Journal Pre-proof

On harmonic oscillator hazard functions

J.A. Christen, F.J. Rubio

Please cite this article as: J.A. Christen and F.J. Rubio, On harmonic oscillator hazard functions. *Statistics and Probability Letters* (2024), doi: https://doi.org/10.1016/j.spl.2024.110304.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

On harmonic oscillator hazard functions

J.A. Christen[1]

Department of Probability and Statistics
Centro de Investigación en Matemáticas (CIMAT-CONAHCYT)
Guanajuato, México

F.J. Rubio[*]

Department of Statistical Science
University College London
London, UK

Abstract

We propose a parametric hazard model obtained by enforcing positivity in the damped harmonic oscillator. The resulting model has closed-form hazard and cumulative hazard functions, facilitating likelihood and Bayesian inference on the parameters. We show that this model can capture a range of hazard shapes, such as increasing, decreasing, unimodal, bathtub, and oscillatory patterns, and characterize the tails of the corresponding survival function. We illustrate the use of this model in survival analysis using real data.

Keywords: Harmonic Oscillator, Hazard Function, ODE, Survival Analysis

1. Introduction

The hazard function plays a key role in the analysis of survival data [1]. For a positive random variable T, with probability density function $f(t)$ and cumulative distribution $F(t)$, the hazard function is defined as:

$$h(t) = \lim_{\Delta t \downarrow 0} \frac{\mathbb{P}\left(t \leq T < t + \Delta t \mid T \geq t\right)}{\Delta t} = \frac{f(t)}{1 - F(t)}. \qquad (1)$$

---
[*]Corresponding author
Email address: `f.j.rubio@ucl.ac.uk` (F.J. Rubio)

Given the intuitive interpretation of the hazard function as the instantaneous failure rate at time $t > 0$, this function serves as basis for defining many survival regression models [2]. Estimation of the hazard function using both parametric and non-parametric methods has received considerable attention. We refer the reader to [3] for a comprehensive overview of these methods, which include the use of parametric distributions, splines, and Bayesian nonparametric methods. In the parametric setting, it is desirable to define models capable of capturing the basic shapes of the hazard function: increasing, decreasing, unimodal (up-then-down), and bathtub (down-then-up). To this end, several generalizations of the Weibull distribution have been proposed, including the generalized gamma distribution, the exponentiated Weibull distribution, and the power generalized Weibull distribution, among others. Similarly, general methods have been proposed to add parameters to a baseline distribution or baseline hazard function. See [4] and [5] for an overview of these methods. More recently, [3] proposed parametrically modeling the hazard function using a system of first-order autonomous ordinary differential equations (ODEs) with positive solutions. This approach offers a general methodology for generating distributions with interpretable parameters and allows for adding flexibility to the resulting solutions by including hidden states.

Following [3], a novel family of hazard functions is derived using a general linear second-order differential equation that represents the classical damped harmonic oscillator. An additional parameter is used to allow the system to reach the equilibrium point at a positive value (in contrast to the classical damped harmonic oscillator, which stabilizes at zero). The resulting model, which can be reduced to three parameters by fixing the initial conditions (as discussed later), contains parameters that influence the shape of the hazard function, resulting in a flexible, tractable, parsimonious (three parameters only) and interpretable model. We characterize the tails of the survival function and the hazard shapes in terms of the parameter values. We present a real data example which shows how the model can capture complex hazard shapes, and that it provides a better fit compared to appropriate competitor models. R code and data are provided in the GitHub repository `https://github.com/FJRubio67/HOH`. Python code may also be obtained from JAC. All proofs are presented in the appendix.

2. Modelling the hazard function with a harmonic oscillator

Let us first introduce some notation and assumptions. Let $\mathbf{o} = (o_1, \ldots, o_n)$ be a sequence of survival times, $c_i \in \mathbb{R}_+$ the right-censoring times, $t_i = \min\{o_i, c_i\}$ be the observed times, and $\delta_i = \mathrm{I}(o_i \leq c_i)$ be the indicator that observation i is uncensored, $i = 1, \ldots, n$. Suppose that the survival times are generated by a continuous probability distribution, with twice continuously differentiable probability density function $f(t)$, cumulative distribution function $F(t) = \int_0^t f(r)dr$, survival function $S(t) = 1 - F(t)$, hazard function $h(t) = \frac{f(t)}{S(t)}$, and cumulative hazard function $H(t) = \int_0^t h(r)dr$.

Based on definition (1), a hazard function for an absolutely continuous, positive, random variable may be defined through a continuous function satisfying $h(t) \geq 0$, and $\int_0^t h(r)dr < \infty$, for $t \geq 0$. Thus, any function satisfying these properties can be used as a hazard function. This allows for defining various types of hazard estimators or hazard-based regression models, such as those that model the hazard function using splines. Building on this idea, we propose modelling the hazard function, $h(t)$, through the linear second order ODE:

$$h''(t) + 2\eta w_0 h'(t) + w_0^2(h(t) - h_b) = 0; \quad h(0) = h_0, \quad h'(0) = r_0. \quad (2)$$

In any case, once h is postulated as a hazard function, this fixes the distribution of the survival times, *i.e.* the probability model involved, and should always be interpreted as in (1). Equation (2) models the acceleration of the hazard function and also has a geometric interpretation, describing the curvature of the hazard function, both at any time $t > 0$. More specifically, this equation can be seen as a shifted version of the *damped harmonic oscillator* [6, 7]. The damped harmonic oscillator is a classical model in Physics which is used to describe the motion of a mass attached to a spring when damping (friction) is present [6, 7]. The damped harmonic oscillator is generally understood as a system that degrades or stabilizes over time due to the combined effects of the restoring force and damping. Thus, the damped harmonic oscillator (2) represents the evolution of the instantaneous failure rate, depicting a system that evolves over time to reach a degraded or stable state. Since the solutions to the damped harmonic oscillator equation can take negative values, and the equilibrium point is exactly 0, we consider

the shifted version (2) of this model to allow for a positive equilibrium point $h_b > 0$. The parameters of this ODE are easily interpreted. The natural frequency, $w_0 > 0$, controls the frequency of oscillation, while the damping coefficient, $\eta > 0$, represents the dissipative forces acting on the system. The shift parameter, $h_b > 0$, represents the stability state or equilibrium level of the solution. In the above parametrization the value of the damping coefficient η determines the system's behavior and, in our case, the shape of the hazard function. The three regimens of the system are the under-damped case ($0 \leq \eta < 1$), the over-damped case ($\eta > 1$), and the critically damped case ($\eta = 1$). We present a more detailed analysis of these cases in the next Sections.

The ODE in (2) defines a parametric hazard function $h(t \mid \boldsymbol{\theta})$ with parameters $\boldsymbol{\theta} = (\eta, w_0, h_b, h_0, r_0)^\top \in \mathbb{R}^4_+ \times \mathbb{R}$. There are parameter values $\boldsymbol{\theta}$ that lead to a negative $h(t \mid \boldsymbol{\theta})$ for some values of t. The reason for this is that the shift parameter $h_b > 0$ only defines the equilibrium position of the hazard function, which may not be sufficient to translate the entire oscillator to the positive quadrant. In such cases, it is necessary to discard parameter values that lead to negative solutions. This approach is related to a method known as *enforcing positivity* of ODE solutions, which is used in several areas [8, 9]. The parameter values leading to positive hazard functions will be referred to as the admissible parameter space, as follows.

Definition 1. *Let* $\boldsymbol{\theta} = (\eta, w_0, h_b, h_0, r_0)^\top \in \mathbb{R}^4_+ \times \mathbb{R}$. *We define the admissible parameter space as:*

$$\Theta_A = \{\boldsymbol{\theta} \in \mathbb{R}^4_+ \times \mathbb{R} : h(t \mid \boldsymbol{\theta}) > 0 \quad \text{for all} \quad t \geq 0\}.$$

In the following Sections, we will present a simple characterization of the admissible parameter space Θ_A based on the parameter values.

2.1. Under-damped case

The analytic solution to (2) in the under-damped case ($0 \leq \eta < 1$) is:

$$h(t \mid \boldsymbol{\theta}) = h_b + A e^{-w_0 \eta t} \sin\left(w_1 t + \phi\right),$$

where $w_1 = w_0 \sqrt{|\eta^2 - 1|}$, A denotes the amplitude of the oscillations, and ϕ is the phase, which represents the position of the hazard function within its cycle of

oscillation. $h(t \mid \boldsymbol{\theta}) - h_b$ tends exponentially to zero as $t \to \infty$, since the term $e^{-w_0\eta t}$ dominates the asymptotic behavior of $h(t \mid \boldsymbol{\theta})$. From the initial conditions $h(0) = h_0$ and $h'(0) = r_0$ one can easily find the value for the amplitude A and the phase ϕ. That is, if $h_b \neq h_0$ then, if $r_0 + w_0\eta(h_0 - h_b) \neq 0$, $A = \dfrac{h_0 - h_b}{\sin(\phi)}$ and $\phi = \tan^{-1}\left(\frac{w_1(h_0 - h_b)}{r_0 + w_0\eta(h_0 - h_b)}\right)$ and, if $r_0 + w_0\eta(h_0 - h_b) = 0$, $\phi = \text{sign}(h_0 - h_b)\frac{\pi}{2}$ and $A = h_0 - h_b$. If $h_0 - h_b = 0$, then $\phi = 0$ and $A = \dfrac{r_0}{w_1}$. The cumulative hazard function can be found by integrating $h(t \mid \boldsymbol{\theta})$ as follows:

$$H(t \mid \boldsymbol{\theta}) = h_b t + \frac{A}{w_0\eta} \frac{\sin(\phi) + \mu\cos(\phi) - e^{-w_0\eta t}(\sin(w_1 t + \phi) + \mu\cos(w_1 t + \phi))}{\mu^2 + 1},$$

where $\mu = \frac{w_1}{w_0\eta} = \frac{\sqrt{|\eta^2 - 1|}}{\eta}$.

If $r_0 < 0$ and/or $h_0 > 2h_b$, $h(t \mid \boldsymbol{\theta})$ may become negative for a region around 0, say $t \in [0, t_M]$. To test if is a set of parameters leads to $h(t \mid \boldsymbol{\theta}) < 0$ for $t \in [0, t_M]$, we can calculate the derivative of the analytic solution to obtain

$$h'(t \mid \boldsymbol{\theta}) = -A e^{-w_0\eta t}\left(w_1\cos(w_1 t + \phi) - w_0\eta\sin(w_1 t + \phi)\right).$$

Solving $h'(t \mid \boldsymbol{\theta}) = 0$ leads to $\tan(w_1 t^* + \phi) = \mu$. This equation has many solutions, but we know that the minimum must be in the first or second critical point after $t = 0$, given that the solution has an envelope function $h_b \pm A e^{-w_0\eta t}$. Therefore, we test these two critical points

$$t^* = \frac{\arctan(\mu) - \phi + (0 \text{ or } \pi)}{w_1},$$

to find the minimum $h(t^* \mid \boldsymbol{\theta})$. If $h(t^* \mid \boldsymbol{\theta}) < 0$, for either solution, then $\boldsymbol{\theta} \notin \Theta_A$, since this implies that $h(t \mid \boldsymbol{\theta}) < 0$ for some region $t \in [0, t_M]$. Otherwise, it follows that $h(t \mid \boldsymbol{\theta}) > 0$ for $t > 0$ and $\boldsymbol{\theta} \in \Theta_A$.

2.2. Over-damped case

The case when $\eta > 1$ is called over-damped, since the state variable (hazard function) $h(t \mid \boldsymbol{\theta})$ does not oscillate and returns, exponentially, to h_b. The analytic solution in this case is

$$h(t \mid \boldsymbol{\theta}) = h_b + e^{-w_0\eta t}\left(\frac{h_0 - h_b + a}{2}e^{w_1 t} + \frac{h_0 - h_b - a}{2}e^{-w_1 t}\right),$$

where $a = \frac{h_0 - h_b}{\mu} + \frac{r_0}{w_1}$ ($w_1 = w_0\sqrt{\eta^2 - 1}$ as above). Since $\sqrt{\eta^2 - 1} < \eta$ then $-w_0\eta \pm w_1 < 0$ and $h(t \mid \boldsymbol{\theta}) - h_b$ also tends exponentially to zero, as $e^{-w_0(\eta - w_1)t}$. The cumulative hazard function is:

$$H(t \mid \boldsymbol{\theta}) = h_b t + \frac{h_0 - h_b + a}{2(-w_0\eta + w_1)}(e^{(-w_0\eta + w_1)t} - 1) + \frac{h_0 - h_b - a}{2(-w_0\eta - w_1)}(e^{(-w_0\eta - w_1)t} - 1).$$

Depending on the value of r_0, $h(t \mid \boldsymbol{\theta})$ may take negative values for a region $t \in$
[0, t_M]. In this case, we know there is at most one critical point, as the solution does not oscillate [6, 7]. By solving $h'(t \mid \boldsymbol{\theta}) = 0$, the critical point is

$$t^* = \frac{1}{2w_1} \log\left(\frac{(h_0 - h_b - a)(w_1 + w_0\eta)}{(h_0 - h_b + a)(w_1 - w_0\eta)}\right),$$

if $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0\eta)} > 0$, otherwise there is no critical point. Moreover, if $r_0 > 0$, the solution is increasing or unimodal (see Proposition 1), which implies that $h(t \mid \boldsymbol{\theta}) > 0$ for $t > 0$, and then $\boldsymbol{\theta} \in \Theta_A$. If $r_0 < 0$, depending on the values of the remaining parameters, the shape of the hazard function can be decreasing or bathtub (see Proposition 1). We can characterize the admissible parameter space as follows. If $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0\eta)} > 0$ and $h(t^* \mid \boldsymbol{\theta}) < 0$, then $\boldsymbol{\theta} \notin \Theta_A$. If $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0\eta)} > 0$ and $h(t^* \mid \boldsymbol{\theta}) > 0$, then $\boldsymbol{\theta} \in \Theta_A$. If $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0\eta)} < 0$, then $\boldsymbol{\theta} \in \Theta_A$.

2.3. Critically-damped case

The case $\eta = 1$ is referred to as critically damped. The analytical solution is

$$h(t \mid \boldsymbol{\theta}) = h_b + \{h_0 - h_b + t(r_0 + w_0(h_0 - h_b))\} e^{-w_0 t}.$$

The corresponding cumulative hazard is:

$$H(t \mid \boldsymbol{\theta}) = h_b t + \frac{(h_0 - h_b)(1 - e^{-w_0 t})}{w_0} + \frac{(r_0 + w_0(h_0 - h_b))}{w_0^2}\left(e^{w_0 t} - w_0 t - 1\right) e^{-w_0 t}.$$

Unfortunately, there is no closed-form expression for the critical point in this case. However, since this is a limit case of the under-damped and over-damped cases, and we will consider continuous prior distributions for η in the Bayesian analysis presented in Section 5, this case will have zero probability. Consequently, it is ignored in our implementation. Nonetheless, the expressions for the hazard and cumulative hazard functions are presented here for completeness.

2.4. Parametrization

Overall, we adopt the parametrization (2), with parameters η, w_0, h_b, h_0, and r_0. Nonetheless, other parametrizations might provide further insights into the model. Looking at the ODE formulation in 2, w_0 is called the "natural frequency" and its units are the inverse of time units. The parameter η has no units, and is referred to as the "damping ratio". Then, $w_1 = w_0\sqrt{|\eta^2 - 1|}$ is the angular frequency in the under-damped case, and has the same units as w_0. This means that if $t \in [0, t_M]$, $t_M > 0$, h will oscillate $\frac{w_1 t_M}{2\pi}$ times in that time window.

Since the hazard function $h(t \mid \boldsymbol{\theta})$ does not contain an explicit scale parameter, it is important to check if one can control the scale of this function through a combination of the parameters. It can be noted that by using the following equivalent parametrization, $(h_b, \eta, w_0, \alpha_0, \theta_0)$ for (2)

$$h''(t) + 2\eta(w_0 h_b)h'(t) + (w_0 h_b)^2(h(t) - h_b) = 0, \quad h(0) = \frac{\alpha_0}{h_b}, \quad h'(0) = \frac{\theta_0}{h_b},$$

h_b becomes a scale parameter in the usual sense. Since both parametrizations are equivalent, the model does not need an additional scale parameter, and which parametrization to use is a matter of convenience. In our applications, we will provide empirical strategies for fixing the initial conditions (see Section 4) and therefore we use the original parametrization (2).

3. Shapes and tail-weight characterization

As previously discussed, one of the desirable properties of parametric survival models is their ability to capture the basic shapes. The following result presents a characterization of the hazard shapes of (2), in terms of the parameter values, which include the basic shapes as well as oscillatory cases. We focus on the under-damped and over-damped cases, which are the main cases of interest for modelling, as discussed in the previous Section. See Figure 1 for examples of the hazard shapes obtained as a solution to (2).

Proposition 1. *Let $T > 0$ be a random variable with hazard function defined by the damped harmonic oscillator model (2), and $\boldsymbol{\theta} \in \Theta_A$. Then, the corresponding hazard function can capture the following shapes:*
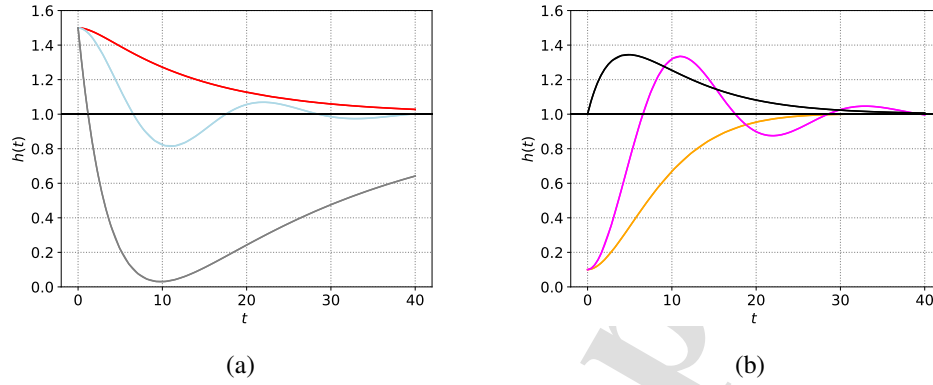
7

Figure 1: Examples of harmonic oscillator hazard functions: (a) Strictly decreasing (red, over-damped), "bathtub" (gray, over-damped) and oscillating (blue, under-damped). (b) Strictly increasing (orange, over-damped), unimodal (black, over-damped) and oscillating (magenta, under-damped). All converge to the asymptotic value h_b ($= 1$).

- *Increasing (monotonic):* $\eta > 1$, $r_0 > 0$, $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0 \eta)} < 0$.

- *Decreasing (monotonic):* $\eta > 1$, $r_0 < 0$, $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0 \eta)} < 0$.

- *Unimodal (up-then-down):* $\eta > 1$, $r_0 > 0$, $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0 \eta)} > 0$.

- *Bathtub (down-then-up):* $\eta > 1$, $r_0 < 0$, $\frac{(h_0 - h_b - a)}{(h_0 - h_b + a)(w_1 - w_0 \eta)} > 0$.

- *Oscillatory (multiple cycles):* $\eta < 1$.

Although some parametric distributions, such as the power generalized Weibull or generalized gamma distributions, can capture the first four shapes, they do not account for oscillatory behaviors. Another important distinction is that the hazard function of the damped harmonic oscillator is always upper-bounded, with its amplitude decaying over time. Consequently, $h(t)$ approaches the stable state h_b as $t \to \infty$. We acknowledge that these are implicit assumptions of the damped harmonic oscillator model, which may not be suitable for all types of data sets. However, we argue that all parametric models have implicit, and often overlooked, limitations. For instance, the Weibull distribution implicitly requires that when the hazard function is increasing, it must satisfy $h(0) = 0$ and remain unbounded, which implies lighter tails than those

8

of the exponential distribution [3]. Conversely, decreasing Weibull hazard functions must satisfy $h(0) = \infty$. The proposed damped harmonic oscillator model allows for a better understanding of the hazard shapes and properties thanks to the interpretability of the parameters in the ODE (2).

180   The following proposition presents a characterization of the right tail of the distribution of a random variable with hazard function defined by the damped harmonic oscillator model (2). As we can see, the shifting parameter plays a key role in controlling the tail-weight in all cases.

Proposition 2. *Let $T > 0$ be a random variable with hazard function defined by the*

185   *damped harmonic oscillator model* (2)*, with $\boldsymbol{\theta} \in \Theta_A$. Then,*

$$\mathbb{P}(T > t \mid \boldsymbol{\theta}) = \mathcal{O}\left(e^{-h_b t}\right).$$

The above result indicates that the distribution associated with the harmonic oscillator hazard function is a sub-exponential distribution [10]. An important difference of model (2) compared to the exponential distribution (which has a constant hazard function) or other parametric sub-exponential distributions [10], is that the proposed

190   model (2) allows for capturing a variety of shapes of the hazard function. This sub-exponential tail-weight characterization also allows for deriving other properties of the distribution induced by (2), such as properties of the moments and moment generating function, as detailed in Chapter 2 of [10], and thus omitted here.

4. Inference

195   The likelihood function for model (2) under right-censoring is fully determined by the hazard and cumulative hazard functions, $h(t \mid \boldsymbol{\theta})$ and $H(t \mid \boldsymbol{\theta})$ as follows,

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^{n} h(t_i \mid \boldsymbol{\theta})^{\delta_i} \exp\left\{-H(t_i \mid \boldsymbol{\theta})\right\}, \quad \boldsymbol{\theta} \in \Theta_A.$$

Since these functions are available in analytic form, the likelihood function can be implemented in any numeric programming language. This also implies that the maximum likelihood estimates can be calculated using general-purpose optimization methods, al-

200   though this may require comparing different methods an initial points since, in general, the concavity of the above likelihood (or log-likelihood) function is not guaranteed.

9

In general, one could consider the initial conditions, h_0 and r_0, as unknown parameters to be estimated. However, estimating these parameters may be challenging in practice for some data sets that do not contain uncensored observations near $t = 0$. Al-

[205] ternatively, one can fix the initial conditions using prior or expert information (see [3] for a discussion on these points). Specifically, setting values for the survival function $S(\Delta t)$ and $S(2\Delta t)$ for some ("small") initial time step Δt, we may approximate the initial condition $h_0 = h(0)$ using

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

Similarly, we may approximate the initial condition $r_0 = h'(0)$ using

$$r_0 = -\frac{S'(0)^2 - S(0)S''(0)}{S(0)^2} \approx \left(\frac{S(\Delta t) - S(0)}{\Delta t S(\Delta t)}\right)^2 - \frac{S(2\Delta t) - 2S(\Delta t) + S(0)}{\Delta t^2 S(\Delta t)}.$$

[210] If there is uncertainty or reservations about these choices, one can choose a prior concentrated on such values. This could serve as an alternative method for estimating the initial conditions or for performing a sensitivity analysis [3]. In the example presented in Section 5 the above method for setting the initial values h_0 and r_0 is used.

5. Real data application

[215] In this Section, we present a real data application in which we analyze the `rotterdam` data set from the `survival` R package. This data set contains information about the survival times of $n = 2,982$ breast cancer patients, of which $1,710$ cases were right-censored. Generally, breast cancer patients have a good prognosis, and the hazard function may begin to increase slowly from the time of diagnosis. Clinically, this indi-

[220] cates that the hazard function starts at a low point and grows slowly at the beginning. Based on these points, and following the discussion presented in Section 4 to fix h_0 and r_0, we assume that $\Delta t = 1/12$ (one month), $S(0) = 1$ (no deaths before the start of follow-up), $S(\Delta t) = 0.999$ and $S(2\Delta t) = 0.998$ (*i.e.* one death per 1000 patients per month, immediately after the start of follow-up). These choices lead to the initial con-

[225] ditions $h_0 = 0.012$ and $r_0 = 0.00014$. Similar values for the initial conditions would be obtained with slight variations in the initial choices, provided they are consistent with the clinical context.

10

We fit the harmonic oscillator model (2), using a Bayesian approach with the fixed initial conditions mentioned above. The three parameters to be estimated are $\boldsymbol{\theta} = (\eta, w_0, h_b)$. We adopt a product prior structure

$$\pi_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \propto \pi(w_0)\pi(\eta)\pi(h_b), \quad \boldsymbol{\theta} \in \Theta_A,$$

where π is a gamma prior with a scale parameter 1000 and shape parameter 0.001. These are "weakly informative" priors which reflect a high degree of prior uncertainty

230 about the parameters. Figure 2 shows the predictive hazard and the predictive survival distributions for these three models, along with the Kaplan-Meier (KM) estimator of survival.

Finally, we compare the fit of the harmonic oscillator model (2) against the Weibull distribution and the power generalized Weibull distribution (PGW), using the Bayesian

235 information criterion (BIC). We choose weakly informative priors for the parameters of these models (gamma priors with scale parameter 1000 and shape parameter 0.001). The Weibull distribution can only capture increasing, decreasing or constant hazard shapes. The PGW distribution is a three-parameter distribution that can capture the basic shapes (except for the oscillatory shape). The BIC for the Weibull model is

240 $9,650.30$, the BIC for the PGW model is $9,590.03$, and the BIC for the harmonic oscillator model is $9,581.037$. Thus, the harmonic oscillator model is clearly favored by the data, based on the BIC.

6. Discussion

We have developed a novel parametric hazard model obtained by enforcing posi-
245 tivity in the damped harmonic oscillator. As shown in the qualitative analysis of the solutions to the corresponding second order ODE, the parameters of this model are interpretable and the model is both tractable and flexible. The real data analysis presented here shows that the proposed model offers competitive performance compared to flexible parametric models commonly used in survival analysis.

250 The focus of this paper has been on survival analysis without covariates. A potential extension involves incorporating covariates into the damped harmonic oscillator hazard
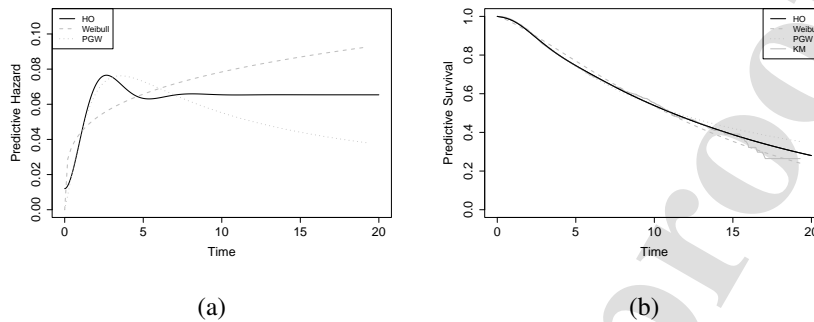
11

(a)

(b)

Figure 2: `rotterdam` data. (a) Predictive hazard functions for the Weibull, PGW, and harmonic oscillator (HO) models; and (b) Predictive survival functions for the Weibull, PGW, and HO models, and Kaplan-Meier (KM) estimator.

function proposed here. This can be achieved by using this model as the baseline hazard in any hazard-based regression model [2], or by introducing covariates through linear predictors on the parameters (similar to distributional regression). These directions will be explored in future work.

Acknowledgements

References

[1] H. Rinne, The Hazard rate: Theory and inference (with supplementary MATLAB-Programs), Justus-Leibig-University, Giessen, Germany, 2014.

[2] F. Rubio, L. Remontet, N. Jewell, A. Belot, On a general structure for hazard-based regression models: an application to population-based cancer research, Statistical Methods in Medical Research 28 (2019) 2404–2417.

[3] J. Christen, F. Rubio, Dynamic survival analysis: modelling the hazard function via ordinary differential equations, Statistical Methods in Medical ResearchIn press.

12

[4] A. Marshall, I. Olkin, A new method for adding a parameter to a family of distributions with application to the exponential and Weibull families, Biometrika 84 (3) (1997) 641–652.

[5] K. Anaya-Izquierdo, M. Jones, A. Davis, A family of cumulative hazard functions and their frailty connections, Statistics & Probability Letters 172 (2021) 109059.

[6] H. Georgi, The Physics of Waves, Prentice Hall, Englewood Cliffs, NJ, 1993.

[7] S. Strogatz, Nonlinear Dynamics and Chaos: with Applications to Physics, Biology, Chemistry, and Engineering, CRC press, Boca Raton, FL, 2018.

[8] L. Shampine, S. Thompson, J. Kierzenka, G. Byrne, Non-negative solutions of ODEs, Applied Mathematics and Computation 170 (1) (2005) 556–569.

[9] S. Blanes, A. I., S. Macnamara, Positivity-preserving methods for ordinary differential equations, ESAIM: Mathematical Modelling and Numerical Analysis 56 (6) (2022) 1843–1870.

[10] R. Vershynin, High-Dimensional Probability: An Introduction with Applications in Data Science, Vol. 47, Cambridge University Press, Cambridge, United Kingdom, 2018.