

Dynamic survival analysis: Modelling the hazard function via ordinary differential equations

Statistical Methods in Medical Research

1–15

© The Author(s) 2024



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/09622802241268504

journals.sagepub.com/home/smmJ Andres Christen¹ and F Javier Rubio² 

Abstract

The hazard function represents one of the main quantities of interest in the analysis of survival data. We propose a general approach for parametrically modelling the *dynamics* of the hazard function using systems of autonomous ordinary differential equations (ODEs). This modelling approach can be used to provide qualitative and quantitative analyses of the evolution of the hazard function over time. Our proposal capitalises on the extensive literature on ODEs which, in particular, allows for establishing basic rules or laws on the dynamics of the hazard function via the use of autonomous ODEs. We show how to implement the proposed modelling framework in cases where there is an analytic solution to the system of ODEs or where an ODE solver is required to obtain a numerical solution. We focus on the use of a Bayesian modelling approach, but the proposed methodology can also be coupled with maximum likelihood estimation. A simulation study is presented to illustrate the performance of these models and the interplay of sample size and censoring. Two case studies using real data are presented to illustrate the use of the proposed approach and to highlight the interpretability of the corresponding models. We conclude with a discussion on potential extensions of our work and strategies to include covariates into our framework. Although we focus on examples of Medical Statistics, the proposed framework is applicable in any context where the interest lies in estimating and interpreting the dynamics of the hazard function.

Keywords

Autonomous ODE, hazard function, ODE solver, ordinary differential equations

1 Introduction

Survival analysis represents a classical area of Statistics which is concerned with the analysis of times to events, potentially subject to censoring. Survival analysis methods have been applied in a number of areas, including medicine, epidemiology, genetics, engineering, and biology, to name but a few. The survival function and the hazard function represent two quantities of interest in this area. The survival function provides information about the probability that an individual or population will survive beyond a certain time point. On the other hand, the hazard function represents the instantaneous failure rate at each time point. From a more practical perspective, the hazard function can be interpreted as a quantification of the risk of the event of interest occurring at a specific time, and the evolution of such risk over time.¹

A number of methods to estimate the survival function have been proposed. From a nonparametric frequentist perspective, the most popular estimators are the Kaplan–Meier estimator,² which aims at estimating the survival function, and the Nelson–Aalen estimator,^{3,4} which aims at estimating the cumulative hazard function (which is the integral of the hazard function between $t = 0$ and $t = t_0$). From a Bayesian nonparametric perspective, several methods for estimating the survival function based on process priors of different types have been developed (see Hjort et al.⁵ for an overview of these methods).

¹Department of Statistics, Centre for Research in Mathematics (CIMAT), Guanajuato, Mexico

²Department of Statistical Science, University College London, London, UK

Corresponding author:

F Javier Rubio, Department of Statistical Science, University College London, London, UK.

Email: f.j.rubio@ucl.ac.uk

Parametric methods, using Bayesian and frequentist estimation methods, have also regained popularity in survival analysis thanks to the developments of flexible parametric distributions, as well as spline-based methods (see Eletti et al.⁶ for a review), and their appealing interpretability.

Several types of estimators of the hazard function have also been proposed. In the frequentist framework, methods based on kernel density estimation and other data-driven smoothing methods have been proposed to estimate the hazard function in the presence of censoring (see Rebora et al.⁷ for a review on this literature).

From a Bayesian nonparametric perspective, several types of estimators of the hazard function based on gamma processes,⁸ infinite mixtures,⁹ piecewise processes¹⁰ have been studied (see Ibrahim et al.¹¹ for an overview of Bayesian survival methods in survival analysis). Although flexible, all of these methods lack interpretability, as the shape of the hazard function depends on the number of knots, number of components in the mixture, or smoothing parameters. Indeed, in many cases the estimated hazard may exhibit a *wiggly* behaviour,^{9,12,13} which complicates the interpretation of such estimates.

Another type of estimators of the hazard function, which enjoys some popularity in practice, are *piecewise constant* estimators.¹⁴ These estimators impose the condition of constant hazard rates within intervals, and their flexibility depends on the number of intervals and their location.

Parametric approaches are often criticised as the shape of the estimated hazard function is restricted to a few possibilities, depending on the model. For example, the hazard function associated with the Weibull distribution can only be increasing, decreasing or flat, while the hazard function associated with the log-normal distribution can only be unimodal (up-then-down). Moreover, the value of the Weibull hazard function at time $t = 0$ is either 0 or ∞ , which is an unrealistic assumption. There exist three-parameter flexible distributions with positive support, such as the power generalised Weibull and generalised gamma, which can capture the basic hazard shapes (increasing, decreasing, unimodal, and bathtub). However, these distributions also impose the condition that the hazard function at time $t = 0$ is restricted to be either 0 or ∞ (see Rubio et al.¹⁵ for a review on distributions with flexible hazard functions).

Modelling the log-hazard (or log-cumulative-hazard) function using splines has also been considered as an alternative method. However, this requires using a large number of parameters to be able to capture complex shapes, and a careful penalised estimation to obtain smooth curves (see Eletti et al.⁶ for an overview on these methods). More recently, Tang et al.¹⁶ proposed a framework for modelling the dynamic change of the cumulative hazard function through an ordinary differential equation, but focused on the inclusion of covariates through this formulation based on known hazard structures. In a more particular setting, Tang et al.¹⁷ consider a general ordinary differential equation (ODE) form for the cumulative hazard function (allowing for the inclusion of covariates), and estimate it through a feed-forward neural network with the cumulative hazard as an input. Related approaches have been recently explored in Danks and Yau,¹⁸ who used neural network-based ODEs for modelling cumulative distribution functions. Indeed, these methods tend to focus on adding flexibility to the model, rather than interpreting the hazard or cumulative hazard functions.

We propose a novel approach for modelling the *dynamics* of the hazard function using systems of first-order ODEs, taking the hazard function as a (positive) state variable of the system of ODEs. In our context, the word *dynamics* refers to the description of the evolution of the hazard function over time through systems of ODEs. This modelling approach can also lead to very flexible hazard functions, depending on the system of ODEs employed to characterise the dynamics of the hazard function. For instance, it is well known that autonomous scalar ODEs have monotonic solutions (Wallach,¹⁹ see also Section 3). Consequently, using a single ODE to model the hazard function can only capture increasing or decreasing shapes of the hazard function. However, by incorporating additional states, it is possible to obtain nonmonotonic hazard shapes, as shown in Section 3. These internal states allow us to account for complex interactions, delays, or memory effects²⁰ that would not be possible to model with a scalar ODE (see Section 7 for an example). Therefore, the choice of the system of ODEs for modelling the hazard function allows for adding flexibility in an interpretable manner, as the additional hidden states not only enrich the family of hazard function models but they also have a clear dynamic influence on the hazard function itself. The proposed framework is of particular interest to users aiming at understanding and interpreting the evolution of the hazard function over time, in contrast to scenarios where the only aim is to estimate the survival function at specific time points. That is, the proposed approach is relevant in cases where the user is interested in a quantitative and qualitative analysis of the dynamics of the hazard function. We capitalise on the vast literature on ODEs, which provides a number of systems of ODEs with interpretable parameters. In particular, we provide examples using autonomous systems of ODEs which facilitate the interpretation of the dynamics of the hazard function and fill a void in the current literature of hazard-based models. For clarity of exposition and to facilitate a comprehensive analysis of the proposed approach, we will focus on the context without covariates, but we conclude with a discussion on general strategies for the inclusion of covariates.

In Section 2, we present the model formulation and discuss the advantages of modelling the hazard function through autonomous systems of ODEs. We present a discussion on modern tools to verify identifiability of hazard models obtained

as a solution of systems of ODEs. We also compare our formulation against previous models for the cumulative hazard function using ODEs.^{16,17} In Section 3, we present two particular models for the hazard function using classical systems of autonomous first-order ODEs. We provide an interpretation of the additional hidden states and the added flexibility obtained through their inclusion. Section 4 presents a discussion on the calculation of the likelihood function associated with the proposed modelling approach, covering the cases where the system of ODEs has an analytic solution or where the use of numerical ODE solvers is required to approximate this solution. Section 5 presents a simulation study that illustrates the ability to recover the true values of the parameters and hazard shapes, as well as the effects of sample size and censoring. Section 6 presents two case studies that illustrate the use of the proposed modelling approach using real data. We conclude with a discussion and possible extensions of this work in Section 7. Software and real data examples can be found at: <https://github.com/FJRubio67/ODESurv> for R code, and <https://github.com/andreschristen/ODESurv> for Python code.

2 Modelling the hazard function through ODEs

Let $\mathbf{o} = \{o_1, \dots, o_n\}$ be a sequence of survival times, $c_i \in \mathbb{R}_+$ be the corresponding right-censoring times, $t_i = \min\{o_i, c_i\}$ be the observed times, and $\delta_i = I(o_i \leq c_i)$ be the indicator that observation i is uncensored, $i = 1, \dots, n$. Suppose that the survival times are generated by an absolutely continuous probability distribution with positive support, and let $f(t)$ denote its probability density function (pdf), $F(t) = \int_0^t f(r)dr$ be the cumulative distribution function (cdf), and $S(t) = 1 - F(t)$ be the survival function. From these functions, we can also derive the hazard function $h(t) = -S'(t)/S(t)$, where $S'(t) = (d/dt)S(t)$, and the cumulative hazard function $H(t) = \int_0^t h(r)dr = -\log S(t)$.

We propose parametrically modelling the hazard function $h(\cdot)$ through a system of first-order ODEs with initial condition as follows. Let $q_j : \mathbb{R}^+ \rightarrow \mathbb{R}$, $j = 1, \dots, m$, be a collection of differentiable functions, and let us denote $\mathbf{Y}(t) = (h(t), q_1(t), \dots, q_m(t))^\top$, $t > 0$. Define the system of ODEs:

$$\begin{cases} \mathbf{Y}'(t) = \boldsymbol{\psi}_\theta(\mathbf{Y}(t), t) \\ H'(t) = Y_1(t) \end{cases} \quad (1)$$

with initial conditions $\mathbf{Y}(0) = \mathbf{Y}_0$ and $H(0) = 0$; taking the initial time $t_0 = 0$ as a simplification, and $Y_1(t) = h(t)$ by definition. Let $D \subset \mathbb{R}^{m+2}$ be a closed rectangle. Assuming that the vector field $\boldsymbol{\psi}_\theta : D \rightarrow \mathbb{R}^{m+1}$ is Lipschitz continuous in \mathbf{y} for every t ($\boldsymbol{\psi}_\theta(\mathbf{y}, t)$), and that \mathbf{Y}_0 is in the interior of D , then there exists a unique solution of the above initial value problem on a vicinity of the initial value (for details, see Po-Fang and Yasutaka²¹ for example). The vector field $\boldsymbol{\psi}_\theta$, its domain D , and the initial condition \mathbf{Y}_0 are such that the solution for $h(t)$ is positive for all $t > 0$, and for any parameter value $\theta \in \Theta \subset \mathbb{R}^d$. That is, (1) is a family of systems of ODEs for which one state variable is the hazard function h , leading to a family of hazard functions defined by θ and \mathbf{Y}_0 . The cumulative hazard function H is also included in the formulation (1), and is thus obtained in the solution of the system. The initial condition must satisfy $h(0) \geq 0$, but it is otherwise arbitrary. The initial conditions in (1) could include $h(0) = h_0 > 0$ indicating that the hazard function takes a nonnegative, finite, value at $t = 0$, and that all individuals are alive at the start of the follow-up.

While there is no general characterisation of ODE systems guaranteeing positive solutions (specifically, for $Y_1(t)$), several strategies can ensure nonnegativity. One approach involves directly verifying nonnegativity through a qualitative analysis of solutions based on the specific structure of the ODE system. This technique is used, for instance, to prove nonnegativity in competitive species models, such as the Lotka–Volterra equations,²² which we study in Section 3. Indeed, a vast literature exists on population dynamics models using ODEs where solutions are naturally positive.²³

There are also particular theoretical characterisations of systems of ODEs with positive solutions, which are often more restrictive. For example, Bernstein and Bhat²⁴ demonstrates that for autonomous systems (see Section 2.1), if the vector field is essentially nonnegative (i.e. entry-wise) and the initial conditions are nonnegative, then the solutions are nonnegative. Another strategy involves modelling the dynamics of a transformed hazard function instead of directly modelling the hazard function, such as $\tilde{h}(t) = \log h(t)$, which allows for using vector fields that can be negative. This approach is also used to improve computational efficiency and stability of ODE solvers, as discussed in Sections 3 and 6.

Finally, another option consists of ‘forcing’ the solution to become positive.²⁵ This can be done theoretically or numerically by adding a parameter-dependent offset that shifts the solution to the positive quadrant or by discarding parameter values that lead to negative solutions altogether.

In any case, establishing the positivity of state variables in a system of ODEs, as described in (1), needs to be done on a case-by-case basis, considering the characteristics of $\boldsymbol{\psi}_\theta$ and the initial conditions \mathbf{Y}_0 .

Regarding the number of state variables q_j ’s, or *hidden states*, there is no unique rule for determining the number of q_j ’s in the ODE system (1). In practice, the number of hidden states to include in (1) may depend on the biological

interpretation of the system. Additional states can also be added to enhance the model's flexibility or incorporate additional effects. In Section 3, we present an example of an ODE system with two states which represent the interaction between the hazard function associated with a disease and the response resulting from an intervention (treatment) on the population. In Section 7, we present an example of a system with three states, where the third state is included to model time-delay effects.

The increased flexibility obtained through additional hidden states is appealing, but as with any parametric model, ensuring parameter identifiability is crucial for making inference about the parameters. Establishing conditions to guarantee identifiability of the solution of systems of ODEs has received a fair amount of attention in the last couple of decades. Studies such as Miao et al.²⁶ and Qiu et al.²⁷ provide conditions to verify identifiability of parameters for families of linear and nonlinear systems of ODEs. Additionally, the Julia package 'StructuralIdentifiability.jl' implements state-of-the-art numerical methods to detect nonidentifiability. Therefore, we encourage users of our framework to take advantage of these tools to verify identifiability of the parameters of any new system of ODEs. Nonetheless, the literature on systems of ODEs is vast, and the identifiability of many of the classical models (such as those presented in Section 3) has already been established.

Before presenting an analysis of this formulation, a remark on the differences of the proposed approach with that proposed in Tang et al.¹⁶ and Tang et al.¹⁷ seems appropriate. As discussed before, Tang et al.¹⁶ proposed modelling the cumulative hazard function via a scalar ODE. That is,

$$\begin{cases} H'(t; \mathbf{x}) = \Psi(H(t; \mathbf{x}), t; \mathbf{x}) \\ H(t_0; \mathbf{x}) = c(\mathbf{x}) \end{cases} \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^p$ are the available covariates. This formulation aims at modelling the dynamics of the cumulative hazard function, which implicitly means modelling the hazard function $h(t; \mathbf{x}) = H'(t; \mathbf{x})$, rather than obtaining it as a solution, in contrast to our formulation (1). Moreover, our formulation (1) produces the hazard function and the cumulative hazard function as solutions of a system of ODEs, in contrast to (2). Finally, formulation (1) allows for including hidden states variables through the inclusion of additional equations in the system.

On the other hand, note that by setting $h(t | \theta) = \Psi(H(t | \theta), t)$ and $\Psi(H(t_0 | \theta), t_0) = \mathbf{Y}_0 = h_0$, we obtain an (theoretical) equivalence between our formulation (1) and formulation (2), restricted to survival models with differentiable hazard functions and models without additional hidden states. More importantly, Tang et al.¹⁷ propose to define Ψ through a neural network, with several parameters to estimate. Then, the formulation (2) is not seen by Tang et al.¹⁷ as an approach to model the (cumulative) hazard function, but simply as a device to use a neural network to construct it.

In sharp contrast, here we focus on *modelling* a family of hazard functions, h , through the general formulation in (1). The modelling will take advantage on the knowledge available on the survival problem at hand, capitalising the vast literature on models using systems of ODEs.

To provide a comprehensive exploration of the proposed modelling approach, we will focus on the case where no covariates are included in the model. Thus, hereafter, we omit the inclusion of covariates in our notation. A comment is included in Section 7 on some strategies for adding covariates to our formulation.

In some cases, the solution to the system of ODEs (1) is analytic, as we will show in our example in Section 3.2. However, this is more the exception than the rule, as many systems of ODEs do not have an analytic solution. In such cases, one can obtain the solution at specific time points $h(t_i)$ by using an appropriate ODE numerical solver, given the vector field ψ_θ (colloquially known as the 'right hand side' or rhs), the initial conditions $\mathbf{Y}(0) = \mathbf{Y}_0, H(0) = 0$ and the time points t_1, \dots, t_n where the solution must be approximated. Fortunately, the theory of the numerical analysis for systems of ODEs, that is, initial value problems, is quite robust (see Butcher²⁸ for example). Solver implementations include sophisticated implicit step algorithms that dynamically choose methods for stiff and nonstiff systems (e.g. the 'Livermore Solver for Ordinary Differential Equations with Automatic method switching'). Moreover, easy-to-use wrappers of these solvers are available both in R (deSolve, Soetaert et al.²⁹), and in Python (scipy.integrate.solve_ivp, Jones et al.³⁰). The examples presented in Section 3 are implemented both in R and in Python using these solvers.

2.1 Autonomous systems of ODEs for modelling the hazard function

If the vector field in the system of ODEs (1) does not explicitly depend on t , it is said to be an autonomous system.²² That is,

$$\begin{cases} \mathbf{Y}'(t) = \psi_\theta(\mathbf{Y}(t)) \\ H'(t) = Y_1(t) \end{cases} \quad (3)$$

with $\mathbf{Y}(0) = \mathbf{Y}_0$ and $H(0) = 0$, and $Y_1(t) = h(t)$. The functions $q_1(t), \dots, q_m(t)$ are hidden state variables, adding flexibility and richness to the modelling of the hazard function h , as shown in the examples in Sections 3 and 7. In fact, Kunze and Vrscay³¹ show how a target smooth function may be arbitrarily close to the solution of an autonomous ODE, using an N degree polynomial $\psi(h)$, for N sufficiently large. This result in turn shows that the ODE system (3) can be used to produce very flexible hazard functions. However, the primary focus of this paper extends beyond simply developing another flexible hazard modeling framework. Instead, it aims to explore the application of autonomous systems of ODEs, as in (3), to *model* the hazard function h in a manner that offers a clear and useful *interpretation* concerning the survival problem under consideration.

As opposed to time-dependent nonautonomous ODEs, autonomous ODEs provide qualitative insights regarding the dynamic evolution of the system. In Section 3, we present an example illustrating a qualitative analysis of a system of ODEs, which shows how the solutions evolve over time by examining their equilibrium points. Since the times of Isaac Newton, autonomous ODEs have been used to model systems by stating principles or laws, namely, the vector field ψ_θ , that remain fixed in time. This powerful tool has been successfully used to model all sorts of phenomena through ‘basic rules’, ranging from a variety of areas such as physics, chemistry, biology, epidemics, and sociology, to name but a few.³² Our goal is to harness this powerful tool to model hazard functions, for specific problems, based on basic rules dictated by the structure of the system of ODEs. In fact, general qualitative descriptions of hazard functions are already common in survival analysis and reliability.^{15,33–35} Here we aim to go one step further, turning qualitative assessments into basic interpretable rules, as a tool for modelling hazard functions, as will be explained in the examples in Section 3.

3 Examples of hazard models defined through ODEs

In this section, we present three examples that we consider of practical interest. In Section 3.1, we show that the ODE formulation (1) can be used to represent common probability distributions with positive support, but also contains extensions, including distributions with a hazard function h with positive and finite initial condition $h(0)$. A particular model without hidden states, obtained by specifying the hazard function as the solution to the logistic-growth model,²² is presented in Section 3.2. This model has been largely used in ecology for modelling population growth, and represents a classical example in most textbooks on ODEs. This model does not include hidden state variables q and has an analytic solution. Section 3.3 introduces a model obtained by defining the hazard function as the solution of a version of the competitive Lotka–Volterra model.^{36,37} The competitive Lotka–Volterra model is often used to model species in competition, but here we provide an alternative use and interpretation of this model to represent the interaction of the hazard associated with a disease and the response coming from a combination of therapeutic interventions and the immune system (at the population level). This model has a hidden state variable q and no analytic solution, so a numerical solver is employed to approximate its solution. This model also illustrates the richness of solutions one can obtain by adding a hidden state to the logistic-growth model.

3.1 Probability distributions defined through a first-order ODE

Note that any given smooth hazard function $h(t)$ may be trivially formulated as the solution of an ODE, with $\psi(h, t) = h'(t)$. The following remark shows that our formulation (1) can be used to represent common absolutely continuous probability distributions with positive support by specifying a Bernoulli-type differential equation.²²

Remark 1. Let $f(t)$ be a differentiable probability density function $f(t)$ with positive support. By differentiating $h(t) = f(t)/S(t)$, it follows that

$$h'(t) = a(t)h(t) + h(t)^2 \quad (4)$$

where $a(t) = f'(t)/f(t) = (d/dt) \log f(t)$. Equation (4) represents a Bernoulli differential equation.³⁸ Equation (4) can also be seen as a Riccati-type differential equation.³⁹

The term $a(t)$ in (4) may explicitly depend on t , which implies that the corresponding ODE is nonautonomous. If the term $a(t)$ in (4) can be written in terms of $h(t)$ only, then the corresponding ODE is autonomous. Thus, the term $a(t)$ in (4) can be interpreted as an ‘autonomy coefficient’.

For example, the Weibull hazard function, $h(t) = \beta\kappa t^{\kappa-1}$, may be obtained as the solution to the ODE (4) with $a(t) = (\kappa - 1)(\beta\kappa/h(t))^{1/(\kappa-1)} - h(t)$, for $\kappa \neq 1$. This is an autonomous ODE. In contrast, for the log-normal distribution $a(t) = (\mu - \sigma^2 - \log(t))/\sigma^2 t$, which implies that the ODE (4) associated with the log-normal distribution is nonautonomous. This is in line with the classical result that, under mild regularity conditions, autonomous scalar ODEs have monotonic solutions.¹⁹ These ideas are summarised in our context in the following remark.

Remark 2. Consider the first-order (one-dimensional) autonomous ODE

$$h'(t) = \psi(h), \quad h(t_0) = h_0$$

where ψ is a continuous function, and $h(t) > 0$ is the corresponding solution. Then, $h(t)$ is monotonic in $t > 0$. Moreover, let $h(t) > 0$ be the solution to a first-order (one-dimensional) ODE

$$h'(t) = \psi(h, t), \quad h(t_0) = h_0 \quad (5)$$

If $h(t)$ is monotonic in $t > 0$, then the ODE (5) is autonomous. That is, $\psi(h, t)$ depends only on h .

This remark also implies that nonmonotonic hazard functions can only be obtained from first-order, one-dimensional, nonautonomous ODEs. As we will show later, one can also obtain nonmonotonic hazard shapes with *systems* of autonomous ODEs with hidden states.

Some references have claimed that equations of type (4) can be seen as a characterisation of absolutely continuous distributions with positive support (see Baran and Coughlin³⁸ and Onose³⁹). However, our formulation (1) indicates otherwise, as formulations with hidden state variables cannot be reduced to a Bernoulli ODE of type (4). Moreover, for most of the commonly used distributions, the corresponding hazard function $h(t)$ is the solution to an equation of type (4), with $\lim_{t \rightarrow 0} h(t)$ restricted to be either 0 or ∞ .

3.2 Logistic growth hazard model

Next, we present an example of (1) without hidden state variables, and with positive and finite initial conditions on h . The logistic-growth hazard model is defined through the following system of ODEs,

$$\begin{cases} h'(t) = \lambda h(t) \left(1 - \frac{h(t)}{\kappa}\right) & h(0) = h_0 \\ H'(t) = h(t), & H(0) = 0 \end{cases} \quad (6)$$

where $\lambda > 0$ represents the intrinsic growth rate of the hazard function, $\kappa > 0$ represents the upper or lower bound of the hazard function, and $h_0 > 0$ is the value of the hazard function at $t = 0$. The logistic-growth ODE is an autonomous Bernoulli ODE of type (4), with $a(t) = \lambda - (1 + \frac{\lambda}{\kappa})h(t)$. This ODE has the following analytic solution:

$$h(t \mid \lambda, \kappa, h_0) = \frac{\kappa h_0 e^{\lambda t}}{\kappa + h_0 (e^{\lambda t} - 1)}$$

The solution $h(t)$, $t \in [0, \infty)$ is positive and finite, and $\lim_{t \rightarrow \infty} h(t) = \kappa$. The corresponding cumulative hazard function is

$$H(t \mid \lambda, \kappa, h_0) = \frac{\kappa}{\lambda} \log \left(\frac{\kappa + h_0 (e^{\lambda t} - 1)}{\kappa} \right)$$

It is straightforward to see that its corresponding pdf is

$$f(t \mid \lambda, \kappa, h_0) = \frac{\kappa^2 h_0 e^{\lambda t - \kappa/\lambda}}{(\kappa + h_0 (e^{\lambda t} - 1))^2}$$

Simulating from this model can be done by inverting the cdf, as follows. Let u be realisation of a $U(0, 1)$ distribution. Then,

$$t = \frac{1}{\lambda} \log \left[1 + \frac{\kappa}{h_0} \left\{ \exp \left(-\frac{\lambda}{\kappa} \log(1 - u) \right) - 1 \right\} \right]$$

is a realisation of the logistic-growth model. An advantage of model (6) is that it does not force the hazard function to be either 0 or ∞ at $t = 0$, in contrast to the Weibull distribution (among other distributions). Moreover, the shapes of the

hazard function are also increasing or decreasing (see Figure 1). The parameters have a clear interpretation as h_0 represents the value of the hazard function at $t = 0$, κ represents the asymptotic point of stability of the hazard function, and λ represents the growth rate of the hazard function. Figure 1 show some examples of the logistic-growth hazard model. The numerical solutions are only presented for illustration, since there exists an analytic solution. The numerical solutions were calculated using the ‘Livermore Solver for Ordinary Differential Equations with Automatic method switching’ (LSODA, SciPy implementation). The analytic solution is plotted for $h_0 = \kappa/10$, for comparison.

In all cases, we can notice how $h(t) \rightarrow \kappa$ as $t \rightarrow \infty$.

3.3 Hazard-response model

The hazard-response model is defined through the system of ODEs:

$$\begin{cases} h'(t) = \lambda h(t) \left(1 - \frac{h(t)}{\kappa}\right) - \alpha q(t)h(t), & h(0) = h_0 \\ q'(t) = \beta q(t) \left(1 - \frac{q(t)}{\kappa}\right) - \alpha q(t)h(t), & q(0) = q_0 \\ H'(t) = h(t), & H(0) = 0 \end{cases} \quad (7)$$

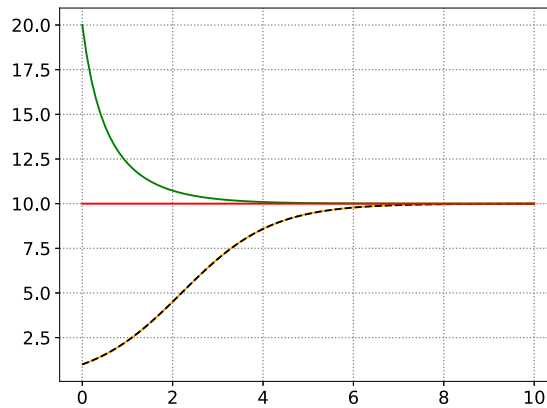


Figure 1. Three examples of the logistic-growth hazard function, $\lambda = 1, \kappa = 10$. Increasing, $h_0 = \kappa/10$ (orange), decreasing, $h_0 = 2\kappa$ (green), and constant, $h_0 = \kappa$ (red). The analytic solution is plotted for $h_0 = \kappa/10$ (dashed black line), for comparison.

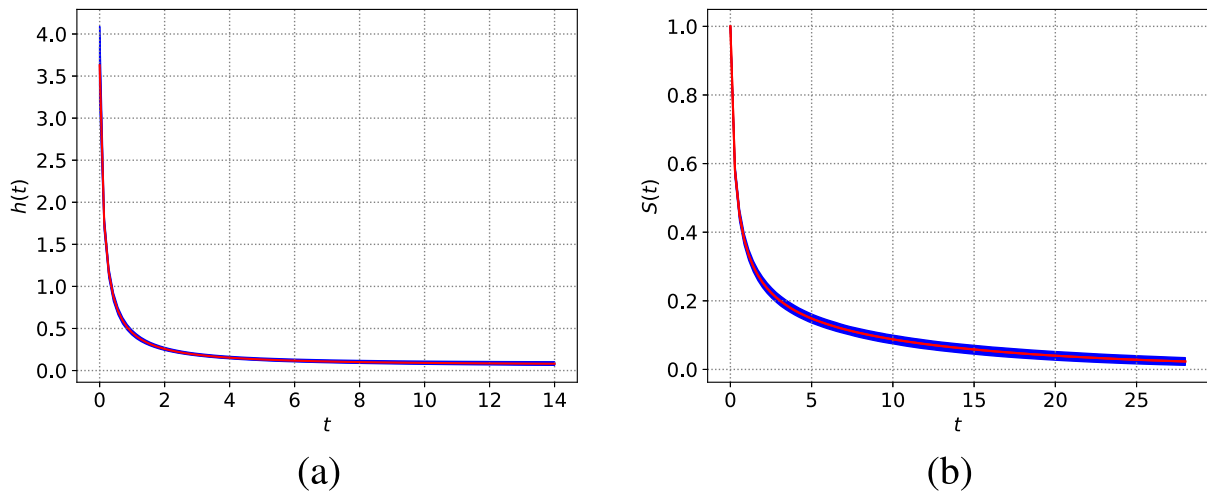


Figure 2. LeukSurv data: (a) posterior predictive hazard function and (b) survival functions.

with $\lambda > 0$, $\alpha \geq 0$, $\beta > 0$, $\kappa > 0$, $h_0 > 0$, and $q_0 > 0$. This model represents a particular version of the competitive Lotka–Volterra model.^{36,37} This model formulation assumes that the hazard function $h(t)$ is in competition with a response $q(t)$ (associated with the immune system and any treatment or intervention at the population level). In the absence of competition ($\alpha = 0$), $h(t)$ follows a logistic growth and reaches its carrying capacity κ as $t \rightarrow \infty$. An analogous argument applies on $q(t)$ for the case when $\alpha = 0$. When $\alpha > 0$, the competition between these two components is modeled through the term $\alpha q(t)h(t)$, and both terms are affected negatively (as the sign is negative in both equations). The term $q(t)$ may be regarded as a hidden, unobserved variable. In that sense, its units are arbitrary and, without loss of generality, we may fix its carrying capacity equal to the carrying capacity of $h(t)$, namely κ . The competition term could include a different coefficient for each equation but, as a parsimonious modelling approach, we use the same coefficient (α) in both equations. The fully parametrised model is studied in Murray³⁷ in the context of ecological modelling as a generalisation of the Lotka–Volterra predator–prey model.

The qualitative analysis of this system is more complicated than that of the simpler logistic-growth model studied in Section 3.2, and may be found in Murray,³⁷ Section 3.5. Briefly, this model has four steady state solutions. These steady states can be found by solving the equation $\psi_\theta(h, q) = \mathbf{0}$, which defines four cases: Case 1, $h = 0, q = 0$; Cases 2 and 3, $h = \kappa, q = 0$ and $h = 0, q = \kappa$, respectively, and Case 4, $h = h^*, q = q^*$, if both positive, where

$$h^* = \kappa \left(\frac{1 - \alpha\kappa\lambda^{-1}}{D} \right) \quad \text{and} \quad q^* = \kappa \left(\frac{1 - \alpha\kappa\beta^{-1}}{D} \right)$$

with $D = 1 - ((\alpha\kappa)^2/\lambda\beta)$. Case 4 is only relevant when $D \neq 0$. The next step is to investigate if these cases are attractors, that is, if $h(t), q(t)$ tend to one of these points as $t \rightarrow \infty$, or not, and under what conditions (see Murray,³⁷ Section 3.5 for details).

Case 1 is not possible as we are assuming $h_0 > 0$ and $q_0 > 0$. It is not possible that $h^* < 0$ and $q^* < 0$ simultaneously. If $\alpha\kappa\lambda^{-1} < 1$ and $\alpha\kappa\beta^{-1} > 1$, then Case 2 is an attractor, that is, the hazard $h(t)$ reaches its carrying capacity as $t \rightarrow \infty$ and the response $q(t)$ ‘loses’ the competition ($q(t) \rightarrow 0$). Conversely, if $\alpha\kappa\lambda^{-1} > 1$ and $\alpha\kappa\beta^{-1} < 1$, then Case 3 is an attractor, that is, the response $q(t)$ ‘wins’ the competition ($q(t) \rightarrow \kappa$) and the hazard ‘loses’ the competition ($h(t) \rightarrow 0$). Case 4 is specially interesting; since both $h^* > 0$ and $q^* > 0$, the hazard remains positive ($h(t) \rightarrow h^* > 0$), although not at its maximum (or minimum), but in an equilibrium with the response ($h^* < \kappa$). However, it is only an attractor when $D > 0$. We will say the hazard-response model is in equilibrium if it is in Case 4 with $D > 0$. If $D < 0$ then Case 2 or Case 3 are the attractors, depending on the initial conditions. We will come back to this classification when analysing the real data using this hazard function in Section 6.

Given that model (7) does not admit an analytic solution, it is not possible to invert the corresponding cdf directly to simulate samples from this model. Alternatively, we propose an algorithm to approximately simulate from this model using the output from an ODE solver. Algorithm 1 shows the steps in the approximate simulation process. The quality of the approximated simulated samples depends on the number of points M , and the choice of an appropriate maximum value t_M .

Algorithm 1. Approximate simulation from the hazard-response model (7)

For given values of the parameters and for a sequence of M time points $\tilde{\mathbf{t}} = \{t_1, \dots, t_M\}$:

1. Obtain a numerical solution of the ODE system (7) at $\tilde{\mathbf{t}}$, using an ODE solver.
 2. Fit a monotonic spline basis with coordinates given by the grid $\tilde{\mathbf{t}}$ and the solutions for the cumulative hazard H at $\tilde{\mathbf{t}}$. This produces an approximation, \tilde{H}^{-1} , of H^{-1} .
 3. Generate $u \sim U(0, 1)$.
 4. Calculate $t^* = \tilde{H}^{-1}(-\log(u))$.
-

4 Inference

Since the proposed hazard models are parametric, the log-likelihood function can be written in terms of the hazard function and the cumulative hazard function, as usual:

$$\ell(\theta, \mathbf{Y}_0) = \sum_{i=1}^n \delta_i \log h(t_i | \theta, \mathbf{Y}_0) - \sum_{i=1}^n H(t_i | \theta, \mathbf{Y}_0)$$

where we are including the initial conditions \mathbf{Y}_0 as unknown parameters in this equation. In some cases, as we will discuss in the simulation study and the applications, in Sections 5–6, it is possible to accurately estimate the initial conditions. However, in other cases, as discussed in Section 6.2, the data contains little information about these parameters, leading to practical nonidentifiability, which is reflected as flat likelihood surfaces.⁴⁰ Indeed, the initial conditions are commonly fixed at reasonable values in many practical applications of ODE models.⁴¹ Once a solution (analytic or numerical) to the system of ODEs (1) has been obtained, we can evaluate this likelihood function at specific values of the parameters θ . In particular, for systems of ODEs without analytic solution and once a solver has been specified, we can retrieve the hazard and cumulative hazard functions, $h(t \mid \theta, \mathbf{Y}_0)$ and $H(t \mid \theta, \mathbf{Y}_0)$, at the sequence of observed time points $\mathbf{t} = \{t_1, \dots, t_n\}$. This allows for calculating the maximum likelihood estimates (MLEs) of the parameters θ and the initial conditions \mathbf{Y}_0 using general-purpose optimisation algorithms (e.g. `optim` or `nlm` in R).

Although it is not possible to come up with general prior choices for the parameters of the hazard models obtained via (1), as different models contain different types of parameters, the ease of interpretation of those parameters facilitates choosing either informative or weakly informative priors. The use of improper priors would require a case-by-case analysis to check the propriety of the posterior distribution.

Since the likelihood function can be evaluated at each parameter value, either for analytic or numerical solutions, any general-purpose sampler can be coupled with the proposed models. These include general Markov chain Monte Carlo (MCMC) samplers such as Metropolis-within-Gibbs sampler (BUGS, `spBayes`), Hamiltonian Monte Carlo (`Stan`), or other ad hoc samplers (`t-walk`, `MCMCPack`).

5 Simulation study

This section presents a simulation study that aims at illustrating the ability to recover the true values of the parameters and hazard shapes, as well as the interplay of sample size and censoring on such aim. We present two simulation scenarios using models (6) and (7), with a combination of sample sizes and censoring rates, as described below.

5.1 Simulation scenarios

For simulation scenario 1, we simulate $M = 250$ samples of sizes $n = 250, 500, 1,000,$ and $5,000$ from the logistic-growth hazard model (6) with parameters $(\lambda, \kappa, h_0) = (0.5, 0.05, 3.5)$. These parameter values produce a decreasing hazard function, which starts at $h_0 = 3.5$, and decreases at a relatively fast rate, $\lambda = 0.5$, to the carrying capacity value $\kappa = 0.05$. These samples are simulated using the procedure described in Section 3.2. Administrative censoring times (i.e. fixed) are used to induce censoring rates of approximately 50% and 25%.

For simulation scenario 2, we simulate $M = 200$ samples of sizes $n = 250, 500, 1,000,$ and $5,000$ from the hazard-response model (7) with parameters $(\lambda, \kappa, \alpha, \beta) = (1.8, 0.1, 6, 4.8)$, using Algorithm 1. The grid $\tilde{\mathbf{t}}$ is defined with equally spaced values on the interval $[0, 150]$, with step size 0.001. In step 2 of Algorithm 1, we use monotone cubic interpolation. We remind the reader that this algorithm produces approximate samples from model (7). Thus, the aim of this scenario is to assess the ability to recover the true values of the parameters and the shape of the hazard using this simulation strategy. Administrative censoring times are used to induce censoring rates of approximately 50% and 25%. Rather than estimating the initial conditions h_0 and q_0 , we fix them at specific values $q_0 = 10^{-2}$ and $q_0 = 10^{-6}$ (in this case, and for illustration purposes, at the true parameter values). These values are inspired by the estimates obtained in case study II, Section 6.2. These parameter values produce a unimodal hazard that stabilises after the mode (stability Case 4). In case study II, Section 6, we will discuss how to use prior information to choose the initial conditions at reasonable values.

For all models, we choose weakly informative priors as follows. For all positive parameters, we adopt gamma priors with scale parameter 2 and shape parameter 2. This prior has mean 4, variance 8, it accumulates 95% of the probability in the interval $(0.5, 11.2)$, and it vanishes at 0 (which helps repelling the MCMC samplers from visiting regions near zero, that may cause numerical problems).

For scenario 1, and for each simulated sample, we obtain a posterior sample of size 1,000 using the `t-walk` sampler⁴² in R, with a burn-in period of 5,000 iterations and a thinning period of 50 iterations (i.e. 55,000 iterations in total). For scenario 2, and for each simulated sample, we obtain a posterior sample of size 1,000 using the adaptive Metropolis-within-Gibbs sampler implemented in the R package `spBayes`, with a burn-in period of 5,000 iterations and a thinning period of 50 iterations, and with approximately 0.44 marginal acceptance rates. In order to obtain a numerically stable solution of the system of ODEs, we first reformulate the system in the logarithmic scale (i.e. in terms of $\log h(t)$, see the online Appendix) and then apply the Livermore Solver for Ordinary Differential Equations (LSODE) solver from the R package `DeSolve`.²⁹ The Jacobian used in the implementation of LSODE is calculated explicitly (see the online Appendix) for a better performance of the ODE solver.

5.2 Results

Tables 1 and 2 in the online Appendix show summaries of the results obtained for simulation scenario 1. These tables show the average posterior mean (across the simulated samples), average posterior median, average posterior standard deviation, average root mean square error (RMSE) with respect to the posterior mean, and coverage of the 95% credible intervals. We can see that it is possible to recover all of the parameters, in the sense that the empirical bias of point Bayes estimates (posterior mean and posterior median) is reduced as the sample size increases or the censoring rate decreases. The empirical coverage is close to the nominal value for all parameters. From these tables, we can also see that, unsurprisingly, the estimation of the initial condition h_0 is the most challenging (see the RMSE and standard deviation values). Figure 3 shows the true hazard function together with the 95% predictive intervals, which are narrow, even for the smallest sample size considered (and while adopting generic weakly informative priors).

Tables 3 and 4 in the online Appendix show the summaries of the results for simulation scenario 2. We can see that the Bayes estimates are close to the true values of the parameters, but exhibit some degree of bias which does not disappear with increasing sample size. This seems to be a result of the step size used in the simulation and the interval length. In order to reduce this bias, it is required to use a smaller step size and larger interval length. On the other hand, we can see

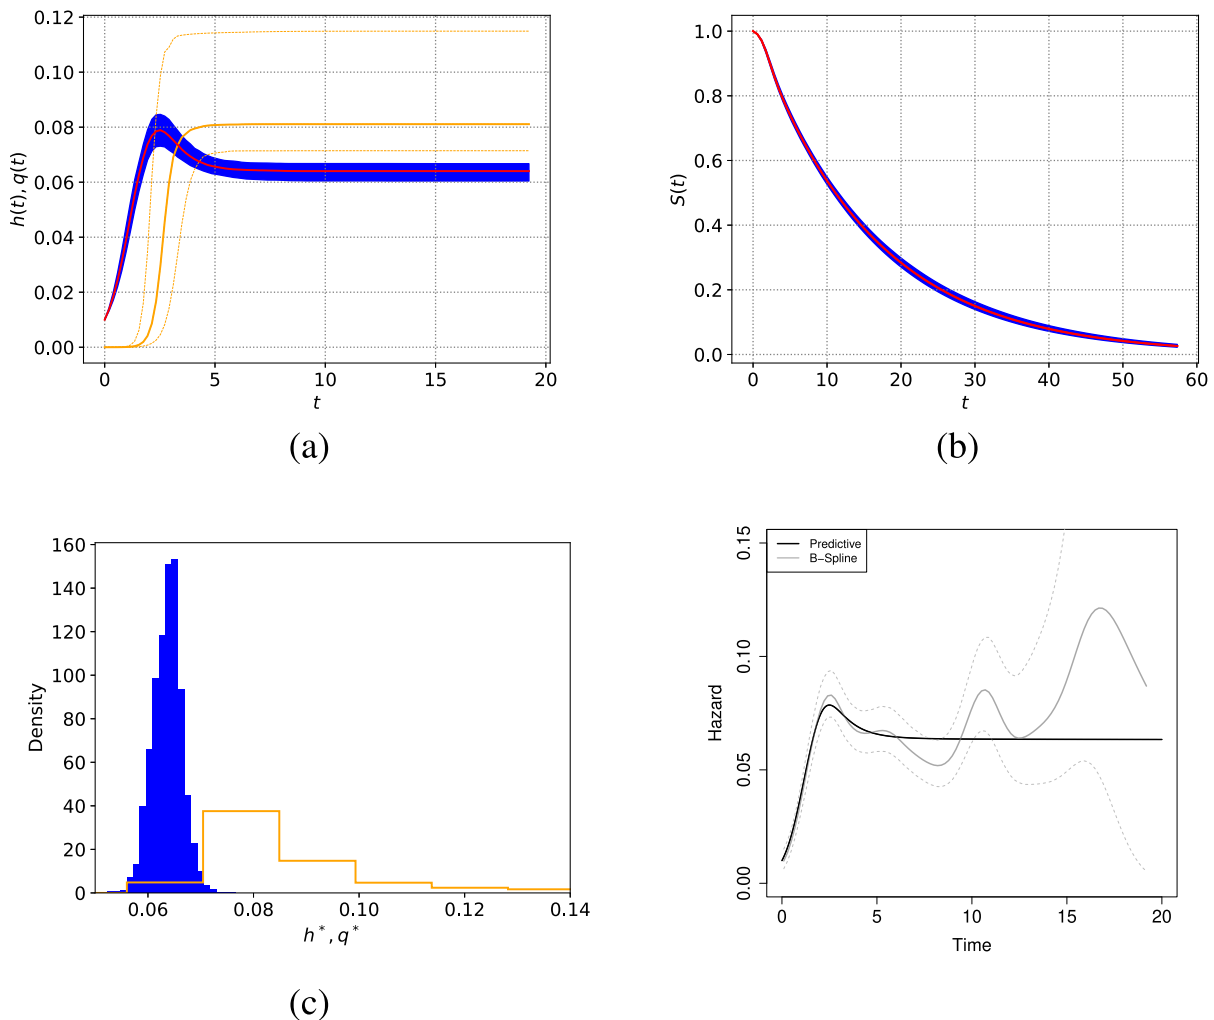


Figure 3. rotterdam data: (a) posterior predictive hazard function (red line) and response function (yellow line); (b) survival function (red line) for the hazard-response model (7), 0.1 to 0.9 quantiles and the median; (c) posterior distribution of h^* (left) and q^* (right); and (d) posterior predictive hazard function (black line) and B-spline estimator (grey line) with 95% confidence interval. The posterior predictive for the 'response' function q is also included (transparent percentile range) along with the hazard function (a). The system is in equilibrium with probability close to 1, and h tends to the asymptotic point h^* .

that this configuration allows one to recover the shape of the hazard with high accuracy, as shown in Figure 4 in the online Appendix, which presents the true hazard function together with the 95% predictive intervals.

For benchmarking purposes, we compare the predictive hazard functions obtained in simulation scenario 2 and a B-spline estimator of the hazard function (using the R package `bshazard`) against the true hazard function. For this purpose, we use the restricted L_1 distance between two hazard functions, h and \tilde{h} , proposed in De Iorio et al.⁴³:

$$d_R(h_0, \tilde{h}) = \int_0^{t^*} |h_0(t) - \tilde{h}(t)| dt$$

We take t^* to be the maximum follow-up time for each case. We calculate this distance in two cases: when \tilde{h} represents the predictive hazard function obtained in simulation scenario 2, and when \tilde{h} is the B-spline estimator of the hazard function, with h_0 being the true hazard function. The corresponding hazard functions are approximated by interpolating the solutions retrieved by the ODE solver and the R package `bshazard` on a grid of 200 equidistant points on $[0, t^*]$ using cubic splines. These approximations are then integrated using the R command `integrate`. Figures 5 and 6 in the online Appendix display the violin plots corresponding to the 250 distances obtained in each case. It is evident that the variability of the B-spline estimator consistently results in higher distances from the true generating model.

6 Real data applications

In this section, we present two case studies using real data on Leukemia and Breast cancer. The first case study concerns the analysis of the `LeukSurv` data set from the R package `spBayesSurv` using the logistic ODE hazard model (6). This represents an example using a hazard model with analytic solution, and where one can estimate all of the model parameters, including the initial conditions. The second case study concerns the analysis of the `rotterdam` breast cancer data set from the R package `survival` using the hazard-response model (7). This represents a case where no analytic solution exists, and thus we need to employ numerical ODE solvers to approximate the solution. Moreover, due to the evolution of breast cancer (the first time-to-event is observed around 6 weeks after the start of follow-up), the data contain no information regarding the initial conditions. We use the interpretation of the model parameters and previous information about this type of cancer in the literature to choose these values.

6.1 Case study I: Leukemia data

In this application, we analyse the `LeukSurv` data set from the `spBayesSurv` R package, using the logistic hazard model (6) presented in Section 3.2. This data set contains information about the survival of $n = 1,043$ patients with acute myeloid leukemia. The maximum follow-up time was 13.6 years for this data set. By the end of the follow-up 879 individuals died and 164 were still alive.

We simulate from the posterior distribution of the parameters (λ, κ, h_0) of the logistic hazard model (6) using the t-walk sampler.⁴² The t-walk is a self-adaptive MCMC algorithm that requires no tuning parameters. We obtain 500,000 posterior samples, and use a burn-in period of 20,000 iterations and a thinning period of 167 iterations. This is the estimated integrated auto-correlation time (see details of output analysis in Molina-Muñoz and Christen⁴⁴), and divided by the dimension (here, three) results in 55.6, leading to an effective sample size of $480,000/167 \approx 2,870$. Figure 2 shows the corresponding posterior predictive hazard and survival functions for this population of cancer patients. The posterior predictive hazard function starts at h_0 (posterior mean ≈ 3.6) and sharply decreases to κ (95% posterior credible interval (0.012, 0.125)) within the first 4 years. These results are in line with the progression of this type of cancer tumours, which tend to develop quickly with only a small portion of the population responding positively to treatments,⁴⁵ leading to poor prognosis and survival. Now, comparing the logistic hazard model against the Weibull distribution using the Bayesian information criterion (BIC):

$$\text{BIC} = k \log(n) - 2\ell(\lambda, \kappa, h_0)$$

where $k = 3$, the total number of unknown parameters, and n is the sample size. the logistic hazard model has a BIC = 1,862.6, while fitting a Weibull distribution to the same data we obtain BIC = 1,898.5. Thus, the data clearly favours the logistic hazard model.

6.2 Case study II: Breast cancer data

In this application, we analyse the `rotterdam` data set⁴⁶ from the `survival` R package. This data set contains information about the survival of $n = 2,982$ breast cancer patients, from which 1,272 died within the maximum follow-up period

(19.3 years). It is known that some of these patients received hormonal treatment, chemotherapy, and/or surgical treatment. We expect that the evolution of the population hazard function over time to depend on the response to the treatments and the natural immunological response. Thus, we model the survival times using the hazard-response model (7). The minimum survival time in this data set is 45 days, thus, we do not expect to have information about the initial conditions h_0 and q_0 . We have verified this assumption by fitting the model where the initial conditions are assumed to be unknown parameters, and we found that the posteriors are virtually the same as the priors, indicating weak identifiability⁴⁰ of h_0 and q_0 . Consequently, we fix these values using prior knowledge. Indeed, fixing the initial conditions in ODE models to some reasonable values is a common practice.⁴¹ It is known that the prognosis of breast cancer is relatively good compared to other cancers, and that the mortality rate during the first few months is very low. Based on this prior information, we assume that the survival probability at one month $\Delta t = 1/12$ is $S(\Delta t) \approx 0.999$. Then, we use the approximation

$$h_0 = h(0) = -\frac{S'(0)}{S(0)} \approx -\frac{S'(\Delta t)}{S(\Delta t)} \approx -\frac{S(\Delta t) - S(0)}{\Delta t S(\Delta t)} \approx 0.01$$

To define the initial condition on q , we use that treatment does not usually start at the beginning of follow-up. Thus, the response on reducing the hazard function should be small at the beginning of follow-up, and we fix this value at $q_0 = 10^{-6}$. One could also assign prior distributions to the initial conditions h_0 and q_0 , as long as they are consistent with these values. As a sensitivity analysis, we provide such implementation in the software provided on our GitHub repositories (<https://github.com/FJRubio67/ODESurv> and <https://github.com/andreschristen/ODESurv>), where we show that very similar results are obtained with both approaches.

Figure 3 shows the posterior predictive hazard, survival, and ‘response’ functions, as well as a comparison against a B-spline estimator of the hazard (using the `bshazard` R package, with three degrees of freedom). Note how the hazard function exhibits an increasing behaviour during the first 2 to 3 years and then decreases and stabilises at a constant value. Interestingly, the equilibrium state is inferred and neither h nor q go to zero. That is, a constant asymptotic hazard h^* is predicted for this data set (see Section 3.3). This coincides nicely with previous studies on the evolution of the population hazard function associated with breast cancer patients.³⁴ The posterior probability of the equilibrium state may be easily calculated, by checking if $h^* > 0$, $q^* > 0$ and $D > 0$ at each iteration of the MCMC. All of our iterations resulted in equilibrium, from an effective sample size of 800. Moreover, the posterior distribution for h^* is presented in Figure 3(c). Figure 3(d) illustrates the large variability of the B-splines estimator of the hazard function compared to the hazard-response model. The ‘wiggleness’ of the B-splines estimator of the hazard function complicates making precise epidemiological statements about the evolution of the disease. Finally, we compare the fit of the hazard-response model against the Weibull and Power Generalised Weibull (PGW, three parameters) distributions using BIC

$$\text{BIC} = k \log(n) - 2\ell(\lambda, \kappa, \alpha, \beta)$$

where $k = 4$, the total number of unknown parameters, and n is the sample size. The corresponding BICs are 9,638.3 and 9,572.0, for the Weibull and PGW distributions, respectively, and 9,561.0 for the hazard-response model. Thus, the BIC clearly favours the proposed ODE model.

7 Discussion

We proposed a novel methodology for modelling the dynamics of the hazard function in survival analysis via systems of ODEs. This framework is particularly useful for researchers interested in qualitative and quantitative analyses of the evolution of the hazard function over time. This modelling approach capitalises on the vast literature on ODEs and systems of ODEs, which have well-understood dynamics and solutions, as well as interpretable parameters. In particular, the models based on autonomous ODEs presented in Section 2.1 represent a novel and interpretable approach for modelling the dynamics of the hazard function. We have presented models with an analytic ODE solution, and consequently their implementation is similar to that of the usual parametric models. We have also shown that numerical ODE solvers can be used in cases where no analytic solution is available, allowing for coupling such models with general-purpose optimisation methods (for likelihood-based inference) as well as MCMC samplers (for Bayesian inference).

The simulation study presented in this work shows it is possible to recover the true parameter values and the shape of the true hazard functions using the proposed modelling approach. This study also provides guidelines about the effect of the sample size and censoring rates on the inference on the parameters, and the increased uncertainty induced by adding more parameters or additional hidden state variables. The case studies presented here illustrate the use of the proposed modelling approach using real data, as well as the qualitative interpretation of the hazard function in the corresponding context. In

particular, Case study II presents an application of the hazard-response model (7), which allows for an interpretation of the competing processes associated with the mortality hazard and the response, at the population level, potentially associated with clinical interventions and the natural immunological response. We conducted model selection using the classical definition of the BIC. Alternatively, one could employ modified versions of BIC tailored for censored samples, or opt for Bayesian methods for model selection, such as the use of Bayes factors or model posterior probabilities.

An important point to consider when choosing an ODE model for the hazard function is the inclusion of a scale parameter. Most classical ODEs already include a scale parameter (potentially under a different parameterisation) to allow for maintaining the same shape under different time scales. However, some textbooks may present ODE formulations without scale parameters, typically to conduct dimensionless analyses. Overall, it is important to keep in mind the need for including scale parameters, as with any other parametric model.

A natural extension of our work consists of using other systems of ODEs to model the hazard function. These include the use of other common population growth models,⁴¹ such as Gompertz or Richard's models, or extensions of competition models. For instance, similar to Farkas,²⁰ we could incorporate the value of the hazard function in the past into the effect on the response (i.e. the response depends on the history and evolution of the hazard) as follows:

$$\begin{cases} h'(t) = \lambda h(t) \left(1 - \frac{h(t)}{\kappa}\right) - \alpha q(t)h(t), & h(0) = h_0 \\ q'(t) = \beta q(t) \left(1 - \frac{q(t)}{\kappa}\right) - \alpha q(t) \int_{-\infty}^t h(\tau)g(t-\tau)d\tau, & q(0) = q_0 \\ H'(t) = h(t), & H(0) = 0 \end{cases}$$

where $g(t) = \eta \exp(-\eta t)$, $\eta > 0$. This represents a Volterra-type integro-differential equation. Since the integral in the previous equation represents a convolution, the integro-differential system can be rewritten in terms of a system of ODEs with an additional state, which avoids the need for numerical integration and allows for applying the methodology proposed in this work

$$\begin{cases} h'(t) = \lambda h(t) \left(1 - \frac{h(t)}{\kappa}\right) - \alpha q(t)h(t), & h(0) = h_0 \\ q'(t) = \beta q(t) \left(1 - \frac{q(t)}{\kappa}\right) - \alpha q(t)\tilde{q}(t), & q(0) = q_0 \\ \tilde{q}'(t) = \eta[h(t) - \tilde{q}(t)], & \tilde{q}(0) = \tilde{q}_0 \\ H'(t) = h(t), & H(0) = 0 \end{cases}$$

This illustrates the variety and richness of models one can obtain by adding hidden states, as well as the interpretability of such additional states.

We have focused on the case without covariates, as we aimed at providing a deeper analysis of the proposed modelling strategy. However, our modelling strategy allows for seamlessly introducing covariates using standard approaches. For instance, one could use the solution to (1) as a baseline hazard in any hazard-based regression model. This includes the proportional hazards, accelerated failure time, accelerated hazards, or general/extended hazards models (see Rubio et al.¹⁵ for a review on classical hazard structures in survival analysis). Another approach that takes advantage of the interpretability of the model parameters consists of modelling them using a transformed linear or additive predictor. These strategies will be explored in future research. The proposed methodology can be extended to other models of interest in survival analysis such as cure models, spatial survival models, relative survival models,¹⁵ and so on.


Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: JAC was partially supported by the ONRG grant N62909-24-1-2016 P01.

ORCID iD

F Javier Rubio  <https://orcid.org/0000-0001-7183-8407>

Supplemental material

Supplemental materials for this article are available online.

References

1. Rinne H. *The hazard rate: Theory and inference (with supplementary MATLAB-programs)*. Giessen, Germany: Justus-Leibig-University, 2014.
2. Kaplan E and Meier P. Nonparametric estimation from incomplete observations. *J Am Stat Assoc* 1958; **53**: 457–481.
3. Aalen O. Nonparametric inference for a family of counting processes. *Ann Stat* 1978; **6**: 701–726.
4. Nelson W. Theory and applications of hazard plotting for censored failure data. *Technometrics* 1972; **14**: 945–966.
5. Hjort N, Holmes C, Müller P, et al. *Bayesian nonparametrics*. New York, USA: Cambridge University Press, 2010.
6. Eletti A, Marra G, Quaresma M, et al. A unifying framework for flexible excess hazard modelling with applications in cancer epidemiology. *J R Stat Soc: Ser C* 2022; **71**: 1044–1062.
7. Reborá P, Salim A and Reilly M. bshazard: A flexible tool for nonparametric smoothing of the hazard function. *R J* 2014; **6**: 114–122.
8. Kalbfleisch J. Non-parametric Bayesian analysis of survival time data. *J R Stat Soc: Ser B (Methodol)* 1978; **40**: 214–221.
9. Kottas A. Nonparametric Bayesian survival analysis using mixtures of weibull distributions. *J Stat Plan Inference* 2006; **136**: 578–596.
10. Arjas E and Gasbarra D. Nonparametric Bayesian inference from right censored survival data, using the Gibbs sampler. *Stat Sin* 1994; **4**: 505–524.
11. Ibrahim J, Chen M and Sinha D. *Bayesian survival analysis*. Vol. 2. New York: Springer, 2004.
12. Hess K, Serachitopol D and Brown B. Hazard function estimators: A simulation study. *Stat Med* 1999; **18**: 3075–3088.
13. McKeague I and Tighiouart M. Bayesian estimators for conditional hazard functions. *Biometrics* 2000; **56**: 1007–1015.
14. Friedman M. Piecewise exponential models for survival data with covariates. *Ann Stat* 1982; **10**: 101–113.
15. Rubio F, Remontet L, Jewell N, et al. On a general structure for hazard-based regression models: an application to population-based cancer research. *Stat Methods Med Res* 2019; **28**: 2404–2417.
16. Tang W, He K, Xu G, et al. Survival analysis via ordinary differential equations. *J Am Stat Assoc* 2023; **118**: 2406–2421.
17. Tang W, Ma J, Mei Q, et al. SODEN: A scalable continuous-time survival model through ordinary differential equation networks. *J Mach Learn Res* 2022; **23**: 1–29.
18. Danks D and Yau C. Derivative-based neural modelling of cumulative distribution functions for survival analysis. In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp.7240–7256.
19. Wallach S. The differential equation $y' = f(y)$. *Am J Math* 1948; **70**: 345–350.
20. Farkas M. Stable oscillations in a predator–prey model with time lag. *J Math Anal Appl* 1984; **102**: 175–188.
21. Po-Fang H and Yasutaka S. *Basic theory of ordinary differential equations*. New York, NY, USA: Springer, 1999.
22. Boyce W, DiPrima R and Meade D. *Elementary differential equations and boundary value problems*. New York: John Wiley & Sons, 2021.
23. Hirsch MW, Smale S and Devaney RL. *Differential equations, dynamical systems, and an introduction to chaos*. New York: Academic Press, 2012.
24. Bernstein D and Bhat S. Nonnegativity, reducibility, and semistability of mass action kinetics. In: *Proceedings of the 38th IEEE Conference on Decision and Control*. IEEE, 1999. vol. 3, pp.2206–2211.
25. Blanes S, Iserles A and Macnamara S. Positivity-preserving methods for ordinary differential equations. *ESAIM: Math Modell Numer Anal* 2022; **56**: 1843–1870.
26. Miao H, Xia X, Perelson A, et al. On identifiability of nonlinear ODE models and applications in viral dynamics. *SIAM Rev* 2011; **53**: 3–39.
27. Qiu X, Xu T, Soltanalizadeh B, et al. Identifiability analysis of linear ordinary differential equation systems with a single trajectory. *Appl Math Comput* 2022; **430**: 127260.
28. Butcher J. *Numerical methods for ordinary differential equations*. 3rd ed. New Jersey, USA: John Wiley & Sons, 2016.
29. Soetaert K, Petzoldt T and Setzer R. Solving differential equations in R: Package deSolve. *J Stat Softw* 2010; **33**: 1–25.
30. Jones E, Oliphant T, Peterson P, et al. SciPy: Open source scientific tools for Python. WebPage, 2001. <http://www.scipy.org/>.
31. Kunze H and Vrscaj E. Solving inverse problems for ordinary differential equations using the Picard contraction mapping. *Inverse Probl* 1999; **15**: 745.
32. Borzi A. *Modelling with ordinary differential equations*. London, England: Chapman & Hall/CRC, 2022.
33. Meeker W, Escobar L and Pascual F. *Statistical methods for reliability data*. 2nd ed. New Jersey, USA: John Wiley & Sons, 2022.
34. Royston P and Parmar M. Flexible parametric proportional-hazards and proportional-odds models for censored survival data, with application to prognostic modelling and estimation of treatment effects. *Stat Med* 2002; **21**: 2175–2197.
35. Simes R and Zelen M. Exploratory data analysis and the use of the hazard function for interpreting survival data: An investigator’s primer. *J Clin Oncol* 1985; **3**: 1418–1431.
36. Kot M. *Elements of mathematical ecology*. Cambridge, UK: Cambridge University Press, 2001.
37. Murray J. *Mathematical biology I: An introduction*. New York, USA: Springer, 2002.
38. Baran R and Coughlin J. An alternative derivation of the hazard rate. *IEEE Trans Reliab* 1987; **36**: 259–260.
39. Onose H. Remarks on the failure rate characterizations. *Bull Fac Sci, Ibaraki Univ Ser A, Math* 1986; **18**: 57–60.

40. Cole D. *Parameter redundancy and identifiability*. Boca Raton, FL: CRC Press, 2020.
41. Simpson M, Browning A, Warne D, et al. Parameter identifiability and model selection for sigmoid population growth models. *J Theor Biol* 2022; **535**: 110998.
42. Christen J and Fox C. A general purpose sampling algorithm for continuous distributions (the t-walk). *Bayesian Anal* 2010; **5**: 263–281.
43. De Iorio M, Johnson W, Müller P, et al. Bayesian nonparametric nonproportional hazards survival modeling. *Biometrics* 2009; **65**: 762–771.
44. Molina-Muñoz J and Christen J. Criterion to determine the sample size in stochastic simulation processes. *Ing Univ* 2022; **26**: 1–21.
45. Shallis R, Wang R, Davidoff A, et al. Epidemiology of acute myeloid leukemia: Recent progress and enduring challenges. *Blood Rev* 2019; **36**: 70–87.
46. Royston P and Altman D. External validation of a Cox prognostic model: principles and methods. *BMC Med Res Methodol* 2013; **13**: 1–15.