Brilleman, S., Crowther, M., May, M. T., Gompels, M., & Abrams, K. (2016). Joint longitudinal hurdle and time-to-event models: an application related to viral load and duration of the first treatment regimen in HIV patients initiating therapy. *Statistics in Medicine*, 35(20), 3583-3594. <https://doi.org/10.1002/sim.6948>

Peer reviewed version

License (if available):
Unspecified

Link to published version (if available):
[10.1002/sim.6948](https://doi.org/10.1002/sim.6948)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the peer reviewed version of the following article: Brilleman, S. L., Crowther, M. J., May, M. T., Gompels, M., and Abrams, K. R. (2016) Joint longitudinal hurdle and time-to-event models: an application related to viral load and duration of the first treatment regimen in patients with HIV initiating therapy. *Statist. Med.*, which has been published in final form at doi: 10.1002/sim.6948. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Joint longitudinal hurdle and time-to-event models: an application related to viral load and duration of the first treatment regimen in HIV patients initiating therapy

Samuel L Brilleman^{1,2 *}

Michael J Crowther^{3,4}

Margaret T May⁵

Mark Gompels⁶

Keith R Abrams³

¹ Department of Epidemiology and Preventive Medicine, Monash University, Alfred Centre, 99 Commercial Road, Melbourne, VIC 3004, Australia

² Victorian Centre for Biostatistics (ViCBiostat), Melbourne, VIC, Australia

³ Department of Health Sciences, University of Leicester, Adrian Building, University Road, Leicester, LE1 7RH, UK

⁴ Department of Medical Epidemiology and Biostatistics, Karolinska Institutet, Box 281, S-171 77 Stockholm, Sweden

⁵ School of Social and Community Medicine, University of Bristol, Canynge Hall, 39 Whatley Road, Bristol BS8 2PS, UK

⁶ North Bristol NHS Trust, Bristol, UK

* Corresponding author: sam.brilleman@monash.edu

Abstract

Shared parameter joint models provide a framework under which a longitudinal response and a time-to-event can be modelled simultaneously. A common assumption in shared parameter joint models has been to assume the longitudinal response is normally distributed. In this paper we instead propose a joint model which incorporates a two-part “hurdle” model for the longitudinal response, motivated in part by longitudinal response data which is subject to a detection limit. The first part of the hurdle model estimates the probability that the longitudinal response is observed above the detection limit, whilst the second part of the hurdle model estimates the mean of the response conditional on having exceeded the detection limit. The time-to-event outcome is modelled using a parametric proportional hazards model, assuming a Weibull baseline hazard. We propose a novel association structure whereby the current hazard of the event is assumed to be associated with the current combined (expected) outcome from the two parts of the hurdle model. We estimate our joint model under a Bayesian framework and provide code for fitting the model using the Bayesian software Stan. We use our model to estimate the association between HIV RNA viral load, which is subject to a lower detection limit, and the hazard of stopping or modifying treatment in HIV patients initiating antiretroviral therapy.

Keywords: joint model; shared parameter model; hurdle model; detection limit; censoring

1. Introduction

In epidemiological or clinical studies it is common for both a longitudinal response and a time-to-event to be outcomes of interest. An emerging class of “shared parameter joint models” provide a framework under which these two outcome types can be modelled simultaneously [1]. A common assumption in shared parameter joint models has been to assume the longitudinal response is normally distributed [2-6]. However in many studies this assumption may not be valid, for example, with longitudinal binary or count data. In other studies the longitudinal response may be a continuous measure but have distributional characteristics such as left- or right-censoring, truncation or extreme skew which make any assumption of normality unreasonable.

In HIV and AIDS research the use of joint models has been particularly prevalent [6-9]. However HIV RNA “viral load”, when measured in HIV positive patients undergoing antiretroviral therapy (ART), is a clinical biomarker which characteristically exhibits non-normality. Viral load is a measure of the number of copies of the HIV virus in a patient’s blood, often quantified as the number of copies per millilitre (copies/ml). As an indicator of clinical efficacy viral load is measured repeatedly over time in patients undergoing ART. However, the assay used to measure viral load generally has a lower detection limit of 40 or 50 copies/ml (depending on the sensitivity of the assay) which results in left-censoring of the observed data. Furthermore the distribution of viral load is usually positively skewed with some patients having extremely high measurements, generally prior to and during the earliest phase of treatment.

In this paper we propose a shared parameter joint model which we use to estimate the association between viral load and the hazard of stopping or modifying treatment in patients on their first ART regimen. To accommodate the distributional characteristics of the viral load data, our joint model incorporates a “hurdle” model for the longitudinal response. Hurdle models consist of two parts and can be used as a method for modelling a single outcome measure as the product of two separate processes [10]. The first part of a hurdle model estimates the probability that an observation is

observed above or below a specified threshold, whilst the second part of the model estimates the mean of observations conditional on them having exceeded the threshold. Traditionally hurdle models have been used for count data with excess zeros (where the threshold distinguishes zero and positive non-zero counts) [10], however they can be equally well applied to continuous data [11, 12]. The first part of our hurdle model consists of a mixed effects logistic regression model and estimates the probability that viral load is “detectable” which we define as $>1.7 \log_{10}$ copies/ml (equivalent to >50 copies/ml). We refer to the threshold of $1.7 \log_{10}$ copies/ml as the “detection limit”. The second part of our hurdle model consists of a mixed effects gamma regression model and estimates – for those observations which are detectable – the mean amount by which they exceed the detection limit. A distinguishing feature of hurdle models is that the second part of the model generally excludes the threshold value from its support; this being in contrast to zero-inflated models.

In the context of joint modelling of longitudinal and time-to-event data there have been very few papers incorporating hurdle models or longitudinal models with two-part processes. Hatfield et al. [13] proposed a joint model which incorporated a zero-augmented beta model for the longitudinal response. Their zero-augmented beta model can be described as a two-part hurdle model in which the first and second parts, respectively, consist of a logistic and a beta mixed effects regression model. They observed longitudinal patient reported outcomes (PROs) on a continuous scale from 0 to 100 from which they derived two measures. Presence was defined as a binary outcome indicating whether the observed PRO was zero or non-zero. Severity was defined as the observed PRO measurement conditional on it being non-zero (rescaled to be between 0 and 1 to ensure it could be accommodated by the support of the beta distribution). They related the presence and severity of PROs to patient survival using a shared parameter joint model. The joint model association structure was based on random intercepts from each part of their zero-augmented beta model. Liu [14] proposed a joint model which consisted of a longitudinal hurdle model for hospital costs and a time-to-event model for deaths. The hurdle model included a logistic mixed effects regression model for the probability of observing a non-zero cost and a linear mixed effects regression model for modelling the non-zero log

costs. Similar to Hatfield et al. [13] the joint model association structure in Liu [14] was based on separate random effects from the two parts of the hurdle model, rather than the combined expected outcome from the hurdle model. Rizopoulos et al. [15] described a joint model which incorporated a two-part model for longitudinal binary data with excess zeros. The first part of their model used a standard logistic regression model to estimate the probability that a patient had all longitudinal measurements equal to zero, whilst the second part used a logistic mixed effects regression model to estimate the probability of a non-zero longitudinal measurement given that the patient had at least one. A common aspect of these previous approaches has been that the association structure in the joint model was based exclusively on subject-specific random effects. Although this potentially provides a computational advantage (since the association structure is not time-dependent) such an association structure may not be meaningful in all contexts. For example, the inclusion of random cubic spline terms may significantly increase the flexibility for modelling subject-specific longitudinal trajectories, but it would render an association structure based purely on random effects somewhat less interpretable.

The joint model presented in this paper differs in several ways to these previous approaches. First, our hurdle model uses restricted cubic splines to flexibly model non-linear changes in the marker over time. This helps to ensure the longitudinal trajectories are appropriately captured, thereby avoiding potential bias in the estimated association parameter [16]. Second, our hurdle model assumes a log link function and gamma error distribution for observations which exceed the detection limit. In our application this has the benefit of allowing for positively skewed viral load data. Third, our joint model uses an association structure based on the combined outcome from the two parts of the hurdle model and to our knowledge this is the first joint model which has been based on this type of association structure. Although the association structure we propose is time-dependent and therefore potentially increases computational complexity, the main motivation for its use is interpretability. Fourth, we estimate our model using the recently developed Bayesian statistical software Stan [17]. In the

supplementary material we discuss some of the benefits of using Stan, as well as providing the code for fitting our joint model.

The remainder of the paper will proceed as follows. In Section 2 we define the joint longitudinal hurdle and time-to-event model. In Section 3 we describe our Bayesian approach to model estimation. In Section 4 we describe the application of our joint model to estimating the association between viral load and the hazard of stopping or modifying a patient's first ART regimen. Section 5 will contain a discussion.

2. Definition of the joint longitudinal hurdle and time-to-event model

The joint longitudinal hurdle and time-to-event models is defined as follows. For the i th patient, $i = 1, \dots, N$, we observe $y_{ij} = y_i(t_{ij})$ as the longitudinal marker measurements taken at some time points $t_{ij}, j = 1, \dots, n_i$. We denote $T_i = \min(T_i^*, C_i)$ as the observed event time for the i th patient, being the minimum of the true event time T_i^* and the censoring time C_i . We define d_i as an indicator variable taking the value 1 if $T_i^* \leq C_i$ and 0 otherwise.

2.1 Hurdle component model

The hurdle component model is used for modelling the longitudinal marker measurements. We assume a threshold value of zero for the hurdle model and define $w_{ij} = I(y_{ij} > 0)$ where $I(\cdot)$ is the indicator function, such that w_{ij} serves as a binary indicator of whether the longitudinal marker measurement for the i th patient at time point t_{ij} is above this threshold.

In our application in Section 4 we follow the common approach of analysing $\log_{10}$ viral load rather than using viral load measurements on the original scale. Since the observed $\log_{10}$ viral load measurements were left-censored at either 1.6 or 1.7 $\log_{10}$ copies/ml (equivalent to 40 or 50 copies/ml) depending on the sensitivity of the assay, but the hurdle model was to be defined for a threshold value of zero (to align with the support of the gamma distribution) we analysed $y_{ij} = \max\{v_{ij} - 1.7, 0\}$ where v_{ij}

is the observed $\log_{10}$ viral load measurement for the i th patient at the j th time point. The longitudinal response y_{ij} therefore represents the amount by which the $\log_{10}$ viral load measurement exceeds the clinically meaningful threshold of $1.7 \log_{10}$ copies/ml which we use to define a “detectable” viral load.

The hurdle model can then be defined in two parts, referred to here as the Bernoulli and gamma parts.

Bernoulli part:

$$w_{ij} \sim \text{Bernoulli}(\mu_{1i}(t_{ij}))$$

$$\eta_{1i}(t_{ij}) = \text{logit}(\mu_{1i}(t_{ij})) = \mathbf{x}'_{1i}(t_{ij})\boldsymbol{\beta}_1 + \mathbf{z}'_{1i}(t_{ij})\mathbf{b}_{1i} + \mathbf{u}'_{1i}\boldsymbol{\delta}_1 \quad (2.1)$$

Gamma part:

$$(y_{ij} \mid w_{ij} = 1) \sim \text{Gamma}(\varphi_i(t_{ij}), \omega)$$

$$\mu_{2i}(t_{ij}) = \frac{\varphi_i(t_{ij})}{\omega} \quad (2.2)$$

$$\eta_{2i}(t_{ij}) = \log(\mu_{2i}(t_{ij})) = \mathbf{x}'_{2i}(t_{ij})\boldsymbol{\beta}_2 + \mathbf{z}'_{2i}(t_{ij})\mathbf{b}_{2i} + \mathbf{u}'_{2i}\boldsymbol{\delta}_2$$

where $\mathbf{x}_{1i}(t_{ij})$, $\mathbf{x}_{2i}(t_{ij})$, $\mathbf{z}_{1i}(t_{ij})$ and $\mathbf{z}_{2i}(t_{ij})$ are vectors of time-dependent covariates with associated vectors of either fixed ($\boldsymbol{\beta}_1, \boldsymbol{\beta}_2$) or random ($\mathbf{b}_{1i}, \mathbf{b}_{2i}$) coefficients and $\mathbf{u}_{1i}$ and $\mathbf{u}_{2i}$ are vectors of time-independent baseline covariates with associated vectors of fixed coefficients $\boldsymbol{\delta}_1$ and $\boldsymbol{\delta}_2$. It is worth noting that the covariates chosen for the Bernoulli and gamma parts (whether time-dependent or time-independent) do not necessarily need to be the same. In fact in many cases they will not be. For example, since the data structure ensures there will be less data available for estimating the distribution of $E(Y_{ij} \mid w_{ij} = 1)$ than $E(W_{ij})$ one may wish to use a simpler approach to modelling changes over time in the gamma part of the model, but we maintain generality in our notation. Flexibility can be incorporated using combinations of splines and/or fractional polynomials in $\mathbf{x}_{1i}(t_{ij})$, $\mathbf{x}_{2i}(t_{ij})$, $\mathbf{z}_{1i}(t_{ij})$ and $\mathbf{z}_{2i}(t_{ij})$.

The random effects from the hurdle model are assumed to be multivariate normally distributed such that

$$\mathbf{b}_i \sim MVN(0, \mathbf{\Sigma}) \quad (2.3)$$

where $\mathbf{b}_i = \{\mathbf{b}_{1i}, \mathbf{b}_{2i}\}$ is the combined vector of random effects (from both parts of the hurdle model) for the i th patient and $\mathbf{\Sigma}$ is the variance-covariance matrix for the random effects. The joint distribution for the random effects is assumed to capture the dependence between the Bernoulli and gamma parts of the hurdle model. This has the potential to alleviate bias in the estimated regression coefficients which might otherwise be induced by assuming the two parts of the hurdle model were independent [18]. Since the two parts of the hurdle model are assumed to be independent conditional on the random effects we can calculate the expected value for the longitudinal marker for patient i at time t , given their vector of random effects $\mathbf{b}_i$, as

$$\mu_i(t) = E(Y_i(t) | \mathbf{b}_i) = \mu_{1i}(t) \times \mu_{2i}(t) \quad (2.4)$$

and this is the quantity used to form the association structure in the joint model. This is discussed in more detail in the next subsection. It is worth noting here, however, that the quantity in (2.4) is the subject-specific mean and not the marginal mean for the hurdle model.

2.2 Time-to-event component model

The time-to-event outcome is modelled using a proportional hazards model of the form

$$h_i(t) = h_0(t) \times \exp(\mathbf{w}_i' \boldsymbol{\gamma} + \alpha \mu_i(t)) \quad (2.5)$$

where $h_i(t)$ is the estimated hazard rate for the event of interest for the i th patient at time t , $\mathbf{w}_i$ is a vector of time-independent baseline covariates with associated vector of fixed effects $\boldsymbol{\gamma}$ and $h_0(t)$ is the baseline hazard evaluated at time t .

We assume the baseline hazard $h_0(t)$ is modelled parametrically. For the application in this paper we use a simple Weibull distribution for the baseline hazard, which appears adequate, however

extensions to more complex shapes for the baseline hazard could be considered (in our discussion section we refer to several alternative choices for the baseline hazard). The regression coefficients γ can be considered as representing the so-called direct association between the baseline covariates $\mathbf{w}_i$ and the log hazard, having adjusted for the underlying level of the longitudinal marker. That is, given that some elements of the baseline covariate vector $\mathbf{w}_i$ may also be present in $\mathbf{x}_{1i}(t_{ij})$ and $\mathbf{x}_{2i}(t_{ij})$, these covariates may also have an indirect effect on the log hazard through their association with the estimated value of the longitudinal marker.

The parameter α is referred to as the association parameter, measuring the strength of association between the hurdle and time-to-event components of the joint model. It can be interpreted as the log hazard ratio associated with a one unit increase in the expected value of the longitudinal marker for a given patient. A key advantage of the approach is that the hazard at time t is able to depend on the expected longitudinal response at time t which itself is not directly observable but likely to be clinically meaningful. It is also worth noting that the form in (2.5) is equivalent to a current value association structure under a setting with just a single continuous longitudinal marker (rather than two outcomes as specified under the hurdle model).

3. Model estimation

3.1 A Bayesian approach

We adopt a Bayesian approach to estimating the joint longitudinal hurdle and time-to-event model as it offers several advantages. Specifying the model is straightforward, particularly when using Bayesian statistical software such as Stan [17]. The use of Markov chain Monte Carlo (MCMC) avoids any need to explicitly integrate over the random effects, potentially providing computational benefits for models with high-dimensional random effects. Derived parameters can also be easily estimated as part of the MCMC sampling. For example, we can easily obtain MCMC estimates for the combined expected outcome from the hurdle model ($\mu_i(t)$) and therefore calculate the uncertainty associated

with its posterior distribution, for example the standard deviation of its posterior distribution or a 95% credible interval. This avoids the need to use, say, bootstrapping to obtain a variance estimate for this parameter.

3.2 Model likelihood and posterior distribution

If $\mathbf{y}_i = \{y_{i1}, \dots, y_{in_i}\}$ represents the combined vector of all observed longitudinal marker measurements for the i th patient then the probability distribution function (up to a proportionality constant) for the full posterior distribution of the joint model can be specified as

$$p(\mathbf{b}_i, \boldsymbol{\theta}_y, \boldsymbol{\theta}_T, \boldsymbol{\Sigma} | \mathbf{y}_i, T_i, d_i) \propto \prod_{i=1}^N \left(\prod_{j=1}^{n_i} p(y_{ij} | \mathbf{b}_i, \boldsymbol{\theta}_y) \right) p(T_i, d_i | \mathbf{b}_i, \boldsymbol{\theta}_T) p(\mathbf{b}_i | \boldsymbol{\Sigma}) p(\boldsymbol{\theta}_y, \boldsymbol{\theta}_T, \boldsymbol{\Sigma}) \quad (3.1)$$

where $\boldsymbol{\theta}_y = \{\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \omega\}$ and $\boldsymbol{\theta}_T = \{\boldsymbol{\gamma}, \alpha\}$ are the combined vectors of fixed effects parameters for the marker and time-to-event components of the joint model respectively. The contribution to the likelihood from the hurdle component of the joint model can be written as

$$p(y_{ij} | \mathbf{b}_i, \boldsymbol{\theta}_y) = \{1 - \mu_{1i}(t_{ij})\}^{1-w_{ij}} \{\mu_{1i}(t_{ij}) p(y_{ij} | w_{ij} = 1; \mathbf{b}_i, \boldsymbol{\theta}_y)\}^{w_{ij}} \quad (3.2)$$

$$p(y_{ij} | w_{ij} = 1; \mathbf{b}_i, \boldsymbol{\theta}_y) = \frac{\omega^{\varphi_i(t_{ij})} y_{ij}^{(\varphi_i(t_{ij})-1)} e^{-\omega y_{ij}}}{\Gamma(\varphi_i(t_{ij}))} \quad (3.3)$$

where $p(y_{ij} | w_{ij} = 1; \mathbf{b}_i, \boldsymbol{\theta}_y)$ corresponds to the gamma probability distribution function for marker measurements which exceed the threshold. Therefore, the use of a different distribution for modelling observations above the threshold would only require changing this element of the full joint model likelihood. The contribution to the likelihood from the time-to-event component of the joint model is

$$p(T_i, d_i | \mathbf{b}_i, \boldsymbol{\theta}_T) = (h_i(T_i | \mathbf{b}_i, \boldsymbol{\theta}_T))^{d_i} \exp\left(-\int_0^{T_i} h_i(u | \mathbf{b}_i, \boldsymbol{\theta}_T) du\right) \quad (3.4)$$

One added complexity of estimating the joint model described here is that the association structure in (2.5) is time-dependent. This means that the resulting likelihood for the time-to-event component

is also time-dependent and therefore the integral in (3.4) has no closed form solution. Therefore we used a Gauss-Kronrod quadrature rule [19] with 15 nodes to numerically evaluate the integral.

4. Application

4.1 Objectives

Major factors potentially impacting on the durability of a first drug regimen in the treatment of HIV include virological treatment failure and intolerance or toxicity. In this study we focus on the former. Virological treatment failure has been defined as “failure to achieve a viral load <50 copies/ml 6 months after commencing ART or following viral suppression to <50 copies/mL a viral load rebound to >400 copies/ml” [20]. However there is also evidence to suggest sustained viral load measurements in the range of 51 to 500 copies/ml are associated with an increased risk of virological treatment failure [21]. The extent by which a patient’s viral load measurements exceed 50 copies/ml is therefore likely to have an influence on the decision to alter their treatment regimen. Using the joint longitudinal hurdle and time-to-event model described in Section 2 we estimate the association between detectable viral load and the hazard of stopping or modifying treatment in patients on their first ART regimen.

4.2 Data description

A total of 780 patients at the North Bristol NHS Trust HIV/AIDS Service started ART between 1st January 2000 and 31st December 2012. Our analyses, of fully anonymised data, included the 550 (71%) patients who had at least two viral load measurements taken during the follow up period. Table 1 shows demographic characteristics of those 550 patients.

A patient’s ART start date was the day on which they were prescribed their first antiretroviral agent. Their first regimen was defined as the combination of drugs they were prescribed within 31 days of the ART start date. A patient’s treatment was considered to have been “stopped or modified” on the

day they were no longer being prescribed at least one of the drugs in their first regimen. Alternatively, a patient's data was censored once five years follow up had been collected or at the study end date of 31st December 2012. There were 387 (70%) patients who switched off their first regimen during a total follow up period of 1,007 person-years. For patients who switched the mean (SD) time-to-switch was 1.49 (1.26) years, whilst for patients who were censored the mean (SD) follow up was 3.25 (1.79) years.

Viral load was measured at approximately 3 monthly intervals during the follow up period (mean = 83 days). During follow up a total of 4,411 viral load measurements were observed with a mean (SD) [range] number of viral load measurements per patient of 8.0 (5.7) [2-29]. There were 1,181 (27%) measurements above the limit of detection ($> 1.7 \log_{10}$ copies/ml) which contributed to estimating the gamma part of the hurdle model; these measurements were observed on 450 patients with a mean (SD) [range] number of measurements per patient of 2.6 (1.8) [1-16]. Figure 1 shows the distribution of $\log_{10}$ viral load measurements which were observed in the detectable range. The distribution is strongly positively skewed.

Additional measured covariates included each patient's age at the date on which they started ART, gender, and the calendar period in which they started ART (referred to as start period). Age was categorised as: 18-29; 30-39; 40-49; and 50+ years. The start period was categorised as: 2000-03; 2004-06; 2007-09; or 2010-12 reflecting changes in availability of ART drugs and recommendations on the threshold below which ART should be started.

4.3 Model specification

The hurdle component model was specified as follows. We used restricted cubic splines to flexibly model changes over time in each patient's log odds of having a detectable viral load. The linear predictor for the Bernoulli part of the hurdle model took the form

$$\eta_{1i}(t_{ij}) = \text{logit}(\mu_{1i}(t_{ij})) = \beta_{10} + \sum_{m=1}^3 \beta_{1m} s_{im}(t_{ij}) + b_{1i0} + b_{1i1} t_{ij} + \mathbf{u}'_{1i} \boldsymbol{\delta}_1 \quad (4.1)$$

where $s_{i1}, \dots, s_{i3}$ are the restricted cubic spline basis functions [22] for knots placed at 0, 3, 12 and 60 months. In other words the model included a fixed intercept and fixed splines with 3 degrees of freedom but only included a random intercept and random linear slope term. We limited the number of individual-level random effects to avoid over-parameterising the model. In the supplementary material we present the results from sensitivity analyses in which we assess the chosen number of degrees of freedom for the spline terms (both fixed and random) as well as the knot locations. The baseline covariates ($\mathbf{u}_{1i}$) included age, gender and start period.

For the gamma part of the model we specified a constant only model with a random intercept

$$\eta_{2i}(t_{ij}) = \log(\mu_{2i}(t_{ij})) = \beta_{20} + b_{2i0} \quad (4.2)$$

Since the linear predictor in (4.2) contains only a fixed and random intercept and does not account for any trends over time, it estimated the subject-specific mean amount by which a patient's $\log_{10}$ viral load measurements exceeded the threshold. We used a simple specification for the gamma part of the hurdle model since there were generally very few observations per patient for estimating this part of the model. However this could be extended to incorporate time trends, both fixed and random.

For the time-to-event component model we assumed a Weibull distribution for the baseline hazard $h_0(t)$. We used the same baseline covariates ($\mathbf{w}_i$) in the time-to-event component model as in the Bernoulli part of the hurdle model, namely age, gender and start period. Details of the model implementation (including specification of the prior distributions and our use of the software Stan [17]) are provided in the supplementary material.

4.4 Results

Table 2 shows the fixed effect parameter estimates from the joint model. The estimated association parameter from the model is 0.44 (95% CrI: 0.06 to 0.78). It is worth noting that a doubling in viral load is equivalent to a 0.3 unit increase in $\log_{10}$ viral load (since $\log_{10} 2 = 0.3$). Therefore our joint model suggests a doubling in the expected detectable viral load for a given individual is associated with an instantaneous hazard of stopping or modifying treatment which is 1.14 (95% CrI: 1.02 to 1.26) times as large (calculated as $\exp(0.3 \times 0.44)$) at time t . The parameter estimates from the time-to-event component of the joint model also suggest demographics may be associated with the hazard. The hazard of stopping or modifying treatment appeared to be highest for patients in the youngest (18-29 years) and oldest (50+ years) age groups, females and those who started ART more recently. The observed outcome data and estimated marker trajectories for 8 selected patients are shown in the supplementary material.

Table 3 shows the standard deviation and correlation parameter estimates for the subject-specific random effects. The estimated correlation between the subject-specific random intercepts in the Bernoulli and Gamma parts of the hurdle model was 0.41 (95% CrI: 0.17 to 0.63), indicating that those patients with a higher probability of having detectable viral load at baseline generally had higher expected mean viral load measurements when they were detectable.

4.5 Comparison with alternative association structures

We also compared our model to three alternative models. The first model assumed a null association structure, in other words, we set the association parameter α equal to zero. Algebraically this is equivalent to fitting the longitudinal hurdle model and the time-to-event model entirely separately. The second model assumed an association structure based on the linear predictor from the Bernoulli part only, in other words, a time-to-event component model of the following form

$$h_i(t) = h_0(t) \times \exp(\mathbf{w}_i' \boldsymbol{\gamma} + \alpha \eta_{1i}(t)) \quad (4.3)$$

such that the association structure is equivalent to a joint model based only on binary marker data. The third model assumed a quadratic association structure based on the expected outcome from the hurdle model, namely

$$h_i(t) = h_0(t) \times \exp\left(\mathbf{w}_i' \boldsymbol{\gamma} + \alpha_1 \mu_i(t) + \alpha_2 (\mu_i(t))^2\right) \quad (4.4)$$

For the alternative models we still estimated all components of the joint hurdle model in Stan (Bernoulli part, gamma part, time-to-event part) but altered only the association structure. Therefore the form for the joint posterior distribution in (3.1) is common to all four models.

We compared our joint hurdle model to these alternative models using the Watanabe-Akaike information criterion (WAIC) [23]. The WAIC provides an indication of model fit (in terms of the expected predictive accuracy of the model) but includes a penalty for model complexity (in terms of the estimated effective number of parameters). As with other information criteria the WAIC can be used to compare models which are non-nested and smaller values indicate a better fitting model. A benefit of the WAIC is that it averages over the posterior distribution for the model parameters and therefore incorporates uncertainty associated with the model parameters [24]. Conversely the widely used deviance information criterion (DIC) [25] is based on a point estimate of the model parameters, namely the posterior means. In Stan we obtain each individual's ($i = 1, \dots, N$) contribution to the total log-likelihood for the model (as defined in Section 3) at each iteration of the MCMC sampling and – following the guidance of Vehtari and Gelman [26] – we use these log-likelihood contributions to calculate the WAIC.

The fixed effect parameter estimates from the model with the quadratic association structure are shown in the supplementary material. The point estimate for α_2 (-0.23) was negative, suggesting that the magnitude of the association between the expected viral load and the hazard of stopping or

modifying treatment diminishes with increasing viral load. However the estimate for α_2 had a wide 95% credible interval (-0.91 to 0.28) incorporating zero, suggesting any evidence of non-linearity was weak. Furthermore, the WAIC values for all four models were in fact very similar. For models with the null, Bernoulli-based, hurdle-based (linear) and hurdle-based (quadratic) association structures the estimated WAICs were 6530.1, 6526.8, 6520.2 and 6525.0 respectively. Therefore, in the context of this application there appeared to be little benefit to allowing for an association between the longitudinal and time-to-event outcomes.

5. Discussion

In this paper we have presented the framework for a joint longitudinal and time-to-event model which incorporates a “hurdle” model for the longitudinal response. A key advantage of the model is it accommodates a skewed and left-censored distribution for the longitudinal response, thereby utilising the full information available in the observed data. Our joint model also demonstrates the use of a novel association structure, based on the combined expected outcome from the longitudinal hurdle model.

In the application we estimated the association between the amount by which $\log_{10}$ viral load exceeds a clinically meaningful detection limit and the hazard of stopping or modifying treatment in patients on their first ART regimen. The association structure we used is meaningful in the context of our application, however for other applications a different parameterisation for the joint model association structure may be necessary. For example, Hatfield et al. [13] used a joint model where the random intercepts from the two parts of the hurdle model were used to form the association structure with the time-to-event model. In the application we found little difference in model fit (based on WAIC) between our model and one which incorporated no association between the longitudinal and time-to-event outcomes. This finding is likely due to the fact that the estimated association parameter α in our model was close to zero (a hazard ratio close to 1) suggesting there is only weak evidence of

an association between the longitudinal marker ($\log_{10}$ viral load) and the event (stopping or modifying treatment) in any case. A similar finding was reported by Thiebaut et al. [9] who used a bivariate joint model for two longitudinal markers (CD4 count and left-censored $\log_{10}$ viral load) and an event (stopping or modifying a first ART regimen). Although the parameterisation of their model was quite different to the model we have presented here, the authors reported that time to treatment modification was more strongly associated with changes in CD4 count and less strongly associated with changes in viral load. We focussed on the association of viral load with switching treatment, although it would be possible to extend the model by including a further longitudinal marker model for CD4 count.

5.1 Limitations of the data

A significant limitation in the application presented here is that the dataset allowed little flexibility for modelling longitudinal trends in the gamma part of the hurdle model. This was due to the fact there were generally few viral load measurements per patient observed above the detection limit. If more data had been available then more complex longitudinal trajectories could have been considered for both the Bernoulli and gamma parts of the hurdle model, which potentially was necessary to realise the underlying association structure with the hazard of modifying or stopping treatment. Nonetheless, with improvements in technology the viral load assays used in clinical practice are becoming increasingly sensitive and, therefore, we would expect the proportion of viral load measurements below the detection limit to be smaller in future studies. Furthermore, in other applications (for example inpatient costs) or other settings (for example viral load measurements taken prior to ART) the censoring rate for the longitudinal measurements may be significantly less.

Another important consideration is that HIV treatment regimens are often relatively complex, generally consisting of three or more drugs at any one time. Treatment modification may be related to just one or more of those drugs, with various potential reasons for switching or stopping each drug. We limited our analysis to each patient's first regimen since this provided some consistency in the

viral load trajectory one might expect. However it is likely that many of the events in our dataset are related to background reasons for switching, unrelated to viral load, thereby making any association between viral load and treatment modification difficult to detect. A study on durability of first ART regimen found that less than 8% of treatment modifications were attributed to virological failure [27]. Other common reasons for modifying treatment include adverse reactions such as intolerance of side-effects, elevated kidney or liver biomarkers, drug-drug interactions in patients developing co-morbidities, patients wanting to change to a newer drug or a single tablet regimen, pregnancy-related issues, and lack of adherence. Furthermore we did not have information on patient deaths, although the total number of deaths was likely to be relatively small in this cohort.

5.2 Model assumptions

Our joint model also depends on several parametric assumptions. We assumed a Weibull distribution for the baseline hazard in the time-to-event model, chosen partly for its simplicity but also because it appeared adequate in exploratory analyses. Extensions to other baseline hazard functions are of course possible, for example the use of other standard distributions such as the log-normal or gamma which can be easily implemented in Stan. A piecewise constant baseline hazard was used by Rizopoulos and Ghosh [28] who provided WinBUGS code for fitting their multivariate joint model. Alternatively restricted cubic splines have been used to estimate flexible non-linear baseline hazard functions in joint models [2, 16].

We have shown our results were slightly sensitive to the number and location of knot points for the restricted cubic splines in the Bernoulli part of the model. In some regard the sensitivity to the number of spline terms is to be expected since the model assumed for the longitudinal marker trajectory will have a direct impact on the association parameter [16]. The sensitivity of the estimated association parameter in our application may also be attributable to a weak “true” underlying association between the marker and event. The knot locations chosen for our final model were guided in part by model fit.

In future work we plan to assess the performance of our model in simulation studies, including a comparison of our model with alternative joint models for censored longitudinal data. For example, the Tobit model has been previously proposed for modelling censored longitudinal data in a joint modelling framework [29]. We also plan to further investigate the role and potential benefits of using Stan for fitting joint models.

Acknowledgements

This paper presents independent research funded by the National Institute for Health Research (NIHR). The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health. SLB was funded in part by an NIHR Methods Fellowship (MET 12-20). MJC was part funded by a NIHR Doctoral Research Fellowship (DRF-2012-05-409). MM was supported by the UK Medical Research Council (MRC) [grant number MR/J002380/1] and the UK Department for International Development (DFID) under the MRC/DFID Concordat agreement and is also part of the EDCTP2 programme supported by the European Union. We would also like to thank the patients and clinical teams for their contribution.

References

1. Lawrence Gould A, Boye ME, Crowther MJ, Ibrahim JG, Quartey G, Micallef S, Bois FY. Joint modeling of survival and longitudinal non-survival data: current methods and issues. Report of the DIA Bayesian joint modeling working group. *Stat Med* 2014.
2. Crowther MJ, Abrams KR, Lambert PC. Flexible parametric joint modelling of longitudinal and survival data. *Stat Med* 2012; **31**: 4456-4471.
3. Henderson R, Diggle P, Dobson A. Joint modelling of longitudinal measurements and event time data. *Biostatistics* 2000; **1**: 465-480.
4. Pantazis N, Touloumi G. Bivariate modelling of longitudinal measurements of two human immunodeficiency type 1 disease progression markers in the presence of informative drop-outs. *Applied Statistics* 2005; **54**: 405-423.
5. Sweeting MJ, Thompson SG. Joint modelling of longitudinal and time-to-event data with application to predicting abdominal aortic aneurysm growth and rupture. *Biom J* 2011; **53**: 750-763.
6. Wulfsohn MS, Tsiatis AA. A joint model for survival and longitudinal data measured with error. *Biometrics* 1997; **53**: 330-339.
7. Wang Y, Taylor JMG. Jointly modeling longitudinal and event time data with application to acquired immunodeficiency syndrome. *Journal of the American Statistical Association* 2001; **96**: 895-905.
8. Guo X, Carlin BP. Separate and joint modeling of longitudinal and event time data using standard computer packages. *Am Stat* 2004; **58**: 16-24.
9. Thiebaut R, Jacqmin-Gadda H, Babiker A, Commenges D, Collaboration C. Joint modelling of bivariate longitudinal data with informative dropout and left-censoring, with application to the

evolution of CD4+ cell count and HIV RNA viral load in response to treatment of HIV infection. *Stat Med* 2005; **24**: 65-82.

10. Mullahy J. Specification and testing of some modified count data models. *J Econometrics* 1986; **33**: 341-365.

11. Cooper NJ, Lambert PC, Abrams KR, Sutton AJ. Predicting costs over time using Bayesian Markov chain Monte Carlo methods: An application to early inflammatory polyarthritis. *Health Econ* 2007; **16**: 37-56.

12. Tooze JA, Grunwald GK, Jones RH. Analysis of repeated measures data with clumping at zero. *Stat Methods Med Res* 2002; **11**: 341-355.

13. Hatfield LA, Boye ME, Hackshaw MD, Carlin BP. Multilevel Bayesian Models for Survival Times and Longitudinal Patient-Reported Outcomes With Many Zeros. *Journal of the American Statistical Association* 2012; **107**: 875-885.

14. Liu L. Joint modeling longitudinal semi-continuous data and survival, with application to longitudinal medical cost data. *Stat Med* 2009; **28**: 972-986.

15. Rizopoulos D, Verbeke G, Lesaffre E, Vanrenterghem Y. A two-part joint model for the analysis of survival and longitudinal binary data with excess zeros. *Biometrics* 2008; **64**: 611-619.

16. Crowther MJ, Andersson TM, Lambert PC, Abrams KR, Humphreys K. Joint modelling of longitudinal and survival data: incorporating delayed entry and an assessment of model misspecification. *Stat Med* 2015 [Epub ahead of print].

17. Stan: A C++ Library for Probability and Sampling, Version 2.6.0, 2014.

18. Su L, Tom BD, Farewell VT. Bias in 2-part mixed models for longitudinal semicontinuous data. *Biostatistics* 2009; **10**: 374-389.

19. Laurie DP. Calculation of Gauss-Kronrod quadrature rules. *Math Comput* 1997; **66**: 1133-1145.

20. Williams I, Churchill D, Anderson J, Boffito M, Bower M, Cairns G, Cwynarski K, Edwards S, Fidler S, Fisher M, Freedman A, Geretti AM, Gilleece Y, Horne R, Johnson M, Khoo S, Leen C, Marshall N, Nelson M, Orkin C, Paton N, Phillips A, Post F, Pozniak A, Sabin C, Trevelion R, Ustianowski A, Walsh J, Waters L, Wilkins E, Winston A, Youle M, Grp W. British HIV Association guidelines for the treatment of HIV-1-positive adults with antiretroviral therapy 2012 (2013 update). *HIV Med* 2014; **15**: 1-85.
21. Greub G, Cozzi-Lepri A, Ledergerber B, Staszewski S, Perrin L, Miller V, Francioli P, Furrer H, Battegay M, Vernazza P, Bernasconi E, Gunthard HF, Hirschel B, Phillips AN, Telenti A, Frankfurt HIVCC, the Swiss HIVCS. Intermittent and sustained low-level HIV viral rebound in patients receiving potent antiretroviral therapy. *AIDS* 2002; **16**: 1967-1969.
22. Durrleman S, Simon R. Flexible regression models with cubic splines. *Stat Med* 1989; **8**: 551-561.
23. Watanabe S. Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. *Journal of Machine Learning Research* 2010; **11**: 3571-3594.
24. Gelman A, Hwang J, Vehtari A. Understanding predictive information criteria for Bayesian models. *Statistics and Computing* 2014; **24**: 997-1016.
25. Spiegelhalter DJ, Best NG, Carlin BR, van der Linde A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society Series B-Statistical Methodology* 2002; **64**: 583-616.
26. Vehtari A, Gelman A. WAIC and cross-validation in Stan. Submitted. In WAIC and cross-validation in Stan. Submitted, Editor (ed)^(eds). City, 2014.
27. Abgrall S, Ingle SM, May MT, Costagliola D, Mercie P, Cavassini M, Reekie J, Samji H, Gill MJ, Crane HM, Tate J, Sterling TR, Antinori A, Reiss P, Saag MS, Mugavero MJ, Phillips A, Manzardo C, Wasmuth JC, Stephan C, Guest JL, Gomez Sirvent JL, Sterne JA, Antiretroviral Therapy Cohort C.

Durability of first ART regimen and risk factors for modification, interruption or death in HIV-positive patients starting ART in Europe and North America 2002-2009. *AIDS* 2013; **27**: 803-813.

28. Rizopoulos D, Ghosh P. A Bayesian semiparametric multivariate joint model for multiple longitudinal outcomes and a time-to-event. *Stat Med* 2011; **30**: 1366-1380.

29. Pike F, Weissfeld L. Joint modeling of censored longitudinal and event time data. *Journal of Applied Statistics* 2013; **40**: 17-27.

Figure 1. Distribution of $\log_{10}$ viral load measurements which were observed above the detection limit, defined as $>1.7 \log_{10}$ copies/ml (equivalent to >50 copies/ml). In total 73% of viral load measurements were below the detection limit and are not shown in the distribution below.

Table 1. Demographic characteristics for the HIV patients used in the analysis

Characteristic	N (%)
Total patients	550 (100%)
Female	192 (34.9%)
Age at start of ART	
18-29 years	101 (18.4%)
30-39 years	239 (43.5%)
40-49 years	140 (25.5%)
50+ years	70 (12.7%)
Ethnicity†	
White background	300 (54.5%)
Black background	209 (38.0%)
Mixed or other background	28 (5.1%)
Route of transmission	
Heterosexual	306 (55.6%)
Male-male sex	202 (36.7%)
Injection drug use	14 (2.5%)
Other	28 (5.1%)
Calendar period for starting ART	
2000-03	97 (17.6%)
2004-06	117 (21.3%)
2007-09	198 (36.0%)
2010-12	138 (25.1%)
Most common ART first regimens	
EFV/TRU	116 (21.1%)
COM/EFV	70 (12.7%)
ATR	49 (8.9%)
ATV/RTV/TRU	49 (8.9%)
DRV/RTV/TRU	45 (8.2%)
COM/NVP	45 (8.2%)
COM/LPV	27 (4.9%)

Notes. †There were 13 (2.4%) individuals with missing ethnicity data.

Abbreviations. ART: antiretroviral therapy. EFV: Efavirenz. TRU: Truvada. COM: Combivir. ATR: Atripla. ATV:Atazanavir. RTV: Ritonavir. DRV:Darunavir. NVP:Nevirapine. LPV:Lopinavir.

Table 2. Fixed effects parameter estimates from the joint longitudinal hurdle and time-to-event model. The Bernoulli part uses a logit link function and the gamma part uses a log link function. The time-to-event model is a proportional hazards model estimated on the log hazard scale.

	Estimated coefficient (95% CrI)	Odds ratio or hazard ratio (95% CrI)
<i>Marker model: Bernoulli part</i>		
Constant	4.58 (3.77 to 5.44)	
Gender (ref: male)		
Female	-1.09 (-1.60 to -0.59)	0.34 (0.20 to 0.56)
Age (ref: 18-29 years)		
30-39 years	0.15 (-0.46 to 0.76)	1.16 (0.63 to 2.13)
40-49 years	0.08 (-0.60 to 0.77)	1.09 (0.55 to 2.16)
50+ years	0.20 (-0.62 to 1.02)	1.22 (0.54 to 2.77)
Start period (ref: 2010-12)		
2000-03	0.02 (-0.67 to 0.72)	1.02 (0.51 to 2.06)
2004-06	-0.78 (-1.45 to -0.13)	0.46 (0.23 to 0.88)
2007-09	-0.63 (-1.24 to -0.03)	0.53 (0.29 to 0.97)
RCS terms		
RCS1	-19.56 (-21.74 to -17.50)	
RCS2	-36.45 (-41.30 to -31.84)	
RCS3	9.31 (8.08 to 10.59)	
<i>Marker model: gamma part</i>		
Constant	0.00 (-0.07 to 0.07)	
Gamma rate parameter	1.18 (1.06 to 1.30)	
<i>Time-to-event model</i>		
Constant	-0.87 (-1.25 to -0.52)	
Gender (ref: male)		
Female	0.31 (0.09 to 0.52)	1.36 (1.10 to 1.68)
Age (ref: 18-29 years)		
30-39 years	-0.30 (-0.57 to -0.04)	0.74 (0.57 to 0.96)
40-49 years	-0.29 (-0.59 to 0.00)	0.75 (0.55 to 1.00)
50+ years	0.02 (-0.35 to 0.38)	1.02 (0.70 to 1.46)
Start period (ref: 2010-12)		
2000-03	-1.13 (-1.51 to -0.75)	0.32 (0.22 to 0.47)
2004-06	-0.32 (-0.64 to 0.00)	0.73 (0.53 to 1.00)
2007-09	0.04 (-0.26 to 0.34)	1.04 (0.77 to 1.40)
Association parameter	0.44 (0.06 to 0.78)	1.55 (1.06 to 2.19)
Weibull shape parameter	1.14 (1.02 to 1.25)	

Abbreviations. CrI: Bayesian credible interval. ref: reference category. RCS: restricted cubic spline.

Table 3. Standard deviation and correlation parameter estimates for the subject-specific random effects from the joint longitudinal hurdle and time-to-event model.

	Estimated coefficient (95% CrI)
Standard deviations	
Random intercept (Bernoulli part), b_{1i0}	2.30 (1.93 to 2.70)
Random linear slope (Bernoulli part), b_{1i1}	0.88 (0.52 to 1.32)
Random intercept (Gamma part), b_{2i0}	0.32 (0.26 to 0.38)
Correlation parameters	
$\rho_{b_{1i0}, b_{1i1}}$	-0.33 (-0.65 to 0.08)
$\rho_{b_{1i0}, b_{2i0}}$	0.41 (0.17 to 0.63)
$\rho_{b_{1i1}, b_{2i0}}$	-0.09 (-0.48 to 0.28)