

## Flexible Bayesian additive joint models with an application to type 1 diabetes research

Meike Köhler <sup>\*,1</sup>, Nikolaus Umlauf <sup>2</sup>, Andreas Beyerlein <sup>1</sup>, Christiane Winkler <sup>1</sup>, Anette-Gabriele Ziegler <sup>1,3</sup>, and Sonja Greven <sup>4</sup>

<sup>1</sup> Institute of Diabetes Research, Helmholtz Zentrum München, and Forschergruppe Diabetes, Klinikum rechts der Isar, Technische Universität München, Ingolstädter Landstraße 1, D-85764 Neuherberg, Germany

<sup>2</sup> Department of Statistics, Faculty of Economics and Statistics Universität Innsbruck, Universitätsstraße 15, A-6020 Innsbruck, Austria

<sup>3</sup> Forschergruppe Diabetes e.V. at the Helmholtz Zentrum München, Ingolstädter Landstraße 1, D-85764 Neuherberg, Germany

<sup>4</sup> Department of Statistics, Ludwig-Maximilians-Universität München, Ludwigstraße 33, D-80539 München, Germany

Received zzz, revised zzz, accepted zzz

The joint modeling of longitudinal and time-to-event data is an important tool of growing popularity to gain insights into the association between a biomarker and an event process. We develop a general framework of flexible additive joint models that allows the specification of a variety of effects, such as smooth nonlinear, time-varying and random effects, in the longitudinal and survival parts of the models. Our extensions are motivated by the investigation of the relationship between fluctuating disease-specific markers, in this case autoantibodies, and the progression to the autoimmune disease type 1 diabetes. By making use of Bayesian P-splines we are in particular able to capture highly nonlinear subject-specific marker trajectories as well as a time-varying association between the marker and the event process allowing new insights into disease progression. The model is estimated within a Bayesian framework and implemented in the R-package `bamlss`.

**Key words:** Anisotropic smoothing; Biomarkers; Longitudinal data; Time-to-event data; P-Splines;

Supporting Information for this article is available under <http://dx.doi.org/10.1022/bimj.XXXXXXX>

### 1 Introduction

The joint modeling of longitudinal biomarkers and the time to disease onset or death offers unique insights into disease progression in various medical domains (Taylor *et al.*, 2013; Gras *et al.*, 2013; Daher Abdi *et al.*, 2013). Depending on the disease and the respective biomarker different challenges have to be faced in joint modeling. In the following, a general framework for the flexible joint modeling of longitudinal data and time-to-event is presented, which was motivated by unique cohort data from studies exploring the development of type 1 diabetes (T1D). The research on T1D underwent a paradigm shift in the past decade, when disease-specific autoantibodies were shown to be diagnostic for the disease before the onset of clinical symptoms and thus paving the way for a pre-clinical diagnosis of T1D (Ziegler *et al.*, 2013; Bonifacio, 2015; Insel *et al.*, 2015). Prior to the onset of clinical symptoms, i.e. the need of insulin substitution, the disease is already progressing and insulin-producing beta-cells in the pancreas are gradually destroyed by

\*Corresponding author: e-mail: meike.koehler@helmholtz-muenchen.de, Phone: +49-(0)89-3068-2917, Fax: +49-(0)89-3187-3144

the body's own immune system. This immune process, leading to an onset of clinical symptoms within months up to more than a decade, can be diagnosed by the emergence of T1D-specific autoantibodies. However, it remains an open question whether the longitudinal patterns of these autoantibodies might be associated with the rate of progression to T1D.

In recent years joint models gained larger popularity in the modeling of associations between time-varying biomarkers and time-to-event. By estimating a submodel for a longitudinal biomarker, usually a mixed model, jointly with the survival submodel of a time-to-event process, one can account for the informative censoring and the within-subject errors in the longitudinal model and can incorporate the longitudinal information, observed only at person-specific discrete timepoints, as a continuous-time covariate in the survival model. Comprehensive overviews on the topic are given in [Tsiatis and Davidian \(2004\)](#), [Rizopoulos \(2012\)](#) and [Gould \*et al.\* \(2015\)](#). In our work we focus on extensions of so-called shared parameter models. These assume that a set of parameters influences both the longitudinal and the survival model, and that there is conditional independence given those parameters.

In T1D research little is known concerning typical trajectories of autoantibodies as biomarkers. At the same time the observed trajectories show highly nonlinear patterns over time and differ strongly between subjects, see [Figure 1a](#). In consequence, a flexible specification of individual trajectories in the longitudinal model is needed in our application.

Much work on joint models has focused on simple parametric longitudinal trajectories, while only few approaches allow for more flexible, potentially non-parametric longitudinal models. [Ding and Wang \(2008\)](#) model mean trajectories by B-splines and allow for one multiplicative random effect per subject. For our application however it remains questionable if such a model is flexible enough to capture the highly different trajectories. Spline based approaches, that allow also the random effects to be non-linear functions in time, are mentioned by [Song and Wang \(2008\)](#) and were employed by [Rizopoulos and Ghosh \(2011\)](#) and [Rizopoulos \*et al.\* \(2014\)](#) as well as [Brown \*et al.\* \(2005\)](#) and [Brown \(2009\)](#). While allowing for flexibility, a disadvantage of all these approaches is finding an optimal number of knots to specify the flexible longitudinal model, e.g. by AIC or DIC. As the number of random effects increases with the number of knots, this number is limited in practice. We aim to avoid the explicit choice of knots and number of basis functions by using a penalized spline approach, where a larger number of knots is specified and smoothness penalties are employed ([Lang and Brezger, 2004](#)). [Tang and Tang \(2015\)](#) also make use of P-Splines in modeling longitudinal trajectories, but do so only in estimating the mean function, whereas we model also the individual trajectories as smooth functions of time. This is similar in spirit to the specification of individual trajectories in [Jiang \*et al.\* \(2015\)](#), however we do not assume an underlying class membership for the random effects.

The estimation of joint models with complex subject-specific trajectories poses a challenge to frequentist estimation approaches due to the necessary integration over potentially high-dimensional random effects distributions. Due to this drawback and further advantages of the Bayesian approach in joint modeling, such as straightforward model assessment and the potential integration of previous knowledge via priors ([Gould \*et al.\*, 2015](#)), many complex joint models, like e.g. the aforementioned models, are specified within a Bayesian framework. The most widely used sampling approach for the parameter distributions in Bayesian joint models is Gibbs Sampling, e.g. [Faucett and Thomas \(1996\)](#); [Guo and Carlin \(2004\)](#); [Brown and Ibrahim \(2003\)](#), also in conjunction with Metropolis-Hastings algorithms ([Tang and Tang, 2015](#)). In addition, the well-established R-package `JMbayes` ([Rizopoulos, 2016a,b](#)) implementing [Rizopoulos \*et al.\* \(2014\)](#) employs a random walk Metropolis-Hastings algorithm. Our Bayesian estimation approach is different as we employ a derivative-based Metropolis-Hastings algorithm, where we draw samples from approximations of the full conditionals using score vectors and Hessians of the parameters. Despite being computationally demanding this algorithm shows a high stability in the model estimation, as we also show in our simulations.

In addition to the need for a flexible longitudinal model, a further generalization of existing joint models seems necessary in our application, namely a time-varying association between the biomarker and the time-to-event. Here, the biomarker indicates an ongoing immune process eventually leading to the destruction

of the insulin-producing beta cells. As the activity of the immune system is constantly regulated, it is plausible that the association between a biomarker and the hazard of T1D varies over time. For example a recent paper by Meyer *et al.* (2016) indicated that patients with an autoimmune disease can also present unique disease-ameliorating autoantibodies. Such a time-varying association has rarely been studied in the context of joint models. Using a discretized time-scale and a probit model for the discrete hazard function, Barrett *et al.* (2015) allow for the association to vary over the discrete time points in their model. However, this flexible specification is not considered in their simulations, the applied examples or the code provided to fit the models. A time-varying coefficient to associate the marker and the event process is the focus of the conditional score estimation approach in Song and Wang (2008). This approach can be seen as a weighted local partial likelihood without any assumptions on the distribution of the random effects. While this approach accounts for measurement error and short-term biological fluctuations in the longitudinal marker when modeling the hazard, it only permits inference on the survival parameters and not on the longitudinal model.

In order to allow for these two extensions, the flexible longitudinal trajectories and a potentially nonlinear time-varying association, both modeled by penalized splines, we develop and implement a highly flexible framework for joint models available within the R-package `bamlss`. As we represent all parts of this flexible joint model as structured additive predictors, which can include linear, parametric but also non-parametric penalized terms, we are able to allow potentially nonlinear, smooth, random, and time-varying effects in both submodels. In consequence the possibilities of this implementation go way beyond the two extensions that originally triggered the development. By applying this flexible model to the combined data set from two German high-risk T1D birth cohorts we aim to shed further light on the complex relationship between T1D-associated autoantibodies and the onset of clinical disease.

The remainder of this paper is structured as follows: The general model structure and potential extensions are outlined in Section 2. In Section 3, details on the Bayesian estimation procedure are given. A thorough testing of the model estimation through simulations is presented in Section 4 and the application to our T1D research question in Section 5. Concluding remarks are given in Section 6 and technical details can be found in the Appendix. The presented model is implemented in the R-package `bamlss` (Umlauf *et al.*, 2016, 2017). Source code to reproduce the simulation results is available as Supporting Information on the journal's web page (<http://onlinelibrary.wiley.com/doi/xx/supinfo>).

## 2 Methods

In the following, the general setup for additive joint models is presented with a special focus on two extensions of existing approaches: the flexible specification of longitudinal trajectories as well as the time-varying association between the longitudinal marker and the event. An overview of potential further model specifications illustrates the flexibility of the presented model family.

### 2.1 General Setup

For every subject  $i = 1, \dots, n$  we observe a potentially right-censored follow-up time  $T_i$  and the event indicator  $\delta_i$  (1 if subject  $i$  experiences the event, 0 if it is censored). We model the hazard of an event at time  $t$  as

$$h_i(t) = \exp\{\eta_i(t)\} = \exp\{\eta_{\lambda_i}(t) + \eta_{\gamma_i} + \eta_{\alpha_i}(t) \cdot \eta_{\mu_i}(t)\} \quad (1)$$

including in the full predictor  $\eta$  a predictor  $\eta_{\lambda}$  for all survival covariates that are time-varying or have a time-varying coefficient including the log baseline hazard, a predictor for baseline survival covariates  $\eta_{\gamma}$  as well as a predictor  $\eta_{\alpha}$  representing the potentially time-varying association between the longitudinal marker  $\eta_{\mu}$  and the hazard.

We also observe a longitudinal response  $\mathbf{y}_i = [y_{i1}, \dots, y_{in_i}]^T$  at the potentially subject-specific ordered

time points  $\mathbf{t}_i = [t_{i1}, \dots, t_{in_i}]^\top$  with  $t_{i1} \leq \dots \leq t_{in_i} \leq T_i$ .  $\mathbf{t} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_n^\top]^\top$  denotes the vector of the  $N = \sum_{i=1}^n n_i$  longitudinal measurement time points of all subjects. The longitudinal response at  $t_{ij}$  with  $j = 1, \dots, n_i$  is modeled as

$$y_{ij} = \eta_{\mu i}(t_{ij}) + \varepsilon_{ij} \quad (2)$$

with independent errors  $\varepsilon_{ij} \sim N(0, \exp[\eta_{\sigma i}(t_{ij})^2])$  allowing to also model the error variance. Thus  $\eta_{\mu i}(t_{ij})$  represents the longitudinally observed marker value without error at timepoint  $t_{ij}$ . This “true” marker value serves as a continuous-time covariate in the hazard in (1) and links the two model equations.

Each predictor  $\eta_{ki}$  with  $k \in \{\lambda, \gamma, \alpha, \mu, \sigma\}$  is a structured additive predictor, i.e. a sum of  $M_k$  functions of covariates  $\mathbf{x}_i$ ,

$$\eta_{ki} = \sum_{m=1}^{M_k} f_{km}(\mathbf{x}_{ki}).$$

Different subsets  $\mathbf{x}_{ki}$  of  $\mathbf{x}_i$  can serve as covariates for the different predictors, with each  $f_{km}$  typically depending on one or two covariates. For time-varying predictors the functions can also depend on time  $\eta_{ki}(t) = \sum_{m=1}^{M_k} f_{km}(\mathbf{x}_{ki}(t), t)$  with a potentially time-varying covariate vector  $\mathbf{x}_{ki}(t)$ . We express the vector of predictors for all subjects as  $\boldsymbol{\eta}_k = [\eta_{k1}, \dots, \eta_{kn}]^\top$ . In the survival part of the model (1) the vectors are of length  $n$  and potentially time-varying where  $\boldsymbol{\eta}_k(t)$  denotes the evaluation at time  $t$ . In the longitudinal part of the model (2) the vector  $\boldsymbol{\eta}_k(\mathbf{t})$  is of length  $N$ , containing entries  $\eta_{ki}(t_{ij})$  for all  $i$  and  $j$ , i.e. evaluations at all observed time points  $\mathbf{t}$  for the corresponding subjects. A general overview and details on the setup of the predictor vectors in the submodels can be found in table 1.

The functions  $f_{km}(\mathbf{x}_{ki})$  can model a variety of effects, such as smooth, spatial, time-varying or random effects terms which can be expressed in a straightforward notation for every term  $m$  of predictor  $k$  by using suitable basis function expansions and corresponding penalties  $\mathbf{P}_{km}$ . In a generic setup we let

$$\mathbf{f}_{km} = \mathbf{X}_{km}\boldsymbol{\beta}_{km} \quad \text{and} \quad \mathbf{P}_{km} = \frac{1}{\tau_{km}^2} \boldsymbol{\beta}_{km}^\top \mathbf{K}_{km} \boldsymbol{\beta}_{km}, \quad (3)$$

with the vector of function evaluations  $\mathbf{f}_{km}$  stacked over subjects, the design matrix  $\mathbf{X}_{km}$ , the coefficient vector  $\boldsymbol{\beta}_{km}$ , the penalty matrix  $\mathbf{K}_{km}$  and the variance parameter  $\tau_{km}^2$  that controls the amount of penalization of the respective term. In the Bayesian setting a penalization is imposed by specifying an appropriate prior distribution for the parameters,  $\boldsymbol{\beta}_{km} \sim N(\mathbf{0}, [\frac{1}{\tau_{km}^2} \mathbf{K}_{km}]^-)$  with  $\mathbf{A}^-$  denoting the generalized inverse of  $\mathbf{A}$ , as presented in more detail in section 3.3. Note that these basic penalties can be extended further as shown in more detail in the next subsection. In analogy to the differences in form in the generic vector of predictors, i.e.  $\boldsymbol{\eta}_k$ ,  $\boldsymbol{\eta}_k(t)$  and  $\boldsymbol{\eta}_k(\mathbf{t})$ , the form of the generic vectors of function evaluations  $\mathbf{f}_{km}$  and the generic design matrices  $\mathbf{X}_{km}$  also differs between predictors and submodels. We refer to table 1 for further details.

We illustrate this setup by two important examples and the exemplary specification of a standard shared parameter joint model. First, smooth functions in time can be modeled using P-splines with a B-spline basis,  $f_{km}(t) = \sum_{d=1}^D \beta_d B_d(t) =: \mathbf{x}_{km}^\top(t) \boldsymbol{\beta}_{km}$ , and corresponding penalty matrix  $\mathbf{K}_{km} = \mathbf{D}_r^\top \mathbf{D}_r$  with the  $r$ -th difference matrix  $\mathbf{D}_r$  of appropriate dimension (Eilers and Marx, 1996). For Bayesian P-Splines, smoothing is induced by appropriate prior specification, where the difference penalties are replaced by their stochastic analogues, i.e. random walks (Lang and Brezger, 2004). Second, random intercepts in the longitudinal part are incorporated by specifying  $\mathbf{X}_{km}$  as an  $N \times n$  indicator matrix, where the  $i$ th column indicates which longitudinal measurements belong to subject  $i$ ,  $\boldsymbol{\beta}_{km} = [\beta_{km1}, \dots, \beta_{kmn}]$  denotes the coefficient vector and an  $n \times n$  identity matrix as penalty  $\mathbf{K}_{km} = \mathbf{I}_n$  ensures  $\beta_{kmi} \sim N(0, \tau_{km}^2)$  independently. Further, although the model allows for highly flexible joint models, a simple shared parameter joint model is contained as a special case with

- P-spline log-baseline hazard  $\eta_\lambda(t) = f_\lambda(t) = \mathbf{X}_\lambda(t)\beta_\lambda$  with  $\mathbf{K}_\lambda = \mathbf{D}_r^\top \mathbf{D}_r$ :  
 $\mathbf{X}_\lambda(t)$  contains  $n$  stacked replications of  $\mathbf{x}_\lambda(t)$ , the vector of  $D$  B-spline basis functions  $B_d(t)$  at  $t$
- parametric effects of baseline survival covariates  $\eta_\gamma = \mathbf{f}_\gamma = \mathbf{X}_\gamma\beta_\gamma$  with  $\mathbf{K}_\gamma = \mathbf{0}$ :  
 $\mathbf{X}_\gamma$  contains the stacked subject-specific covariate vectors  $\mathbf{x}_{\gamma i}$
- time-constant association between longitudinal and survival model  $\eta_\alpha = \mathbf{f}_\alpha = \mathbf{1}_n\beta_\alpha$  with  $\mathbf{K}_\alpha = \mathbf{0}$
- longitudinal model with a random intercept  $\eta_\mu(\mathbf{t}) = \mathbf{f}_{\mu 1}(\mathbf{t}) + \mathbf{f}_{\mu 2}(\mathbf{t}) = \mathbf{X}_{\mu 1}(\mathbf{t})\beta_{\mu 1} + \mathbf{X}_{\mu 2}(\mathbf{t})\beta_{\mu 2}$   
with  $\mathbf{K}_{\mu 1} = \mathbf{0}$ ,  $\mathbf{K}_{\mu 2} = \mathbf{I}_n$ :  
 $\mathbf{X}_{\mu 1}(\mathbf{t})$  is the fixed effects design matrix potentially including a parametric effect of time and  $\mathbf{X}_{\mu 2}(\mathbf{t})$  the indicator matrix for a random intercept
- constant error variance  $\eta_\sigma(\mathbf{t}) = \mathbf{f}_\sigma(\mathbf{t}) = \mathbf{1}_N\beta_\sigma$  with  $\mathbf{K}_\sigma = \mathbf{0}$

where  $\mathbf{1}_n$  and  $\mathbf{1}_N$  are vectors of ones of length  $n$  and  $N$ , respectively, and  $\mathbf{0}$  is a zero matrix. Note that we drop the index  $m$  for predictors which consist of only one term.

## 2.2 Important extensions of current models

A special focus in our joint model approach lies on the flexibility of the longitudinal predictor  $\eta_\mu$ . We model the trajectory for every subject as the sum of fixed covariate effects, a smooth function of time, a random intercept as well as smooth subject-specific deviations from this function over time,

$$\eta_{\mu i}(t) = f_{\mu 1}(t) + f_{\mu 2}(i) + f_{\mu 3}(t, i) + \sum_{m=4}^{M_\mu} f_{\mu m}(\mathbf{x}_{\mu m}). \quad (4)$$

In this parameterization  $f_{\mu 1}(t)$  is a smooth effect of time and  $f_{\mu 2}(i)$  is a random intercept. The term  $f_{\mu 3}(t, i)$  denotes the smooth subject-specific deviations from the global time effect using functional random intercepts (Scheipl *et al.*, 2015). Additionally linear or parametric effects, including a global intercept, as well as further smooth effects of covariates can be represented by an extra term in  $\sum_{m=4}^{M_\mu} f_{\mu m}(\mathbf{x}_{\mu m})$ . The basis for the functional random intercepts can be specified within the basis function approach as row tensor products of the marginal basis of a random intercept, marked by the subscript  $s$ , and the marginal basis for a smooth effect of time, marked by the subscript  $t$ . We denote the vector of function evaluations at every observed longitudinal time point in  $\mathbf{t}$  for the corresponding subjects in  $\mathbf{i} = [1, \dots, n]^\top$  as

$$f_{\mu 3}(\mathbf{t}, \mathbf{i}) = (\mathbf{X}_{\mu 3s} \odot \mathbf{X}_{\mu 3t})\beta_{\mu 3} = \mathbf{X}_{\mu 3}\beta_{\mu 3}, \quad (5)$$

where  $\mathbf{X}_{\mu 3s}$  is an  $N \times n$  indicator matrix as the basis for a random intercept as specified for  $\mathbf{X}_{\mu 2}$  in the previous sub-section,  $\mathbf{X}_{\mu 3t}$  is an  $N \times D$  matrix of evaluations of a marginal spline basis at  $\mathbf{t}$  and  $\mathbf{X}_{\mu 3}$  is the  $N \times nD$  basis matrix resulting from the row tensor product. The row tensor product  $\odot$  of a  $p \times a$  matrix  $\mathbf{A}$  and a  $p \times b$  matrix  $\mathbf{B}$  is defined as the  $p \times ab$  matrix  $\mathbf{A} \odot \mathbf{B} = (\mathbf{A} \otimes \mathbf{1}_b^\top) \cdot (\mathbf{1}_a^\top \otimes \mathbf{B})$  with  $\cdot$  denoting element-wise multiplication and  $\otimes$  denoting the Kronecker product.

The corresponding penalty term is constructed from the marginal penalty matrices:

$$\mathbf{P}_{\mu 3} = \beta_{\mu 3}^\top \left( \frac{1}{\tau_{\mu 3s}^2} \mathbf{K}_{\mu 3s} \otimes \mathbf{I}_t + \frac{1}{\tau_{\mu 3t}^2} \mathbf{I}_s \otimes \mathbf{K}_{\mu 3t} \right) \beta_{\mu 3} = \beta_{\mu 3}^\top \left( \frac{1}{\tau_{\mu 3s}^2} \tilde{\mathbf{K}}_{\mu 3s} + \frac{1}{\tau_{\mu 3t}^2} \tilde{\mathbf{K}}_{\mu 3t} \right) \beta_{\mu 3}, \quad (6)$$

where  $\mathbf{K}_{\mu 3s} = \mathbf{I}_n$  is the penalty matrix for the random effect and  $\mathbf{K}_{\mu 3t}$  is an appropriate penalty matrix for the smooth effect of time such as a difference penalty for B-splines. The enlarged penalty matrices  $\tilde{\mathbf{K}}_{\mu 3s}$  and  $\tilde{\mathbf{K}}_{\mu 3t}$  yield a penalization for every subject, resulting in a random effects structure and a smoothness penalization across time for each subject. Note that by specifying two variance parameters,  $\tau_{\mu 3s}^2$  and  $\tau_{\mu 3t}^2$ ,

the amount of penalization can differ in the direction of time and across subjects, resulting in an anisotropic penalty. This specification allows for a highly flexible modeling of individual trajectories over time.

Given the specification of a separate global intercept and subject-specific random intercepts, the constraints  $\int f_{\mu 1}(t)dt = 0$  and  $\int f_{\mu 3}(t, i)dt = 0$  for every  $i$  are set in order to ensure identifiability. The necessary linear constraint  $\int f_{\mu 1}(t)dt = 0$  is implemented for B-splines by transforming the marginal basis  $\mathbf{X}_{\mu 3t}$  into an  $N \times (D - 1)$  matrix  $\check{\mathbf{X}}_{\mu 3t}$  for which it holds that  $\check{\mathbf{X}}_{\mu 3t} \mathbf{1}_{D-1} = \mathbf{0}$  as shown in Wood (2006, chapter 1.8), and adjusting the penalty accordingly. Constructing the row tensor product  $\mathbf{X}_{\mu 3}$  using the transformed marginal basis matrix  $\check{\mathbf{X}}_{\mu 3t}$  with correspondingly adjusted marginal penalty ensures that the identification constraint  $\int f_{\mu 3}(t, i)dt = 0$  for every  $i$  is also fulfilled.

As a second extension to existing shared-parameter models we also specify the association between the longitudinal and the survival model as a structured additive predictor  $\eta_{\alpha}$ . In consequence, this predictor can be modeled as a function of time and/or other covariates. Motivated by our applied research questions we model  $\eta_{\alpha}(t) = f_{\alpha}(t)$  as a smooth function of time by using penalized splines, as specified for the baseline hazard. This allows us to find patterns beyond the standard joint model specification to explain the relationship between the longitudinal marker and the survival process. These patterns could for example be critical time windows in which a non-zero effect of  $\eta_{\alpha}$  is present or a potential change in the direction of the association  $\eta_{\alpha}$  over time.

### 2.3 Further potential specifications

The presented general framework of structured additive joint models allows for a variety of different effect specifications by making use of the flexibility of Bayesian structured additive regression models as well as adding functional extensions. Besides the presented smooth, time-varying, random effects and functional random intercept terms, a variety of further effects can be incorporated. Table 2 gives an overview of possible terms. All these terms can be specified by formulating the desired effect in a basis function representation with an appropriate penalty term. For details on the specification of such effects please refer to Fahrmeir *et al.* (2004); Scheipl *et al.* (2015); Wood (2006). Further details on the practical aspects within our implementation are given in section 3.4.

## 3 Estimation

We estimate the model in a Bayesian framework using Newton-Raphson and Markov chain Monte Carlo (MCMC) algorithms.

### 3.1 Likelihood

Under the assumption of conditional independence of the survival outcomes  $[T_i, \delta_i]$  and the longitudinal outcome  $\mathbf{y}_i$ , given the random effects, the likelihood of the specified joint model is the product of the two submodel likelihoods  $L^{\text{surv}}$  and  $L^{\text{long}}$  for the survival and the longitudinal model.

The log-likelihood of the survival part is

$$\ell^{\text{surv}}[\boldsymbol{\eta}_{\lambda}(\mathbf{T}), \boldsymbol{\eta}_{\gamma}, \boldsymbol{\eta}_{\alpha}(\mathbf{T}), \boldsymbol{\eta}_{\mu}(\mathbf{T}) | \mathbf{T}, \boldsymbol{\delta}] = \boldsymbol{\delta}^{\top} \boldsymbol{\eta}(\mathbf{T}) - \mathbf{1}_n^{\top} \boldsymbol{\Lambda}(\mathbf{T}), \quad (7)$$

where  $\mathbf{T} = [T_1, \dots, T_n]^{\top}$  and  $\boldsymbol{\delta} = [\delta_1, \dots, \delta_n]^{\top}$ . Here  $\boldsymbol{\Lambda}(\mathbf{T}) = [\Lambda_1(T_1), \dots, \Lambda_n(T_n)]^{\top}$  is the vector of the cumulative hazard rates  $\Lambda_i(T_i) = \exp(\eta_{\gamma i}) \int_0^{T_i} \exp[\eta_{\lambda i}(u) + \eta_{\alpha i}(u) \cdot \eta_{\mu i}(u)] du$  and  $\boldsymbol{\eta}(\mathbf{T}) = [\eta_1(T_1), \dots, \eta_n(T_n)]$  denotes the vector of the full predictors evaluated at the subject-specific survival times. The additive predictors implicitly also depend on covariates and model parameters. The log-likelihood of the longitudinal part of the model is

$$\ell^{\text{long}}[\boldsymbol{\eta}_{\mu}(\mathbf{t}), \boldsymbol{\eta}_{\sigma}(\mathbf{t}) | \mathbf{y}] = -\frac{N}{2} \log(2\pi) - \mathbf{1}_N^{\top} \boldsymbol{\eta}_{\sigma}(\mathbf{t}) - \frac{1}{2} (\mathbf{y} - \boldsymbol{\eta}_{\mu}(\mathbf{t}))^{\top} \mathbf{R}^{-1} (\mathbf{y} - \boldsymbol{\eta}_{\mu}(\mathbf{t})). \quad (8)$$

$\boldsymbol{\eta}_\mu(\mathbf{t})$  and  $\boldsymbol{\eta}_\sigma(\mathbf{t})$  are the predictor vectors of length  $N$  corresponding to the longitudinal response  $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top]^\top$  and  $\mathbf{R} = \text{blockdiag}(\mathbf{R}_1, \dots, \mathbf{R}_n)$ , where  $\mathbf{R}_i$  can reflect the error structure of interest. In our case, we assume  $\mathbf{R}_i = \text{diag}(\exp[\eta_{\sigma i}(t_{i1})]^2, \dots, \exp[\eta_{\sigma i}(t_{in_i})]^2)$  so that  $\mathbf{R}$  reduces to a diagonal matrix.

### 3.2 Priors and Posterior

In this general framework above, a variety of terms (cf. Table 2) can be specified using corresponding priors. For linear or parametric terms we use vague normal priors on the vectors of the regression coefficients, e.g.  $\boldsymbol{\beta}_{km} \sim N(0, 1000^2)$ , approximately corresponding to the precision matrices  $\mathbf{K}_{km} = \mathbf{0}$  as presented above. Smooth and random effect terms are regularized by placing suitable multivariate normal priors on the coefficients

$$p(\boldsymbol{\beta}_{km} | \tau_{km}^2) \propto \left( \frac{1}{\tau_{km}^2} \right)^{\frac{\text{rank}(\mathbf{K}_{km})}{2}} \exp \left( -\frac{1}{2\tau_{km}^2} \boldsymbol{\beta}_{km}^\top \mathbf{K}_{km} \boldsymbol{\beta}_{km} \right)$$

with precision matrix  $\mathbf{K}_{km}$  as specified in the penalty (3). We use independent inverse Gamma hyperpriors  $\tau_{km}^2 \sim IG(0.001, 0.001)$  to obtain an inverse Gamma full conditional for the variance parameters. In addition to the inverse gamma distribution, different priors are possible for the variance parameters in our implementation, such as Half-Cauchy and Half-normal distributions. The variance parameters  $\tau_{km}^2$  control the trade-off between flexibility and smoothness in the nonlinear modeling of effects. As such they can be interpreted analogous to inverse smoothing parameters in a frequentist approach.

For anisotropic smooths, when multiple variance parameters  $\tau_{km}^2 = (\tau_{kms}^2, \tau_{kmt}^2)$  are involved as in (6), we use the prior

$$p(\boldsymbol{\beta}_{km} | \tau_{km}^2) \propto \left| \frac{1}{\tau_{kms}^2} \tilde{\mathbf{K}}_{kms} + \frac{1}{\tau_{kmt}^2} \tilde{\mathbf{K}}_{kmt} \right|^{\frac{1}{2}} \exp \left( -\frac{1}{2} \boldsymbol{\beta}_{km}^\top \left[ \frac{1}{\tau_{kms}^2} \tilde{\mathbf{K}}_{kms} + \frac{1}{\tau_{kmt}^2} \tilde{\mathbf{K}}_{kmt} \right] \boldsymbol{\beta}_{km} \right). \quad (9)$$

The resulting posterior of the model is

$$p(\boldsymbol{\theta} | \mathbf{T}, \boldsymbol{\delta}, \mathbf{y}) \propto L^{\text{surv}} [\boldsymbol{\eta}_\lambda(\mathbf{T}), \boldsymbol{\eta}_\gamma, \boldsymbol{\eta}_\alpha(\mathbf{T}), \boldsymbol{\eta}_\mu(\mathbf{T}) | \mathbf{T}, \boldsymbol{\delta}] \cdot L^{\text{long}} [\boldsymbol{\eta}_\mu(\mathbf{t}), \boldsymbol{\eta}_\sigma(\mathbf{t}) | \mathbf{y}] \cdot \prod_{k \in \{\lambda, \gamma, \alpha, \mu, \sigma\}} \prod_{m=1}^{M_k} [p(\boldsymbol{\beta}_{km} | \tau_{km}^2) p(\tau_{km}^2)],$$

where  $\boldsymbol{\theta}$  is the vector of all parameters in the model and  $p(\boldsymbol{\beta}_{km} | \tau_{km}^2)$  and  $p(\tau_{km}^2)$  are the priors for the regression coefficients and variance parameters for each term  $m$  and predictor  $k$ , respectively.

### 3.3 Bayesian Estimation

Point estimates of  $\boldsymbol{\theta}$  can be obtained by posterior mode and posterior mean estimation. We estimate the posterior mode by maximizing the log-posterior of the model using a Newton-Raphson procedure, the posterior mean is obtained via derivative-based Metropolis-Hastings sampling and thus computationally demanding. We therefore recommend to use posterior mode estimates for a first quick assessment of the model and in order to obtain starting values for the posterior mean sampling.

In the maximization of the log-posterior to obtain the posterior mode, we update blockwise each term  $m$  of predictor  $k$  in each iteration  $l$  as

$$\boldsymbol{\beta}_{km}^{[l+1]} = \boldsymbol{\beta}_{km}^{[l]} - \nu_{km}^{[l]} \mathbf{H} \left( \boldsymbol{\beta}_{km}^{[l]} \right)^{-1} \mathbf{s} \left( \boldsymbol{\beta}_{km}^{[l]} \right)$$



with potentially varying steplength  $\nu_{km}^{[l]}$  and with the score vector  $s(\beta_{km})$  and the Hessian  $\mathbf{H}(\beta_{km})$ , which can be found in the Appendix. We optimize the variance parameters in each updating step to minimize the corrected AIC (AICc, Hurvich *et al.*, 1998), which showed good performance in smoothing parameter estimation in Belitz and Lang (2008). Additionally we optimize the steplength  $\nu_{km}^{[l]}$  over  $(0, 1]$  in each step to maximize the log-posterior. We assume the coefficients to have an approximately normal posterior distribution and derive credibility intervals from  $N(\hat{\beta}_{km}, [-\mathbf{H}(\hat{\beta}_{km})]^{-1})$  for quick approximate inference.

For the posterior mean sampling we construct approximate full conditionals  $\pi(\beta_{km}|\cdot)$  based on a second order Taylor expansion of the log-posterior centered at the last state  $\beta_{km}^{[l]}$ , similar to Fahrmeir *et al.* (2004), Klein *et al.* (2015a) and Klein *et al.* (2015b) and as shown in more detail in Umlauf *et al.* (2017). The proposal density from this approximate full conditional is proportional to a multivariate normal distribution with the precision matrix  $(\Sigma_{km}^{[l]})^{-1} = -\mathbf{H}(\beta_{km}^{[l]})$  and the mean  $\mu_{km}^{[l]} = \beta_{km}^{[l]} - \mathbf{H}(\beta_{km}^{[l]})^{-1}s(\beta_{km}^{[l]})$ . In each iteration  $l$  of the Metropolis-Hastings sampler and for updating block  $km$  a candidate  $\beta_{km}^*$  is drawn from the proposal density  $q(\beta_{km}^*|\beta_{km}^{[l]}) = N(\mu_{km}^{[l]}, \Sigma_{km}^{[l]})$ . By drawing candidates from a close approximation of the full conditional, we approximate a Gibbs sampler and achieve high acceptance rates and good mixing.

For the sampling of the variance parameters  $\tau_{km}^2$  Gibbs sampling is employed, as the full conditionals  $\pi(\tau_{km}^2|\cdot)$  follow an inverse Gamma distribution, if inverse Gamma hyperpriors are used. Slice sampling is employed when no simple closed-form full conditional can be obtained as for example in the sampling of variance parameters for anisotropic smooths (9) or for other hyperpriors.

Model selection can be conducted via DIC which is readily available within our implementation.

### 3.4 Implementation details

The model estimation is implemented within R (R Core Team, 2016) in the package `bamlss` (Umlauf *et al.*, 2016, 2017) that allows the Bayesian estimation of a variety of models within the framework of Bayesian Additive Models for Location, Scale and Shape. The specification of appropriate design matrices and penalties for the desired effects is conducted internally via the R-package `mgcv` (Wood, 2011). In consequence the full range of implemented smoothing approaches, such as P-splines, thin-splite regression splines, random effects, and Markov Random Fields, can be used within our implementation. We refer to Wood (2006) and Wood *et al.* (2016) for further information on model terms, bases and penalties. In our model specification in the simulations and the application we make use of Bayesian P-splines (Lang and Brezger, 2004) to model smooth effects. As the integrals in the survival likelihood as well as in the respective scores and Hessians have no analytical solution, they are approximated numerically using the trapezoidal rule and a fixed number of 25 integration points. Starting values for the posterior mean sampling are obtained by estimating the posterior mode of the joint model. For estimating the highly flexible functional random intercepts, the implementation makes use of the blockdiagonal structure of the design matrices to increase computation speed. The posterior mean sampling is implemented to potentially run in parallel on a number of specified cores on Linux systems. More details can be found in the documentation of the `bamlss` R-package.

## 4 Simulation

We assess the estimation of our model by means of a simulation study with focus on two aspects: First, comparing our results with the established joint model implementation in `JMbayes` (Rizopoulos, 2016a) for models with time-constant  $\eta_\alpha$ . Second, we want to assess the ability to model highly complex longitudinal trajectories as well as a time-varying effect of  $\eta_\alpha(t)$ , the two important new extensions within our framework. With this simulation we also aim to gain insights into the estimation quality of the model when applied to real data sets of T1D cohorts that motivated our methods development. Therefore we simulate two differing data situations, mimicking real cohort data. The first simulated data setting, corresponding



to the German cohort data presented in the Application Section, has less subjects, at more variably spaced time points but with a longer follow up, than the other. Finally we aim to assess how well the posterior mode estimation can approximate the effects in comparison with the posterior mean estimates.

#### 4.1 Simulation design

For every setting we generate longitudinal measurements for  $n$  subjects at a fixed grid of time points  $\mathcal{P}$  based on a true longitudinal model  $\eta_\mu(t)$  as specified in (4) with the time effect  $f_{\mu 1}(t) = 0.1(t + 2) \exp(-0.075t)$ , the random intercepts  $f_{\mu 2}(i) = r_i$  where  $r_i \sim N(0, 0.25)$ , the functional random intercepts  $f_{\mu 3}(t, i) = \mathbf{X}_{\mu 3} \beta_{\mu 3}$ , and the global intercept and covariate effect  $f_{\mu 4}(\mathbf{x}_{\mu i}) = 0.5$  and  $f_{\mu 5}(\mathbf{x}_{\mu i}) = 0.6 \sin(x_{2i})$  with  $x_{2i} \sim \text{unif}(-3, 3)$ . We simulate the functional random intercepts flexibly by P-Splines where we draw the true vector of spline-coefficients for all subjects from  $\beta_{\mu 3} \sim N(\mathbf{0}, [(1/\tau_s^2) \tilde{\mathbf{K}}_s + (1/\tau_t^2) \tilde{\mathbf{K}}_t]^{-1})$  as in (6) with  $\mathbf{K}_t = \mathbf{D}_2^T \mathbf{D}_2$ ,  $\tau_s^2 = 1$  and  $\tau_t^2 = 0.2$ . The hazard  $h_i(t)$  for every subject is calculated according to (1) using the true survival predictor functions  $\eta_\lambda(t) = 1.4 \log((t + 10)/1000)$ ,  $\eta_\gamma = 0.5 \sin(x_{1i})$ , with  $x_{1i} \sim \text{unif}(-3, 3)$  and  $\eta_\alpha(t)$  varying for the two simulation settings. Based on  $h_i(t)$ , survival times are generated for every subject as described in Bender *et al.* (2005) and Crowther and Lambert (2013). Every subject is censored after  $\max(\mathcal{P})$  and we additionally apply uniform censoring  $U(0, 1.5 \cdot \max(\mathcal{P}))$  to the survival times. In order to mimic missing measurements in the real data,  $p\%$  of the remaining longitudinal data are randomly set to missing. Longitudinal observations are obtained from  $\eta_{\mu i}(t)$  by adding independent errors  $\epsilon_{ij} \sim N(0, 0.3^2)$  for each  $t_{ij}$  in  $\mathbf{t}$ . The influence of different data structures on the estimation is assessed by simulating two different data settings in each of the two simulation settings. In the smaller data setting,  $a$ , observations for  $n_a = 150$  subjects are generated at the measurements points  $\mathcal{P}_a = [0, 1, \dots, 120]$  where  $p_a = 75\%$  of the longitudinal measurements are missing and on average 108 (72 %) events occur, compared to  $n_b = 300$  subjects at the time points  $\mathcal{P}_b = [0, 3, \dots, 72]$  with  $p_b = 10\%$  missings and 165 (55 %) events in the larger data setting,  $b$ .

In each data and simulation setting we draw  $Q = 200$  samples. To ensure convergence, we run the model estimation with 23000 samples, a burn-in of 3000 and a thinning of 20, yielding 1000 samples, as assessed in preliminary simulations. For each estimated model  $q$  within a simulation setting we assess bias, mean-squared error (MSE) and frequentist coverage of the 95% credibility intervals, defined by the 2.5th and the 97.5th percentiles of the MCMC samples for the posterior mean and the approximate normal intervals for the posterior mode. We evaluate bias, MSE and coverage both averaged over all time points and averaged per time point. For the predictors in the longitudinal model, i.e.  $k \in \{\mu, \sigma\}$ , the average bias in each sample  $q$  is  $B_k^q = \frac{1}{N} \sum_{i=1}^n \sum_{j=1}^{n_i} [\hat{\eta}_{ki}^q(t_{ij}) - \eta_{ki}^q(t_{ij})]$  where  $\hat{\eta}_{ki}$  denotes the estimate. To assess the model fit over time we also evaluate the bias per timepoint  $B_k^q(t) = \frac{1}{n} \sum_{i=1}^n [\hat{\eta}_{ki}^q(t) - \eta_{ki}^q(t)]$  for all  $t$  in  $\mathcal{P}$ . The computations for MSE and coverage are analogous. For the survival predictors, i.e.  $k \in \{\gamma, \lambda, \alpha\}$ , the average bias is  $B_k^q = \frac{1}{n} \sum_{i=1}^n [\hat{\eta}_{ki}^q(T_i) - \eta_{ki}^q(T_i)]$  using evaluations at the subject's event times. The bias of the time-varying survival predictor  $\eta_\lambda$ , and for setting 2 also  $\eta_\alpha$ , is additionally evaluated at the fixed grid of time points  $t$  in  $\mathcal{P}$  as above with MSE and coverage computed accordingly. These error measures are then averaged over all  $Q$  samples per setting.

For the comparison with the joint model implementation in `JMbayes` in settings 1a and 1b, data is generated with  $\eta_\alpha(t) = 1$  as time-constant. In our implementation we model the longitudinal submodel by P-splines with cubic B-splines, a second order difference penalty and 12 knots (4 internal knots), for both the overall mean as well as the individual trajectories. After application of the constraints this yields  $7 \cdot n$  basis functions. For the time-varying effect of the baseline hazard,  $\eta_\lambda$ , as well as the nonlinear effect in  $\eta_\gamma$  we use 10 knots (2 internal knots) resulting in 5 basis functions per effect after application of the constraints. In order to achieve a comparable model in the package `JMbayes` we model nonlinear effects in the longitudinal submodel and survival covariate effects by B-splines and determine the number of knots to minimize the DIC in preliminary simulations. Details on the inclusion of nonlinear effects in both submodels can be found in the source code of the Supporting Information. As a result we model the longitudinal part by cubic B-splines for both the fixed and random effects with 1 internal knot for the

larger data setting and without internal knots for the smaller data setting, resulting in 4 and 3 basis functions for both the fixed and random effects of time, respectively. As prior simulations had shown convergence issues using an unrestricted covariance matrix of the random effects, we restrict it to be diagonal, resulting in independent random effects. Also based on DIC from preliminary simulations we specify the effect in  $\eta_\gamma$  in the survival part with cubic B-splines with 3 internal knots using 5 basis functions. We model the baseline hazard with P-splines using the default settings from `JMbayes`, i.e. a cubic B-spline basis with 17 basis functions and a second order difference penalty. For the MCMC procedure we also use the default settings of 20000 iterations, including a burn-in of 3000 and a thinning such that 2000 samples are kept.

In our second simulation, i.e. settings 2a and 2b, we specify the longitudinal trajectories as before but generate data using a time-varying association predictor  $\eta_\alpha(t) = \cos((t - 33)/33)$  for data in *a* and  $\eta_\alpha(t) = \cos((t - 20)/20)$  for *b* in order to achieve a similar shape despite a differing time scale. We fit the model using the same specification as in setting 1. Additionally  $\eta_\alpha$  is modeled as a P-spline with 10 knots (2 internal knots) resulting in 5 basis functions after application of the constraints.

## 4.2 Simulation results

The focus of the first simulation is the comparison with the package `JMbayes` regarding the accuracy of the modeling of the longitudinal trajectories and the time-constant association parameter  $\eta_\alpha$  in settings 1a and 1b. Table 3 shows the MSE, bias and coverage for the estimation of  $\eta_\alpha$ .

For both methods  $\eta_\alpha$  is estimated more precisely and with a higher coverage in the larger data setting *b* compared to *a*. In both data settings `bamlss` achieves lower MSE, less bias and a higher coverage in the estimation of the association compared to `JMbayes`. For `JMbayes` the coverage for  $\eta_\alpha$  is not satisfactory in both settings (0.840 and 0.890). As the further survival predictors,  $\eta_\gamma$  and  $\eta_\lambda$ , are parameterized differently in the two estimation methods with regard to the intercept term and sum-to-zero constraints, we assess only the prediction quality of  $\eta_\lambda + \eta_\gamma$ . We observe that `JMbayes` shows a higher bias in the estimation of the sum of these two predictors. Regarding the longitudinal submodel for  $\eta_\mu$  both methods are fairly equal regarding the average MSE over the larger data setting (`bamlss`: 0.028 vs. `JMbayes`: 0.029), but our approach seems to be more precise in the smaller data setting (`bamlss`: 0.022 vs. `JMbayes`: 0.031). To further understand the cause of this difference we look at the bias in the estimation of  $\eta_\mu$  over the whole observed time course for the smaller data setting. As shown in Figure 2, `JMbayes` seems to underestimate some nonlinearity of the true predictor. Both methods show higher uncertainty for later time points when, due to censoring and the occurrence of events, less information is available. For the longitudinal predictors we were not able to calculate credibility intervals in `JMbayes`. Finally, the estimation of the error variance is more precise in `bamlss`.

There are large runtime differences where `JMbayes` models took on average 9 minutes and 13 minutes for data setting *a* and *b*, respectively, and the implementation `bamlss`, due to the more flexible functional random effects specification, took on average 2.5 hours and 5.1 hours to run on a single core of a 2.60 GHz Intel Xeon Processor E5-2650.

The aim of the second simulation setting is to shed light on the precision of the estimation of all predictors in the model with a special focus on the estimation of  $\eta_\alpha$ , which is nonlinear in time. Additionally, we also compare the precision of the posterior mode to the posterior mean estimation. Table 4 gives an overview of the estimation precision of all predictors.

Similarly to setting 1 we observe an effect of sample size: All survival predictors ( $\eta_\lambda, \eta_\gamma, \eta_\alpha$ ) show a smaller MSE for data setting *b* compared to *a* probably due to the higher number of events. In contrast, the MSE is smaller for the estimation of  $\eta_\mu$  in data setting *a* compared to *b* potentially due to the longer follow-up and a slightly higher number of longitudinal observations per subject. Whereas the precision of the point estimates is overall similar or only slightly worse for the posterior mode compared to the posterior mean estimation, the coverage is not acceptable for the posterior mode but close to 95% for the posterior mean. The only exception is the estimation of  $\eta_\sigma$ , where the coverage is somewhat lower for

the posterior mean. As  $\eta_\sigma$  is very precisely estimated and formal inference is usually not of interest for this predictor, we do not rate this under-coverage as too problematic. In order to illustrate the precision in the time-varying effect estimates and to assess the cause of differences in MSE, Figure 3 displays the true and estimated predictors  $\eta_\lambda(t)$ ,  $\eta_\alpha(t)$  and their joint effect  $\eta_\lambda(t) + \eta_\alpha(t) \cdot \frac{1}{n} \sum_{i=1}^n \eta_{\mu i}(t)$  evaluated at the mean of the estimated effect of  $\eta_\mu(t)$  for all time-points. Overall the estimated predictors match the true functions quite well. For the smaller data sets there is more uncertainty in the estimation, especially at later time points, when less subjects are still observed. For example in setting *a* at a late time point ( $t = 110$ ) on average only 6 subjects remain in the risk set as compared to 37 subjects at a late timepoint ( $t = 65$ ) in setting *b*. Therefore the estimates for later timepoints are highly dependent on the remaining subjects. This uncertainty is not only visible in the single predictors but also in their combined effect, see the right panel of Figure 3 (later time points in column a).

With on average only 10 and 18 minutes to run for data setting *a* and *b* respectively, the posterior mode estimation has clear advantages in computation time over the more precise posterior mean estimation with 2.6 and 5.2 hours on average in this setting.

In conclusion, our simulations show that the estimation of models with constant associations between marker and event performs well, even outperforming the implementation in `JMbayes` in some aspects. The estimation of more flexible models that are newly covered by our approach in contrast to existing implementations, i.e. with a time-varying association parameter and the specification of flexible trajectories, is equally satisfactory. While the more precise posterior mean estimation is time-consuming, the posterior mode offers a computationally efficient way to quickly assess the point estimates in a given model specification, even though credibility bands are only approximate.

## 5 Application

In order to gain insights into our motivating research question we apply the model to a combined data set of two ongoing German T1D risk cohorts to investigate whether longitudinal trajectories of insulin autoantibodies (IAA) are associated with the rate of progression to T1D. Whereas different autoantibodies are diagnostic for a preclinical stage of the disease, our focus lies on the analysis of the levels of IAA as a marker from the time when it first exceeded a specific threshold, called seroconversion, to the onset of T1D or loss to follow-up. The marker IAA is most often the first autoantibody to appear (Ziegler *et al.*, 1993, 1999; Hummel *et al.*, 2004a). Both its initial value at seroconversion and its mean over time have been shown to be positively associated with the emergence of T1D and negatively related to the age at T1D diagnosis (Steck *et al.*, 2011, 2015). As seroconversion defines the beginning of a pre-clinical, symptom-free stage of T1D (Insel *et al.*, 2015) we chose this timepoint as the starting point for our analysis instead of modeling the trajectories from birth and accounting for zeros up to the timepoint of seroconversion (Liu, 2009; Hatfield *et al.*, 2012; Rizopoulos *et al.*, 2008).

The BABYDIAB and BABYDIET studies, both prospective birth cohorts with a joint study protocol, aim to investigate the natural history of T1D development, i.e. explore environmental and genetic factors associated with the progression to T1D. In the BABYDIAB study 1650 offspring of patients with T1D were followed from birth to the development of T1D or loss to follow-up for up to 21 years (median: 15.7 years) (Ziegler *et al.*, 1993, 1999; Hummel *et al.*, 2004b). Using the BABYDIAB study protocol, the BABYDIET study subsequently recruited 791 additional offspring or siblings of patients with T1D (Hummel *et al.*, 2011). In both studies, autoantibody measurements were taken at age 9 months and 2, 5, 8, 11, 14 and 17 years and additionally every 6 months after positive islet autoantibodies had emerged. The exact age at the emergence of clinical T1D was assessed also between study visits.

In our joint model we use data of  $n = 127$  children who developed IAA during follow-up, censored at 15 years after seroconversion due to the extremely low sample size at later time points, of which 69 (54%) subjects progressed to T1D (see also WebFigure 1 in the Supporting Information). The median progression

time after seroconversion for IAA was 4.7 years (range: 0, 14.5 years; 25th and 75th percentile: 1.9, 6.7) with a median event-free follow-up of 5.9 years (range: 0, 15 years; 25th and 75th percentile: 3.4, 8.6). Subjects seroconverted for IAA at a median age of 2.1 years (range: 0.6, 18.1 years) with 86% of the subjects having IAA as the first autoantibody to appear. In total  $N = 894$  longitudinal measurements of IAA after seroconversion (median number of measurements per subject: 6, range: 1, 24) were used and log-transformed as  $\log(IAA + 1)$  for the analysis. We model subject's transformed autoantibody levels using functional random intercepts and two further covariates which were chosen a priori based on medical knowledge to fully capture the pattern of autoantibody levels. First, the age at seroconversion is included as a linear effect and second a binary variable indicates whether the autoantibody was among the first autoantibodies to appear. We model the association between marker and event,  $\eta_\alpha(t)$ , to be a non-linear function of time. Further we allow the covariates in the longitudinal model to also influence the survival process directly and expect a positive association between the age at seroconversion and the time to T1D (Steck *et al.*, 2011; Ziegler *et al.*, 2013). In our Bayesian model estimation we sample for 33000 iterations with a burnin of 3000 and thinning of 15 to obtain 2000 samples, with starting values for the posterior mean estimation obtained from the posterior mode estimates. Convergence is assessed by the inspection of traceplots, of which a subset is presented in the Supporting Information. In order to assess the sensitivity of the results to the number of knots we specify three models with differing numbers of knots. We specify two models using either 12 (i.e. 4 internal) knots or 20 (i.e. 8 internal) knots for the overall mean as well as the individual trajectories in the functional random intercepts and 10 (i.e. 2 internal) knots in the survival submodel. Additionally we specify a model with 20 (i.e. 8 internal) knots for nonlinear terms in both, the longitudinal and the survival submodels.

The results from the three specified models in our sensitivity analysis are highly similar for all predictors with regard to mean estimates and the credibility intervals. However, we observe lower DIC for the models with more knots in the functional random intercepts along with a closer fit of the individual trajectories and more narrow credibility intervals for the estimated association  $\eta_\alpha(t)$  (cf. WebFigure 2). Using more knots in the survival submodel results in a better mixing in the traceplots and a slightly lower DIC. Hence we assume the results to be robust regarding the exact number of knots and present results of the model with the lowest DIC in the following.

As shown in Figure 1b for 5 randomly selected subjects, we are able to closely approximate the individual non-linear trajectories of IAA. The association between the marker and the onset of clinical T1D is estimated as stable over time with an average slope of -0.01 [95% credibility interval: -0.09, 0.07]. The average slope was defined as the mean over the first derivative of the association  $\eta_\alpha(t)$  evaluated at all observed event and follow-up times  $\mathbf{T}$ , and its posterior distribution can be easily obtained by numerically deriving  $\eta'_\alpha(t)$  in every sample. The credibility interval for the estimated association is above 0 for the first 2 years after seroconversion and there is more uncertainty at later timepoints, i.e. when less event and follow-up times are observed and when less subjects remain in the risk set, as indicated by the credibility intervals (Figure 4). In the longitudinal submodel we observe that trajectories have a lower level, if subjects seroconverted at an older age (in years,  $\beta_{\mu 4} = -0.07$ ; 95% credibility interval: [-0.14, -0.004]) and a higher level if IAA was amongst the first markers to appear ( $\beta_{\mu 5} = 0.89$ ; [0.23, 1.53]). In the survival submodel the log-hazard is decreased if IAA was amongst the first markers to appear ( $\beta_{\gamma 2} = -0.94$ ; [-1.73, -0.12]). In sum if IAA is amongst the first markers to appear the log-hazard is reduced by 0.7. This net effect can be derived as the sum of the direct effect in  $\eta_\gamma$  and the indirect effect in  $\eta_\alpha \cdot \eta_\mu$  with an average association of  $\frac{1}{n} \sum_i \eta_\alpha(T_i) = 0.28$ . Additionally we do not observe a direct effect of the age at seroconversion ( $\beta_{\gamma 3} = -0.09$ ; [-0.19, 0.01]).

As the estimated association  $\eta_\alpha(t)$  showed no clear time trend we further compared this main model with a simpler model assuming a time-constant association using both our joint model implementation in `bamlss` and the implementation in `JMbayes`. In `bamlss` the same model was specified as in our main analysis with the exception that only an intercept was allowed in the predictor  $\eta_\alpha$ . In `JMbayes` individual trajectories were modeled by cubic B-splines (see section 4). The models resulted in highly similar effect

estimates for the covariate effects and the intercept of the association effect (cf. Figure 4b). The effects on the log-hazard of the association and the covariate indicating that IAA was the first marker to appear were estimated slightly smaller in the `JMbayes` model. Despite more flexibility in the modeling, our main model achieved a slightly lower DIC than the time-constant `bamlss` fit (DIC: 2645.6 vs. 2648.9). As the DIC values are random and show small variations depending on the posterior sample this small difference in DIC should not be overinterpreted. The DIC of the `JMbayes` estimation was not directly comparable due to a different likelihood formulation. We further conducted a sensitivity analysis regarding the prior specification of the variance parameters as well as the parametric terms. Basing on our main model we estimated 3 additional models (a) using  $\tau_{km}^2 \sim IG(0.001, 1)$ , i.e. a distribution closer to a uniform for the variance parameters, (b) a Half-Cauchy distribution for the variance parameters, and (c) using a slightly more informative prior  $\beta_{km} \sim N(0, 50^2)$  for all linear terms. Again all estimated effects were highly similar (cf. WebFigure 6).

In line with previous findings (Steck *et al.*, 2011, 2015) these results indicate that the quantitative levels of the marker IAA are informative for the rate of progression to T1D in the first years after seroconversion with higher levels increasing the hazard of T1D. The direct relationship between the hazard and age at seroconversion is not supported by the model, suggesting that the previously established influence of this baseline covariate on T1D progression may be mediated by the marker levels, i.e. the effect in the respective log-hazard is reduced if the marker levels over time are taken into account as in our flexible parameterization. The association between IAA and the hazard of T1D over time was not found to be strongly time-varying in this relatively small study despite a slightly smaller DIC for the time-varying model. There is much uncertainty around the nonlinear time-varying estimate of the association  $\eta_\alpha(t)$ . This uncertainty is potentially a result of the flexibility in the estimation in combination with the amount of data in the survival part. Furthermore the precision in the assessment of the exact age at seroconversion and therefore the timepoint 0 per subject is limited by the timing of the follow-up visits, potentially inducing further noise in the assessment of the time-varying association between the biomarker and the time to clinical T1D. All estimated effects were robust regarding different model specifications, i.e. varying number of knots, differently specified associations, as well as different prior distributions.

## 6 Discussion and Outlook

We presented a flexible joint model that allows to fit a broad range of joint model specifications using structured additive predictors for all model components. The approach is fully implemented in the R-package `bamlss`. While the framework is very flexible, as illustrated by Table 2, the focus in this work lies on the flexible modeling of individual trajectories and the specification of a time-varying association between marker and event. The proposed model shows satisfactory performance in various simulation settings and has the potential to offer new insights into complex relationships between biomarkers and time-to-event processes.

Our methods development was motivated by a specific research question from T1D studies and data from two cohorts, the presented combined German BABYDIAB/BABYDIET cohort as well as a multinational cohort. We aimed at assessing the potentially nonlinear time-varying association between a highly variable disease-specific marker and the time to T1D onset. We saw that even by using the combined BABYDIAB/BABYDIET cohort, the sample size of the data set considered in the application in Section 5 is at the lower limit for the complexity of our model, as indicated by our simulation study and by the width of the credibility intervals in the applied results. Nevertheless we found a positive association between a disease-related biomarker and the occurrence of clinical T1D. Although our model allows for a time-varying association between the biomarker and the event process, at least in this small data set it was estimated to be roughly constant with a DIC comparable to that of a simpler constant model and a slightly larger effect early after seroconversion. In consequence our flexible model can also be used to check the



modeling assumptions of simpler models that are commonly used. We aim to further explore the relationship between T1D-specific autoantibodies and the progression to T1D in the larger data set from the multinational T1D cohort (with sample size exceeding data setting  $b$  in our simulations) as in Steck *et al.* (2015). We note that these results, as the majority of findings on pre-clinical T1D, can only be generalized to subjects with increased risk of T1D and not to the general population.

Due to the complexity of the model and its estimation, the computation speed is still a drawback in our implementation. Hence we are constantly working on speeding up the computations further. As shown in simulation 2, the posterior mode estimation offers a computationally efficient way to obtain point estimates from a flexible joint model before starting the full MCMC sampling. These posterior mode estimates show a precision similar to that of the posterior mean estimates. However, the credibility intervals obtained from posterior modes are not wide enough, potentially due to the fact that the uncertainty around the variance parameters  $\tau_{km}^2$  is not included in the credibility intervals. In consequence, only the credibility intervals of the posterior mean estimates should be used for inference.

As is well-known in the survival context, the number of potential parameters in the model is limited by the number of observed events (Harrell *et al.*, 1996). This also holds in our approach for the predictors in the survival part of the model,  $\eta_\lambda$ ,  $\eta_\gamma$ , and  $\eta_\alpha$ . We achieve to alleviate this issue to some extent by the penalized approach, which decreases the effective number of degrees of freedom and thus allows for a richer model than would be possible without a penalty. Still, we recommend to model only those effects as non-linear functions, where a strong indication for non-linearity is given.

Within the framework of the presented additive joint model several further extensions are possible. So far model selection is conducted via DIC. We note that more advanced model selection techniques such as Bayesian Lasso selection (Tang *et al.*, 2017) or boosting (Waldmann *et al.*, 2017) have been developed. Including these techniques into the presented framework are topics for future work. Regarding potential effects to be specified we aim to extend the model by including the derivative of the longitudinal trajectories to model the event process similar to Ye *et al.* (2008), Brown (2009) and Rizopoulos *et al.* (2014), allowing to model the potentially time-varying association between changes in the marker and the hazard. Further, functional historical effects of the trajectories, including information on the history of the marker (Malfait and Ramsay, 2003; Gellar *et al.*, 2014), could potentially offer additional insights into complex relationships between markers and event processes.

**Acknowledgements** We thank Lorenz Lachmann, Claudia Matzke, Joanna Stock, Stephanie Krause, Annette Knopff, Florian Haupt, Maren Pflüger, Marlon Scholz and Anita Gavrisan (all: Institute for Diabetes Research, Helmholtz Zentrum München) for data collection and expert technical assistance, Ramona Puff (Institute for Diabetes Research, Helmholtz Zentrum München) for laboratory management, and Peter Achenbach (Institute for Diabetes Research, Helmholtz Zentrum München) and Ezio Bonifacio (Center for Regenerative Therapies Dresden and Paul Langerhans Institute Dresden, Technische Universität Dresden) for overseeing antibody measurement and for fruitful discussions and advice on modeling T1D-specific autoantibodies. We also thank all pediatricians and family doctors in Germany for participating in the BABYDIAB Study. Furthermore we thank Fabian Scheipl (Ludwig-Maximilians-Universität München) for advice on modeling functional random intercepts and the reviewers as well as the associate editor for their constructive comments which helped to improve the paper. The work was supported by the JDRF (JDRF-2-SRA-2015-13-Q-R) and by grants from the German Federal Ministry of Education and Research (BMBF) to the German Center for Diabetes Research (DZD e.V.). Meike Köhler's work was supported by a grant from the Helmholtz International Research Group (HIRG-0018) and Sonja Greven acknowledges funding from the German research foundation (DFG) through Emmy Noether grant GR 3793/1-1.

### Conflict of Interest

*The authors have declared no conflict of interest.*



## Appendix

We derive score vectors and Hessians for the regression coefficients of every predictor. We introduce some further notation to formulate these derivatives. For the time-varying predictors of the survival part  $k \in \{\lambda, \alpha, \mu\}$  the design matrix  $\mathbf{X}_k(\mathbf{T})$  denotes the  $n \times p_k$  matrix of evaluations at the vector of survival times  $\mathbf{T}$ . For the time-varying predictors of the longitudinal part  $k \in \{\mu, \sigma\}$  the  $N \times p_k$  design matrix  $\mathbf{X}_k(\mathbf{t})$  contains the evaluations at all observed subject-specific timepoints  $\mathbf{t}$ . The column vector  $\mathbf{x}_{ki}$  is the  $i$ -th row of the respective design matrix  $\mathbf{X}_k$  for  $k \in \{\gamma, \lambda, \alpha, \mu, \sigma\}$  and  $\boldsymbol{\beta}_k = [\beta_{k1}, \dots, \beta_{kM_k}]^\top$ . Let  $\ell$  denote the log-likelihood, i.e. the sum of the contributions of the longitudinal and survival submodels defined in (7) and (8). In more detail, the full log-likelihood is

$$\begin{aligned} \ell[\boldsymbol{\theta}|\mathbf{T}, \boldsymbol{\delta}, \mathbf{y}] &= \boldsymbol{\delta}^\top [\mathbf{X}_\lambda(\mathbf{T})\boldsymbol{\beta}_\lambda + \mathbf{X}_\gamma\boldsymbol{\beta}_\gamma + \mathbf{X}_\alpha(\mathbf{T})\boldsymbol{\beta}_\alpha \cdot \mathbf{X}_\mu(\mathbf{T})\boldsymbol{\beta}_\mu] \\ &\quad - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \boldsymbol{\beta}_\gamma) \int_0^{T_i} \exp[\mathbf{x}_{\lambda i}^\top(u) \boldsymbol{\beta}_\lambda + \mathbf{x}_{\alpha i}^\top(u) \boldsymbol{\beta}_\alpha (\mathbf{x}_{\mu i}^\top(u) \boldsymbol{\beta}_\mu)] du \\ &\quad - \frac{N}{2} \log(2\pi) - \mathbf{1}_N^\top \mathbf{X}_\sigma(\mathbf{t}) \boldsymbol{\beta}_\sigma - \frac{1}{2} (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \boldsymbol{\beta}_\mu)^\top \mathbf{R}^{-1} (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \boldsymbol{\beta}_\mu) \end{aligned}$$

resulting in the log-posterior

$$\log p(\boldsymbol{\theta}|\mathbf{T}, \boldsymbol{\delta}, \mathbf{y}) \propto \ell[\boldsymbol{\theta}|\mathbf{T}, \boldsymbol{\delta}, \mathbf{y}] + \sum_{k \in \{\lambda, \gamma, \alpha, \mu, \sigma\}} \sum_{m=1}^{M_k} [\log p(\boldsymbol{\beta}_{km} | \boldsymbol{\tau}_{km}^2) + \log p(\boldsymbol{\tau}_{km}^2)].$$

The full score vectors  $\mathbf{s}(\boldsymbol{\beta}_k)$  and Hessians  $\mathbf{H}(\boldsymbol{\beta}_k)$  of the respective predictors are computed as the sum of the score  $\mathbf{s}^*(\boldsymbol{\beta}_k)$  and Hessians  $\mathbf{H}^*(\boldsymbol{\beta}_k)$  based on the log-likelihood function and the score and Hessian of the respective log-prior densities. For the multivariate normal prior as specified in 3.2 these are  $-\frac{1}{\tau_{km}^2} \mathbf{K}_{km} \boldsymbol{\beta}_{km}$  and  $-\frac{1}{\tau_{km}^2} \mathbf{K}_{km}$ , respectively.

### 6.1 Score Vectors

$$\begin{aligned} \mathbf{s}^*(\boldsymbol{\beta}_\mu) &= \frac{\partial \ell}{\partial \boldsymbol{\beta}_\mu} = \mathbf{X}_\mu(\mathbf{t})^\top \mathbf{R}^{-1} (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \boldsymbol{\beta}_\mu) + \mathbf{X}_\mu^\top(\mathbf{T}) \text{diag}(\boldsymbol{\delta}) [\mathbf{X}_\alpha(\mathbf{T}) \boldsymbol{\beta}_\alpha] \\ &\quad - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \boldsymbol{\beta}_\gamma) \int_0^{T_i} \omega_i(u) \mathbf{x}_{\alpha i}^\top(u) \boldsymbol{\beta}_\alpha \mathbf{x}_{\mu i}(u) du \\ \mathbf{s}^*(\boldsymbol{\beta}_\gamma) &= \frac{\partial \ell}{\partial \boldsymbol{\beta}_\gamma} = \boldsymbol{\delta}^\top \mathbf{X}_\gamma - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \boldsymbol{\beta}_\gamma) \mathbf{x}_{\gamma i} \int_0^{T_i} \omega_i(u) du \\ \mathbf{s}^*(\boldsymbol{\beta}_\alpha) &= \frac{\partial \ell}{\partial \boldsymbol{\beta}_\alpha} = \mathbf{X}_\alpha^\top(\mathbf{T}) \text{diag}(\boldsymbol{\delta}) [\mathbf{X}_\mu(\mathbf{T}) \boldsymbol{\beta}_\mu] - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \boldsymbol{\beta}_\gamma) \int_0^{T_i} \omega_i(u) \mathbf{x}_{\alpha i}(u) (\mathbf{x}_{\mu i}^\top(u) \boldsymbol{\beta}_\mu) du \\ \mathbf{s}^*(\boldsymbol{\beta}_\lambda) &= \frac{\partial \ell}{\partial \boldsymbol{\beta}_\lambda} = \boldsymbol{\delta}^\top \mathbf{X}_\lambda(\mathbf{T}) - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \boldsymbol{\beta}_\gamma) \int_0^{T_i} \omega_i(u) \mathbf{x}_{\lambda i}(u) du \\ \mathbf{s}^*(\boldsymbol{\beta}_\sigma) &= \frac{\partial \ell}{\partial \boldsymbol{\beta}_\sigma} = -\mathbf{X}_\sigma(\mathbf{t})^\top \mathbf{1}_N + [\mathbf{X}_\sigma(\mathbf{t}) \odot (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \boldsymbol{\beta}_\mu)]^\top \mathbf{R}^{-1} (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \boldsymbol{\beta}_\mu) \end{aligned}$$

## 6.2 Hessian

$$\begin{aligned}
\mathbf{H}^*(\beta_\mu) &= \frac{\partial^2 \ell}{\partial \beta_\mu \partial \beta_\mu^\top} = -\mathbf{X}_\mu(\mathbf{t})^\top \mathbf{R}^{-1} \mathbf{X}_\mu(\mathbf{t}) \\
&\quad - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \beta_\gamma) \int_0^{T_i} \omega_i(u) [\mathbf{x}_{\alpha i}^\top(u) \beta_\alpha]^2 \mathbf{x}_{\mu i}(u) \mathbf{x}_{\mu i}^\top(u) du \\
\mathbf{H}^*(\beta_\gamma) &= \frac{\partial^2 \ell}{\partial \beta_\gamma \partial \beta_\gamma^\top} = - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \beta_\gamma) \mathbf{x}_{\gamma i} \mathbf{x}_{\gamma i}^\top \int_0^{T_i} \omega_i(u) du \\
\mathbf{H}^*(\beta_\alpha) &= \frac{\partial^2 \ell}{\partial \beta_\alpha \partial \beta_\alpha^\top} = - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \beta_\gamma) \int_0^{T_i} \omega_i(u) [\mathbf{x}_{\mu i}^\top(u) \beta_\mu]^2 \mathbf{x}_{\alpha i}(u) \mathbf{x}_{\alpha i}^\top(u) du \\
\mathbf{H}^*(\beta_\lambda) &= \frac{\partial^2 \ell_i}{\partial \beta_\lambda \partial \beta_\lambda^\top} = - \sum_{i=1}^n \exp(\mathbf{x}_{\gamma i}^\top \beta_\gamma) \int_0^{T_i} \omega_i(u) \mathbf{x}_{\lambda i}(u) \mathbf{x}_{\lambda i}^\top(u) du \\
\mathbf{H}^*(\beta_\sigma) &= \frac{\partial^2 \ell}{\partial \beta_\sigma \partial \beta_\sigma^\top} = -2 [\mathbf{X}_\sigma(\mathbf{t}) \odot (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \beta_\mu)]^\top \mathbf{R}^{-1} [\mathbf{X}_\sigma(\mathbf{t}) \odot (\mathbf{y} - \mathbf{X}_\mu(\mathbf{t}) \beta_\mu)]
\end{aligned}$$

where  $\omega_i(u) = \exp[\mathbf{x}_{\lambda i}^\top(u) \beta_\lambda + \mathbf{x}_{\alpha i}^\top(u) \beta_\alpha (\mathbf{x}_{\mu i}^\top(u) \beta_\mu)]$  and  $\mathbf{R} = \text{diag}(\exp[\mathbf{X}_\sigma(\mathbf{t}) \beta_\sigma]^2)$ .

## References

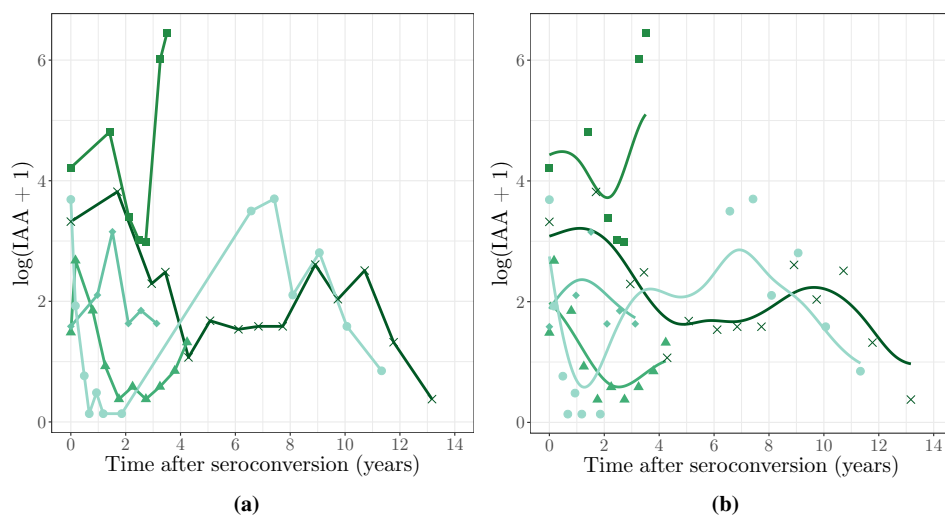
- Barrett, J., Diggle, P., Henderson, R., and Taylor-Robinson, D. (2015). Joint modelling of repeated measurements and time-to-event outcomes: flexible model specification and exact likelihood inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **77**, 131–148.
- Belitz, C. and Lang, S. (2008). Simultaneous selection of variables and smoothing parameters in structured additive regression models. *Computational Statistics & Data Analysis* **53**, 61–81.
- Bender, R., Augustin, T., and Blettner, M. (2005). Generating survival times to simulate Cox proportional hazards models. *Statistics in Medicine* **24**, 1713–1723.
- Bonifacio, E. (2015). Predicting Type 1 Diabetes Using Biomarkers. *Diabetes Care* **38**, 989–996. URL <http://care.diabetesjournals.org/content/38/6/989>.
- Brown, E. R. (2009). Assessing the association between trends in a biomarker and risk of event with an application in pediatric HIV/AIDS. *The Annals of Applied Statistics* **3**, 1163–1182.
- Brown, E. R. and Ibrahim, J. G. (2003). A Bayesian semiparametric joint hierarchical model for longitudinal and survival data. *Biometrics* **59**, 221–228.
- Brown, E. R., Ibrahim, J. G., and DeGruttola, V. (2005). A flexible B-spline model for multiple longitudinal biomarkers and survival. *Biometrics* **61**, 64–73.
- Crowther, M. J. and Lambert, P. C. (2013). Simulating biologically plausible complex survival data. *Statistics in Medicine* **32**, 4118–4134.
- Daher Abdi, Z., Essig, M., Rizopoulos, D., Le Meur, Y., PrÁl'maud, A., *et al.* (2013). Impact of longitudinal exposure to mycophenolic acid on acute rejection in renal-transplant recipients using a joint modeling approach. *Pharmaceutical Research* **72**, 52–60. URL <http://www.sciencedirect.com/science/article/pii/S1043661813000583>.
- Ding, J. and Wang, J.-L. (2008). Modeling longitudinal data with nonparametric multiplicative random effects jointly with survival data. *Biometrics* **64**, 546–556.
- Eilers, P. H. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical Science* **11**, 89–102. URL <http://www.jstor.org/stable/2246049>.
- Fahrmeir, L., Kneib, T., and Lang, S. (2004). Penalized structured additive regression for space-time data: a Bayesian perspective. *Statistica Sinica* **14**, 731–762.

- Faucett, C. L. and Thomas, D. C. (1996). Simultaneously modelling censored survival data and repeatedly measured covariates: A Gibbs sampling approach. *Statistics in Medicine* **15**, 1663–1685.
- Gellar, J. E., Colantuoni, E., Needham, D. M., and Crainiceanu, C. M. (2014). Variable-Domain Functional Regression for Modeling ICU Data. *Journal of the American Statistical Association* **109**, 1425–1439. URL <http://dx.doi.org/10.1080/01621459.2014.940044>.
- Gould, A. L., Boye, M. E., Crowther, M. J., Ibrahim, J. G., Quartey, G., *et al.* (2015). Joint modeling of survival and longitudinal non-survival data: current methods and issues. Report of the DIA Bayesian joint modeling working group. *Statistics in Medicine* **34**, 2181–2195.
- Gras, L., Geskus, R. B., Jurriaans, S., Bakker, M., van Sighem, A., *et al.* (2013). Has the Rate of CD4 Cell Count Decline before Initiation of Antiretroviral Therapy Changed over the Course of the Dutch HIV Epidemic among MSM? *PLoS ONE* **8**. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3664616/>.
- Guo, X. and Carlin, B. P. (2004). Separate and joint modeling of longitudinal and event time data using standard computer packages. *The American Statistician* **58**, 16–24.
- Harrell, F. E., Lee, K. L., and Mark, D. B. (1996). Multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine* **15**, 361–387.
- Hatfield, L. A., Boye, M. E., Hackshaw, M. D., and Carlin, B. P. (2012). Multilevel Bayesian Models for Survival Times and Longitudinal Patient-Reported Outcomes With Many Zeros. *Journal of the American Statistical Association* **107**, 875–885. URL <http://www.tandfonline.com/doi/abs/10.1080/01621459.2012.664517>.
- Hummel, M., Bonifacio, E., Schmid, S., Walter, M., Knopff, A., *et al.* (2004a). Brief communication: early appearance of islet autoantibodies predicts childhood type 1 diabetes in offspring of diabetic parents. *Annals of Internal Medicine* **140**, 882–886.
- Hummel, M., Bonifacio, E., Schmid, S., Walter, M., Knopff, A., *et al.* (2004b). Brief Communication: Early Appearance of Islet Autoantibodies Predicts Childhood Type 1 Diabetes in Offspring of Diabetic Parents. *Annals of Internal Medicine* **140**, 882–886. URL <http://dx.doi.org/10.7326/0003-4819-140-11-200406010-00009>.
- Hummel, S., Pflüger, M., Hummel, M., Bonifacio, E., and Ziegler, A.-G. (2011). Primary Dietary Intervention Study to Reduce the Risk of Islet Autoimmunity in Children at Increased Risk for Type 1 Diabetes. *Diabetes Care* **34**, 1301–1305. URL <http://care.diabetesjournals.org/content/34/6/1301>.
- Hurvich, C. M., Simonoff, J. S., and Tsai, C.-L. (1998). Smoothing parameter selection in nonparametric regression using an improved Akaike information criterion. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60**, 271–293.
- Insel, R. A., Dunne, J. L., Atkinson, M. A., Chiang, J. L., Dabelea, D., *et al.* (2015). Staging Presymptomatic Type 1 Diabetes: A Scientific Statement of JDRF, the Endocrine Society, and the American Diabetes Association. *Diabetes Care* **38**, 1964–1974. URL <http://care.diabetesjournals.org/content/38/10/1964>.
- Jiang, B., Wang, N., Sammel, M. D., and Elliott, M. R. (2015). Modelling short- and long-term characteristics of follicle stimulating hormone as predictors of severe hot flashes in the Penn Ovarian Aging Study. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **64**, 731–753.
- Klein, N., Kneib, T., Klasen, S., and Lang, S. (2015a). Bayesian Structured Additive Distributional Regression for Multivariate Responses. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **64**, 569–591.
- Klein, N., Kneib, T., Lang, S., and Sohn, A. (2015b). Bayesian Structured Additive Distributional Regression with an Application to Regional Income Inequality in Germany. *Annals of Applied Statistics* **9**, 1024–1052.
- Lang, S. and Brezger, A. (2004). Bayesian P-Splines. *Journal of Computational and Graphical Statistics* **13**, 183–212.
- Liu, L. (2009). Joint modeling longitudinal semi-continuous data and survival, with application to longitudinal medical cost data. *Statistics in Medicine* **28**, 972–986. URL <http://doi.wiley.com/10.1002/sim.3497>.
- Malfait, N. and Ramsay, J. O. (2003). The historical functional linear model. *Canadian Journal of Statistics* **31**, 115–128. URL <http://onlinelibrary.wiley.com/doi/10.2307/3316063/full>.
- Meyer, S., Woodward, M., Hertel, C., Vlaicu, P., Haque, Y., *et al.* (2016). AIRE-Deficient Patients Harbor Unique High-Affinity Disease-Ameliorating Autoantibodies. *Cell* **166**, 582–595. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4967814/>.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.

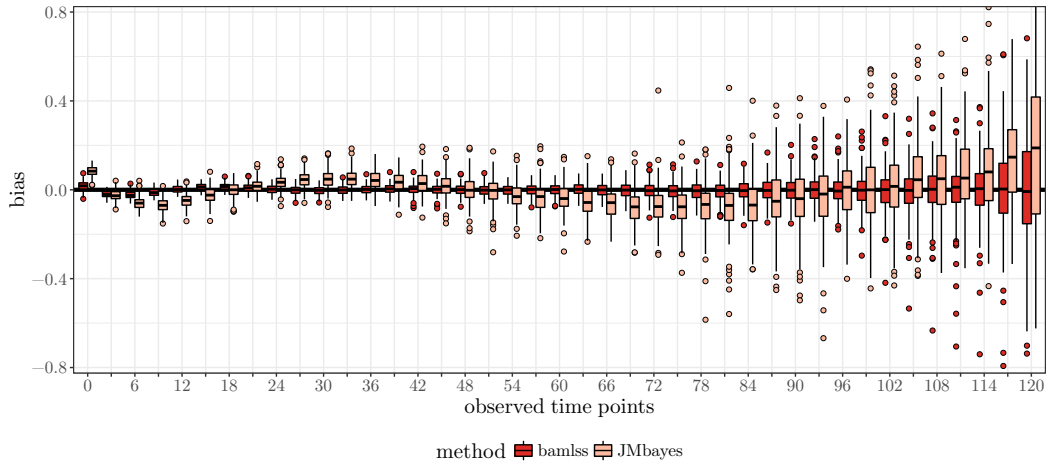
- Rizopoulos, D. (2012). *Joint Models for Longitudinal and Time-to-event Data, with Applications in R*. Chapman and Hall/CRC, Boca Raton, Florida.
- Rizopoulos, D. (2016a). *JMbayes: Joint Modeling of Longitudinal and Time-to-Event Data under a Bayesian Approach*. URL <https://CRAN.R-project.org/package=JMbayes>, R package version 0.7-9.
- Rizopoulos, D. (2016b). The R Package JMbayes for fitting joint models for longitudinal and time-to-event data using MCMC. *Journal of Statistical Software* **72**, 1–46. URL <https://www.jstatsoft.org/index.php/jss/article/view/v072i07>.
- 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. URL <http://onlinelibrary.wiley.com/doi/10.1002/sim.4205/abstract>.
- Rizopoulos, D., Hatfield, L. A., Carlin, B. P., and Takkenberg, J. J. M. (2014). Combining dynamic predictions from joint models for longitudinal and time-to-event data using Bayesian model averaging. *Journal of the American Statistical Association* **109**, 1385–1397.
- Rizopoulos, D., Verbeke, G., Lesaffre, E., and Vanrenterghem, Y. (2008). A Two-Part Joint Model for the Analysis of Survival and Longitudinal Binary Data with Excess Zeros. *Biometrics* **64**, 611–619. URL <http://onlinelibrary.wiley.com.amedien.ub.uni-muenchen.de/doi/10.1111/j.1541-0420.2007.00894.x/abstract>.
- Scheipl, F., Staicu, A.-M., and Greven, S. (2015). Functional additive mixed models. *Journal of Computational and Graphical Statistics* **24**, 477–501.
- Song, X. and Wang, C. Y. (2008). Semiparametric approaches for joint modeling of longitudinal and survival data with time-varying coefficients. *Biometrics* **64**, 557–566.
- Steck, A. K., Johnson, K., Barriga, K. J., Miao, D., Yu, L., *et al.* (2011). Age of islet autoantibody appearance and mean levels of insulin, but not GAD or IA-2 autoantibodies, predict age of diagnosis of type 1 diabetes: Diabetes Autoimmunity Study in the Young (TEDDY). *Diabetes Care* **34**, 1397–1399. URL <http://care.diabetesjournals.org/cgi/doi/10.2337/dc10-2088>.
- Steck, A. K., Vehik, K., Bonifacio, E., Lernmark, A., Ziegler, A.-G., *et al.* (2015). Predictors of Progression From the Appearance of Islet Autoantibodies to Early Childhood Diabetes: The Environmental Determinants of Diabetes in the Young (TEDDY). *Diabetes Care* **38**, 808–813. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4407751/>.
- Tang, A.-M. and Tang, N.-S. (2015). Semiparametric Bayesian inference on skew-normal joint modeling of multivariate longitudinal and survival data. *Statistics in Medicine* **34**, 824–843.
- Tang, A.-M., Zhao, X., and Tang, N.-S. (2017). Bayesian variable selection and estimation in semiparametric joint models of multivariate longitudinal and survival data. *Biometrical Journal* **59**, 57–78. URL <http://onlinelibrary.wiley.com.amedien.ub.uni-muenchen.de/doi/10.1002/bimj.201500070/abstract>.
- Taylor, J. M. G., Park, Y., Ankerst, D. P., Proust-Lima, C., Williams, S., *et al.* (2013). Real-time individual predictions of prostate cancer recurrence using joint models. *Biometrics* **69**, 206–213.
- Tsiatis, A. A. and Davidian, M. (2004). Joint modeling of longitudinal and time-to-event data: an overview. *Statistica Sinica* **14**, 809–834.
- Umlauf, N., Klein, N., and Zeileis, A. (2017). BAMLSS: Bayesian Additive Models for Location, Scale and Shape (and Beyond). Working Paper, Faculty of Economics and Statistics, University of Innsbruck. URL <http://econpapers.repec.org/paper/innwpaper/2017-05.htm>.
- Umlauf, N., Klein, N., Zeileis, A., and Köhler, M. (2016). *bamlss: Bayesian Additive Models for Location Scale and Shape (and Beyond)*. R package version 0.1-1.
- Waldmann, E., Taylor-Robinson, D., Klein, N., Kneib, T., Pressler, T., *et al.* (2017). Boosting joint models for longitudinal and time-to-event data. *Biometrical Journal*, n/a–n/a URL <http://onlinelibrary.wiley.com.amedien.ub.uni-muenchen.de/doi/10.1002/bimj.201600158/abstract>.
- Wood, S. N. (2006). *Generalized additive models: an introduction with R*. Chapman & Hal/CRC, Boca Raton, Florida.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**, 3–36.
- Wood, S. N., Pya, N., and Säfken, B. (2016). Smoothing Parameter and Model Selection for General Smooth Models. *Journal of the American Statistical Association* **111**, 1548–1563. URL <https://www.tandfonline.com/doi/full/10.1080/01621459.2016.1180986>.

- Ye, W., Lin, X., and Taylor, J. M. G. (2008). Semiparametric Modeling of Longitudinal Measurements and Time-to-Event Data: A Two-Stage Regression Calibration Approach. *Biometrics* **64**, 1238–1246. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1541-0420.2007.00983.x/abstract>.
- Ziegler, A. G., Hillebrand, B., Rabl, W., Mayrhofer, M., Hummel, M., *et al.* (1993). On the appearance of islet associated autoimmunity in offspring of diabetic mothers: a prospective study from birth. *Diabetologia* **36**, 402–408.
- Ziegler, A.-G., Hummel, M., Schenker, M., and Bonifacio, E. (1999). Autoantibody appearance and risk for development of childhood diabetes in offspring of parents with type 1 diabetes: the 2-year analysis of the German BABY-DIAB Study. *Diabetes* **48**, 460–468. URL <http://diabetes.diabetesjournals.org/content/48/3/460.short>.
- Ziegler, A. G., Rewers, M., Simell, O., Simell, T., Lempainen, J., *et al.* (2013). Seroconversion to multiple islet autoantibodies and risk of progression to diabetes in children. *JAMA* **309**, 2473–2479.


## Figures and Tables



**Figure 1** Longitudinal marker values of  $\log(IAA + 1)$  for five randomly selected subjects in the BABY-DIAB/BABYDIET study. (a) Observed values (points) and linear interpolation (lines); (b) Observed values (points) and estimated trajectories (lines).

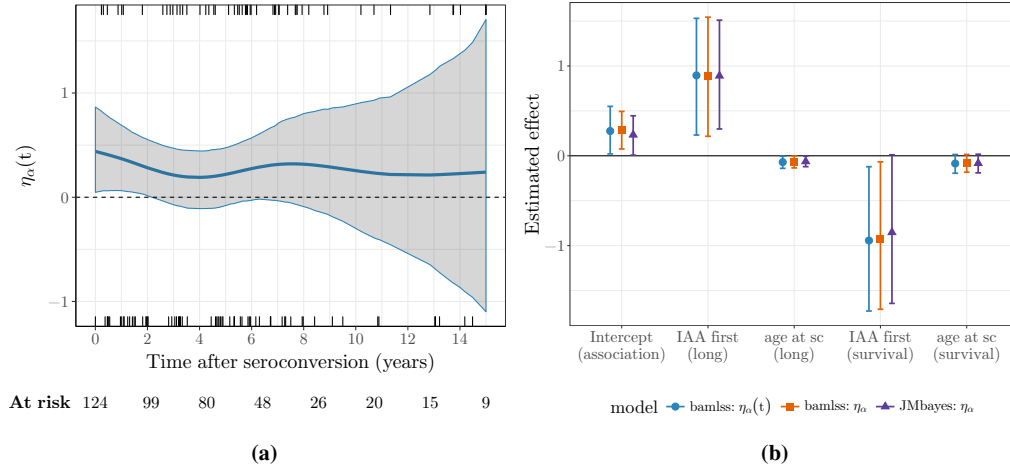


**Figure 2** Comparison of the bias over time for  $\eta_\mu(t)$  estimates from `bamlss` and `JMbayes` in setting 1a.



**Figure 3** True (black) and estimated (grey) predictors from posterior mean and posterior mode estimation for small (a) and large (b) data sets in simulation setting 2. Left:  $\eta_\lambda(t)$ ; middle:  $\eta_\alpha(t)$ ; right:  $\eta_\lambda(t) + \eta_\alpha(t) \cdot \frac{1}{n} \sum_{i=1}^n \eta_{\mu i}(t)$ .





**Figure 4** Estimated effects in the BABYDIAB/BABYDIET data. (a): Estimated posterior mean of  $\eta_\alpha(t)$  with 95% pointwise credibility bands (shaded area), observed event times (rugs bottom), and censoring times (rugs top). (b): linear effect estimates from the survival and longitudinal submodel from model fits based on (i) *bamlss* with time-varying  $\eta_\alpha(t)$ , (ii) *bamlss* with time-constant  $\eta_\alpha$  and (iii) *JMBayes*. *association* stands for effects in  $\eta_\alpha$ , *long* for the longitudinal submodel, *sc* indicates seroconversion.

**Table 1** Overview of the predictor vectors, function evaluations and design matrices in the survival and longitudinal submodel.

	predictor vector	function evaluation	design matrix
survival model			
$k \in \{\gamma\}$	$\boldsymbol{\eta}_k = [\eta_{ki}]^\top$ $n \times 1$	$\boldsymbol{f}_k = [f_k(\boldsymbol{x}_{ki})]^\top$ $n \times 1$	$\mathbf{X}_k$ $n \times p_k$
$k \in \{\lambda, \alpha, \mu\}$	$\boldsymbol{\eta}_k(t) = [\eta_{ki}(t)]^\top$ $n \times 1$	$\boldsymbol{f}_k(t) = [f_k(\boldsymbol{x}_{ki}(t), t)]^\top$ $n \times 1$	$\mathbf{X}_k(t)$ $n \times p_k$
longitudinal model			
$k \in \{\mu, \sigma\}$	$\boldsymbol{\eta}_k(\mathbf{t}) = [\eta_{ki}(\mathbf{t}_i)]^\top$ $N \times 1$	$\boldsymbol{f}_k(\mathbf{t}) = [f_k(\boldsymbol{x}_{ki}(\mathbf{t}_i), \mathbf{t}_i)]^\top$ $N \times 1$	$\mathbf{X}_k(\mathbf{t})$ $N \times p_k$

For ease of notation we denote the vector  $\mathbf{a}^\top = [a_1, \dots, a_n]$  as  $[a_i]$  for  $i = 1, \dots, n$  and drop the subscript  $m$  for the different terms per predictor in this illustration.

**Table 2** Effects  $f_{km}(\mathbf{x}_{ki})$  that can be specified within a predictor  $\eta_k$  in structured additive joint models; modified from a similar table in Scheipl *et al.* (2015).

covariate (subset of $\mathbf{x}$ )	$f_{km}(\mathbf{x}_k)$ constant over $t$	$f_{km}(\mathbf{x}_k)$ varying over $t$
no covariate	scalar intercept $1 \cdot \beta$	smooth effect of time $f(t)$
scalar covariate $z$	linear effect $z \cdot \beta$ smooth effect $f(z)$	linear effect varying over time $z \cdot f(t)$ smooth effect over time $f(z, t)$
spatial covariate(s) $s$	spatial effect $f(s)$	spatial effect over time $f(s, t)$
grouping variable $g$	random intercept $\beta_g$	functional random intercept $f_g(t)$
scalar and grouping variable	random slope $z \cdot \beta_g$	functional random slope $z \cdot f_g(t)$
vector of scalars $[z_1, z_2]$	linear interaction $z_1 \cdot z_2 \cdot \beta$ varying coefficient $z_1 \cdot f(z_2)$ smooth effect $f(z_1, z_2)$	linear interaction over time $z_1 \cdot z_2 \cdot f(t)$

**Table 3** Posterior mean simulation results from `bamlss` and results from `JMbayes` from setting 1 (time-constant  $\eta_\alpha$ ) for small (*a*) and large (*b*) data sets.

		MSE		bias		coverage	
		a	b	a	b	a	b
$\eta_\alpha$	<code>bamlss</code>	0.032	0.016	0.003	-0.009	0.925	0.970
	<code>JMbayes</code>	0.049	0.021	0.100	0.048	0.840	0.890
$\eta_\gamma + \eta_\lambda$	<code>bamlss</code>	0.127	0.077	-0.007	0.011	0.935	0.946
	<code>JMbayes</code>	0.155	0.101	-0.095	-0.048	0.743	0.742
$\eta_\mu$	<code>bamlss</code>	0.022	0.028	0.001	0.000	0.944	0.942
	<code>JMbayes</code>	0.031	0.029	-0.001	0.008	*	*
$\eta_\sigma$	<code>bamlss</code>	0.001	0.001	0.009	0.014	0.940	0.875
	<code>JMbayes</code>	0.007	0.002	0.080	0.039	*	*

\* No credibility intervals and thus no coverage could be calculated for these predictors.

**Table 4** Posterior mode and posterior mean simulation results for setting 2 (time-varying  $\eta_\alpha(t)$ ) for small (*a*) and large (*b*) data sets.

		MSE		bias		coverage	
		a	b	a	b	a	b
$\eta_\alpha$	mean	0.171	0.078	0.007	0.002	0.940	0.961
	mode	0.177	0.117	0.058	0.069	0.608	0.593
$\eta_\gamma$	mean	0.097	0.062	-0.035	-0.032	0.931	0.948
	mode	0.089	0.059	0.022	-0.001	0.804	0.795
$\eta_\lambda$	mean	0.083	0.065	0.000	0.000	0.945	0.957
	mode	0.101	0.082	0.000	0.000	0.592	0.549
$\eta_\mu$	mean	0.022	0.028	0.000	0.000	0.943	0.942
	mode	0.025	0.031	0.000	0.000	0.882	0.865
$\eta_\sigma$	mean	0.001	0.001	0.009	0.015	0.905	0.855
	mode	0.004	0.004	-0.057	-0.057	0.175	0.045