Dynamic prediction in event history analysis

Grand, M.K.

5

Dynamic prediction with a joint model

Uveitis is characterised as a recurrent inflammation of the eye and an ongoing inflammation can have severe impact on the visual acuity of the patient. The Rotterdam Eye Hospital has been collecting data on every uveitis patient visiting the hospital since 2000. We propose a joint model for the inflammation and visual acuity with the purpose of making dynamic predictions. Dynamic prediction models allow predictions to be updated during the follow-up of the patient based on the patient's disease history.

The joint model consists of a submodel for the inflammation, the event history outcome, and one for the visual acuity, the longitudinal outcome. The inflammation process is described with a two state reversible multi-state model, where transition times are interval censored. Correlated log-normal frailties are included in the multi-state model to account for the within eye and within patient correlation. A linear mixed model is

used for the visual acuity. The joint model is fitted in a two-stage procedure and we illustrate how the model can be used to make dynamic predictions. The performance of the method was investigated in a simulation study. The novelty of the proposed model includes the extension to a multi-state outcome, whereas previously the standard has been to consider survival or competing risk outcomes. Furthermore, it is usually the case that the longitudinal outcome affects the event history outcome, but in this model the relation is reversed.

5.1 Introduction

Uveitis is an intraocular inflammation of the uvea, which typically is episodic. An active inflammation can be very painful for the patient. After the onset of the disease it is vital that the patient is provided with proper treatment to keep the inflammation under control. An untreated inflamed eye will over time progress towards poorer visual acuity, but correct treatment can suppress the inflammation, and the eye may over time recover and regain visual acuity. In $35 - 50\%$ of cases, there is no known cause[31] and the interplay between the two eyes is unresolved. Unlike other eye diseases, that usually affect the elderly, uveitis affects all ages. Accurate assessment of the risk of inflammation and poor visual acuity is highly relevant for these patients as uveitis is one of the leading causes of preventable legal blindness in developed countries[31].

We propose a joint model for dynamic prediction of visual acuity and inflammation for patients with uveitis. The data that motivated the joint model was collected at the Rotterdam Eye Hospital, and it is comprised of uveitis patients that started visiting the hospital in the period from 2000 to 2014. Most previous studies on uveitis have been cross-sectional, so the longitudinal data collected in Rotterdam offers unique possibilities to understand how different risk factors affect the disease progression.

Early key papers on joint models for event history and longitudinal outcomes include Faucett and Thomas[28] and Wulfsohn and Tsiatis[99], and joint models have since been an increasingly popular research field. A somewhat recent overview can be found in Diggle

et al[27]. The classical example is when a biomarker is measured repeatedly over time which may be related to a time to event outcome such as death. There are three main objectives for employing a joint model. The objective can be to analyse either the time to event outcome or the longitudinal outcome or to study the relationship between the two. Our main objective is to analyse the longitudinal outcome, since visual acuity is what ultimately matters for the patients. However, the inflammation process, which we consider as a time to event outcome, is also of secondary interest. Often when the longitudinal outcome is the object of interest the joint model approach is used to correct for informative censoring[98,84]. This is however not the case here, since the changes in the inflammation do not terminate the measurements of the visual acuity, which would otherwise be the case if the time to event outcome where death. Instead, our motive to employ a joint model is based on clinical considerations; that the time spent with an active inflammation or the time spent in recovery is what drives the progression of the visual acuity[31]. However, the exact time of transition from one state to another is interval censored in our data, since the inflammation status is only observed at the visits to the hospital. Other examples of joint models for an interval censored time to event outcome can be found in Gueorguieva et al[44] and Rouanet et al[79]. We used random effects both to account for the dependence of observations within an individual and within an eye as well as to allow for individualised predictions. Using correlated random effects, rather than just one shared random effect, has been a popular way of connecting the longitudinal and time to event outcomes[43,87]. Given the complexity of the joint model, particularly the random effects structure, we employed a two-stage approach to estimating the parameters of the joint model. Two-stage approaches has been criticized as being subject to possible bias and poor coverage[86]. Nevertheless our two-stage approach differs from the conventional approach in a number of ways and we conducted a simulation study to evaluate the performance of the proposed estimation procedure.

Joint models can be used for dynamic prediction[70,76], where predictions are updated based on the information that is available on the patient at a given time during follow-

up. Predictions may change over the follow-up due to changes in the patient's covariates or due to changes in the effect of the covariates or the baseline. Early work on dynamic prediction used a Cox model with time-varying covariates [47,95], and van Houwelingen [92] proposed to use landmarking [11]. Although joint models are usually more complex than the alternatives, they may also provide more insight, as both the longitudinal and time to event outcome are modelled.

We start by describing the data from the Rotterdam Eye Hospital in Section 5.2. In Section 5.3 we describe the joint model, how the estimation is carried out and how the joint model can be used to make dynamic prediction of both outcomes. To evaluate the performance of the proposed estimation procedure we conducted a simulation study described in Section 5.4. In Section 5.5 we show the results of fitting the joint model to the uveitis data along with the results from a sensitivity analysis of the assumptions. Section 5.6 is devoted to discussion. Additional results from the uveitis data and the simulation study are provided in the Supporting Information.

5.2 Uveitis data

The data consists of 366 uveitis patients that started frequenting the Rotterdam Eye Hospital in the period from 2000 to 2014. These patients contributed with data on 714 eyes and 10816 observations, with a mean follow-up time of 2.5 years and the mean number of visits was 15. The visits were in principle prescheduled, and patients would only be discharged from the hospital after five years without any inflammation episodes. At each visit information was collected on the inflammation status, visual acuity and covariates. The inflammation status was either observed to be active (present) or quiescent (inflammation free). However, the exact transition times are unknown, since the inflammation status was only observed at the visits. The total number of observed transitions was 980 to quiescent and 657 to active. The visual acuity was measured on the Snellen scale, where an eye with normal vision would score $20/20 = 1$ and a completely blind eye would score 0.

Data collected on three patients are shown in Figure 5.1. It shows the inflammation status and visual acuity measured at every visit since the patients' first visit to the hospital. Patient A has close to two years of follow-up, where the left eye started out with inflammation and declining visual acuity, but after a while the eye turned quiescent and the visual acuity recovered somewhat. The right eye only had one visit with an active inflammation, and the visual acuity did not change as much as it did in the left eye. Patient B is an example of a patient where only one eye seemed to be affected by the disease. In contrast to patient B, patient C is an example of a patient where the visual acuity and to some degree the inflammation on both eyes followed similar patterns. Patients B and C illustrate that for most patients uveitis takes on a chronic nature and in these cases only proper treatment may help suppress future episodes. Furthermore, Figure 5.1 illustrates that there is a high level of heterogeneity between these patients and that the visual acuity is affected by the status of the inflammation process.

The covariates include age, early onset, treatments, surgeries and complications. Table 5.1 contains a summary of the covariates in the data set. The few missing values ($< 5\%$) have been replaced by the value at the previous visit. The patient level covariates are also baseline covariates. Age is defined as the patient's age at the first visit to the hospital. Visual acuity is expected to decline with age in the general population. Early onset denotes the patients that had more than six weeks between the onset of the first complaints and the first visit to the hospital. Although the number of patients with early onset is small, it is believed to be an important predictor of the outcomes, as early treatment of uveitis is considered to be crucial for future recovery. The eye level covariates are also time-varying, and they are therefore presented on an aggregated level. The patients could receive a whole range of treatments in the form of eye drops, pills or injections in various combinations and with varying intensities. A high intensity treatment increases the suppression of the inflammation, but it also increases the risk of adverse events. All the treatments have been grouped according to intensity as either maintenance or active treatment, i.e. medium or high intensity. The surgeries that were considered clinically relevant for the inflammation were phaco, YAG and vitrec-

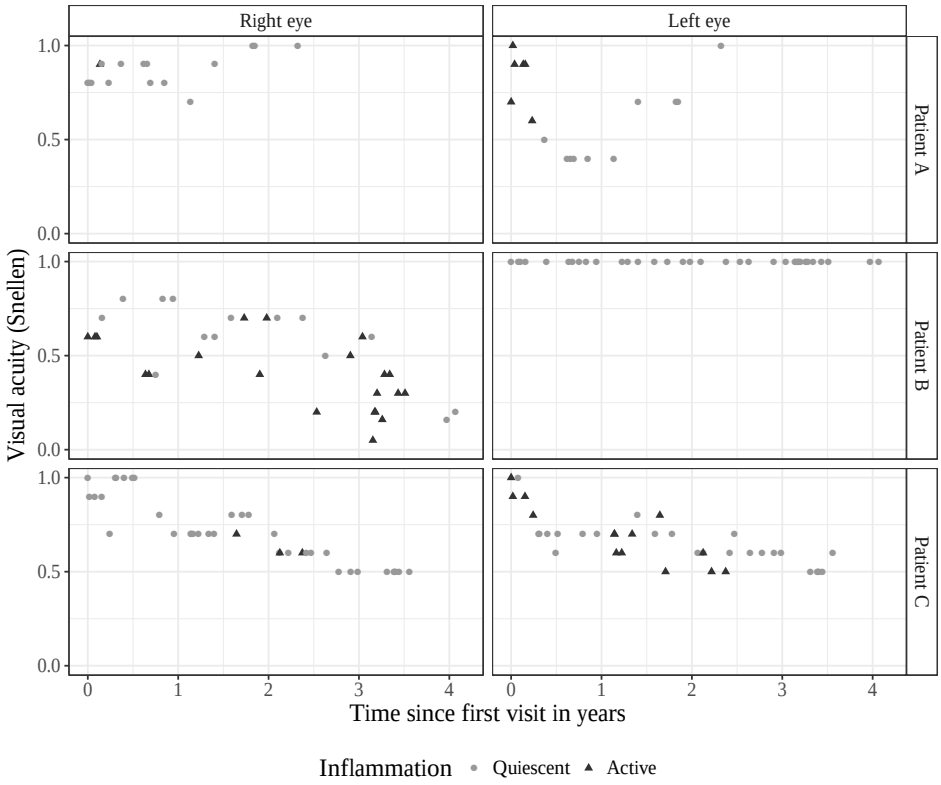

Figure 5.1: Illustration of data collected on both eyes from three selected uveitis patients. The x-axis is time since the first visit, where the patients came to the Rotterdam Eye Hospital with complaints. Every dot is a visit and the observed inflammation status is represented by shape and colour. The visual acuity is depicted on the y-axis, where normal vision is 1 and blind is 0.

Table 5.1: Summary of the covariates in the uveitis data set. All the patient level covariates are also baseline covariates. The eye level covariates are time-varying, and they are therefore presented on an aggregated level.

Covariate	Number	%
Patient level		
Age (mean,sd)	45	18
Early onset		
no	358	98
yes	8	2
Eye level		
Treatment		
no	122	17
maintenance	43	6
active	547	77
missing	2	0
Surgery		
no	598	84
yes	116	16
Complication		
no	456	63
yes	244	34
missing	14	2

tomy surgery. The complications that were considered relevant were macular edema, macular pucker, atrophy, choroidal neovascularization and retinal detachment, and the presence of either one was recoded at each visit. Table 5.1 shows the number of eyes that never received any treatment (no) and how many that had treatment at least once during follow-up (maintenance or active). It also shows how many eyes had at least one of the relevant surgeries performed during follow up (yes) and at least one of the relevant complications (yes).

5.3 Method

Let n be the number of subjects in the sample and let $v_0, \ldots, v_{N_i}$ denote the N_i+1 visit times for the ith patient. The visit times are not necessarily the same for every patient. The time scale is time since the first visit to the hospital, which for most patients is the same as the onset of the disease (Table 5.1). The information collected at the time of the first visit, $v_0 = 0$, is used as baseline information. At each visit we observe the inflammation status $X_{il}(t)$, the visual acuity $Y'_{il}(t)$ and the covariates $Z_{il}(t)$ on both eyes $l \in \{R, L\}$. Throughout $Z_{il}(t)$ will denote the value of the covariates just prior to time t. The joint model consists of two parts; a model for the inflammation and a model for the visual acuity.

5.3.1 Models

Inflammation model

The inflammation process $X_{il}(t)$ can be described by the multi-state model in Figure 5.2. The process can move back and forth between the two states quiescent 1 and active 2. We assume that the transition intensity for making a transition into state g, for eye l of subject i, takes the form

$$\lambda_g(t|Z_{il}(t), b_{ilg}) = \lambda_{g,0} \exp(Z_{il}(t)\beta_g + b_{ilg}) \text{ for } g \in \{1, 2\} \ . \tag{5.1}$$

The baseline transition intensity $\lambda_{g,0}$ is assumed to be constant, which is considered to be reasonable in view of the chronic nature of the disease. The smaller the transition intensity the longer time the process will spend in the current state. The effect β_g of the time-varying covariates $Z_{il}(t)$ is assumed to be time-constant. The eye and subject specific frailty is denoted by b_{ilg}. It is expected that the frailties between the two transitions will be negatively correlated. The frailties are therefore assumed to be multivariate normal, which unlike the gamma distribution also allows the correlation to be negative. It would however be too ambitious to attempt to estimate all variance

and correlation parameters in an unstructured covariance matrix, so we impose some structure. We assume that the vector of frailties b_i for subject i can be decomposed into a component that is common for both eyes and a component that is unique for each eye. Let $b'_i \sim N_2(0, \Sigma_{b'})$ denote the common component and let $b'_{il} \sim N_2(0, \Sigma_{b''})$ denote the component that is unique for eye l. We assume that b'_i, b'_{iR} and b'_{iL} are independent. As a result we have that

$$b_i = \begin{bmatrix} b_{iR1} \\ b_{iR2} \\ b_{iL1} \\ b_{iL2} \end{bmatrix} = \begin{bmatrix} b'_i + b'_{iR} \\ b'_i + b'_{iL} \end{bmatrix} \sim N_4(0, \Sigma_b) \ ,$$

where the variance matrix, due to independence, can be decomposed as

$$\Sigma_b = \begin{bmatrix} \Sigma_{b'} + \Sigma_{b''} & \Sigma_{b'} \\ \Sigma_{b'} & \Sigma_{b'} + \Sigma_{b''} \end{bmatrix} \ .$$

We assume that the inflammation status can change at most once between two visits. In this way we are certain whether or not there was a transition between two visits. So if the inflammation status between two visits was unchanged, then we assume that there were no transitions. If there was a change, then we assume that only one transition took place. Let $T_{il1}, \ldots, T_{ilM_{il}}$ denote the M_{il} unobserved transition times. The first period between the first visit and the first transition will be referred to as spell 0, and the period between the first and the second transition will be referred to as spell 1 etc. Hence, with M_{il} transitions we will have $M_{il} + 1$ spells.

Visual acuity model

The visual acuity is first transformed from the Snellen scale y' to a new scale y given by

$$y = \log\left(\frac{y' + \epsilon_1}{1 - y' + \epsilon_2}\right) \ ,$$

where $\epsilon_1, \epsilon_2 > 0$ are small. The reasoning behind the transformation is that y' is on the Snellen scale, which is bounded and in order to ensure that predictions will stay within the range of the visual acuity scale we transform it to an unbounded scale. Furthermore, the model assumption about normality is more appropriate after the transformation. The visual acuity on the new scale is assumed to follow a linear mixed model

$$Y_{il}(v_j) = \mu_{il}(v_j) + \epsilon_{ilj} \quad \text{for } l \in \{R, L\} \text{ and } j \in \{1, \ldots, N_i\} \ ,$$

where $Y_{il}(t)$ is the visual acuity on the transformed scale at visit time v_j and $\mu_{il}(t)$ is its expectation given random effects, which will be specified in a moment. The error terms ϵ_{ilj} are assumed to be independent and identically distributed with $N(0, \sigma_\epsilon^2)$.

The key motivation for the joint model, and hence the visual acuity model, is that the time that the eye spent with a quiescent or active inflammation, is the driving force behind changes in the visual acuity[31]. We therefore assume that $\mu_{il}(t)$ is a linear function of the time that the eye has spent in the quiescent and active inflammation state.

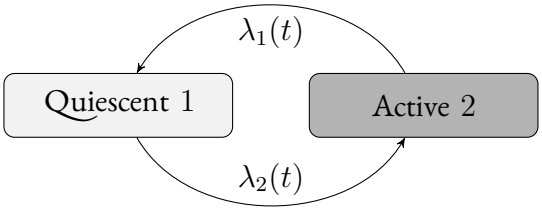

Figure 5.2: Multi-state model describing the inflammation process within the eye.

For now we will carry on as if the transition times of the inflammation process were known. We will discuss later how to incorporate the inherent uncertainty arising from the fact that the transition times are unobserved. Let t_1 and t_2 denote the time that the eye has spent in the quiescent and active inflammation state up until time t, such that $t = t_1 + t_2$. The part of $\mu_{il}(t)$ that does not depend on t_1 or t_2 is referred to as the intercept and the part that does is referred to as the progression of $\mu_{il}(t)$.

The progression part of $\mu_{il}(t)$ is given by

$$\big(Z_{il}^\top(t)\alpha_1 + a_{ilm1}\big)t_1 + \big(Z_{il}^\top(t)\alpha_2 + a_{ilm2}\big)t_2 \quad \text{for } m \in \{0, \ldots, M_{il}\} \ ,$$

where the vectors α_1 and α_2 are the fixed effects of the covariates $Z_{il}(t)$ on the progression part. Hence, $Z_{il}^\top(t)\alpha_1$ and $Z_{il}^\top(t)\alpha_2$ are the fixed effect slopes for time spent in the quiescent or active state. They depend on the covariates, since the presence of complications is expected to have an effect on the slopes. Furthermore, a_{ilm1} and a_{ilm2} denote the random effect part of the slopes. They also depend on time as they are spell-specific and m indicates what spell the eye is in at time t.

The intercept of $\mu_{il}(t)$ is given by

$$Z_{il}^\top(t)\alpha_0 + a_{ilm0}$$

where α_0 is a vector of fixed effect of the covariates $Z_{il}^\top(t)$ and a_{ilm0} denotes the random intercept for spell m and eye l. Since the random intercept is spell-specific, the model allow for discontinuities at the transition times between spells.

Similar to the inflammation model, we also simplify the random effect structure in the visual acuity model by decomposing it into a part that is common within the eye and one that is specific for each spell, as we assume that the random effects between the two eyes are independent. Let $a_{ilm} = [a_{ilm0}, a_{ilm1}, a_{ilm2}]^\top$ denote the vector of the spell specific random effects for eye l on subject i. The vector of all random effects $a_{il} = [a_{il0}^\top, \ldots, a_{ilM_{il}}^\top]^\top$ for eye l on subject i can be decomposed into a contribution from the eye $a'_{il} \sim N_3(0, \Sigma_{a'})$ and from the spells $a'_{ilm} \sim N_3(0, \Sigma_{a''})$. We assume

that $a'_{il}, a'_{il0}, \ldots, a'_{ilM_{il}}$ are independent. As a result we have that

$$a_{il} = \begin{bmatrix} a_{il0} \\ \vdots \\ a_{ilM_{il}} \end{bmatrix} = \begin{bmatrix} a'_{il} + a'_{il0} \\ \vdots \\ a'_{il} + a'_{ilM_{il}} \end{bmatrix} \sim N_{3(M_{il}+1)}(0, \Sigma_a) \ .$$

Hence, the intercept and slopes between spells on the same eye are allowed to be dependent. As mentioned earlier, the random effects between the two eyes on the same subject, a_{iR} and a_{iL}, are assumed to be independent. Furthermore, the random effects from the inflammation model are assumed to be independent from the random effects and error terms from the visual acuity model. The visual acuity model could be simplified by assuming that the random effects are the same for all spells within an eye, and thus that there is only one random intercept and slope for each eye. We explore this later in Section 5.5.

Joint model

An illustration of the dependence between the variables and the random effects in the joint model is shown in Figure 5.3. It includes both the unobserved (circles) and observed variables (squares). It illustrates that any correlation between two eyes' visual acuity is induced by the frailty term in the inflammation model. The joint model relies on a number of assumptions, and we list the essential ones below:

- The visit times are non-informative.
- Missing values are missing at random[80].
- The inflammation process changes at most once between two visit times.
- The baseline transition intensities are constant.
- Given the inflammation status the visual acuity processes from the two eyes are independent.

- Expected visual acuity, on the new scale, is a linear function of time spent with and without inflammation.

- Censoring is independent of the inflammation and visual acuity processes.

Most of these assumptions are based on clinical input. Nonetheless it is important, if possible, to verify them from the data or conduct sensitivity analyses. To this end, we performed a sensitivity analysis of the first assumption in Section 5.5 and the rest is left for the discussion.

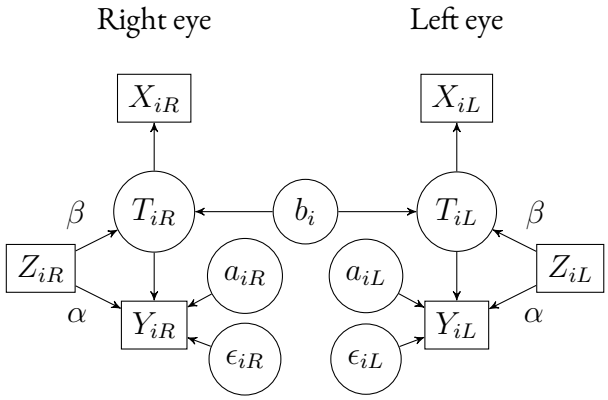

Figure 5.3: Illustration of the dependencies in the joint model between unobserved (circles) and observed variables (squares).

5.3.2 Estimation

Let X, Y and $Z = (Z_Y, Z_X)$ denote the observed data, i.e. the status of the inflammation, the visual acuity and the covariates, which are all observed at every visit. Let T denote the unobserved transition times and let a, b denote the unobserved random effects of the visual acuity and inflammation model. Let $\theta = (\theta_Y, \theta_X) = \big((\alpha, \Sigma_a, \sigma_\epsilon), (\beta, \Sigma_b)\big)$ be the collection of all the parameters in the joint model. The observed data

likelihood, conditional on the covariates, can be decomposed as

$$\begin{aligned} L^*(\theta|Y,X,Z) &= P(Y,X|Z,\theta) \\ &= P(Y|X,Z,\theta)\, P(X|Z,\theta) \\ &= E\Big(P(Y|T,Z_Y,\theta_Y)\Big|X,Z,\theta\Big)\, P(X|Z_X,\theta_X)\ . \end{aligned}$$

Maximization of the observed likelihood is complicated as Y depends on the unobserved transition times. Furthermore, the observed data likelihood consists of integrals which have no closed form solution and thus would need to be approximated, which is computationally intractable with the available software. For joint models with random effects the expectation maximisation (EM) algorithm has proven to be a convenient estimation approach[99], as the random effects can be considered as missing data. However, in our setting we have both unobserved random effects and transition times, which make a classic EM algorithm approach intractable. Instead the joint model is fitted in two steps. First the parameters of the inflammation model are estimated and the output, along with its uncertainty, is used to estimate the parameters of the visual acuity model.

Inflammation model

The parameters of the inflammation model θ_X are estimated using Poisson regression with random effects. Poisson regression with random effects can be performed in `R` using `glmer` from the package `lme4`[13] when the transition times are known. We use an EM type algorithm where we consider the unobserved transition times as missing data, calculate their expectations using current values of the estimates (E-step), then use these expectations to obtain updated estimates of the parameters using `glmer` (M-step).

More specifically, for the E-step we use the empirical Bayes estimates of the random effects for each subject to calculate the expected transition time within each interval where a transition took place. We obtain the empirical Bayes estimates via `ranef` func-

tion, which calculates the conditional mode given by

$$\bar{b}_i = \text{argmax}_b \log\big(f(b|T_i, X_i, Z_i)\big) \ ,$$

where $f(b|T_i, X_i, Z_i)$ is the conditional density of the random effect. To calculate the expected transition times we assume that a transition took place between two visit times v_j and v_{j+1} if $X_{il}(v_j) \neq X_{il}(v_{j+1})$. Let $\Delta_j = v_{j+1} - v_j$ denote the length of the interval and define the intensity in the interval as

$$\gamma_{ilj} = \begin{cases} \lambda_1(v_j|Z_{il}(v_j), b_{il1}) & \text{for } X_{il}(v_j-) = 2 \\ \lambda_2(v_j|Z_{il}(v_j), b_{il2}) & \text{for } X_{il}(v_j-) = 1 \end{cases}$$

where $X_{il}(v_j-)$ is the value of the inflammation just prior to time v_j. The expectation of the unobserved transition time T given the observed data and the frailties is given by

$$\begin{aligned} & E\big(T|X_{il}(v_j), X_{il}(v_{j+1}), Z_{il}(v_j), b_i\big) \\ &= E\Big(TI_{(v_j,v_{j+1})}(T)|Z_{il}(v_j), b_i\Big) \Big/ P\Big(T \in \big(v_j, v_{j+1}\big)\Big) \\ &= \int_{v_j}^{v_{j+1}} s \exp\big(-\gamma_{ilj}s\big)\gamma_{ilj}ds \Big/ \int_{v_j}^{v_{j+1}} \exp\big(-\gamma_{ilj}s\big)\gamma_{ilj}ds \\ &= \Big(v_j + \tfrac{1}{\gamma_{ilj}} - \Big(v_{j+1} + \tfrac{1}{\gamma_{ilj}}\Big)\exp\big(-\gamma_{ilj}\Delta_j\big)\Big) \Big/ \Big(1 - \exp\big(-\gamma_{ilj}\Delta_j\big)\Big) \ , \end{aligned} \tag{5.2}$$

which is straightforward to calculate given β and b_i. Thus, the expectations of the unobserved transition times are estimated by plugging in $\hat{\beta}$ and $\bar{b}_i$.

Visual acuity model

Once the inflammation model has been fitted we use the estimated parameters as input to estimate the parameters of the visual acuity model. Rather than using the estimated parameters from the inflammation model to obtain the unobserved transition times we use multiple imputation. Hence, we start by imputing the transitions times given

the observed data and the estimated parameters from the inflammation model. That is, between two visit times v_j and v_{j+1} with a transition we impute the unobserved event time T by drawing a $p \sim \text{uniform}[0,1]$ and letting

$$T_{ilj} = -\tfrac{1}{\gamma_{ilj}} \log\Big(\exp(-\gamma_{ilj} v_j) - \big(\exp(-\gamma_{ilj} v_j) - \exp(-\gamma_{ilj} v_{j+1}) \big) p \Big) \ .$$

From the imputed transition times we can calculate the time each eye has spent with an quiescent or active inflammation prior to each visit time. Let $\delta_{ilj} = I(X_{il}(v_j) \neq X_{il}(v_{j+1}))$ denote the indicator for a transition between visit v_j and v_{j+1}. We compute the time eye l has spent in state g prior to visit v_j by

$$\begin{aligned} t_{ilg}(v_j) = \ & \textstyle\sum_{k=1}^{j-1} \Big((1 - \delta_{ilk}) I(X_{il}(v_k) = g) \Delta_k \\ & + \ \delta_{ilk} \big(I(X_{il}(v_k) = g)(T_{ilk} - v_k) + I(X_{il}(v_k) \neq g)(v_{k+1} - T_{ilk}) \big) \Big) \ , \end{aligned}$$

and for short we use t_g to denote $t_{ilg}(v_j)$ for $g = 1, 2$. After calculating t_1 and t_2 for each visit it is straightforward to estimate the parameters of the visual acuity model by maximising $P(Y|T, Z, \theta_Y)$, as it is a standard linear mixed model. This procedure is repeated a number of times and the estimated parameters are then pooled. The pooled estimate of the parameters in α are obtained by taking the mean of the estimate of α obtained in each imputation.

Variance estimation

The estimated standard errors obtained within the fitting procedure do not account for the two-stage estimation of the parameters and are therefore most likely too small. The standard errors of the estimates in the joint model are therefore obtained by bootstrapping. A bootstrap sample is obtained by sampling from the pool of subjects with replacement until the sample has the same number of subjects as in the original data set. Hence, the same subject can appear more than once and the number of observations is not necessarily the same as in the original sample. The bootstrap sample is then used

to re-estimate the model parameters. This is repeated a large number of times and the variance of the estimates are calculated as the variance of the estimated parameters in the bootstrap samples.

5.3.3 Dynamic prediction

Here we describe how we use the joint model to make dynamic predictions by simulation. Consider a patient i with a current follow-up time of s years after the first visit to the hospital. For this patient we wish to predict the inflammation status and visual acuity for the lth eye up until a horizon τ. We first estimate the expected transition times in the past and then simulate the future transitions times up until τ. Both the past and future transition times are then used to predict the visual acuity from s up until τ. All time-dependent covariates $Z_{il}(t)$ need to be specified beforehand. In other words, the predictions will be for a predetermined set of treatment decisions etc., which will typically be taken as constant and in what follows we describe them as constant.

First the empirical Bayes estimates of the frailties $\bar{b}_i = (\bar{b}_{iR1}, \bar{b}_{iR2}, \bar{b}_{iL1}, \bar{b}_{iL2})$ are calculated. The expected transition times in the past are obtained by using equation (5.2). For the future transition times, the transition intensities for eye l on subject i are obtained by replacing the parameters with their estimates $\hat{\beta} = (\hat{\beta}_1, \hat{\beta}_2)$ and $\bar{b}_i$ into

$$\hat{\lambda}_{ilg} = \hat{\lambda}_{g,0} \exp\left(Z_{il}(s)\hat{\beta}_g + \bar{b}_{ilg}\right) \text{ for } g \in \{1, 2\} \ .$$

We can then simulate the time to the next transition by drawing a $u \sim \text{uniform}[0, 1]$ and letting

$$\Delta T = \frac{-\log(u)}{\exp(\hat{\lambda}_{ilg})} \ , \tag{5.3}$$

where g is determined by what state the previous transition was made from. The kth

transition time after time s is given by

$$\hat{T}_{ilk} = s + \sum_{j=1}^{k} \Delta T_j \ .$$

This is repeated until $\hat{T}_{ilk} > \tau$. The time the eye will spend in either the quiescent t_1 or active state t_2 up until τ is calculated from the past and future transition times. The procedure generates a single trajectory for the inflammation process.

After generating the transition times for a single trajectory of the inflammation process, we generate a single trajectory from the predictive distribution of $Y_{il}(t)$, for $s < t \leq \tau$ given these transition times. Let $t_1(s)$ and $t_2(s)$ denote the time the eye has spent in quiescent and active state up until time s and let m denote the current spell the eye is in at time s. First the empirical Bayes estimates of the decomposed eye specific random effects $\bar{a}'_{il} = [\bar{a}'_{il0}, \bar{a}'_{il1}, \bar{a}'_{il2}]^\top$ and current spell specific random effects $\bar{a}'_{ilm} = [\bar{a}'_{ilm0}, \bar{a}'_{ilm1}, \bar{a}'_{ilm2}]^\top$ are calculated. Then we determine the current true value of the visual acuity

$$\begin{array}{rcllllll} \hat{Y}_{il}(s) & = & Z_{il}^\top(s)\hat{\alpha}_0 & + & Z_{il}^\top(s)\hat{\alpha}_1 t_1(s) & + & Z_{il}^\top(s)\hat{\alpha}_2 t_2(s) & \text{(fixed)} \\ & + & \bar{a}'_{il0} & + & \bar{a}'_{il1} t_1(s) & + & \bar{a}'_{il2} t_2(s) & \text{(eye)} \\ & + & \bar{a}'_{ilm0} & + & \bar{a}'_{ilm1} t_1(s) & + & \bar{a}'_{ilm2} t_2(s) & \text{(spell)} \ . \end{array}$$

Subsequently, we predict $\hat{Y}_{il}(t)$ until the first transition time after s as a straight line with slope $Z_{il}^\top(s)\hat{\alpha}_1 + \bar{a}'_{il1} + \bar{a}'_{ilm1}$ if the current state is quiescent, or with $Z_{il}^\top(s)\hat{\alpha}_2 + \bar{a}'_{il2} + \bar{a}'_{ilm2}$ if the current state is active. Every time a transition time $\hat{T}_{ilk}$ is encountered a new set of spell specific random effects are drawn from the estimated distribution. The new set of spell specific random effects $\bar{a}'_{il(m+k)}$ replaces the set from the previous spell. Using the updated spell specific random effects, $t_1(\hat{T}_{ilk})$ and $t_2(\hat{T}_{ilk})$, we can determine the true value of the visual acuity at the transition time $\hat{Y}_{il}(\hat{T}_{ilk})$ as before. The visual acuity $\hat{Y}_{il}(t)$ is then predicted as a straight line with slope $Z_{il}^\top(s)\hat{\alpha}_1 + \bar{a}'_{il1} + \bar{a}'_{il(m+k)1}$ if the state is quiescent, or with $Z_{il}^\top(s)\hat{\alpha}_2 + \bar{a}'_{il2} + \bar{a}'_{il(m+k)2}$ if the state is

active, until the next transition time is encountered. The procedure is repeated until the horizon τ.

After having generated a number of trajectories, the results are gathered, and the mean and 2.5% and 97.5% percentiles are used to obtain a point prediction and 95% prediction interval.

5.4 Simulations

In order to evaluate the performance of the proposed estimation procedure we conducted a simulation study. The main objective were to evaluate the estimates of the fixed effects and the variance of the random effects under the assumption that the model is correctly specified in a scenario resembling the uveitis data. In addition, we also looked at the performance of the estimates when the model would be misspecified to not take into account the dependence between the eyes.

5.4.1 Setup

To generate the data for a single subject i, we first generated three patient level baseline covariates, where Z_{1i} and Z_{2i} are binary and each level were sampled with equal probability and $Z_{3i} \sim \mathrm{N}(0, 15)$. Then we generated the subject $b_i' \sim N_2(0, \Sigma_{b'})$ and eye $b_{il}' \sim N_2(0, \Sigma_{b''})$ specific frailty components. The specific parameters values are reported in the Supporting Information. The parameters were chosen such that the simulated data resembled the uveitis data. The transition intensities for subject i were assumed to be given by

$$\lambda_g(t|Z_{1i}, b_{ilg}) = \lambda_{g,0} \exp(Z_{1i}\beta_g + b_{ilg}) \text{ for } g \in \{1, 2\} \ .$$

The initial states were also random, that is there was a 50% chance that one of the eyes were inflamed, a 40% chance that both were inflamed and a 10% chance that none of the eyes had an active inflammation at time 0. Given the random effects and the

baseline information we then generated the time to the next transition by employing the same strategy as in (5.3). New transition times were generated until the sum reached the time horizon of 5 years for each eye. In order to induce the interval censoring of the event times, we simulated a number of prescheduled visit times $N_i \sim \text{Poisson}(\lambda_N)$ with equal distance between time 0 and 5 years. It was possibly for the subject to receive extra visits to ensure that every transition was observed. Every subject had a minimum of two visits. We generated the eye $a'_{il} \sim N_3(0, \Sigma_{a'})$ and spell $a'_{ilm} \sim N_3(0, \Sigma_{a''})$ specific random effect components, where the number of spells was determined by the simulated transitions. The transformed visual acuity was simulated at each visit time based on the true history of the inflammation process according to the model

$$Y_{il}(v_j) = \alpha_0 + \alpha_1 t_1 + \alpha_2 t_2 + \alpha_3 Z_{2i} + \alpha_4 Z_{3i} + a_{ilm0} + a_{ilm1} t_1 + a_{ilm2} t_2 + \epsilon_{ilj} \ ,$$

where $\epsilon_{ilj} \sim N(0, 0.36)$. We generated data with sample sizes of 100 or 300 with an average of 15 or 30 visits per subject and repeated the simulations 1000 times. In scenario A we analysed the simulated interval censored data using the correct submodels for the inflammation and visual acuity, but in scenario B we used an inflammation model which assumed that the eyes were independent. The performance of the estimation procedure was evaluated by calculating the bias, variance, root mean squared error (RMSE) of the fixed effects, along with the coverage rate of the 95% confidence intervals based on the variance estimate with or without bootstrap. Due to computation time, the bootstrapped coverage rates were only based on 100 repetitions and not 1000. We also calculated the bias, variance and RMSE of the variance estimates of the random effects.

5.4.2 Results

The results from the simulation study can be found in the Supporting Information. In scenario A with sample sizes of 100 or 300 and an average number of visits of 15, the bias of the fixed effects in both submodels was overall of a reasonable size compared to the true effect size even with a sample size of 100. In general the bias, variance and

RMSE improved with an increase in sample size, although the improvement in bias was less for the inflammation model parameters. The coverage rate without bootstrapping the variance was lower than the nominal 95% for the inflammation, but it was adequate for the visual acuity. The bootstrapped confidence intervals in the inflammation model had a somewhat better coverage rate. The conclusion for the variance of the random effects is broadly the same as for the fixed effects.

In scenario A it was found that an increase in the average number of visit times improved the bias for both the fixed effects and the variance of the random effects, but had less of an impact on the variance and RMSE. Even with an increase in visits the coverage rate without bootstrapping was still too low. The boostrapped coverage rate performed reasonable, although somewhat variable probably due to the low number of repetitions.

In scenario B the misspecification of the inflammation model lead to an increased bias of both the fixed and random effects in the inflammation model, but did not have a sizeable effect on the estimation of the visual acuity model parameters.

All in all the simulation study suggests that the two-stage estimation procedure with multiple imputation performed satisfactory.

5.5 Uveitis results

The joint model was applied to the uveitis data using early onset, treatment and surgery for the inflammation model and patient age, centred at age 43, and complications for the visual acuity model. The estimated fixed effects for the two submodels can be found in Table 5.2. The estimated baseline transition intensities $\lambda_{1,0}$ and $\lambda_{2,0}$, the Intercept in Table 5.2, imply that the eyes in general move quicker to the quiescent state than to the active state. Moreover, since the baseline is time-constant, it also implies that the reference group is expected to spend $1/\lambda_{2,0} \approx 1$ years in the active state and $1/\lambda_{1,0} \approx 3$ years in the quiescent state. Maintenance or active treatments increase the transition intensity to the quiescent state considerably, but they did not have a significant effect at

the 5% level on the transitions to the active state. Surgery increases the transition intensity to active inflammation, which was expected since surgery may distress the eye and thereby cause more inflammation. The estimated slopes in the visual acuity submodel imply that the visual acuity improves with time, although the time spent with an active inflammation was not found to be significant. The explanation for the increase over time may be due to our relatively young population. At baseline the older ages have a lower intercept and we investigated if there was an interaction between time spent in quiescent or active state and age, but it was found not to be significant. The presence of complications had a significant negative impact on the intercept and a nonsignificant negative impact on the progression of visual acuity over time.

The estimated variances of the random effects are shown in Table 5.3. From $\hat{\Sigma}_{b'}$ and $\hat{\Sigma}_{b''}$ we can see that there is a negative correlation between the two transitions both within the patient and within an eye. Furthermore, the variance is larger for transitions to active than to quiescent. In addition, we can see that there largely is a negative correlation between the intercept and the two slopes in the visual acuity model, and that the variance of the slope for time spent with inflammation is larger than of the slope for time spent without inflammation.

We investigated whether it would be sufficient to have random effects in the visual acuity model on the eye level, instead of a set for each spell. All the same, a likelihood ratio test strongly suggested that the more complex model was preferable. The estimates from the model without spell specific random effects can be found in the Supporting Information.

Figure 5.4 shows the model estimates of the inflammation and visual acuity for the three patients from Figure 5.1. The y-axis depicts the visual acuity on the new scale used in the model, instead of the Snellen scale. On the new scale higher values correspond to better visual acuity and lower values to poorer visual acuity. The model estimates are depicted as lines either with or without the empirical Bayes estimates of the random effects. The fixed effect estimates are straight lines with an intercept and slope that depends on

Table 5.2: Estimates of the fixed effects parameters in the inflammation and visual acuity model for the uveitis data with 95% bootstrapped confidence intervals (CI).

Inflammation					
Transitions to	Quiescent		Active		
Covariates	$\exp(\beta_1)$	CI	$\exp(\beta_2)$	CI	
Intercept	0.984	(0.71, 1.37)	0.303	(0.24, 0.38)	
Early onset					
yes	1.233	(0.77, 1.96)	1.394	(0.77, 2.51)	
Treatment					
maintenance	3.429	(2.33, 5.04)	1.287	(0.93, 1.79)	
active	4.425	(3.02, 6.48)	1.35	(0.9, 2.02)	
Surgery					
yes	0.787	(0.11, 5.74)	5.695	(3.33, 9.75)	

Visual acuity		
Covariates	α	CI
Intercept	1.111	(0.96, 1.27)
t_1	0.055	(0.00, 0.11)
t_2	0.071	(−0.19, 0.33)
Age	−0.019	(−0.02,−0.01)
Age^2	−0.001	(0.00, 0.00)
Complication		
yes	−0.525	(−0.74,−0.31)
t_1	0.077	(0.00, 0.15)
t_2	−0.082	(−0.23, 0.07)

Table 5.3: Estimates of the variance of the frailties and random effects in the joint model.

$$\hat{\Sigma}_{b_1'} = \begin{bmatrix} 0.35 & -0.01 \\ -0.01 & 0.63 \end{bmatrix} \qquad \hat{\Sigma}_{b_2'} = \begin{bmatrix} 0.001 & -0.02 \\ -0.02 & 0.37 \end{bmatrix}$$

$$\hat{\Sigma}_{a_1'} = \begin{bmatrix} 1.77 & 0.01 & -0.60 \\ 0.01 & 0.15 & -0.12 \\ -0.60 & -0.12 & 2.28 \end{bmatrix} \quad \hat{\Sigma}_{a_2'} = \begin{bmatrix} 0.40 & -0.05 & -0.13 \\ -0.05 & 0.03 & -0.05 \\ -0.13 & -0.05 & 0.24 \end{bmatrix}$$

the estimated time spent in the two states. The lines are not continuous, because the time-varying covariates can modify both the intercept and the slopes. The lines where the empirical Bayes estimates are included allow the lines to be even more discontinuous, as the inclusion of spell specific intercepts allow the visual acuity to jump at the transition times.

An illustration of dynamic predictions of inflammation and visual acuity based on data from the three patients can be found in the Supporting Information.

We also conducted a sensitivity analysis of the assumption of non-informative visits. The assumption was based on the input that visits were prescheduled. However, since the inflammation can be very painful, it is possible that some patients requested an earlier appointment due to an onset of an inflammation episode. In order to address this concern, we refitted the model under the assumption that the onset of an inflammation episode happened exactly at the visit time, where an onset was registered. The offset of an inflammation episode was still assumed to be subject to interval censoring. It was simple to implement, as the only thing that changed in the estimation procedure in Section 5.3.2, was that only the transition times going from active to quiescent needed to be updated. The estimates of the fixed effects are given in Table 5.4. In the inflammation model the baseline transition intensity for transitions to quiescent went from

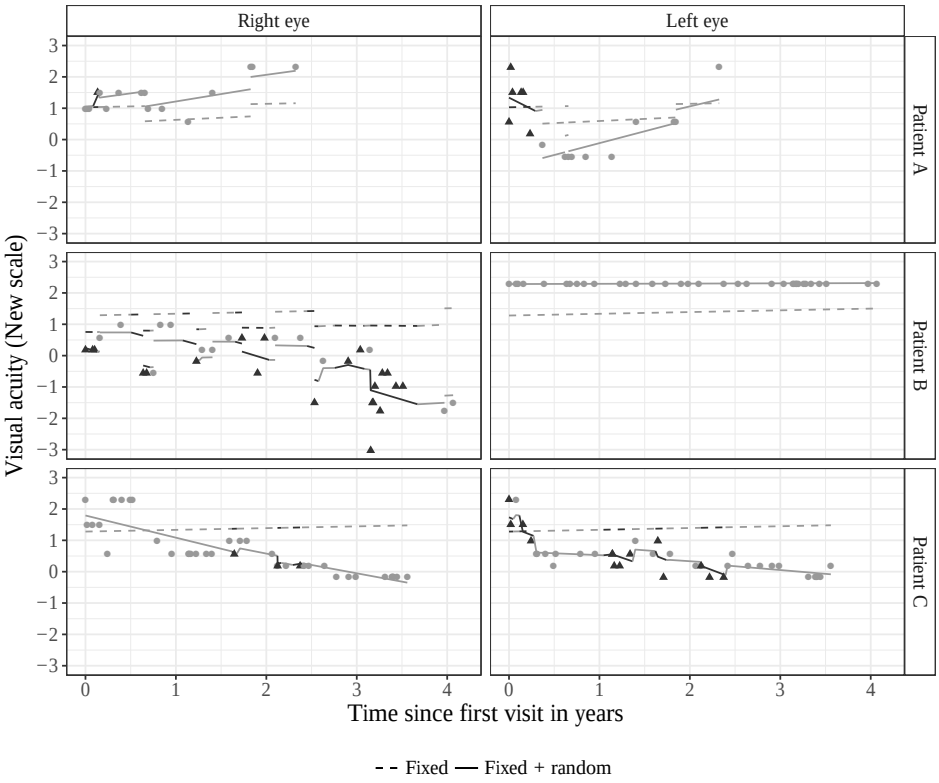

Figure 5.4: Illustration of the model applied to data from three patients. The observed inflammation (colour) and visual acuity on the model scale (y-axis) is indicated with transparent dots. The lines depicts the model estimates of the inflammation and visual acuity either with (Fixed + random) or without (Fixed) the empirical Bayes estimates of the random effects.

0.984 to 3.599 under the new assumption. Considering that the assumption leads to less time being spent in the active inflammation state, this is not surprising. However, what is surprising is that the treatment effects on transitions to quiescent were also noticeably reduced. In the visual acuity model results were qualitatively the same, except for the effect of time in the quiescent state, which increased from 0.055 to 0.1.

To evaluate the model predictions we looked at the Brier Score for the inflammation model, and for the visual acuity we looked at the bias and root mean squared error (RMSE). The evaluation measures were calculated for three different time points during the patients' follow-up at 0, 1 and 2 years. Using the data that were available at a given follow-up time point predictions were assessed at 1 and 3 years ahead in time. We compared predictions from three joint models. The first model (Model 1.a) did not include any covariates in the two submodels and it was fitted under the assumption that the transition times of the inflammation process were interval censored. The second model (Model 1.b) is the one that was reported in Table 5.2. It was fitted under the same assumption, but it included covariates. The last model (Model 2) was reported in Table 5.4. The model included covariates, but it was fitted under the assumption that the onset of inflammation episodes were observed and happened at the visit time. For each patient 50 simulated predictions were obtained from each of the three models. Although it arguably is an imperfect solution, we compared the mean of the 50 predictions to the last observed value of either the inflammation or the visual acuity. Figure 5.5 show the results of the evolutions. In general Model 1.a has the lowest Brier Score and Model 2 has the lowest RMSE. In terms of bias there is no one model that performs better than the others. Since the primary concern of the patients is their visual acuity we tend to favour the models that perform better on the visual acuity scale. For this reason we ultimately decided to favour Model 2 and furthermore the implied assumption about the interval censoring seems reasonable for the uveitis data.

Albeit, the predictions were evaluated on the same data that were used to estimate the models' parameters, it is still reasonably to compare the models based on their predictive

Table 5.4: Estimates of the fixed effects parameters in the inflammation and visual acuity model for the uveitis data with 95% bootstrapped confidence intervals (CI), where it was assumed that onset of inflammation happened at the visit time.

Inflammation					
Transitions to	Quiescent		Active		
Covariates	$\exp(\beta_1)$	CI	$\exp(\beta_2)$	CI	
Intercept	3.599	(2.53, 5.12)	0.274	(0.22, 0.34)	
Early onset					
yes	1.262	(0.73, 2.19)	1.394	(0.81, 2.39)	
Treatment					
maintenance	1.629	(1.09, 2.43)	1.419	(1.06, 1.90)	
active	1.263	(0.88, 1.82)	1.544	(1.11, 2.15)	
Surgery					
yes	1.244	(0.1, 15.69)	5.375	(3.46, 8.34)	

Visual acuity		
Covariates	α	CI
Intercept	1.135	(0.99, 1.28)
t_1	0.100	(0.04, 0.16)
t_2	−0.035	(−0.10, 0.03)
Age	−0.018	(−0.02,−0.01)
Age^2	−0.001	(0.00, 0.00)
Complication		
yes	−0.561	(−0.77,−0.35)
t_1	0.120	(0.03, 0.21)
t_2	0.002	(−0.08, 0.08)

performance. Nonetheless, the final model should be evaluated on a new data set to avoid overoptimism.

5.6 Discussion

We have proposed a joint model for dynamic prediction of visual acuity and inflammation in uveitis patients, which accounts for the special features of the data. The proposed joint model distinguishes itself by dealing with an episodic interval censored multi-state outcome. In addition it is unusual in that the multi-state outcome affects the longitudinal outcome and not the other way around.

The joint model is complicated by the need to account for the special dependence structure in the uveitis data. We employed random effects both to account for the dependence structure and to obtain subject-specific predictions. However, a classic criticism of random effect models is that the assumed distribution is difficult to verify from the data, and instead the choice is often based on what is computationally convenient. For other applications the structure could be simplified, which would reduce the dimensions of the random effects, and such a model would probably prove easier to estimate in one step instead of two. The current estimation procedure could be improved by finding a way to directly estimate all the parameters in one step. Ways to solve the problem of computational intractability could be to use Laplace approximations[102] or adaptive Gaussian quadratures, which would likely be much faster than using an EM algorithm. Due to the complexity of the uveitis data and consequently the joint model, we employed a two-stage estimation procedure. Although two-stage procedures can lead to bias and loss of efficiency compared to other procedures[3,86,101], our simulation study showed that the estimation procedure performed satisfactory when the model was correctly specified. It is however a disadvantage of the two-stage procedure that we cannot compute a full likelihood.

The model relies on a number of assumptions, which were largely motivated by clinical insight. One of the assumptions was that the baseline transition intensities

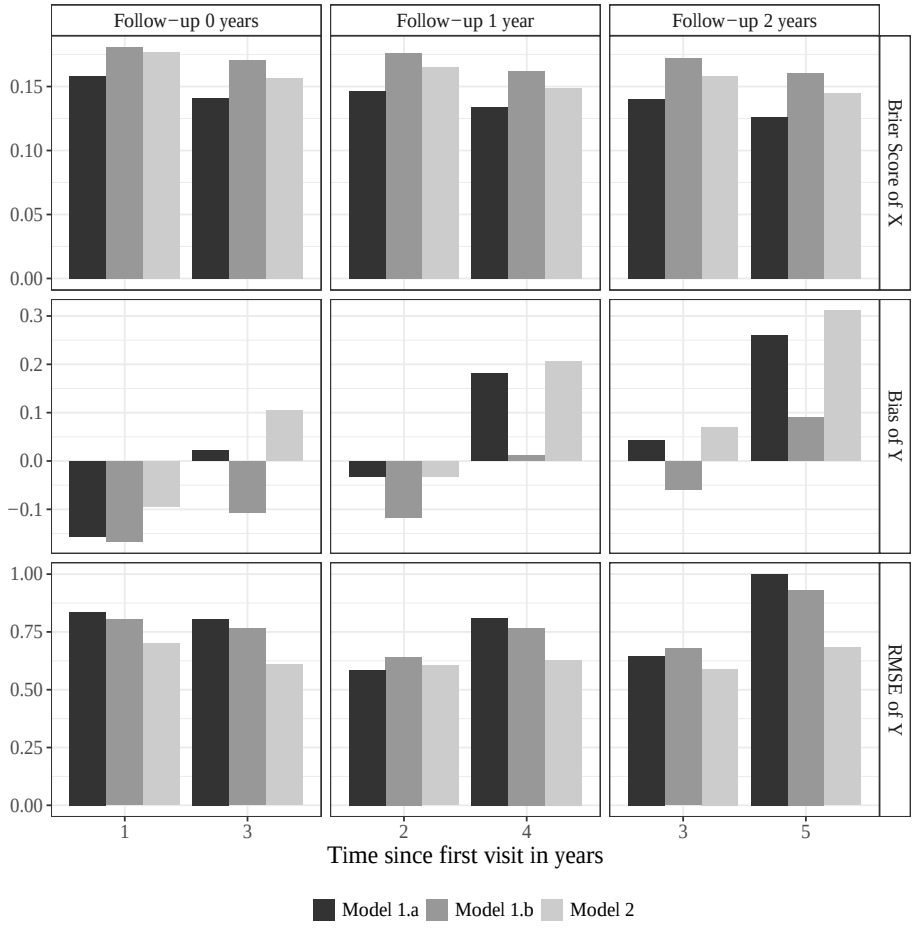

Figure 5.5: Evaluations of the dynamic predictions from three models: Model 1.a and 1.b refers to the joint model where it is assumed that the inflammation is interval censored. Model 2 assumes that onset of inflammation happens at the visit time. Model 1.b and 2 do in include covariates, whereas 1.a does not. The predictions from the inflammation submodel are evaluated with the Brier Score and predictions form the visual acuity submodel are evaluated in terms of bias and RMSE. The evaluation measures are calculated at three follow-up time points and at 1 or 3 years ahead in time.

were constant. This was believed to be reasonable for this application, since uveitis is a chronic disease in most cases. However, for other applications it would be a natural extension of the model to allow the baseline to be time-varying. This could also be a way of confirming the assumption about the constant baseline. A necessary assumption to fit the model was that all transitions between the quiescent and active state were observed. The assumption is believed to be reasonably for the uveitis data, however it could be an issue if the assumption is violated. Dropout from the study could be a cause for concern as well for the missing at random assumption. The standard procedure at the Rotterdam Eye Hospital was to only discharge a patient after five years without any inflammation episodes, unfortunately information about discharges was not available to us. In addition, it is also imaginable that patients could have neglected to turn up for the appointments if their eyes had been improving over a longer period of time. We investigated the consequences of one of the other assumptions in a sensitivity analysis. There we either assumed that the onset of an inflammation episode was always observed or interval censored. It turned out to result in a higher baseline transition rate to quiescent and smaller treatment effects. It is likely that the truth is somewhere in between.