$ and $\mathbf{z}_i(t)$ denote the time-dependent design vectors for the fixed-effects $\boldsymbol{\beta}$ and for the random effects \mathbf{b}_i , respectively, and $\varepsilon_i(t)$ the corresponding error terms that are assumed independent of the random effects, and $\text{cov}\{\varepsilon_i(t), \varepsilon_i(t')\} = 0$ for $t' \neq t$. For the survival process, we assume that the risk for an event depends on the true and unobserved value of the marker at time t , denoted by $m_i(t)$ in (1). More specifically, we have

$$\begin{aligned} h_i(t \mid \mathcal{M}_i(t), \mathbf{w}_i) &= \lim_{s \rightarrow 0} \Pr\{t \leq T_i^* < t + s \mid T_i^* \geq t, \mathcal{M}_i(t), \mathbf{w}_i\} / s \\ &= h_0(t) \exp\{\boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha m_i(t)\}, \quad t > 0, \end{aligned} \tag{2}$$

where $\mathcal{M}_i(t) = \{m_i(s), 0 \leq s < t\}$ denotes the history of the true unobserved longitudinal process up to t , $h_0(\cdot)$ denotes the baseline hazard function, and \mathbf{w}_i is a vector of baseline covariates with corresponding regression coefficients $\boldsymbol{\gamma}$. Parameter α quantifies the association between the true value of the marker at t and the hazard for an event at the same time point. To complete the specification of the survival process we need to make appropriate assumptions for the baseline hazard function $h_0(\cdot)$. Typically in survival analysis this is left unspecified. However, since our aim is to make subject-specific predictions of survival probabilities, we use a smoother assumption for this function. We achieve this, while still allowing for flexibility in the specification of $h_0(\cdot)$ by using a B-splines approach. In particular, the

log baseline hazard function is expressed as

$$\log h_0(t) = \gamma_{h_0,0} + \sum_{q=1}^Q \gamma_{h_0,q} B_q(t, \mathbf{v}), \quad (3)$$

where $B_q(t, \mathbf{v})$ denotes the q -th basis function of a B-spline with knots v_1, \dots, v_Q and γ_{h_0} the vector of spline coefficients. Increasing the number of knots Q increases the flexibility in approximating $\log h_0(\cdot)$; however, we should balance bias and variance and avoid overfitting. A standard rule of thumb is to keep the total number of parameters, including the parameters in the linear predictor in (2) and in the model for $h_0(\cdot)$, between 1/10 and 1/20 of the total number of events in the sample (Harrell 2001, Section 4.4). After the number of knots has been decided, their location is based on percentiles of the observed event times T_i or of the true event times $\{T_i : T_i^* \leq C_i, i = 1, \dots, n\}$ to allow for more flexibility in the region of greatest density.

3.2 Estimation

For the estimation of the joint model's parameters we follow the Bayesian paradigm and derive posterior inferences using a Markov chain Monte Carlo algorithm (MCMC). The likelihood of the models is derived under the assumption that the vector of time-independent random effects \mathbf{b}_i accounts for all interdependencies between the observed outcomes. That is, given the random effects, the longitudinal and event time process are assumed independent, and in addition, the longitudinal responses of each subject are assumed independent. Formally we have,

$$p(\mathbf{y}_i, T_i, \delta_i \mid \mathbf{b}_i, \boldsymbol{\theta}) = p(\mathbf{y}_i \mid \mathbf{b}_i, \boldsymbol{\theta}) p(T_i, \delta_i \mid \mathbf{b}_i, \boldsymbol{\theta}), \quad (4)$$

$$p(\mathbf{y}_i \mid \mathbf{b}_i, \boldsymbol{\theta}) = \prod_l p(y_{il} \mid \mathbf{b}_i, \boldsymbol{\theta}), \quad (5)$$

where $\boldsymbol{\theta}^\top = (\boldsymbol{\theta}_t^\top, \boldsymbol{\theta}_y^\top, \boldsymbol{\theta}_b^\top)$ denotes the full parameter vector, with $\boldsymbol{\theta}_t$ denoting the parameters for the event time outcome, $\boldsymbol{\theta}_y$ the parameters for the longitudinal outcomes, and $\boldsymbol{\theta}_b$ the

unique parameters of the random-effects covariance matrix, and $p(\cdot)$ denotes an appropriate probability density function. In addition, we assume that given the observed history of longitudinal responses up to time s , the censoring mechanism and the visiting process are independent of the true event times and future longitudinal measurements at $t > s$. Under these assumptions, the likelihood contribution for the i -th subject conditional on the parameters and random effects takes the form:

$$\begin{aligned}
p(\mathbf{y}_i, T_i, \delta_i \mid \boldsymbol{\theta}, \mathbf{b}_i) &= \prod_{l=1}^{n_i} p(y_{il} \mid \mathbf{b}_i; \boldsymbol{\theta}_y) p(T_i, \delta_i \mid \mathbf{b}_i; \boldsymbol{\theta}_t, \boldsymbol{\beta}) p(\mathbf{b}_i; \boldsymbol{\theta}_b) \\
&\propto \left[(\sigma^2)^{-n_i/2} \exp \left\{ - \sum_l (y_{il} - \mathbf{x}_{il}^\top \boldsymbol{\beta} - \mathbf{z}_{il}^\top \mathbf{b}_i)^2 / 2\sigma^2 \right\} \right. \\
&\times \left[\exp \left\{ \sum_q \gamma_{h_0,q} B_q(T_i, \mathbf{v}) + \boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha m_i(T_i) \right\} \right]^{\delta_i} \\
&\times \exp \left[- \exp(\boldsymbol{\gamma}^\top \mathbf{w}_i) \int_0^{T_i} \exp \left\{ \sum_q \gamma_{h_0,q} B_q(s, \mathbf{v}) + \alpha m_i(s) \right\} ds \right] \\
&\times \det(\mathbf{D})^{-1/2} \exp(-\mathbf{b}_i^\top \mathbf{D}^{-1} \mathbf{b}_i / 2) \Big],
\end{aligned} \tag{6}$$

where the intercept term $\gamma_{h_0,0}$ from the definition of the baseline risk function (3) has been incorporated into the design vector \mathbf{w}_i . The integral in the definition of the survival function

$$S_i(t \mid \mathcal{M}_i(t), \mathbf{w}_i) = \exp \left\{ - \int_0^t h_0(s) \exp \{ \boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha m_i(s) \} ds \right\}, \tag{7}$$

does not have a closed-form solution, and thus a numerical method must be employed for its evaluation. Here we use a 15-point Gauss-Kronrod quadrature rule. For the parameters $\boldsymbol{\theta}$ we take standard prior distributions. In particular, for the vector of fixed effects of the longitudinal submodel $\boldsymbol{\beta}$, for the regression parameters of the survival model $\boldsymbol{\gamma}$, for the vector of spline coefficients for the baseline hazard $\boldsymbol{\gamma}_{h_0}$, and for the association parameter α we use independent univariate diffuse normal priors. For the variance of the error terms σ^2 we take an inverse-Gamma prior, while for covariance matrices we assume an inverse Wishart prior.

3.3 Dynamic Individualized Predictions

Under the Bayesian specification of the joint model, presented in Section 3, we can derive subject-specific predictions for either the survival or longitudinal outcomes (Yu et al. 2008; Rizopoulos 2011, 2012). To put it more formally, based on a joint model fitted in a sample $\mathcal{D}_n = \{T_i, \delta_i, \mathbf{y}_i; i = 1, \dots, n\}$ from the target population, we are interested in deriving predictions for a new subject j from the same population that has provided a set of longitudinal measurements $\mathcal{Y}_j(t) = \{y_j(s); 0 \leq s \leq t\}$, and has a vector of baseline covariates \mathbf{w}_j . The fact that biomarker measurements have been recorded up to t , implies survival of this subject up to this time point, meaning that it is more relevant to focus on the conditional subject-specific predictions, given survival up to t . In particular, for any time $u > t$ we are interested in the probability that this new subject j will survive at least up to u , i.e.,

$$\pi_j(u | t) = \Pr(T_j^* \geq u | T_j^* > t, \mathcal{Y}_j(t), \mathbf{w}_j, \mathcal{D}_n).$$

Similarly, for the longitudinal outcome we are interested in the predicted longitudinal response at u , i.e.,

$$\omega_j(u | t) = E\{y_j(u) | T_j^* > t, \mathcal{Y}_j(t), \mathcal{D}_n\}.$$

The time-dynamic nature of both $\pi_j(u | t)$ and $\omega_j(u | t)$ is evident because when new information is recorded for patient j at time $t' > t$, we can update these predictions to obtain $\pi_j(u | t')$ and $\omega_j(u | t')$, and therefore proceed in a time dynamic manner.

Under the joint modeling framework of Section 3, estimation of either $\pi_j(u | t)$ or $\omega_j(u | t)$ is based on the corresponding posterior predictive distributions, namely

$$\pi_j(u | t) = \int \Pr(T_j^* \geq u | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathcal{D}_n) d\boldsymbol{\theta},$$

for the survival outcome, and analogously

$$\omega_j(u | t) = \int E\{y_j(u) | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}\} p(\boldsymbol{\theta} | \mathcal{D}_n) d\boldsymbol{\theta},$$

for the longitudinal one. The calculation of the first part of each integrand takes full advantage of the conditional independence assumptions (4) and (5). In particular, we observe that the first factor of the integrand of $\pi_j(u | t)$ can be rewritten by noting that:

$$\begin{aligned} \Pr(T_j^* \geq u | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) &= \int \Pr(T_j^* \geq u | T_j^* > t, \mathbf{b}_j, \boldsymbol{\theta}) p(\mathbf{b}_j | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) d\mathbf{b}_j \\ &= \int \frac{S_j\{u | \mathcal{M}_j(u, \mathbf{b}_j), \boldsymbol{\theta}\}}{S_j\{t | \mathcal{M}_j(t, \mathbf{b}_j), \boldsymbol{\theta}\}} p(\mathbf{b}_j | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) d\mathbf{b}_j, \end{aligned}$$

whereas for $\omega_j(u | t)$ we similarly have:

$$\begin{aligned} E\{y_j(u) | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}\} &= \int E\{y_j(u) | \mathbf{b}_j, \boldsymbol{\theta}\} p(\mathbf{b}_j | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) d\mathbf{b}_j \\ &= \mathbf{x}_j^\top(u) \boldsymbol{\beta} + \mathbf{z}_j^\top(u) \bar{\mathbf{b}}_j^{(t)}, \end{aligned}$$

with

$$\bar{\mathbf{b}}_j^{(t)} = \int \mathbf{b}_j p(\mathbf{b}_j | T_j^* > t, \mathcal{Y}_j(t), \boldsymbol{\theta}) d\mathbf{b}_j.$$

Combining these equations with the MCMC sample from the posterior distribution of the parameters for the original data \mathcal{D}_n , we can devise a simple simulation scheme to obtain Monte Carlo estimates of $\pi_j(u | t)$ and $\omega_j(u | t)$. More details can be found in [Yu et al. \(2008\)](#) and [Rizopoulos \(2011, 2012\)](#).

4 Functional Form

In ordinary proportional hazards models it has been long recognized that the functional form of time-varying covariates influences the derived inferences; see, for example, [Fisher and Lin \(1999\)](#) and references therein. In the joint modeling framework however, where the longitudinal outcome plays the role of a time-dependent covariate for the survival process, this topic has received much less attention. The two main functional forms that have been primarily used so far in joint models include in the linear predictor of the relative risk model (2) either the subject-specific means $m_i(t)$ from the longitudinal submodel or just the

random effects b_i (Henderson et al. 2000; Rizopoulos and Ghosh 2011). Nonetheless, in our setting, where our primary interest is in producing accurate predictions, we expect that the link between the longitudinal and event time processes to be important, and therefore it is relevant to investigate if and how the accuracy of predictions is influenced by the assumed functional form. This is motivated by the fact that there could be other characteristics of the patients’ longitudinal profiles that are more strongly predictive for the risk of an event, such as the rate of increase/decrease of the biomarker’s levels or a suitable summary of the whole longitudinal trajectory.

In this section, we present a few examples of alternative formulations for the association structure between the longitudinal outcome and the risk for an event. These different parameterizations can be seen as special cases of the following general formulation of the relative risk model:

$$h_i(t) = h_0(t) \exp[\boldsymbol{\gamma}^\top \mathbf{w}_i + \boldsymbol{\alpha}^\top f\{t, \mathbf{b}_i, \mathcal{M}_i(t)\}],$$

where $f(\cdot)$ is a time-dependent function that may depend on the random effects and on the true longitudinal history, as approximated by the mixed-effects model, and $\boldsymbol{\alpha}$ is a vector of association parameters.

4.1 Time-Dependent Slopes

The standard formulation (2) postulates that the risk for an event at time t is associated with parameter α measuring the strength of this association. Even though this is a very intuitively appealing parameterization with a clear interpretation for α , it cannot distinguish between patients who, for instance, at a specific time point have equal true marker levels, but they may differ in the rate of change of the marker, with one patient having an increasing trajectory and the other a decreasing one. An extension of (2) to capture such a setting has been considered by Ye et al. (2008) who posited a joint model in which the risk depends on both the current true value of the trajectory and the slope of the true trajectory at time t .

More specifically, the relative risk survival submodel takes the form,

$$h_i(t) = h_0(t) \exp\{\boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha_1 m_i(t) + \alpha_2 m'_i(t)\}, \quad (8)$$

where $m'_i(t) = d\{\mathbf{x}_i^\top(t)\boldsymbol{\beta} + \mathbf{z}_i^\top(t)\mathbf{b}_i\}/dt$. The interpretation of parameter α_1 remains the same as in the standard parameterization (2). Parameter α_2 measures the association between the value of the slope of the true longitudinal trajectory at time t and the risk for an event at the same time point, provided that $m_i(t)$ remains constant.

4.2 Cumulative Effects

A common characteristic of both (2) and (8) is that the risk for an event at any time t is assumed to be associated with features of the longitudinal trajectory at the same time point. However, in ordinary time-dependent Cox models several authors have argued that this assumption is over-simplistic, and in many real settings we may benefit from allowing the risk to depend on a more elaborate function of the history of the time-varying covariate (Sylvestre and Abrahamowicz 2009). Extending this concept in the context of joint models, we will account for the cumulative effect of the longitudinal outcome by including in the linear predictor of the relative risk submodel the integral of the longitudinal trajectory from baseline up to time t . More specifically, the survival submodel takes the form

$$h_i(t) = h_0(t) \exp\left\{\boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha \int_0^t m_i(s) ds\right\}, \quad (9)$$

where for any particular time point t , α measures the strength of the association between the risk for an event at time point t and the area under the longitudinal trajectory up to the same time t , with the area under the longitudinal trajectory taken as a suitable summary of the whole marker history $\mathcal{M}_i(t) = \{m_i(s), 0 \leq s < t\}$.

A feature of (9) is that it assigns the same weight on all past values of the longitudinal trajectory. This may not be reasonable when biomarker values closer to t are considered more relevant. An extension that allows placing different weights at different time points

is to multiply $m_i(t)$ with an appropriately chosen weight function $\varpi(\cdot)$ that places different weights at different time points, i.e.,

$$h_i(t) = h_0(t) \exp\left\{\boldsymbol{\gamma}^\top \mathbf{w}_i + \alpha \int_0^t \varpi(t-s)m_i(s) ds\right\}. \quad (10)$$

A possible family of functions with this property are probability density functions of known parametric distributions, such as the normal, the Student's- t and the logistic. The scale parameter in these densities and the degrees of freedom parameter in the Student's- t density can be used to tune the relative weights of more recent marker values compared to older ones.

4.3 Shared Random Effects

As mentioned earlier, one of the standard formulations of joint models includes in the linear predictor of the risk model only the random effects of the longitudinal submodel, that is,

$$h_i(t) = h_0(t) \exp(\boldsymbol{\gamma}^\top \mathbf{w}_i + \boldsymbol{\alpha}^\top \mathbf{b}_i), \quad (11)$$

where $\boldsymbol{\alpha}$ denotes a vector of association parameters each one measuring the association between the corresponding random effect and the hazard for an event. This parameterization is more meaningful when a simple random-intercepts and random-slopes structure is assumed for the longitudinal submodel, in which case the random effects express subject-specific deviations from the average intercept and average slope. Under this setting this parameterization postulates that patients who have a lower/higher level for the longitudinal outcome at baseline (i.e., intercept) or who show a steeper increase/decrease in their longitudinal trajectories (i.e., slope) are more likely to experience the event. In that respect, this formulation shares also similarities with the time-dependent slopes formulation (8).

A computational advantage of formulation (11) is that it is time-independent, and therefore leads to a closed-form solution (under certain baseline risk functions) for the integral in the definition of the survival function (7). This facilitates computations since we do

not have to numerically approximate this integral. However, an important disadvantage of (11) emerges when polynomials or splines are used to capture nonlinear subject-specific evolutions, in which case the random effects do not have a straightforward interpretation, complicating in turn the interpretation of α . Nonetheless, in our setting, we are primarily interested in predictions and not that much in interpretation, and thus we can consider (11) even under an elaborate mixed model.

5 Bayesian Model Averaging

The previous section demonstrated that there are several choices for the link between the longitudinal and event time outcomes. The common practice in prognostic modeling is to base predictions on a single model that has been selected based on an automatic algorithm, such as, backward, forward or stepwise selection, or on likelihood-based information criteria, such as, AIC, BIC, DIC and their variants. However, what is often neglected in this procedure is the issue of model uncertainty. For example, in the scenario that two models are correct, model selection forces us to choose one of the models even if we are not certain which model is true. In addition, with respect to predictions, there could be several competing models that could offer almost equally good predictions or even that some models may produce more accurate predictions for some subjects with specific longitudinal profiles, while other models may produce better predictions for subjects whose profiles have other features. Here we follow another approach and we explicitly take into account model uncertainty by combining predictions under different association structures using Bayesian model averaging (BMA) (Hoeting et al. 1999). We should stress that in our setting there is no concern in using BMA, because we do not average parameters that possibly have different interpretations under the different association structures, but rather predictions that maintain the same interpretation whatever the chosen functional form.

Due to space limitations, we will only focus here on dynamic BMA predictions of survival probabilities. BMA predictions for the longitudinal outcome can be produced with similar

methodology. Following the definitions of Section 3.3, we assume that we have available data $\mathcal{D}_n = \{T_i, \delta_i, \mathbf{y}_i; i = 1, \dots, n\}$ based on which we fit M_1, \dots, M_K joint models with different association structures. Interest is in calculating predictions for a new subject j from the same population who has provided a set of longitudinal measurements $\mathcal{Y}_j(t) = \{y_j(s); 0 \leq s \leq t\}$, and has a vector of baseline covariates \mathbf{w}_j . We let $\mathcal{D}_j(t) = \{\mathcal{Y}_j(t), T_j^* > t, \mathbf{w}_j\}$ denote the available data for this subject. The model-averaged probability of subject j surviving time $u > t$, given her survival up to t is given by the expression:

$$\Pr(T_j^* > u \mid \mathcal{D}_j(t), \mathcal{D}_n) = \sum_{k=1}^K \Pr(T_j^* > u \mid M_k, \mathcal{D}_j(t), \mathcal{D}_n) p(M_k \mid \mathcal{D}_j(t), \mathcal{D}_n). \quad (12)$$

The first term in the right-hand side of (12) denotes the model-specific survival probabilities, derived as in Section 3.3, and the second term denotes the posterior weights of each of the competing joint models. The unique characteristic of these weights is that they depend on the observed data of subject j , in contrast to classic applications of BMA where the model weights are the same for all subjects. This means that, in our case, the model weights are both subject- and time-dependent, and therefore, for different subjects, and even for the same subject but at different times points, different models may have higher posterior probabilities. Hence, this framework is capable of better tailoring predictions to each subject than standard prognostic models, because at any time point we base risk assessments on the models that are more probable to describe the association between the observed longitudinal trajectory of a subject and the risk for an event.

For the calculation of the model weights we observe that these are written as:

$$p(M_k \mid \mathcal{D}_j(t), \mathcal{D}_n) = \frac{p(\mathcal{D}_j(t) \mid M_k) p(\mathcal{D}_n \mid M_k) p(M_k)}{\sum_{\ell=1}^K p(\mathcal{D}_j(t) \mid M_\ell) p(\mathcal{D}_n \mid M_\ell) p(M_\ell)},$$

where

$$p(\mathcal{D}_j(t) \mid M_k) = \int p(\mathcal{D}_j(t) \mid \boldsymbol{\theta}_k) p(\boldsymbol{\theta}_k \mid M_k) d\boldsymbol{\theta}_k$$

and $p(\mathcal{D}_n \mid M_k)$ is defined analogously. The likelihood part $p(\mathcal{D}_n \mid \boldsymbol{\theta}_k)$ is given by (6),

whereas $p(\mathcal{D}_j(t) \mid \boldsymbol{\theta}_k)$ equals

$$p(\mathcal{D}_j(t) \mid \boldsymbol{\theta}_k) = p(\mathcal{Y}_j(t) \mid \mathbf{b}_i, \boldsymbol{\theta}_k) S_j(t \mid \mathbf{b}_i, \boldsymbol{\theta}_k) p(\mathbf{b}_j \mid \boldsymbol{\theta}_k).$$

Thus, the subject-specific information in the model weights at time t comes from the available longitudinal measurements $\mathcal{Y}_j(t)$ but also from the fact that this subject has survived up to t . Closed-form expressions for the marginal densities $p(\mathcal{D}_n \mid M_k)$ and $p(\mathcal{D}_j(t) \mid M_k)$ are obtained by means of Laplace approximations (Tierney and Kadane 1986) performed in two-steps, namely, first integrating out the random effects \mathbf{b}_i and then the parameters $\boldsymbol{\theta}_k$. The details are given in the supplementary material. Finally, a priori we assume that all models are equally probable, i.e., $p(M_k) = 1/K$, for all $k = 1, \dots, K$.

6 Analysis of the Aortic Valve Dataset

We return to the Aortic Valve dataset introduced in Section 2. Our aim here is to derive a prediction tool that will utilize all recorded information of a patient to provide accurate individualized predictions for both future aortic gradient levels and the risk of re-operation-free survival. Following the discussion of Sections 4 and 5, we will compare predictions under different association structures between the longitudinal and survival outcomes, with the BMA predictions that are based on all association structures simultaneously.

We start by defining the set of joint models from which predictions will be calculated. For the longitudinal process we allow a flexible specification of the subject-specific square root aortic gradient trajectories using natural cubic splines of time. More specifically, the linear mixed model takes the form

$$\begin{aligned} y_i(t) = & \beta_1 \mathbf{SI}_i + \beta_2 \mathbf{RR}_i + \beta_3 \{B_1(t, \lambda) \times \mathbf{SI}_i\} + \beta_4 \{B_1(t, \lambda) \times \mathbf{RR}_i\} \\ & + \beta_5 \{B_2(t, \lambda) \times \mathbf{SI}_i\} + \beta_6 \{B_2(t, \lambda) \times \mathbf{RR}_i\} \\ & + \beta_7 \{B_3(t, \lambda) \times \mathbf{SI}_i\} + \beta_8 \{B_3(t, \lambda) \times \mathbf{RR}_i\} \\ & + b_{i0} + b_{i1} B_1(t, \lambda) + b_{i2} B_2(t, \lambda) + b_{i3} B_3(t, \lambda) + \varepsilon_i(t), \end{aligned}$$

where $B_n(t, \lambda)$ denotes the B-spline basis for a natural cubic spline with boundary knots at baseline and 19 years and internal knots at 2.1 and 5.5 years (i.e., the 33.3% and 66.6% percentiles of observed follow-up times), **SI** and **RR** are the dummy variables for the sub-coronary implantation and root replacement groups, respectively, $\varepsilon_i(t) \sim \mathcal{N}(0, \sigma^2)$ and $\mathbf{b}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{D})$. For the survival process we consider five relative risk models, each positing a different association structure between the two processes, namely:

$$\begin{aligned}
M_1 : \quad & h_i(t) = h_0(t) \exp\{\gamma_1 \mathbf{RR}_i + \gamma_2 \mathbf{Age}_i + \gamma_3 \mathbf{Female}_i + \alpha_1 m_i(t)\}, \\
M_2 : \quad & h_i(t) = h_0(t) \exp\{\gamma_1 \mathbf{RR}_i + \gamma_2 \mathbf{Age}_i + \gamma_3 \mathbf{Female}_i + \alpha_1 m_i(t) + \alpha_2 m'_i(t)\}, \\
M_3 : \quad & h_i(t) = h_0(t) \exp\{\gamma_1 \mathbf{RR}_i + \gamma_2 \mathbf{Age}_i + \gamma_3 \mathbf{Female}_i + \alpha_1 \int_0^t m_i(s) ds\}, \\
M_4 : \quad & h_i(t) = h_0(t) \exp\{\gamma_1 \mathbf{RR}_i + \gamma_2 \mathbf{Age}_i + \gamma_3 \mathbf{Female}_i + \alpha_1 \int_0^t \phi(t-s) m_i(s) ds\}, \\
M_5 : \quad & h_i(t) = h_0(t) \exp(\gamma_1 \mathbf{RR}_i + \gamma_2 \mathbf{Age}_i + \gamma_3 \mathbf{Female}_i + \alpha_1 b_{i0} + \alpha_2 b_{i1} + \alpha_3 b_{i2} + \alpha_4 b_{i3}),
\end{aligned}$$

where the baseline hazard is approximated with splines, as described in Section 3.1, **Female** denotes the dummy variable for females, and $\phi(\cdot)$ denotes the probability density function of the standard normal distribution. We fitted each of these joint models using a single chain of 115000 MCMC iterations from which we discarded the first 15000 samples as burn-in. All computations have been performed in R (version 2.15.2) and JAGS (version 3.3.0). Trace and auto-correlations plots did not show any alarming indications of convergence failure. Tables 1 and 2 in the supplementary material show estimates and the corresponding 95% credible intervals for the parameters in the longitudinal and survival submodels, respectively. We observe that the parameter estimates in the relative risk models show much greater variability between the posited association structures than the parameters in the linear mixed models. However, we should note that the interpretation of the regression coefficients γ is not the same in all five models because we condition on different components of the longitudinal process.

We continue with the calculation of dynamic predictions based on these five joint models. To mimic the real-life use of a prognostic tool, we calculate predictions for three patients,

namely Patients 20, 22 and 81, who have been excluded from the dataset we used to fit the joint models. Patient 20 is a female of 64 years old and has received a sub-coronary implantation, Patient 22 is a male of 53 and has also received a sub-coronary implantation, and Patient 81 is a male of 39 and has undertaken a root replacement. The longitudinal trajectories of square root aortic gradient for these patients are illustrated in Figure 3, from which we see that Patients 20 and 22 show similar profiles with an initial drop in their aortic gradient levels after the operation and up to about five years followed by a stable increase for the next five years and a drop again. On the contrary, Patient 81 showed a steady increase of aortic gradient for the duration of his follow-up. For each one of those subjects we calculate dynamic predictions (i.e., a new prediction after each aortic gradient measurement has been recorded) for both the longitudinal and survival outcomes based on each of the five fitted joint models. In addition, following the methodology of Section 5, we also calculate BMA dynamic predictions for each patient by taking weighted averages of the predictions of the five joint models. We show the predictions of re-operation-free survival probabilities for Patients 20 and 81 in Figures 4 and 5, respectively. The predictions for Patient 22, and the predictions of future square root aortic gradient levels are presented in Figures 1–4 in the supplementary material. Similarly to the results in Tables 1 and 2 in the supplementary material, we observe that the predictions of survival probabilities seem to be more sensitive to the assumed association structure than the predictions of future square root aortic gradient levels. Much greater differences can be seen for Patient 81 who showed a steeper increase in his aortic gradient levels than the other two patients. Regarding the BMA predictions, it is interesting to note that they show some variability and they are not always dictated by a single joint model. In particular, Table 1 presents the time-dependent subject-specific BMA weights, from which some interesting observations can be made. First, only models M_2 and M_3 contribute in the BMA predictions of these three patients with the weights of the other three models being practically zero. Second, we indeed observe that the choice of the most appropriate model changes in time; for example, for Patient 20 and using only the first measurement we observe that there is little to choose between models M_2

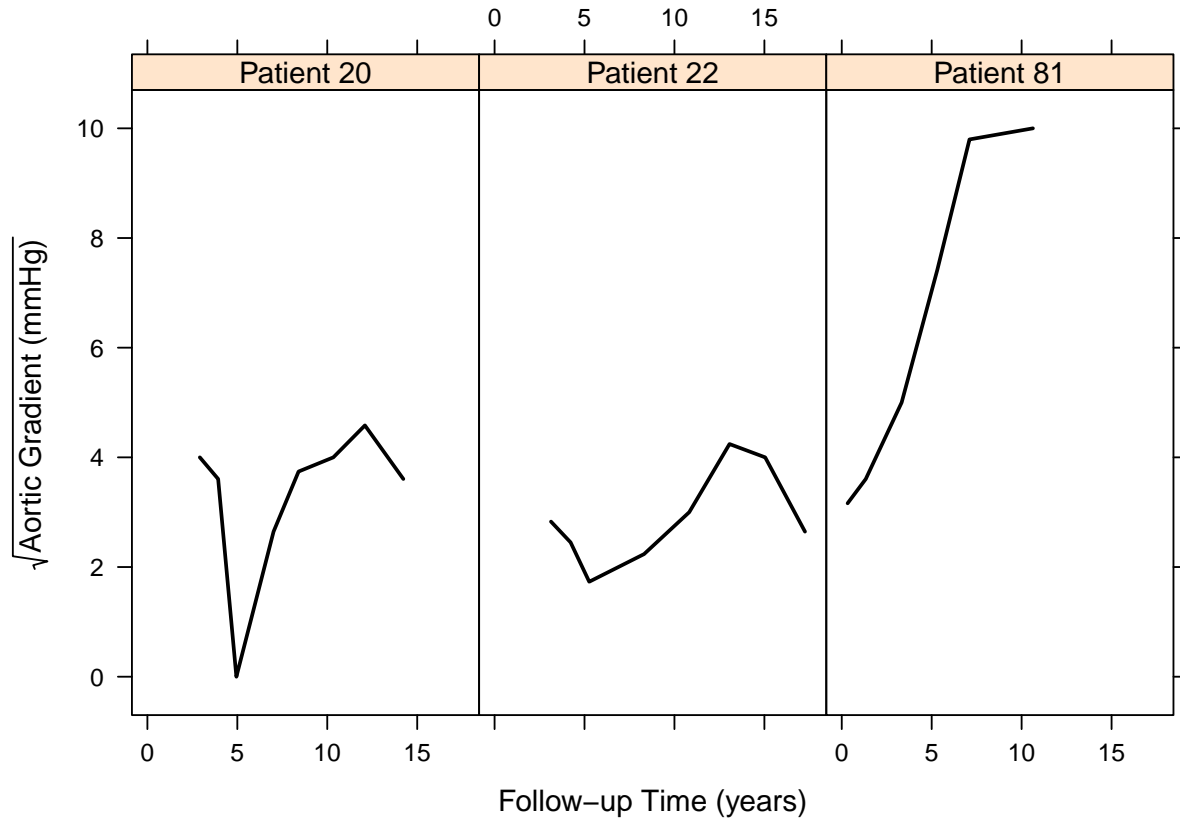


Figure 3: Longitudinal trajectories of square root aortic gradient for Patients 20, 22 and 81 who have been excluded from the analysis dataset, and for who we calculate predictions.

and M_3 , when the second measurement is observed then M_2 dominates, whereas when the third measurement is recorded prediction is based solely on model M_3 (though predictions based on the first two measurements are very similar from all five models, more variability is observed for latter time points). Similar behavior is observed for the other two patients as well. This convincingly demonstrates that it would be not optimal to base predictions on a single model for all patients and during the whole follow-up.

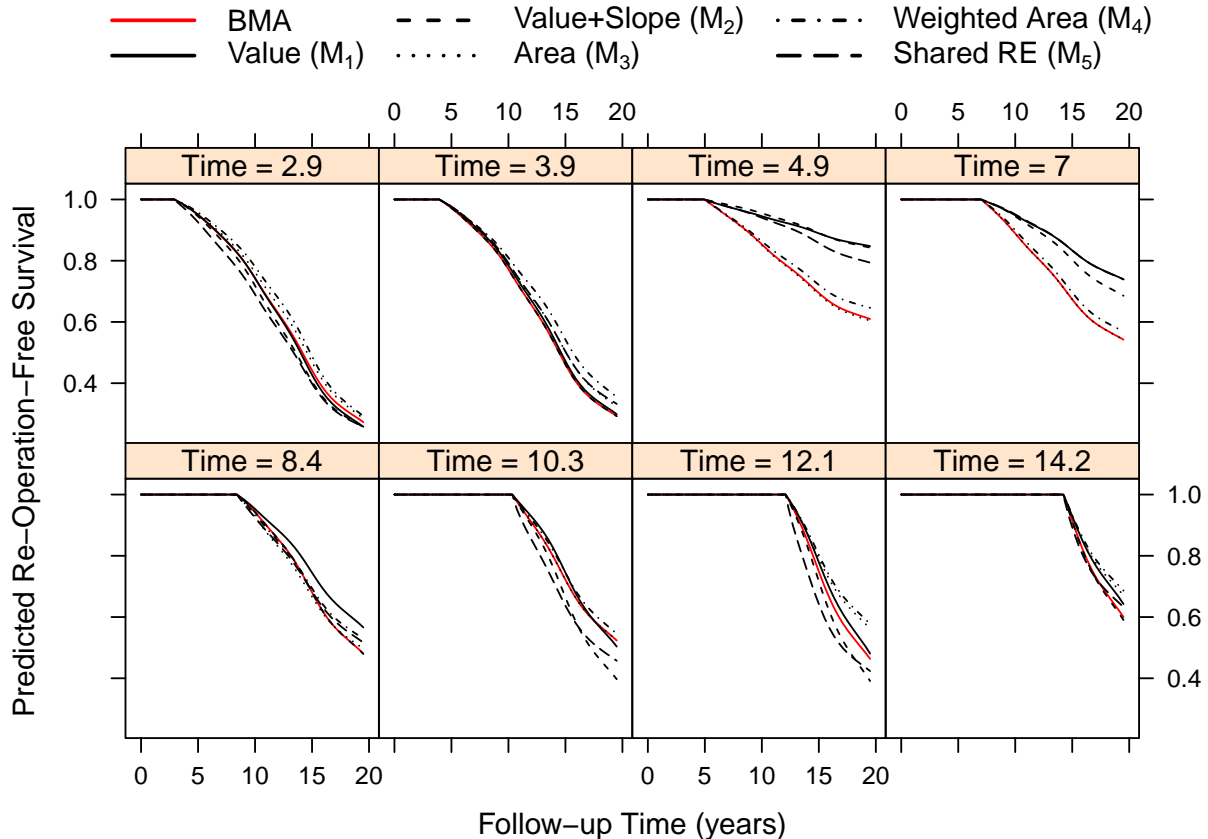


Figure 4: Dynamic predictions of re-operation-free survival for Patient 20 under the five joint models along with the BMA predictions. Each panel shows the corresponding conditional survival probabilities calculated after each of her longitudinal measurements have been recorded.

7 Simulations

7.1 Design

We performed a series of simulations to evaluate the finite sample performance of the BMA subject-specific predictions. The design of our simulation study is motivated by the set of joint models fitted to the Aortic Valve dataset in Section 6. In particular, we assume 300 patients who have been followed-up for a period of 19 years, and were planned to provide longitudinal measurements at baseline and afterwards at nine random follow-up times. For the longitudinal process, and similarly to the model fitted in the Aortic Valve dataset, we used natural cubic splines of time with two internal knots placed at 2.1 and 5.5 years, and

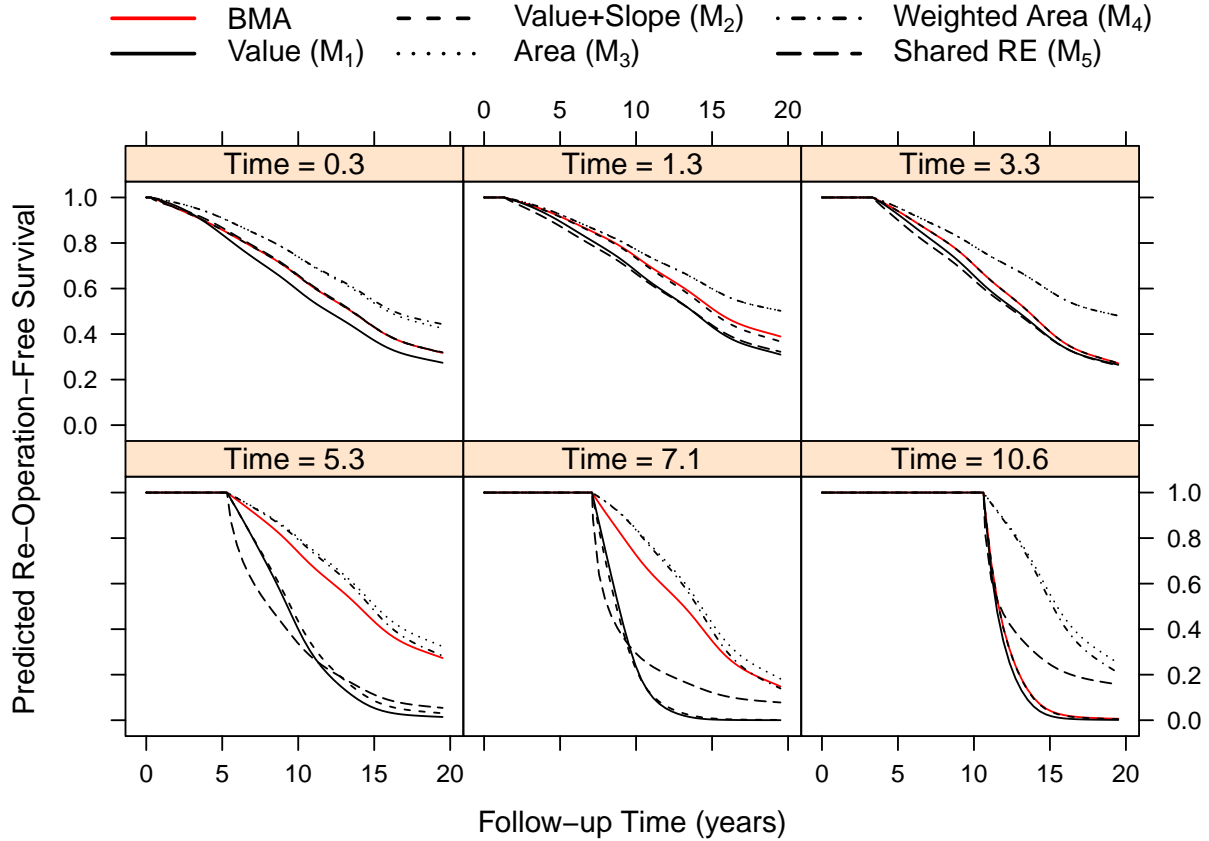


Figure 5: Dynamic predictions of re-operation-free survival for Patient 81 under the five joint models along with the BMA predictions. Each panel shows the corresponding conditional survival probabilities calculated after each of his longitudinal measurements have been recorded.

boundary knots placed at baseline and 19 years, i.e., the form of the model is as follows

$$\begin{aligned}
y_i(t) = & \beta_1 \text{Trt}0_i + \beta_2 \text{Trt}1_i + \beta_3 \{B_1(t, \lambda) \times \text{Trt}0_i\} + \beta_4 \{B_1(t, \lambda) \times \text{Trt}1_i\} \\
& + \beta_5 \{B_2(t, \lambda) \times \text{Trt}0_i\} + \beta_6 \{B_2(t, \lambda) \times \text{Trt}1_i\} \\
& + \beta_7 \{B_3(t, \lambda) \times \text{Trt}0_i\} + \beta_8 \{B_3(t, \lambda) \times \text{Trt}1_i\} \\
& + b_{i0} + b_{i1} B_1(t, \lambda) + b_{i2} B_2(t, \lambda) + b_{i3} B_3(t, \lambda) + \varepsilon_i(t),
\end{aligned}$$

where $B_n(t, \lambda)$ denotes the B-spline basis for a natural cubic spline, $\text{Trt}0$ and $\text{Trt}1$ are the dummy variables for the two treatment groups, $\varepsilon_i(t) \sim \mathcal{N}(0, \sigma^2)$ and $\mathbf{b}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{D})$ with \mathbf{D} taken to be diagonal.

Subject	Year	(Aort.Grad.) ^{1/2}	M ₁	M ₂	M ₃	M ₄	M ₅
20	2.9	4.0	0.00	0.43	0.57	0.00	0.00
	3.9	3.6	0.00	1.00	0.00	0.00	0.00
	4.9	0.0	0.00	0.02	0.98	0.00	0.00
	7.0	2.6	0.00	0.00	1.00	0.00	0.00
	8.4	3.7	0.00	0.99	0.01	0.00	0.00
	10.3	4.0	0.00	0.00	1.00	0.00	0.00
	12.1	4.6	0.00	0.57	0.43	0.00	0.00
	14.2	3.6	0.00	0.88	0.12	0.00	0.00
22	3.1	2.8	0.00	0.95	0.05	0.00	0.00
	4.2	2.4	0.00	1.00	0.00	0.00	0.00
	5.3	1.7	0.00	0.00	1.00	0.00	0.00
	8.3	2.2	0.00	1.00	0.00	0.00	0.00
	10.8	3.0	0.00	1.00	0.00	0.00	0.00
	13.1	4.2	0.00	0.01	0.99	0.00	0.00
	15.0	4.0	0.00	1.00	0.00	0.00	0.00
	17.3	2.6	0.00	0.67	0.33	0.00	0.00
81	0.3	3.2	0.00	1.00	0.00	0.00	0.00
	1.3	3.6	0.00	0.84	0.16	0.00	0.00
	3.3	5.0	0.00	1.00	0.00	0.00	0.00
	5.3	7.4	0.00	0.17	0.83	0.00	0.00
	7.1	9.8	0.00	0.19	0.81	0.00	0.00
	10.6	10.0	0.00	0.99	0.01	0.00	0.00

Table 1: BMA posterior weights for the five joint models for each subject and after each measurement.

For the survival process, we have assumed four scenarios, each one corresponding to a different functional form for the association structure between the two processes. More specifically,

$$\text{Scenario I: } h_i(t) = h_0(t) \exp\{\gamma_0 + \gamma_1 \text{Trt} \mathbf{1}_i + \alpha_1 m_i(t)\},$$

$$\text{Scenario II: } h_i(t) = h_0(t) \exp\{\gamma_0 + \gamma_1 \text{Trt} \mathbf{1}_i + \alpha_1 m_i(t) + \alpha_2 m'_i(t)\},$$

$$\text{Scenario III: } h_i(t) = h_0(t) \exp\{\gamma_0 + \gamma_1 \text{Trt} \mathbf{1}_i + \alpha_1 \int_0^t m_i(s) ds\},$$

$$\text{Scenario IV: } h_i(t) = h_0(t) \exp(\gamma_0 + \gamma_1 \text{Trt} \mathbf{1}_i + \alpha_1 b_{i0} + \alpha_2 b_{i1} + \alpha_3 b_{i2} + \alpha_4 b_{i3}),$$

with $h_0(t) = \sigma_i t^{\sigma_i - 1}$, i.e., the Weibull baseline hazard. The values for the regression coefficients in the longitudinal and survival submodels, the variance of the error terms of the

mixed model, the covariance matrix for the random effects, and the scale of the Weibull baseline risk function are given in Section 3 of the supplementary material. Censoring times were simulated from a uniform distribution in $(0, t_{max})$ with t_{max} set to result in about 45% censoring in each scenario. For each scenario we simulated 200 datasets.

7.2 Results

Motivated by the Aortic Valve dataset to produce accurate risk predictions, but also from previous experience regarding the robustness of predictions for the longitudinal outcome on the assumed association structure, we have focused on dynamic predictions for the survival outcome. More specifically, under each scenario and for each simulated dataset, we randomly excluded ten subjects whose event times were censored (it is more meaningful to calculate predictions for those individuals since they have not experienced the event yet), and in the remaining patients, we fitted five joint models. The longitudinal submodel was the same as the one we simulated from. For the survival process we assumed relative risk submodels with: the current value term (as in Scenario I), the current value and current slope terms (as in Scenario II), the cumulative effect (as in Scenario III), the weighted cumulative effect (10) with weight function the probability density function of the standard normal distribution, and only including the random effects (as in Scenario IV). Because the aim of this simulation study was to investigate how the different association structures affect predictions, we also assumed the same Weibull baseline hazard function from which we simulated the data and not the B-spline approximated baseline hazard (3) presented in Section 3.1.

Based on the five fitted joint models, we calculated predictions for each of the ten subjects we have originally excluded, at ten equidistant time points between their last available longitudinal measurement and the end of follow-up. These predictions were calculated under the true model for each scenario, BMA including all five models, and BMA based on four models without including the true one for the specific scenario. These predicted survival probabilities were then compared with the gold standard survival probabilities, calculated as $S_j\{u \mid \mathcal{M}_j(u, \mathbf{b}_j), \boldsymbol{\theta}\} / S_j\{t \mid \mathcal{M}_j(t, \mathbf{b}_j), \boldsymbol{\theta}\}$, using the true parameter values and the true

values of the random effects for the subjects we excluded. In each simulated dataset and for each of the ten subjects, we calculated root mean squared prediction errors (RMSEs) between the gold standard survival probabilities and the predictions under the three methods. The RMSEs over all the subjects from the 200 datasets are shown in Figure 6. For all scenarios

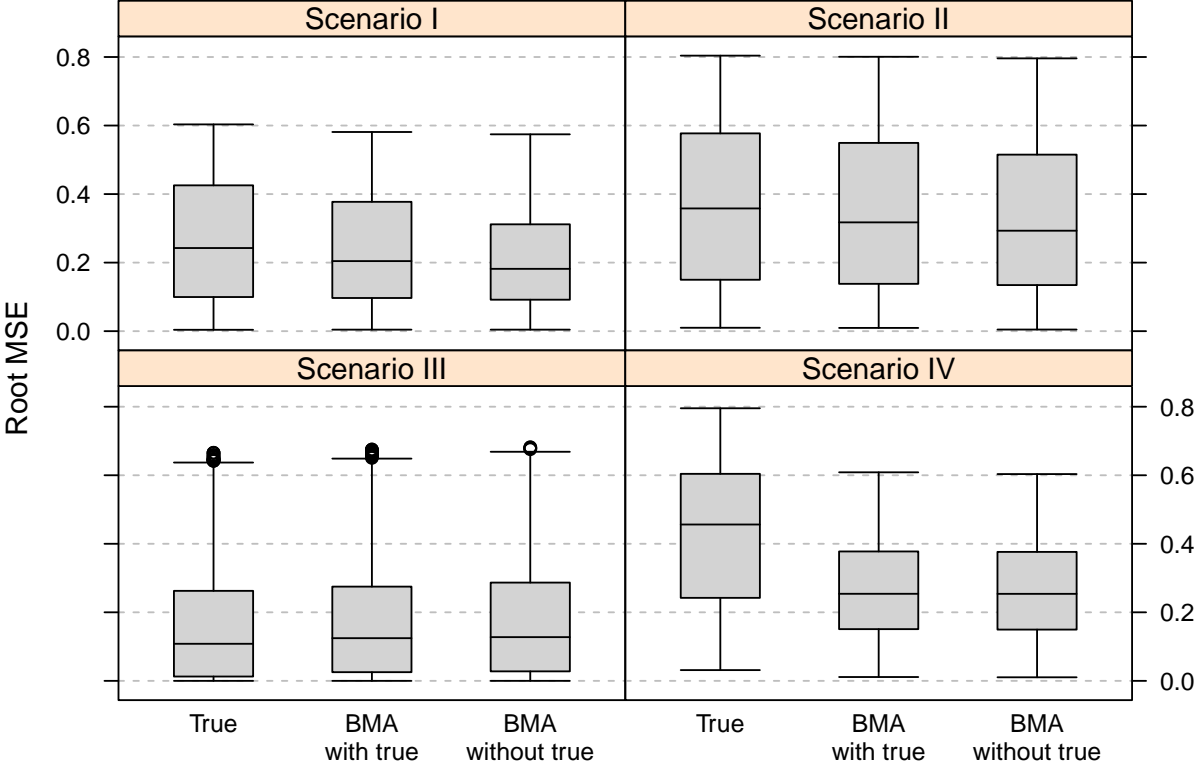


Figure 6: Simulation results under the four scenarios based on 200 datasets. Each boxplot shows the distribution of the root mean squared predictions error of the corresponding method to compute predictions (i.e., based on the true model, BMA including the true model and BMA without the true model) versus the gold standard.

we observe that in practice the BMA subject-specific predictions perform very well against the corresponding predictions from the true model. The greatest differences are observed in Scenario IV in which the BMA predictions seem to considerably outperform the true model predictions. A more careful examination of this scenario showed that this behavior was due to the fact that all models produced on average accurate predictions, but the predictions from the true model were much more variable than ones from the other models. In addition,

for the scenarios we considered in this study, it seems that BMA works equally well whether or not the true model is included in the list of the models that are averaged. This is a promising result because it suggests that BMA-based predictions could ‘protect’ against a misspecification of the association structure between the two processes.

8 Discussion

This paper illustrated how dynamic predictions from joint models with different association structures can be optimally combined using Bayesian model averaging. The novel feature of this approach is that the weights for combining predictions depend on the recorded information for the subject for whom predictions are of interest. Thus, for different subjects and even for the same subject but at different follow-up times, different models may have higher weights. This explicitly accounts for model uncertainty and acknowledges that a single prognostic model may not be adequate for quantifying the risk of all patients. Our simulation study showed that BMA predictions perform very well in comparison with predictions from the true model, even if the true model is not included in the list of models that are averaged. This gives us more confidence in trusting BMA for deriving predictions for future patients from the Aortic Valve study population based on the five joint models we have considered.

In our developments we have considered a simple joint model for a single longitudinal outcome and one time-to-event. However, often in longitudinal studies several outcomes are recorded on each patient during follow-up. For example, for the patients enrolled in the Aortic Valve study also aortic regurgitation was recorded, which is another measure of valve function. Hence it would be of interest to investigate whether by considering both longitudinal biomarkers we could improve the accuracy of predictions. In addition, in our analysis we have considered the composite event re-operation or death (whatever comes first), but for the treating physicians it could be of interest to have separate risk estimates for the two events. Based on recent advances in joint modeling that include joint models for multiple markers and multiple event times ([Huang et al. 2011](#); [Rizopoulos and Ghosh 2011](#);

Liu and Huang 2009), our ideas could be relatively easily extended to these more elaborate cases. The challenge in such settings will be the high dimensionality of the model space. This is because the number of possible combinations of association structures between the longitudinal and the survival processes grows exponentially with the number of outcomes.

A topic that we have partially addressed in this paper, but which is of high relevance for the practicality of prognostic models concerns the external validation of the derived predictions in terms of both discrimination and calibration. In particular, from our simulations we saw that the BMA predictions perform satisfactorily compared to the predictions using the true model, but this does not answer the question of how accurately the longitudinal outcome can predict the survival one or if it can discriminate between subjects of low and high risk. In standard survival analysis there has been a lot of research at these two fronts (see e.g., Zheng and Heagerty 2007; Gerds and Schumacher 2006, and references therein), but within the joint modeling framework relatively little work has been done (Rizopoulos 2011; Proust-Lima and Taylor 2009). Theoretically, all standard measures for calibration (e.g., Brier score) and discrimination (time-dependent sensitivity, specificity and ROC curves) can be defined for joint models, but their estimation is more challenging.

9 Supplementary Material

Supplementary material are available in `SuppPredParam.pdf`, and include **Section 1**: Figures and Tables with results from the analysis of the Aortic Valve dataset; **Section 2**: Details on the Laplace approximations to calculate $p(\mathcal{D}_n | M_k)$ and $p(\mathcal{D}_j(t) | M_k)$; **Section 3**: Details on the simulation settings.

References

Bekkers, J., Klieverik, L., Raap, G., Takkenberg, J., and Bogers, A. (2011), “Re-operations for aortic allograft root failure: Experience from a 21-year single-center prospective follow-up study,” *European Journal of Cardio-Thoracic Surgery*, 40, 35–42.

- Faucett, C. and Thomas, D. (1996), “Simultaneously modelling censored survival data and repeatedly measured covariates: A Gibbs sampling approach,” *Statistics in Medicine*, 15, 1663–1685.
- Fisher, L. and Lin, D.-Y. (1999), “Time-dependent covariates in the Cox proportional-hazards regression model,” *Annual Review of Public Health*, 20, 145–157.
- Gerds, T. and Schumacher, M. (2006), “Consistent estimation of the expected Brier score in general survival models with right-censored event times,” *Biometrical Journal*, 48, 1029–1040.
- Guo, X. and Carlin, B. (2004), “Separate and joint modeling of longitudinal and event time data using standard computer packages,” *The American Statistician*, 58, 16–24.
- Harrell, F. (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, New York: Springer-Verlag.
- Henderson, R., Diggle, P., and Dobson, A. (2000), “Joint modelling of longitudinal measurements and event time data,” *Biostatistics*, 1, 465–480.
- Hoeting, J., Madigan, D., Raftery, A., and Volinsky, C. (1999), “Bayesian model averaging: A tutorial,” *Statistical Science*, 14, 382–417.
- Huang, X., Li, G., Elashoff, R., and Pan, J. (2011), “A general joint model for longitudinal measurements and competing risks survival data with heterogeneous random effects,” *Lifetime Data Analysis*, 17, 80–100.
- Liu, L. and Huang, X. (2009), “Joint analysis of correlated repeated measures and recurrent events processes in the presence of death, with application to a study on acquired immune deficiency syndrome,” *Journal of the Royal Statistical Society, Series C*, 58, 65–81.
- Proust-Lima, C. and Taylor, J. (2009), “Development and validation of a dynamic prognostic tool for prostate cancer recurrence using repeated measures of posttreatment PSA: a joint modeling approach,” *Biostatistics*, 10, 535–549.

- Rizopoulos, D. (2011), “Dynamic predictions and prospective accuracy in joint models for longitudinal and time-to-event data,” *Biometrics*, 67, 819–829.
- (2012), *Joint Models for Longitudinal and Time-to-Event Data, with Applications in R*, Boca Raton: Chapman & Hall/CRC.
- Rizopoulos, D. and Ghosh, P. (2011), “A Bayesian semiparametric multivariate joint model for multiple longitudinal outcomes and a time-to-event,” *Statistics in Medicine*, 30, 1366–1380.
- Sylvestre, M.-P. and Abrahamowicz, M. (2009), “Flexible modeling of the cumulative effects of time-dependent exposures on the hazard,” *Statistics in Medicine*, 28, 3437–3453.
- Tierney, L. and Kadane, J. (1986), “Accurate approximations for posterior moments and marginal densities,” *Journal of the American Statistical Association*, 81, 82–86.
- Tsiatis, A. and Davidian, M. (2004), “Joint modeling of longitudinal and time-to-event data: An overview,” *Statistica Sinica*, 14, 809–834.
- Wulfsohn, M. and Tsiatis, A. (1997), “A joint model for survival and longitudinal data measured with error,” *Biometrics*, 53, 330–339.
- Ye, W., Lin, X., and Taylor, J. (2008), “Semiparametric modeling of longitudinal measurements and time-to-event data – a two stage regression calibration approach,” *Biometrics*, 64, 1238–1246.
- Yu, M., Taylor, J., and Sandler, H. (2008), “Individualized prediction in prostate cancer studies using a joint longitudinal-survival-cure model,” *Journal of the American Statistical Association*, 103, 178–187.
- Zheng, Y. and Heagerty, P. (2007), “Prospective Accuracy for Longitudinal Markers,” *Biometrics*, 63, 332–341.