



HAL
open science

Revisiting Whittaker-Henderson Smoothing

Guillaume Biessy

► **To cite this version:**

| Guillaume Biessy. Revisiting Whittaker-Henderson Smoothing. 2024. hal-04124043v3

HAL Id: hal-04124043

<https://hal.science/hal-04124043v3>

Preprint submitted on 15 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Revisiting Whittaker-Henderson Smoothing

Guillaume Biessy*, PhD, LinkPact[†] and Sorbonne Université[‡]

October 15, 2024

Introduced nearly a century ago, Whittaker-Henderson smoothing is still widely used by actuaries for constructing one-dimensional and two-dimensional experience tables for mortality, disability and other Life Insurance risks. Our paper reframes this smoothing technique within a modern statistical framework and addresses six questions of practical relevance regarding its use.

Firstly, we adopt a Bayesian view of this smoothing method to build credible intervals. Next, we shed light on the choice of the observation and weight vectors to which the smoothing should be applied by linking it to a maximum likelihood estimator introduced in the context of duration models. We then enhance the precision of the smoothing by relaxing an implicit asymptotic normal approximation on which it relies. Afterward, we select the smoothing parameters based on maximizing a marginal likelihood function. We later improve numerical performance in the presence of many observation points and consequently parameters. Finally, we extrapolate the results of the smoothing while preserving, through the use of constraints, consistency between estimated and predicted values.

Keywords: smoothing, duration models, experience tables, maximum likelihood, generalized additive models, empirical Bayes approach, marginal likelihood, extrapolation.

Notations

In this paper, vector names are written in bold characters and matrix names in uppercase letters. If \mathbf{y} is a vector and A is a matrix, $\text{Var}(\mathbf{y})$ denotes the variance-covariance matrix associated with \mathbf{y} , $\mathbf{diag}(A)$ represents the diagonal of matrix A , and $\text{Diag}(\mathbf{y})$ is the diagonal matrix such that $\mathbf{diag}(\text{Diag}(\mathbf{y})) = \mathbf{y}$. The sum of the diagonal elements of A is denoted as

*guillaume.biessy78@gmail.com

[†]LinkPact, 80 boulevard Pasteur, 75015 Paris, France

[‡]Sorbonne Université, CNRS, Laboratoire de Probabilités, Statistique et Modélisation, LPSM, 4 place Jussieu, 75005 Paris, France

$\text{tr}(A)$ and its transpose as A^T . In the case where A is invertible, A^{-1} denotes its inverse, A^{-T} is the inverse of its transpose, and $|A|$ is the product of the eigenvalues of A . For a non-invertible matrix A , A^- refers to the Moore-Penrose pseudoinverse of A , and $|A|_+$ denotes the product of the non-zero eigenvalues of A . By writing the eigendecomposition as $A = U\Sigma V^T$, where U and V are two orthogonal matrices and Σ is a diagonal matrix containing the eigenvalues of A , and by denoting Σ^- as the matrix obtained by replacing the non-zero eigenvalues in Σ with their inverses while keeping the zero eigenvalues unchanged, the pseudoinverse is given by $A^- = V\Sigma^-U^T$. The Kronecker product of two matrices A and B is denoted as $A \otimes B$, and their Hadamard product, or element-wise product, is denoted as $A \odot B$. $\text{vec}(A)$ represents the vector obtained by stacking the columns of matrix A together. Finally, the symbol \propto indicates a proportional relationship between the expressions on both sides of it.

1 Introduction

Whittaker-Henderson smoothing is a graduation method aimed at correcting the effect of sampling fluctuations on a vector of evenly-spaced discrete observations. Initially proposed by Whittaker (1922) for graduating mortality tables and further improved by Henderson (1924), it quickly became one of the most popular graduation method among actuaries, especially in North America. Extending to the two-dimensional case, Whittaker-Henderson (WH) smoothing may be used to build experience tables for a wide range of Life insurance risks including but not limited to disability, long-term care, lapse, mortgage default, and unemployment. We first give a brief reminder of the method then introduce the objectives and plan of the paper.

1.1 A brief reminder of Whittaker-Henderson mathematical formulation

1.1.1 The one-dimensional case

Let \mathbf{y} be a vector of observations and \mathbf{w} a vector of positive weights, both of size n . The estimator associated with Whittaker-Henderson smoothing is given by:

$$\hat{\mathbf{y}} = \underset{\boldsymbol{\theta}}{\text{argmin}}\{F(\mathbf{y}, \mathbf{w}, \boldsymbol{\theta}) + R_{\lambda,q}(\boldsymbol{\theta})\} \quad (1)$$

where:

- $F(\mathbf{y}, \mathbf{w}, \boldsymbol{\theta}) = \sum_{i=1}^n w_i (y_i - \theta_i)^2$ represents a fidelity criterion to the observations,
- $R_{\lambda,q}(\boldsymbol{\theta}) = \lambda \sum_{i=1}^{n-q} (\Delta^q \boldsymbol{\theta})_i^2$ represents a smoothness criterion.

In the latter expression, $\lambda \geq 0$ is a smoothing parameter and Δ^q denotes the forward difference operator of order q , such that for any $i \in [1, n - q]$:

$$(\Delta^q \boldsymbol{\theta})_i = \sum_{k=0}^q \binom{q}{k} (-1)^{q-k} \theta_{i+k}.$$

Let us define $W = \text{Diag}(\mathbf{w})$, the diagonal matrix of weights, and $D_{n,q}$ as the order q difference matrix of dimensions $(n - q) \times n$, such that $(D_{n,q} \boldsymbol{\theta})_i = (\Delta^q \boldsymbol{\theta})_i$ for all $i \in [1, n - q]$. The most commonly used difference matrices of order 1 and 2 have the following forms:

$$D_{n,1} = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \quad \text{and} \quad D_{n,2} = \begin{bmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 1 & -2 & 1 \end{bmatrix}.$$

The fidelity and smoothness criteria can be rewritten with matrix notations as:

$$F(\mathbf{y}, \mathbf{w}, \boldsymbol{\theta}) = (\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta}) \quad \text{and} \quad R_{\lambda,q}(\boldsymbol{\theta}) = \lambda \boldsymbol{\theta}^T D_{n,q}^T D_{n,q} \boldsymbol{\theta}.$$

The associated estimator for smoothing hence becomes:

$$\hat{\mathbf{y}} = \underset{\boldsymbol{\theta}}{\text{argmin}} \left\{ (\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta}) + \boldsymbol{\theta}^T P_\lambda \boldsymbol{\theta} \right\} \quad (2)$$

where $P_\lambda = \lambda D_{n,q}^T D_{n,q}$.

1.1.2 The two-dimensional case

In the two-dimensional case, let us consider a matrix Y of observations and a matrix Ω of non-negative weights, both of dimensions $n_x \times n_z$. The estimator associated with the Whittaker-Henderson smoothing can be written as:

$$\hat{Y} = \underset{\Theta}{\text{argmin}} \{ F(Y, \Omega, \Theta) + R_{\lambda,q}(\Theta) \}$$

where:

- $F(Y, \Omega, \Theta) = \sum_{i=1}^{n_x} \sum_{j=1}^{n_z} \Omega_{i,j} (Y_{i,j} - \Theta_{i,j})^2$ represents a fidelity criterion to the observations,
- $R_{\lambda,q}(\Theta) = \lambda_x \sum_{j=1}^{n_z} \sum_{i=1}^{n_x - q_x} (\Delta^{q_x} \Theta_{\bullet,j})_i^2 + \lambda_z \sum_{i=1}^{n_x} \sum_{j=1}^{n_z - q_z} (\Delta^{q_z} \Theta_{i,\bullet})_j^2$ is a smoothness criterion with $\lambda = (\lambda_x, \lambda_z)$.

This latter criterion can be written as the sum of two one-dimensional regularization criteria, with orders q_x and q_z , weighted by non-negative smoothing parameters λ_x and λ_z , and applied respectively to all rows and all columns of Θ . It is also possible to adopt matrix notations by defining $\mathbf{y} = \text{vec}(Y)$, $\mathbf{w} = \text{vec}(\Omega)$, and $\boldsymbol{\theta} = \text{vec}(\Theta)$ as the vectors obtained by stacking the columns of the matrices Y , Ω , and Θ , respectively. Additionally, let us denote $W = \text{Diag}(\mathbf{w})$ and $n = n_x \times n_z$. The fidelity and smoothness criteria can then be rewritten as linear combinations of the vectors \mathbf{y} , \mathbf{w} , and $\boldsymbol{\theta}$:

$$\begin{aligned} F(\mathbf{y}, \mathbf{w}, \boldsymbol{\theta}) &= (\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta}) \\ R_{\lambda,q}(\boldsymbol{\theta}) &= \boldsymbol{\theta}^T (\lambda_x I_{n_z} \otimes D_{n_x, q_x}^T D_{n_x, q_x} + \lambda_z D_{n_z, q_z}^T D_{n_z, q_z} \otimes I_{n_x}) \boldsymbol{\theta}. \end{aligned}$$

and the associated estimator also takes the form of Equation 2 except in this case

$$P_\lambda = \lambda_x I_{n_z} \otimes D_{n_x, q_x}^T D_{n_x, q_x} + \lambda_z D_{n_z, q_z}^T D_{n_z, q_z} \otimes I_{n_x}.$$

1.1.3 An explicit solution

Equation 2 has an explicit solution, given by:

$$\hat{\mathbf{y}} = (W + P_\lambda)^{-1} W \mathbf{y}. \quad (3)$$

Indeed, as a minimum, $\hat{\mathbf{y}}$ satisfies:

$$0 = \left. \frac{\partial}{\partial \boldsymbol{\theta}} \right|_{\hat{\mathbf{y}}} \left\{ (\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta}) + \boldsymbol{\theta}^T P_\lambda \boldsymbol{\theta} \right\} = -2W(\mathbf{y} - \hat{\mathbf{y}}) + 2P_\lambda \hat{\mathbf{y}}.$$

It follows that $(W + P_\lambda)\hat{\mathbf{y}} = W\mathbf{y}$, and if $W + P_\lambda$ is invertible, then $\hat{\mathbf{y}}$ is indeed a solution of Equation 3. When $\lambda \neq 0$, $W + P_\lambda$ is invertible as long as \mathbf{w} has q non-zero elements in the one-dimensional case, and Ω has at least $q_x \times q_z$ non-zero elements distributed over q_x different rows and q_z different columns in the two-dimensional case. Those conditions are always satisfied in practice.

1.2 Plan for the paper

Even after a century, Whittaker-Henderson smoothing remains widely used by actuaries in many countries, notably France and the U.S.A. Other non-parametric smoothing methods have emerged, in particular spline-based methods (Reinsch 1967), which became even more popular after the introduction of P-Splines smoothing (Eilers and Marx 1996). When used on evenly-spaced discrete observations, this last method and WH smoothing are actually very similar. Indeed, they both balance a fidelity criterion with a smoothness criterion, the latter being based on difference matrices. The key difference is that P-spline smoothing uses a low-rank basis of cubic splines while WH smoothing relies on a full-rank basis whose parameters are directly the fitted values of the model (the associated model matrix being the identity

matrix). This elegant simplicity has been for much in the success of WH smoothing and gives it a special status among smoothing methods. The obvious drawback is that its large number of parameters makes it more computationally intensive than low-rank based smoothers. For an overview of alternative smoothers, the interested reader may find Chapter 5 of Wood (2006) interesting.

Whittaker-Henderson smoothing was first introduced as an empirical method and an alternative to weighted average or polynomial regression. While the statistical theory about smoothing methods has had a lot of interesting development, notably within the framework of generalized additive models (Hastie and Tibshirani 1990), this has not fundamentally changed the way actuaries view Whittaker-Henderson smoothing. The aim of this paper is to bridge the gap between this empirical view of Whittaker-Henderson smoothing and the recent statistical theory about non-parametric smoothing methods. To do so, we reframe Whittaker-Henderson within a modern statistical framework, which allows us to tackle 6 operational questions that arise from the use of this method, each of which being covered in a dedicated Section.

1.2.1 How to measure uncertainty in smoothing results?

We propose to measure the uncertainty in smoothing results based on the available observation volume. In a Frequentist approach, the estimator associated with WH smoothing is biased, which prevents the construction of unbiased finite-size sample confidence intervals. However, smoothing may under some conditions be naturally interpreted in a Bayesian framework, allowing the construction of credible intervals. The connection between Whittaker-Henderson smoothing and Bayesian approaches had already been hinted by Whittaker himself who referred to the smoothness criterion as the logarithm of an *antecedent probability* (Whittaker and Robinson 1924). This connection was latter identified by Kimeldorf and Jones (1967) and made more explicit by Taylor (1992).

1.2.2 Which observation and weight vectors to use?

For the Bayesian interpretation of Whittaker-Henderson smoothing to hold, it needs to be applied to a vector \mathbf{y} of independent observations with known variances and a vector of weights \mathbf{w} containing the inverses of those variances. In the framework of duration models and under the assumption of piecewise constant transition intensities, we show that the maximum likelihood estimator of crude rates provides vectors (\mathbf{y}, \mathbf{w}) that asymptotically satisfies those conditions.

1.2.3 How to improve the accuracy of smoothing with limited data volume?

The previous theoretical framework relies on a 2-step procedure where a crude rates estimator is first derived, which proves asymptotically normally distributed, and Whittaker-Henderson smoothing is then applied to this estimator rather than to the original data. This approach

however is not well founded in practical situations where the available data volume is limited. We propose a generalization of Whittaker-Henderson smoothing where this 2-step approach is replaced by the maximization of a penalized log-likelihood involving the aggregated event and exposure counts directly. This maximization is carried out by Newton method. We compare those two approaches using several simulated datasets whose size and characteristics are close to those of real actuarial datasets and show that the normal approximation in the 2-step method sometimes cause significant bias in the results.

1.2.4 How to choose the smoothing parameter(s)?

We then address the crucial choice of the smoothing parameter λ . For many years, it has been left to the hand of the actuary, using graphical validation. Giesecke and Center (1981) sets the smoothing parameter so that the variance in smoothing results correspond to the average variance for a Chi-square statistic but consider the associated degrees of freedom to be $n - q$, neglecting the reduction in effective dimension caused by the smoothness criterion. Brooks et al. (1988) uses the global cross-validation criteria introduced by Wahba (1980). However, as shown in Wood (2011), this criteria may sometimes lead to severe under-smoothing. We propose instead to rely on the maximization of a marginal likelihood function, an approach proposed by Patterson H. D. (1971) and first used for smoothing parameter selection by Anderssen and Bloomfield (1974). It fits nicely in the Bayesian framework previously introduced and possess good finite-size sample properties as shown by Reiss and Todd Ogden (2009). Maximization of the marginal likelihood function relies on the optimization methods of Brent (1973) and John A. Nelder and Mead (1965), employed respectively in the one-dimensional and two-dimensional case. These methods are implemented in the `optimize` and `optim` functions in the `stats` library of the statistical programming language R. In the context of the proposed generalization of smoothing, where solving likelihood equations and selecting the smoothing parameter are performed through nested iterations, we compare the precision and computational time of two strategies called *outer iteration* and *performance iteration*, the latter approach having been introduced by Gu (1992) as a faster alternative to the former. These strategies differ in the order in which the iterations are nested, and we demonstrate that the latter approach indeed significantly reduces computation time at the cost of negligible loss of precision.

1.2.5 How to improve smoothing performance with a large number of data points?

We propose a practical solution to alleviate the computational burden when the number of observations - and thus parameters - is substantial. This situation is common in practice, especially for two-dimensional tables used in the modelling of disability and long-term care risks. To address this issue, an eigendecomposition of the one-dimensional penalization matrices involved in smoothing is used. This decomposition provides a new interpretation of smoothing and opens the door to the use of a rank reduction method to reduce the number of parameters

and consequently the computational time, while incurring minimal loss of precision. This approach is inspired by Wood (2020) (p. 311).

1.2.6 How to extrapolate smoothing results?

Finally, we cover the extrapolation of smoothing results. The fact that results from models such as WH smoothing or P-splines smoothing may, as is the case for parametric models, be naturally extrapolated beyond the range of the initial data is frequently overlooked by actuaries. Indeed, the literature on this issue is sparse and focused on forecasting mortality rates for mortality modelling. Currie, Durban, and Eilers (2004) uses P-splines to fit and forecast mortality rates by treating the extrapolated positions as 0-weight observations. In a valuable contribution, Carballo, Durban, and Lee (2021) shows that while this solution works well in one-dimensional case, in the two-dimensional case adding 0-weight observations alter the results compared to working with the initial observations alone. To remediate this issue, the authors enforce constraints to maintain the values that would have been obtained using the initial observations alone. However, the confidence intervals given in Carballo, Durban, and Lee (2021) for the extrapolation of the smoothing do not account for the innovation error which means it implicitly assumes that, while the underlying law of the observations has some degree of wigglyness in the range of the initial data, it is perfectly smooth outside this range. We provide a refined expression for credible intervals that also accounts for the innovation error.

2 How to measure uncertainty in smoothing results?

The explicit solution given by Equation 3 indicates that $\mathbb{E}(\hat{\mathbf{y}}) = (W + P_\lambda)^{-1}W\mathbb{E}(\mathbf{y}) \neq \mathbb{E}(\mathbf{y})$ when $\lambda \neq 0$. This implies that penalization introduces a smoothing bias, which prevents the construction of a finite-size sample confidence interval centered on $\mathbb{E}(\mathbf{y})$. Therefore, in this Section, we turn to a Bayesian approach where smoothing can furthermore be interpreted more naturally.

2.1 Maximum a posteriori estimate

Let us suppose that $\mathbf{y}|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}, W^-)$ and $\boldsymbol{\theta} \sim \mathcal{N}(0, P_\lambda^-)$. The Bayes formula allows us to express the posterior likelihood $f(\boldsymbol{\theta}|\mathbf{y})$ associated with these choices in the following form:

$$f(\boldsymbol{\theta}|\mathbf{y}) \propto f(\mathbf{y}|\boldsymbol{\theta})f(\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\left[(\mathbf{y} - \boldsymbol{\theta})^T W(\mathbf{y} - \boldsymbol{\theta}) + \boldsymbol{\theta}^T P_\lambda \boldsymbol{\theta}\right]\right).$$

Hence the mode of the posterior distribution, $\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}}[f(\boldsymbol{\theta}|\mathbf{y})]$, also known as the maximum a posteriori (MAP) estimate, coincides with the solution $\hat{\mathbf{y}}$ from Equation 2, whose explicit form is given by Equation 3.

2.2 Posterior distribution of $\boldsymbol{\theta}|\mathbf{y}$

A second-order Taylor expansion of the log-posterior likelihood around $\hat{\mathbf{y}} = \hat{\boldsymbol{\theta}}$ gives us:

$$\ln f(\boldsymbol{\theta}|\mathbf{y}) = \ln f(\hat{\boldsymbol{\theta}}|\mathbf{y}) + \left. \frac{\partial \ln f(\boldsymbol{\theta}|\mathbf{y})}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}^T (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) + \frac{1}{2} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T \left. \frac{\partial^2 \ln f(\boldsymbol{\theta}|\mathbf{y})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \quad (4)$$

where $\left. \frac{\partial \ln f(\boldsymbol{\theta}|\mathbf{y})}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = 0$ and $\left. \frac{\partial^2 \ln f(\boldsymbol{\theta}|\mathbf{y})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = -(W + P_\lambda)$.

As this last derivative no longer depends on $\boldsymbol{\theta}$, higher-order derivatives beyond two are all zero. The Taylor expansion allows for an exact computation of $\ln f(\boldsymbol{\theta}|\mathbf{y})$. By substituting the result back into Equation 4, we obtain:

$$\begin{aligned} f(\boldsymbol{\theta}|\mathbf{y}) &\propto \exp \left[\ln f(\hat{\boldsymbol{\theta}}|\mathbf{y}) - \frac{1}{2} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T (W + P_\lambda) (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \right] \\ &\propto \exp \left[-\frac{1}{2} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T (W + P_\lambda) (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \right] \end{aligned}$$

which can immediately be recognized as the density of the $\mathcal{N}(\hat{\boldsymbol{\theta}}, (W + P_\lambda)^{-1})$ distribution.

2.3 Consequence for Whittaker-Henderson smoothing

The assumption $\boldsymbol{\theta} \sim \mathcal{N}(0, P_\lambda^-)$ corresponds to a simple Bayesian formalization of the use of the smoothness criterion. It reflects an (improper) prior belief of the modeller about the underlying distribution of the observation vector \mathbf{y} .

The use of Whittaker-Henderson smoothing in the Bayesian framework and the construction of credible intervals are conditioned on the validity of the assumption $\mathbf{y}|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}, W^-)$, *i.e.* the components of the observation vector should be independent and have known variances. In addition, the weight vector \mathbf{w} may not be chosen empirically but should contain the inverses of these variances. Under these assumptions, $100(1 - \alpha)\%$ credible intervals for smoothing can be constructed and take the form:

$$\mathbb{E}(\mathbf{y})|\mathbf{y} \in \left[\hat{\mathbf{y}} \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\text{diag} \{ (W + P_\lambda)^{-1} \}} \right] \quad (5)$$

where $\hat{\mathbf{y}} = (W + P_\lambda)^{-1} W \mathbf{y}$ and Φ denotes the cumulative distribution function of the standard normal distribution. According to Marra and Wood (2012), these credible intervals provide satisfactory coverage of the corresponding Frequentist confidence intervals and may therefore be used as practical substitutes.

3 Which observation and weight vectors to use?

Section 2 highlighted the need to apply the Whittaker-Henderson smoothing to independent observation vectors \mathbf{y} and to weight vectors \mathbf{w} corresponding to the inverses of the variances of the components of \mathbf{y} in order to obtain a measure of uncertainty in the results. In this Section, we propose, within the framework of duration models used for constructing experience tables for life insurance risks, vectors \mathbf{y} and \mathbf{w} that satisfy these conditions.

3.1 Duration models framework: one-dimensional case

Consider the observation of m individuals in a longitudinal study subject to the phenomena of left truncation and right censoring. Suppose one wants to estimate a distribution that depends on only one continuous explanatory variable, denoted by x . For example, one may think of a mortality distribution with the explanatory variable of interest x representing age. Such a distribution is fully characterized by one of the following quantities:

- the cumulative distribution function $F(x)$ or its complement, the survival function $S(x) = 1 - F(x)$,
- the associated probability density function $f(x) = -\frac{d}{dx}S(x)$,
- the instantaneous hazard function $\mu(x) = -\frac{d}{dx} \ln S(x)$.

Suppose that the considered distribution depends on a vector of parameters $\boldsymbol{\theta}$ that one wants to estimate using maximum likelihood. The likelihood associated with the observation of the individuals can be written as follows:

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^m \left[\frac{f(x_i + t_i, \boldsymbol{\theta})}{S(x_i, \boldsymbol{\theta})} \right]^{\delta_i} \left[\frac{S(x_i + t_i, \boldsymbol{\theta})}{S(x_i, \boldsymbol{\theta})} \right]^{1-\delta_i} \quad (6)$$

where x_i represents the age at the start of observation, t_i represents the observation duration, i.e., the time elapsed between the starting and ending dates of observation, and δ_i is the indicator of event observation, which takes the value 1 if the event of interest is observed and 0 if the observation is censored. We will not go into the details of how these three quantities are derived, however they should take into account individual-specific information such as the subscription date, lapse date if applicable, as well as the global characteristics of the product such as the presence of a waiting period or medical selection phenomenon, and the choice of a restricted observation period due to delays in the reporting of event of interests. These factors typically lead to a narrower observation period than the actual period where individuals are present in the portfolio. The various quantities introduced above are related by the following relationships:

$$S(x) = \exp \left(- \int_{u=0}^x \mu(u) du \right) \quad \text{and} \quad f(x) = \mu(x)S(x).$$

The log-likelihood associated with Equation 6 can be rewritten using only the instantaneous hazard function (also known as force of mortality in the case of the death risk):

$$\ell(\boldsymbol{\theta}) = \sum_{i=1}^m \left[\delta_i \ln \mu(x_i + t_i, \boldsymbol{\theta}) - \int_{u=0}^{t_i} \mu(x_i + u, \boldsymbol{\theta}) du \right] \quad (7)$$

While it is possible to base model estimation on direct maximisation of Equation 7, this approach does not scale well with the number of individuals m and requires the approximation of the integral through quadrature, except for some simple parametric models. We instead discretize the problem by assuming that the mortality rate is piecewise constant over one-year intervals between two integer ages or more formally $\mu(x + \epsilon) = \mu(x)$ for all $x \in \mathbb{N}$ and $\epsilon \in [0, 1[$. The advantages of this discretization are underlined for example in Gschlössl, Schoenmaekers, and Denuit (2011) but the idea goes back to at least Hoem (1971). Let us further note that, if $\mathbf{1}$ denotes the indicator function, then for any $x_{\min} \leq a < x_{\max}$, we have $\sum_{x=x_{\min}}^{x_{\max}} \mathbf{1}(x \leq a < x + 1) = 1$, where $x_{\min} = \min(\mathbf{x})$ and $x_{\max} = \max(\mathbf{x})$. Equation 7 may therefore be rewritten as:

$$\ell(\boldsymbol{\theta}) = \sum_{i=1}^m \left[\sum_{x=x_{\min}}^{x_{\max}} \delta_i \mathbf{1}(x \leq x_i + t_i < x + 1) \ln \mu(x_i + t_i, \boldsymbol{\theta}) - \int_{u=0}^{t_i} \sum_{x=x_{\min}}^{x_{\max}} \mathbf{1}(x \leq x_i + u < x + 1) \mu(x_i + u, \boldsymbol{\theta}) du \right].$$

The assumption of piecewise constant mortality rate implies that:

$$\begin{aligned} \mathbf{1}(x \leq x_i + t_i < x + 1) \ln \mu(x_i + t_i, \boldsymbol{\theta}) &= \mathbf{1}(x \leq x_i + t_i < x + 1) \ln \mu(x, \boldsymbol{\theta}) \quad \text{and} \\ \mathbf{1}(x \leq x_i + u < x + 1) \mu(x_i + u, \boldsymbol{\theta}) &= \mathbf{1}(x \leq x_i + u < x + 1) \mu(x, \boldsymbol{\theta}). \end{aligned}$$

It is then possible to interchange the two summations to obtain the following expressions:

$$\begin{aligned} \ell(\boldsymbol{\theta}) &= \sum_{x=x_{\min}}^{x_{\max}} [\ln \mu(x, \boldsymbol{\theta}) d(x) - \mu(x, \boldsymbol{\theta}) e_c(x)] \quad \text{where} \\ d(x) &= \sum_{i=1}^m \delta_i \mathbf{1}(x \leq x_i + t_i < x + 1) \quad \text{and} \\ e_c(x) &= \sum_{i=1}^m \int_{u=0}^{t_i} \mathbf{1}(x \leq x_i + u < x + 1) du = \sum_{i=1}^m [\min(t_i, x - x_i + 1) - \max(0, x - x_i)]^+ \end{aligned}$$

by denoting $a^+ = \max(a, 0)$, where $d(x)$ and $e_c(x)$ correspond to the number of observed deaths between ages x and $x + 1$ and the sum of observation durations of individuals between these ages, respectively (the latter quantity is also known as central exposure to risk).

3.2 Extension to the two-dimensional case

The extension of the proposed approach to the two-dimensional framework requires only minor adjustments to the previous reasoning. Let $z_{\min} = \min(\mathbf{z})$ and $z_{\max} = \max(\mathbf{z})$. The piecewise constant assumption for the mortality rate needs to be extended to the second dimension. Formally, one now assumes that $\mu(x + \epsilon, z + \xi) = \mu(x, z)$ for all pairs $x, z \in \mathbb{N}$ and $\epsilon, \xi \in [0, 1[$. The sums involving the variable x are then replaced by double sums considering all combinations of x and z . The log-likelihood is given by:

$$\begin{aligned} \ell(\boldsymbol{\theta}) &= \sum_{x=x_{\min}}^{x_{\max}} \sum_{z=z_{\min}}^{z_{\max}} [\ln \mu(x, z, \boldsymbol{\theta}) d(x, z) - \mu(x, z, \boldsymbol{\theta}) \mathbf{e}_c(x, z)] \quad \text{where} \\ d(x, z) &= \sum_{i=1}^m \delta_i \mathbf{1}(x \leq x_i + t_i < x + 1) \mathbf{1}(z \leq z_i + t_i < z + 1) \quad \text{and} \\ \mathbf{e}_c(x, z) &= \sum_{i=1}^m \int_{u=0}^{t_i} \mathbf{1}(x \leq x_i + u < x + 1) \mathbf{1}(z \leq z_i + u < z + 1) du \\ &= \sum_{i=1}^m [\min(t_i, x + 1 - x_i, z + 1 - z_i) - \max(0, x - x_i, z - z_i)]^+ \end{aligned}$$

3.3 Likelihood equations

The choice $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbf{exp}(\boldsymbol{\theta})$, which includes one parameter per observation, allows us to relate to the Whittaker-Henderson smoothing. Using the exponential function ensures positive values for the estimated mortality rate. The expressions of the likelihood in the one-dimensional or two-dimensional case can then be written in a common vectorized form:

$$\ell(\boldsymbol{\theta}) = \boldsymbol{\theta}^T \mathbf{d} - \mathbf{exp}(\boldsymbol{\theta})^T \mathbf{e}_c \quad (8)$$

where \mathbf{d} and \mathbf{e}_c represent the vectors of observed deaths and exposures to risk, respectively. The derivatives of the likelihood function for this model are given by:

$$\frac{\partial \ell}{\partial \boldsymbol{\theta}} = [\mathbf{d} - \mathbf{exp}(\boldsymbol{\theta}) \odot \mathbf{e}_c] \quad \text{and} \quad \frac{\partial^2 \ell}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} = -\text{Diag}(\mathbf{exp}(\boldsymbol{\theta}) \odot \mathbf{e}_c). \quad (9)$$

Note that these likelihood equations are exactly what we would obtain by assuming that the observed numbers of deaths, conditional on the observed exposures to risk \mathbf{e}_c , follow Poisson distributions with parameters $\boldsymbol{\mu}(\boldsymbol{\theta}) \odot \mathbf{e}_c$. The model presented here has many similarities with a Poisson GLM (John Ashworth Nelder and Wedderburn 1972) although the initial assumptions are not the same for both models.

These likelihood equations have an explicit solution given by $\hat{\boldsymbol{\theta}} = \ln(\mathbf{d}/\mathbf{e}_c)$. This model, which treats each age independently, is known as the crude rates estimator. The properties of the

maximum likelihood estimator imply that asymptotically $\hat{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\theta}, W_{\hat{\boldsymbol{\theta}}}^{-1})$, where $W_{\hat{\boldsymbol{\theta}}}$ is a diagonal matrix with elements $\mathbf{exp}(\hat{\boldsymbol{\theta}}) \odot \mathbf{e}_c = (\mathbf{d}/\mathbf{e}_c) \odot \mathbf{e}_c = \mathbf{d}$. It should be noted that the asymptotic nature and hence the validity of this approximation are conditioned by the number of individuals m in the portfolio and not the size n of the aggregated vectors \mathbf{d} and \mathbf{e}_c .

Thus, we have shown that in the framework of duration models, using the crude rates estimator, asymptotically $\ln(\mathbf{d}/\mathbf{e}_c) \sim \mathcal{N}(\ln \boldsymbol{\mu}, W^{-1})$, where $W = \text{Diag}(\mathbf{d})$. This justifies applying the Whittaker-Henderson smoothing to the observation vector $\mathbf{y} = \ln(\mathbf{d}/\mathbf{e}_c)$ and weight vector $\mathbf{w} = \mathbf{d}$. According to the results from Section 2, we obtain the following associated credible intervals:

$$\ln \boldsymbol{\mu} | \mathbf{d}, \mathbf{e}_c \in \left[\hat{\boldsymbol{\theta}} \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\mathbf{diag} \{ (\text{Diag}(\mathbf{d}) + P_{\lambda})^{-1} \}} \right]$$

where $\hat{\boldsymbol{\theta}} = (\text{Diag}(\mathbf{d}) + P_{\lambda})^{-1} \text{Diag}(\mathbf{d}) [\ln(\mathbf{d}) - \ln \mathbf{e}_c]$. Confidence intervals for $\boldsymbol{\mu}$ may then be obtained directly by exponentiating the above expression.

4 How to improve the accuracy of smoothing with limited data volume?

4.1 Generalized Whittaker-Henderson smoothing

The approach described in Section 3.3 relies on the asymptotic properties of the maximum likelihood estimator, which provide a theoretical framework for applying smoothing using the crude rates estimator. However, the validity of these asymptotic properties in practice, where we have a limited number of observations, is questionable, as will be shown in this Section.

Another approach is to apply the Bayesian reasoning presented in Section 2 directly to the likelihood of Equation 8. Let us assume again that $\boldsymbol{\theta} \sim \mathcal{N}(0, P_{\lambda}^{-})$ and write, using Bayes' theorem:

$$f(\boldsymbol{\theta} | \mathbf{d}, \mathbf{e}_c) \propto f(\mathbf{d}, \mathbf{e}_c | \boldsymbol{\theta}) f(\boldsymbol{\theta}) \propto \exp \left[\ell(\boldsymbol{\theta}) - \frac{1}{2} \boldsymbol{\theta}^T P_{\lambda} \boldsymbol{\theta} \right].$$

The quantity $\ell_P(\boldsymbol{\theta}) = \ell(\boldsymbol{\theta}) - \frac{1}{2} \boldsymbol{\theta}^T P_{\lambda} \boldsymbol{\theta}$ will be referred to as the penalized likelihood. The maximum a posteriori $\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\text{argmax}} f(\boldsymbol{\theta} | \mathbf{d}, \mathbf{e}_c)$ corresponds, once again, to the maximum of the penalized likelihood $\tilde{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\text{argmax}} \ell_P(\boldsymbol{\theta})$.

A second-order Taylor expansion of the posterior log-likelihood around $\hat{\boldsymbol{\theta}}$ yields:

$$\ln f(\boldsymbol{\theta} | \mathbf{d}, \mathbf{e}_c) \simeq \ln f(\hat{\boldsymbol{\theta}} | \mathbf{d}, \mathbf{e}_c) + \left. \frac{\partial \ln f(\boldsymbol{\theta} | \mathbf{d}, \mathbf{e}_c)}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}^T (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) + \frac{1}{2} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T \left. \frac{\partial^2 \ln f(\boldsymbol{\theta} | \mathbf{d}, \mathbf{e}_c)}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \quad (10)$$

$$\text{with } \left. \frac{\partial \ln f(\boldsymbol{\theta}|\mathbf{d}, \mathbf{e}_c)}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = 0 \quad \text{and} \quad \left. \frac{\partial^2 \ln f(\boldsymbol{\theta}|\mathbf{d}, \mathbf{e}_c)}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = -(W_{\hat{\boldsymbol{\theta}}} + P_\lambda).$$

where $W_{\hat{\boldsymbol{\theta}}} = \text{Diag}(\mathbf{exp}(\hat{\boldsymbol{\theta}}) \odot \mathbf{e}_c)$.

Unlike the normal case studied in Section 2, the higher-order derivatives of the posterior log-likelihood are not zero, and the Equation 10 represents an approximation of the posterior log-likelihood known as the Laplace approximation, which becomes more accurate as the number of observations increases. Asymptotically:

$$\begin{aligned} f(\boldsymbol{\theta}|\mathbf{d}, \mathbf{e}_c) &\propto \exp \left[\ell_P(\hat{\boldsymbol{\theta}}) - \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T (W_{\hat{\boldsymbol{\theta}}} + P_\lambda)(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \right] \\ &\propto \exp \left[-\frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T (W_{\hat{\boldsymbol{\theta}}} + P_\lambda)(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \right] \end{aligned}$$

and the posterior distribution is thus equivalent to $\mathcal{N}(\hat{\boldsymbol{\theta}}, (W_{\hat{\boldsymbol{\theta}}} + P_\lambda)^{-1})$. Since $\boldsymbol{\theta} = \ln \boldsymbol{\mu}$, this result allows for the construction of $100(1 - \alpha)\%$ asymptotic credible intervals for $\ln \boldsymbol{\mu}|\mathbf{d}, \mathbf{e}_c$, of the form:

$$\ln \boldsymbol{\mu}|\mathbf{d}, \mathbf{e}_c \in \left[\hat{\boldsymbol{\theta}} \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\mathbf{diag} \{ (W_{\hat{\boldsymbol{\theta}}} + P_\lambda)^{-1} \}} \right].$$

Unlike Equation 9, it is not possible here to explicitly determine the maximum of the penalized likelihood $\hat{\boldsymbol{\theta}}$. However, the Newton algorithm allows for a numerical solution of the likelihood equations by constructing a sequence of estimators $(\boldsymbol{\theta}_k)_{k \geq 0}$ that converges to $\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\text{argmax}} \ell_P(\boldsymbol{\theta})$.

These estimators are recursively defined as:

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \boldsymbol{\theta}_k - \left(\left. \frac{\partial^2 \ell_P}{\partial \boldsymbol{\theta}^2} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_k} \right)^{-1} \left. \frac{\partial \ell_P}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_k} \\ &= \boldsymbol{\theta}_k + (W_k + P_\lambda)^{-1} (\mathbf{d} - \mathbf{exp}(\boldsymbol{\theta}_k) \odot \mathbf{e}_c - P_\lambda \boldsymbol{\theta}_k) \\ &= (W_k + P_\lambda)^{-1} W_k \mathbf{z}_k \end{aligned}$$

by denoting $W_k = \text{Diag}(\mathbf{exp}(\boldsymbol{\theta}_k) \odot \mathbf{e}_c)$ and $\mathbf{z}_k = \boldsymbol{\theta}_k + W_k^{-1}[\mathbf{d} - \mathbf{exp}(\boldsymbol{\theta}_k) \odot \mathbf{e}_c]$. An interesting initialization choice for the algorithm starts with the crude rates estimator, setting $\boldsymbol{\theta}_0 = \ln(\mathbf{d}/\mathbf{e}_c)$, from which we derive $W_0 = \text{Diag}(\mathbf{d})$ and $z_0 = \ln(\mathbf{d}/\mathbf{e}_c)$.

With this choice, the result of the first iteration $\boldsymbol{\theta}_1 = (W_0 + P_\lambda)^{-1} W_0 z_0$ coincides with the solution of the original WH smoothing, while subsequent iterations may be seen as successive applications of the original smoothing to pseudo-observation vectors \mathbf{z}_k and associated weight vectors $w_k = \mathbf{diag}(W_k)$, adjusted at each iteration. The proposed approach can thus be considered as an iterative variant of the original Whittaker-Henderson smoothing, where the choice of observation and weight vectors is refined at each step. The connection between

the approach presented in this Section and the original WH smoothing is very similar to the one between linear model and *generalized linear models*. Therefore, we refer to this approach as *generalized Whittaker-Henderson smoothing* through the rest of the paper. Algorithm 1 provides an implementation of this generalization of the initial WH smoothing.

4.2 Impact of the normal approximation made by the original Whittaker-Henderson smoothing

We have shown that the original Whittaker-Henderson smoothing can be interpreted as an approximation of a penalized likelihood maximization problem, based on a crude rate estimator which has an asymptotically normal distribution. To assess the impact of this approximation in practice, we rely on 6 simulated datasets that capture the characteristics of real biometric risks :

- The first 3 datasets mimic the behaviour of three annuity portfolios with 20,000, 100,000 and 500,000 policyholders respectively. The only covariate considered in this setting is age.
- The subsequent 3 datasets mimic the behaviour of three annuity portfolios of policyholders in a long-term care situation, with 1,000, 5,000, and 25,000 dependent policyholders respectively. Modelling of long-term care typically relies on a multi-state model with 3 states (*active*, *disabled* and *dead*) and in this case we only focus on the transition between the *disabled* and *dead* states. We consider both *age* and *time already spent in the long-term care situation* as relevant covariates in this setting.

Those 6 datasets initially consist of individual longitudinal data from which we compute observed event counts \mathbf{d} as well as central exposure to risk \mathbf{e}_c for each combination of discretized explanatory variables of interest: age x for the first 3 portfolios and combinations (x, z) of age and long-term care duration for the subsequent portfolio. Further information about how those datasets were generated is provided in the appendices. We then apply Whittaker-Henderson smoothing to the observation vectors $\mathbf{y} = \ln(\mathbf{d}/\mathbf{e}_c)$ and weight vectors $\mathbf{w} = \mathbf{d}$ on one hand, or directly to the vectors \mathbf{d} and \mathbf{e}_c within the framework of *generalized WH smoothing* proposed in Section 4 on the other hand. The parameter λ used is determined for each of the 6 portfolios by using a method that will be presented in Section 5. For now, let us just point out that for each portfolio, the same parameter λ , and therefore the same prior belief on the vector $\boldsymbol{\theta} = \ln \boldsymbol{\mu}$, is used for both approaches.

The two approaches estimate the same quantity $\boldsymbol{\theta} = \ln \boldsymbol{\mu}$ hence they may be directly compared through their respective estimates $\hat{\boldsymbol{\theta}}_{\text{norm}}$ and $\hat{\boldsymbol{\theta}}_{\text{ML}}$. Since the first estimator can be seen as an approximate version of the second, we propose to use as error measure:

$$\Delta(\boldsymbol{\theta}) = \frac{\ell_P(\hat{\boldsymbol{\theta}}_{\text{ML}}) - \ell_P(\boldsymbol{\theta})}{\ell_P(\hat{\boldsymbol{\theta}}_{\text{ML}}) - \ell_P(\hat{\boldsymbol{\theta}}_{\infty})} \quad (11)$$

Table 1: Impact of the normal approximation associated with Whittaker smoothing on a selection of 6 fictional insurance portfolios.

Portfolio	Head count	Relative error on penalized deviance
One-dimensional	20 000	7,51%
One-dimensional	100 000	0,23%
One-dimensional	500 000	0,09%
Two-dimensional	1 000	1 784,61%
Two-dimensional	5 000	122,89%
Two-dimensional	25 000	6,95%

where $\hat{\theta}_\infty$ corresponds to the parameter vector that maximizes the penalized likelihood with an infinite penalty. It will become apparent in Section 6.2 that this choice corresponds to the polynomial function of degree $q - 1$ that maximizes the likelihood. By definition, $\hat{\theta}_{\text{ML}} = \underset{\theta}{\operatorname{argmax}} \ell_P(\theta)$, and thus $\Delta(\theta) \geq 0$ for any vector θ . Furthermore, $\Delta(\hat{\theta}_{\text{ML}}) = 0$ and $\Delta(\hat{\theta}_\infty) = 1$. A model for which $\Delta(\theta) > 1$ can be considered as having no practical interest as, according to the prior belief, it is less probable than a simple polynomial fit. The indicator Δ constructed in this way allows interpreting the quality of the used approximation within certain limits. Going forward, we make the arbitrary choice that a difference of less than 1% is considered negligible, while a difference greater than 10% is considered prohibitive. Differences between 1% and 10% will be interpreted in a more nuanced way.

The discrepancies between the two approaches obtained for each of the 6 portfolios studied are presented in Table 1. These discrepancies naturally decrease with the size of the portfolio, as the validity of the normal approximation increases. In the one-dimensional case, the discrepancy is modest for the two largest portfolios but significant for the smallest one. In the two-dimensional case, the discrepancy is significant for the largest portfolio and prohibitive for the two smallest ones. This is because using the observed death vector as the weight vector introduces a bias by overweighting observations with the highest crude rate and underweighting those with the lowest crude rate. In light of these results, it seems appropriate to favour generalized *Whittaker-Henderson smoothing*, especially since its implementation does not pose significant practical difficulties.

5 How to choose the smoothing parameter(s)?

5.1 Impact of the smoothing parameter(s)

This Section is dedicated to the choice of the smoothing parameter λ in the one-dimensional case or the pair $\lambda = (\lambda_x, \lambda_z)$ in the two-dimensional case. Those smoothing parameters control the relative importance of the fidelity and smoothness criterion in Equation 1. Figure 1 provides an illustration of one-dimensional smoothing on the annuity dataset based on 100,000 policyholders

introduced in Section 4.2, using three different values of the smoothing parameter. The effective degrees of freedom shown in this figure is computed by summing the diagonal values of H , the *hat matrix* of the model. The hat matrix satisfies $\hat{\mathbf{y}} = H\mathbf{y}$ and can be identified in Equation 3 as $H = (W + P_\lambda)^{-1}W$. These effective degrees of freedom serve as a non-parametric equivalent of the number of independent parameters in parametric models but can take non-integer values. The concept of degrees of freedom will be further discussed in Section 6 of the paper. It can be observed here that the result of smoothing is highly sensitive to the chosen value of the smoothing parameter. The choice of $\lambda = 10^1$ leads to a highly volatile fit that largely reproduces the sampling fluctuations present in the data. On the other hand, the choice of $\lambda = 10^7$ seems too rigid to capture the underlying pattern of the data. The choice of $\lambda = 10^4$ appears, at first glance, as a satisfactory compromise between those two extremes.




Figure 1: Whittaker-Henderson smoothing applied to a synthetic portfolio of 100,000 policyholders for 3 choices of the smoothing parameter. The points represent the initial observations, the curves depict the result of the smoothing, and the shaded areas represent the associated credibility intervals. edf refers to the effective degrees of freedom of the model.

Figure 2 illustrates the application of smoothing to the simulated annuity dataset with 5,000 policyholders in situation of long-term care also presented in Section 4.2. Different combinations of the two smoothing parameters in this model are used to produce the 9 plots represented in this figure. While some combinations of smoothing parameters may be deemed implausible, such as those associated with the choice $\lambda_x = 10^0$, it is challenging in the two-dimensional case to determine graphically which combination of parameters is most suitable. This is why the selection of the smoothing parameter should rather be based on the optimization of a statistical criterion.

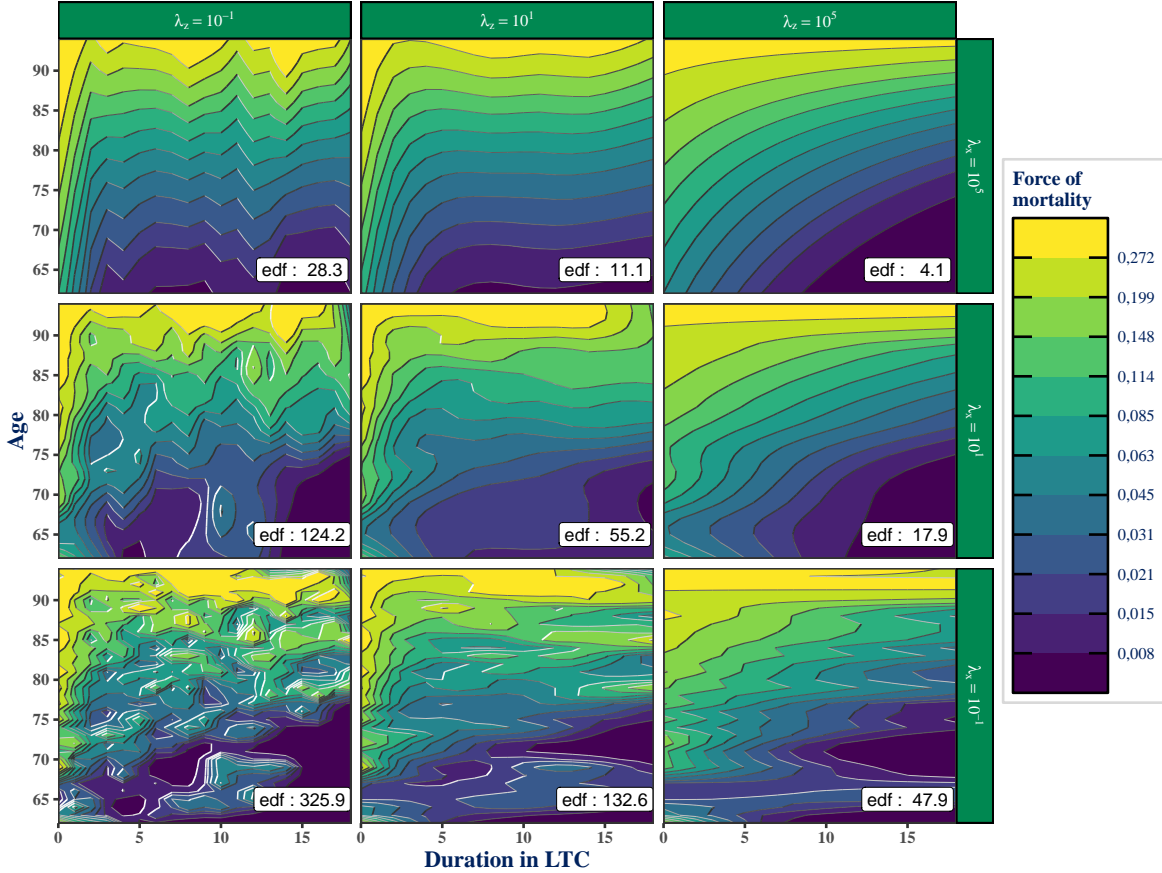


Figure 2: Whittaker-Henderson smoothing applied to a synthetic portfolio of 5,000 policyholders in long-term care situation for 9 combinations of smoothing parameters. The contour lines and colors depict the surface of the smoothed force of mortality as a function of the age of the policyholders and the duration in the long-term care state.

5.2 Statistical criteria for parameter selection

Statistical criteria used for the choice of smoothing parameters typically belong to one of two major families. On one hand, there are criteria based on the minimization of the prediction error for the model, among which the Akaike information criterion (AIC, Akaike 1973) and the generalized cross-validation (GCV, Wahba 1980). On the other hand, it is possible to rely on the maximisation of a marginal likelihood function. It was first introduced by Patterson H. D. (1971) in the Gaussian case, initially under the name of restricted likelihood (REML), and applied by Anderssen and Bloomfield (1974) for the selection of smoothing parameters. Wahba (1985) and Kauermann (2005) show that criteria minimizing prediction error have the best asymptotic performance, but their convergence to the optimal smoothing parameters is slower. For finite-size samples, criteria based on the maximization of a likelihood function prove a more

robust choice, as shown in Reiss and Todd Ogden (2009) or Wood (2011).

Figure 3 represents the values of GCV and marginal likelihood as a function of the smoothing parameter, based on the synthetic datasets used to produce Figure 1. While the marginal likelihood exhibits a clear maximum, GCV has two minima, one of which coincides with the marginal likelihood maximum. The second and global minimum results in a model with nearly 45 degrees of freedom, which clearly is not a plausible choice for the underlying mortality curve. For these reasons, we favour the marginal likelihood as the selection criterion, which also naturally fits within the Bayesian framework introduced in Sections 2 and 4. We first discuss smoothing parameter selection in the case of the original Whittaker-Henderson smoothing and then turn to the generalized case introduced in Section 4, introducing two competing approaches that are subsequently compared.




Figure 3: Comparison in the context of one-dimensional Whittaker-Henderson smoothing of two criteria for selecting the smoothing parameter: the Generalized Cross-Validation (GCV) criterion and the marginal likelihood

5.3 Selection in the context of the original smoothing

Let us start from the notations and assumptions from Section 2, namely $\mathbf{y}|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}, W^-)$ and $\boldsymbol{\theta}|\lambda \sim \mathcal{N}(0, P_\lambda^-)$. In a purely Bayesian approach, it would be necessary to define a prior distribution on λ and then estimate the posterior distribution of each parameter vector using methods such as Markov Chain Monte Carlo (see for example Gilks, Richardson, and Spiegelhalter 1995). The empirical Bayesian approach we adopt rather seeks to find the value of λ that maximizes the marginal likelihood:

$$\mathcal{L}_{\text{norm}}^m(\lambda) = f(\mathbf{y}|\lambda) = \int f(\mathbf{y}, \boldsymbol{\theta}|\lambda) d\boldsymbol{\theta} = \int f(\mathbf{y}|\boldsymbol{\theta}) f(\boldsymbol{\theta}|\lambda) d\boldsymbol{\theta}.$$

This corresponds to the maximum likelihood method applied to the smoothing parameter. Let us explicitly rewrite the expressions of $f(\mathbf{y}|\boldsymbol{\theta})$ and $f(\boldsymbol{\theta}|\lambda)$ introduced in Section 2:

$$\begin{aligned} f(\mathbf{y}|\boldsymbol{\theta}) &= \sqrt{\frac{|W|_+}{(2\pi)^{n_*}}} \exp\left(-\frac{1}{2}(\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta})\right) \\ f(\boldsymbol{\theta}|\lambda) &= \sqrt{\frac{|P_\lambda|_+}{(2\pi)^{p-q}}} \exp\left(-\frac{1}{2}\boldsymbol{\theta}^T P_\lambda \boldsymbol{\theta}\right) \end{aligned}$$

where $|A|_+$ denotes the product of the non-zero eigenvalues of A , n_* is the number of non-zero diagonal elements of W , and q is the number of zero eigenvalues of P_λ ($q = q_x \times q_z$ in the two-dimensional case). Based on the Taylor expansion used in Section 2, let us recall that:

$$\ln f(\mathbf{y}, \boldsymbol{\theta}|\lambda) = \ln f(\mathbf{y}, \hat{\boldsymbol{\theta}}_\lambda|\lambda) + \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)^T (W + P_\lambda) (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)$$

which leads to:

$$\begin{aligned} \mathcal{L}_{\text{norm}}^m(\lambda) &= \int \exp[\ln f(\mathbf{y}, \boldsymbol{\theta}|\lambda)] d\boldsymbol{\theta} \\ &= f(\mathbf{y}, \hat{\boldsymbol{\theta}}_\lambda|\lambda) \int \exp\left[-\frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)^T (W + P_\lambda) (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)\right] d\boldsymbol{\theta} \\ &= f_{\mathbf{y}}(\mathbf{y}|\hat{\boldsymbol{\theta}}_\lambda) f_{\boldsymbol{\theta}}(\hat{\boldsymbol{\theta}}_\lambda|\lambda) \sqrt{\frac{(2\pi)^p}{|W + P_\lambda|}} \\ &= \sqrt{\frac{|W|_+ |P_\lambda|_+}{(2\pi)^{n_*-q} |W + P_\lambda|}} \exp\left(-\frac{1}{2} \left[(\mathbf{y} - \hat{\boldsymbol{\theta}}_\lambda)^T W (\mathbf{y} - \hat{\boldsymbol{\theta}}_\lambda) + \hat{\boldsymbol{\theta}}_\lambda^T P_\lambda \hat{\boldsymbol{\theta}}_\lambda \right]\right). \end{aligned}$$

The associated log-likelihood can be expressed as follows:

$$\begin{aligned} \ell_{\text{norm}}^m(\lambda) &= -\frac{1}{2} \left[(\mathbf{y} - \hat{\boldsymbol{\theta}}_\lambda)^T W (\mathbf{y} - \hat{\boldsymbol{\theta}}_\lambda) + \hat{\boldsymbol{\theta}}_\lambda^T P_\lambda \hat{\boldsymbol{\theta}}_\lambda \right. \\ &\quad \left. - \ln |W|_+ - \ln |P_\lambda|_+ + \ln |W + P_\lambda| + (n_* - q) \ln(2\pi) \right]. \end{aligned} \tag{12}$$

The lack of an explicit solution to Equation 12 forces us to resort to numerical methods to find $\hat{\lambda}_{\text{norm}}$ that maximizes ℓ_{norm}^m . The Newton algorithm could once again be employed here and is a robust choice. This approach was notably adopted by Wood (2011). However, explicitly calculating the derivatives of the likelihood ℓ_{norm}^m is rather difficult from an operational perspective. Instead, we rely on the general heuristics provided by Brent (1973) and John A. Nelder and Mead (1965), which are applicable to any sufficiently regular function. They do not require prior computation of derivatives and are implemented in the `optimize` and `optim` functions of the statistical programming language R.

Computing the marginal likelihood using Equation 12 requires the estimation of various intermediate quantities, the most demanding being the estimation of $\hat{\boldsymbol{\theta}}_\lambda$. This estimation is based on Equation 3 and ideally involves the inversion of the symmetric matrix $W + P_\lambda$ through a

Cholesky decomposition. The calculation of $\ln |W + P_\lambda|$ is immediate for the triangular matrices resulting from this decomposition. Note that since the matrix P_λ is a linear combination of the matrix $D_{n,q}^T D_{n,q}$ in the one-dimensional case, and the matrices $I_{n_z} \otimes D_{n_x,q_x}^T D_{n_x,q_x}$ and $D_{n_z,q_z}^T D_{n_z,q_z} \otimes I_{n_x}$ in the two-dimensional case, it is only necessary to form these matrices once. Furthermore, the calculation of $\ln |P_\lambda|_+$ can be efficiently performed by computing once the eigendecomposition of the matrix $D_{n,q}^T D_{n,q}$ (or the matrices $D_{n_x,q_x}^T D_{n_x,q_x}$ and $D_{n_z,q_z}^T D_{n_z,q_z}$ in the two-dimensional case). This calculation then only requires multiplying these eigenvalues by the corresponding element of λ and taking the logarithm. Finally, since the terms $\ln |W|_+$ and $(n_* - q) \ln(2\pi)$ do not depend on λ , they can be ignored altogether.

5.4 Selection in the generalized smoothing framework: *outer iteration* approach

It is possible to extend the previous approach to the framework of penalized likelihood introduced in Section 4. The Taylor expansion used in Section 5.3 can be applied in this context and yields:

$$\begin{aligned} \mathcal{L}_{\text{ML}}^m(\lambda) &= \int \exp[\ln f(\mathbf{d}, \mathbf{e}_c, \boldsymbol{\theta}|\lambda)] d\boldsymbol{\theta} \\ &\simeq f(\mathbf{d}, \mathbf{e}_c, \hat{\boldsymbol{\theta}}_\lambda|\lambda) \int \exp\left(-\frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)^T (W_{\hat{\boldsymbol{\theta}}_\lambda} + P_\lambda)(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\lambda)\right) d\boldsymbol{\theta} \\ &\simeq l(\hat{\boldsymbol{\theta}}_\lambda) f_{\boldsymbol{\theta}}(\hat{\boldsymbol{\theta}}_\lambda|\lambda) \sqrt{\frac{(2\pi)^p}{|W_{\hat{\boldsymbol{\theta}}_\lambda} + P_\lambda|}} \\ &\simeq l(\hat{\boldsymbol{\theta}}_\lambda) \sqrt{\frac{(2\pi)^q |P_\lambda|_+}{|W_{\hat{\boldsymbol{\theta}}_\lambda} + P_\lambda|}} \exp\left(-\frac{1}{2} \hat{\boldsymbol{\theta}}_\lambda^T P_\lambda \hat{\boldsymbol{\theta}}_\lambda\right). \end{aligned}$$

This gives us the expression for the likelihood in this framework:

$$\ell_{\text{ML}}^m(\lambda) \approx l(\hat{\boldsymbol{\theta}}_\lambda) - \frac{1}{2} \left[\hat{\boldsymbol{\theta}}_\lambda^T P_\lambda \hat{\boldsymbol{\theta}}_\lambda - \ln |P_\lambda|_+ + \ln |W_{\hat{\boldsymbol{\theta}}_\lambda} + P_\lambda| - q \ln(2\pi) \right]. \quad (13)$$

Unlike in the normal case, the Taylor expansion does not allow for an exact calculation of the marginal likelihood but provides an approximation known as the Laplace approximation, whose validity depends on the number of available observations. Equation 13 does not have an explicit solution and requires numerical resolution. The calculation of the Laplace approximation of the marginal likelihood is done in a similar manner to Equation 12 from Section 5.3, with the only notable difference being that the vector $\hat{\boldsymbol{\theta}}_\lambda$ must be estimated iteratively using Algorithm 1. This leads to Algorithm 2 in the appendices for selecting the smoothing parameter which includes two nested iterative calculations: the calculation of $\hat{\boldsymbol{\theta}}_\lambda$, for a fixed λ , using the Newton algorithm, and the search for the maximum of ℓ_{ML}^m using heuristics such as the Brent and Nelder-Mead algorithms in the one-dimensional and two-dimensional case, respectively. Since the selection of λ corresponds to the outer iterative loop here, and to distinguish it from the approach introduced in the next Section, we will refer to it as the *outer iteration* approach.

5.5 Selection in the generalized smoothing framework: *performance iteration* approach

An alternative to the approach developed in Section 5.4 is to start from Algorithm 1 and notice that at each step of the algorithm, the equations solved coincide with those obtained by assuming that $\mathbf{z}_k|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}, W_k^{-1})$ and $\boldsymbol{\theta} \sim \mathcal{N}(0, P_\lambda^-)$. Therefore, it is possible to rely on the methodology proposed in Section 5.3 and numerically estimate, at the beginning of each iteration of Algorithm 1, the smoothing parameter that maximizes the marginal likelihood ℓ_{norm}^m in which \mathbf{y} would be replaced by \mathbf{z}_k and W by W_k . This is just what Algorithm 3 in the appendix proposes. A justification for this approach is given by (Wood 2006, p149). From a practical point of view, this approach conceptually reverses the nesting order of the two types of iterations and may, in some cases, provide a considerable reduction in computation time. It was hence introduced by Gu (1992) as the *performance (oriented) iteration* approach. However, unlike the *outer iteration* approach, the convergence of the *performance iteration* approach cannot be guaranteed. Indeed, the penalized likelihood calculated in Algorithm 3 is not directly comparable between iterations because it is based on a different smoothing parameter and, therefore, a different prior. There are situations in which this algorithm might not converge, although for the optimization problem addressed in this paper convergence was always achieved.

5.6 Comparison of *outer iteration* and *performance iteration* approaches

In Sections 5.4 and 5.5, we introduced two alternatives for the selection of parameters in generalized Whittaker-Henderson smoothing. This Section aims to study the consequences of choosing one approach over the other in terms of both result accuracy and computational time. While the objective of Section 4.2 was to highlight a bias in the asymptotic approximation associated with the original Whittaker-Henderson smoothing for small portfolios, here we aim to measure these impacts more precisely. We limit our analysis to intermediate-sized portfolios as presented in Section 4.2 (i.e., 100,000 rows for the one-dimensional case and 5,000 rows for the two-dimensional case), however we use 100 replicates of each portfolio, generated from the same mortality and censoring laws, to increase the robustness of our analysis.

The *outer iteration* and *performance iteration* approaches can be directly compared based on the selected parameter λ and using the marginal likelihood $\ell_{\text{ML}}^m(\lambda)$. We define the error criterion in a similar way to Section 4.2, :

$$\Delta(\lambda) = \frac{\ell_{\text{ML}}^m(\hat{\lambda}_{\text{outer}}) - \ell_{\text{ML}}^m(\lambda)}{\ell_{\text{ML}}^m(\hat{\lambda}_{\text{outer}}) - \ell_{\text{ML}}^m(\infty)} \quad (14)$$

where $\ell_{\text{ML}}^m(\infty)$ corresponds to the approximation of the marginal likelihood associated with the choice of an infinite smoothing parameter. Again, the defined Δ satisfies the properties $\Delta(\lambda_{\text{outer}}) = 0$ and $\Delta(\infty) = 1$. This criterion is interpreted in the same way as the one given by Equation 11.

Figure 4 represents the empirical distribution of the difference between the *outer iteration* and *performance iteration* approaches, and Figure 5 represents the distribution of the associated computation time for each approach. The use of the *performance iteration* approach results in a difference below 1% for the 100 simulated datasets, both in the one-dimensional and two-dimensional case. The associated computation time is reduced by a factor of 2 in the one-dimensional case and nearly 4 in the two-dimensional case. Considering the significant computation times involved in the two-dimensional case, the use of the *performance iteration* approach therefore presents a real advantage from an operational perspective.




Figure 4: Distribution of the difference between the two presented approaches for selecting the smoothing parameter, for 100 simulated portfolios, in the one-dimensional and two-dimensional cases.

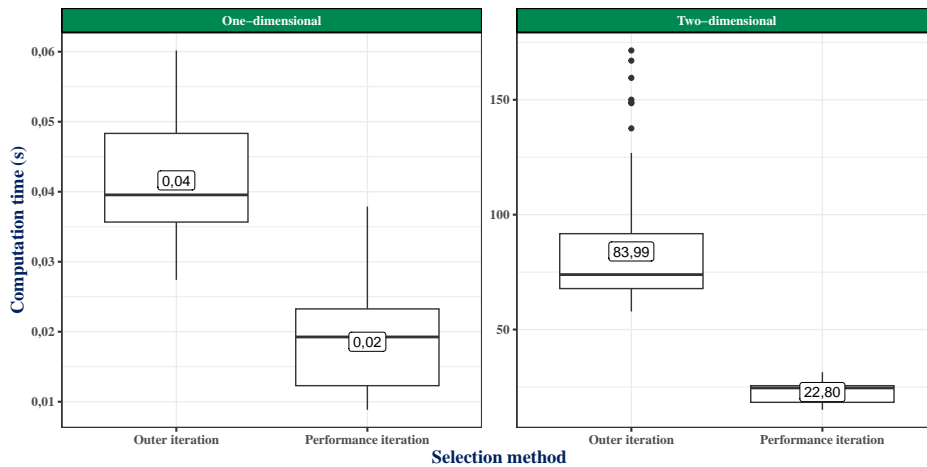


Figure 5: Distribution of the computation time associated with the two presented approaches for selecting the smoothing parameter, for 100 simulated portfolios, in the one-dimensional and two-dimensional cases.

6 How to improve smoothing performances with a large number of data points?

6.1 Motivation

Whittaker-Henderson smoothing is a full-rank smoothing as it contains as many parameters as there are observation points. This characteristic allows it to faithfully reproduce any input signal as long as a sufficient number of observations are supplied. More formally, the estimator associated with smoothing is asymptotically unbiased since:

$$\mathbb{E}(\hat{\mathbf{y}}) = (W + P_\lambda)^{-1}W\mathbb{E}(\mathbf{y}) \xrightarrow{m \rightarrow \infty} \mathbb{E}(\mathbf{y})$$

where m represents the number of observed individuals on which the matrix W depends. The downside is that WH smoothing may become impractical in presence of many observations points n . Indeed, the algorithms presented in Section 5 require the inversion of the $n \times n$ matrix $W + P_\lambda$, an operation with a time complexity of $O(n^3)$, which needs to be repeated at each iteration of the algorithm, regardless of the chosen method.

For biometric risks that depend solely on age, discretized on an annual basis, the number of observations rarely exceeds 100, and computation time is not a significant concern. However, in the two-dimensional case, the number of observations can take much larger values in several practical cases, including:

- For the disability risk in France, it is necessary to construct disability survival tables for entry ages ranging from 18 to 61 and exit ages ranging from the entry age to 62. This represents $[(62 - 18) \times (62 - 18 + 1)]/2 = 990$ observations.
- Also for the disability risk in France, it is necessary to construct transition tables from incapacity to disability for entry ages ranging from 18 to 67 and monthly incapacity durations ranging from 0 to 36 months. This represents $(67 - 18) \times (36 - 0) = 1,764$ observations.
- For the long-term care risk, as the coverage is lifelong, it is necessary to construct tables for all ages and years spent in long-term care encountered in the data. Assuming observation ages ranging from 50 to 110 and durations ranging from 0 to 20 years, this represents $(110 - 50) \times (20 - 0) = 1,200$ observations.

The practical applications described above thus require computation times of at least several minutes for each smoothing application. When such smoothing is used repeatedly, for example when combined with simulations, reducing the computation time becomes a vital goal.

6.2 Smoothing and eigendecomposition

The eigendecomposition of the penalization matrix is key to a better understanding of smoothing and allows for an approximation of the smoothing problem with $p < n$ parameters. Let us consider the one-dimensional case and write the decomposition for the symmetric matrix $D_{n,q}^T D_{n,q}$. It takes the form $D_{n,q}^T D_{n,q} = U \Sigma U^T$, where Σ is a diagonal matrix containing the eigenvalues of $D_{n,q}^T D_{n,q}$ and U is an orthogonal matrix such that $U^T U = U U^T = I_n$. Let us perform the reparameterization $\boldsymbol{\beta} = U^T \boldsymbol{\theta} \Leftrightarrow \boldsymbol{\theta} = U \boldsymbol{\beta}$. The smoothness criterion becomes:

$$\boldsymbol{\theta}^T P_\lambda \boldsymbol{\theta} = (U \boldsymbol{\beta})^T P_\lambda (U \boldsymbol{\beta}) = \lambda \boldsymbol{\beta}^T U^T U \Sigma U^T U \boldsymbol{\beta} = \lambda \boldsymbol{\beta}^T \Sigma \boldsymbol{\beta}.$$

and Equation 2 may be rewritten as:

$$\hat{\mathbf{y}} = U \hat{\boldsymbol{\beta}} \quad \text{where} \quad \hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \left\{ (\mathbf{y} - U \boldsymbol{\beta})^T W (\mathbf{y} - U \boldsymbol{\beta}) + \lambda \boldsymbol{\beta}^T \Sigma \boldsymbol{\beta} \right\}. \quad (15)$$

In the original formulation of smoothing, the parameters directly corresponds to the smoothed values. With the new parameterization of Equation 15, $\boldsymbol{\beta}$ can be interpreted as a vector of coordinates in the basis of eigenvectors of P_λ , providing a decomposition of the signal into components that are more or less smooth according to the difference matrix. Figure 6 represents 8 of the eigenvectors associated with $q = 2$ for a basis of size $n = 74$, which corresponds to the number of observations in Figure 1. The eigenvalues associated with the first q eigenvalues are zero. This can easily be seen by noting that $D_{n,q}$ is a matrix of dimensions $(n - q) \times n$ and of rank $n - q$.

By using the fact that $U^{-1} = U^T$ and making the connection with Equation 3, we obtain the explicit solution:

$$\hat{\mathbf{y}} = U (U^T W U + S_\lambda)^{-1} U^T W \mathbf{y} \quad \text{where} \quad S_\lambda = \lambda \Sigma \quad (16)$$

In order to interpret Equation 16, let us consider the special case where all weights are equal to 1 and therefore:

$$\hat{\mathbf{y}} = U (U^T U + \lambda \Sigma)^{-1} U^T \mathbf{y} = U (I_n + \lambda \Sigma)^{-1} U^T \mathbf{y}$$

The transformation from \mathbf{y} to $\hat{\mathbf{y}}$ can then be seen as a 3-step process, reading the equation from right to left:

1. Decomposition of the signal \mathbf{y} in the basis of eigenvectors through the left multiplication by U^T .
2. Attenuation of the signal components based on the eigenvalues associated with these components. If we denote $s = \operatorname{diag}(\Sigma)$, then $(I_n + \lambda \Sigma)^{-1} = \operatorname{Diag}[1/(1 + \lambda s)]$. After the left multiplication by $(I_n + \lambda \Sigma)^{-1}$, each component is hence divided by a factor




Figure 6: Graphical representation of a subset of columns from a basis of eigenvectors associated with $n = 74$ and $q = 2$ in the one-dimensional case. The 8 vectors in the subset are displayed in order of increasing penalization.

$1 + \lambda s \geq 1$. This coefficient increases linearly with λ , but at different rates depending of the associated eigenvalue.

3. Recomposition of the attenuated signal in the canonical basis through the left multiplication by U .

In the presence of non-unit weights, things are not as straightforward since $U^T W U$ is no longer a diagonal matrix. However, it is still possible to interpret the effect of smoothing thanks to the matrix $F = (U^T W U + S_\lambda)^{-1} U^T W U$. Indeed:

$$U^T \hat{\mathbf{y}} = U^T U \hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}} = (U^T W U + S_\lambda)^{-1} U^T W \mathbf{y} = (U^T W U + S_\lambda)^{-1} U^T W U U^T \mathbf{y} = F U^T \mathbf{y}$$

Since the vectors $U^T \mathbf{y}$ and $\hat{\boldsymbol{\beta}} = U^T \hat{\mathbf{y}}$ represent the coordinates of \mathbf{y} and $\hat{\mathbf{y}}$ in the basis of eigenvectors of $D_{n,q}^T D_{n,q}$, F corresponds to a coordinate transformation matrix playing a similar role for the parameters as the hat matrix $H = U(U^T W U + S_\lambda)^{-1} U^T W$ does for the observations. The diagonal values of F can be interpreted as the effective degrees of freedom associated with each eigenvector after smoothing. It can be verified that:

$$\text{tr}(F) = \text{tr}[(U^T W U + S_\lambda)^{-1} U^T W U] = \text{tr}[U(U^T W U + S_\lambda)^{-1} U^T W] = \text{tr}(H)$$

which means that the sum of the effective degrees of freedom remains the same whether it is counted per observation or per parameter.

Figure 7 represents the effective degrees of freedom per parameter in the previous illustration of smoothing. The first q eigenvectors are never penalized, so their effective degrees of freedom are always equal to 1, regardless of the smoothing parameter used. The other eigenvectors have strictly decreasing effective degrees of freedom with λ . These degrees of freedom are generally decreasing with increasing eigenvalues of $D_{n,q}^T D_{n,q}$, although in the presence of non-unit weights and for small values of λ , this may not always be the case.




Figure 7: Distribution of residual degrees of freedom per eigenvector after applying the Whittaker-Henderson smoothing, in the one-dimensional case, for different values of the smoothing parameter.

6.3 Extension to the two-dimensional case

In the two-dimensional case, we have $P_\lambda = \lambda_x I_{n_x} \otimes D_{n_x, q_x}^T D_{n_x, q_x} + \lambda_z D_{n_z, q_z}^T D_{n_z, q_z} \otimes I_{n_x}$. Similar to the one-dimensional case, we can perform the eigendecomposition of the matrices $D_{n_x, q_x}^T D_{n_x, q_x}$ and $D_{n_z, q_z}^T D_{n_z, q_z}$, yielding $D_{n_x, q_x}^T D_{n_x, q_x} = U_x \Sigma_x U_x^T$ and $D_{n_z, q_z}^T D_{n_z, q_z} = U_z \Sigma_z U_z^T$. Let us define $U = U_z \otimes U_x$ and perform the reparameterization $\beta = U^T \theta \Leftrightarrow \theta = U \beta$. By leveraging the properties of the Kronecker product, we can rewrite the smoothness criterion in a simplified form:

$$\theta^T P_\lambda \theta = (U \beta)^T P_\lambda (U \beta) = \beta^T (\lambda_x I_{n_z} \otimes \Sigma_x + \lambda_z \Sigma_z \otimes I_{n_x}) \beta.$$

This leads to an alternative formulation of the optimization problem:

$$\hat{y} = U \hat{\beta} \quad \text{where} \quad \hat{\beta} = \underset{\beta}{\operatorname{argmin}} \left\{ (\mathbf{y} - U \beta)^T W (\mathbf{y} - U \beta) + \lambda \beta^T (\lambda_x I_{n_z} \otimes \Sigma_x + \lambda_z \Sigma_z \otimes I_{n_x}) \beta \right\}. \quad (17)$$

The solution to the smoothing problem, as in the one-dimensional case, is given by:

$$\hat{\mathbf{y}} = U(U^T W U + S_\lambda)^{-1} U^T W \mathbf{y} \quad \text{where} \quad S_\lambda = \lambda_x I_{n_z} \otimes \Sigma_x + \lambda_z \Sigma_z \otimes I_{n_x}. \quad (18)$$

Figure 8 represents the residual degrees of freedom associated with each parameter after applying the smoothing, in the two-dimensional case, for different combinations of the smoothing parameters. Similar to the one-dimensional case, these degrees of freedom decrease as the smoothing parameters increase and are particularly small for higher eigenvalues. The eigenvectors are sorted in ascending order of eigenvalues for each one-dimensional penalty matrix $D_{n_x, q_x}^T D_{n_x, q_x}$ and $D_{n_z, q_z}^T D_{n_z, q_z}$.




Figure 8: Distribution of residual degrees of freedom of the model by eigenvector after applying Whittaker-Henderson smoothing in the two-dimensional case, for different combinations of smoothing parameters.

6.4 Eigendecomposition and rank reduction

In addition to providing a more intuitive interpretation of the smoothing effect, the eigendecomposition of the penalty matrix is key to reducing the dimension of the optimization problem associated with it. Figures 7 and 8 show that eigenvectors corresponding to higher eigenvalues are more severely penalized by the smoothing, to the extent that a large majority of the eigenvectors represented in these figures have residual degrees of freedom very close to 0 for most values of λ . This suggests that if we simply remove the parameters associated with these eigenvectors from the model, the results of smoothing would hardly be affected. We propose setting the coordinates associated with the $n - p$ largest eigenvectors to 0, thus retaining only a reduced number $p < n$ of parameters to estimate. In the one-dimensional case, this translates to replacing $\hat{\mathbf{y}}$ by the approximate estimator:

$$\hat{\mathbf{y}}_p = U_p(U_p^T W U_p + \lambda \Sigma_p)^{-1} U_p^T W \mathbf{y} \quad (19)$$

where the matrix U_p contains the first p columns of U , and Σ_p contains the rows and columns associated with the p smallest eigenvalues of Σ . In the two-dimensional case, we use:

$$\hat{\mathbf{y}}_{p_x, p_z} = U_{p_x, p_z} (U_{p_x, p_z}^T W U_{p_x, p_z} + \lambda_x I_{p_x} \otimes \Sigma_{x, p_x} + \lambda_z \Sigma_{z, p_z} \otimes I_{p_x})^{-1} U_{p_x, p_z}^T W \mathbf{y} \quad (20)$$

where $U_{p_x, p_z} = U_{z, p_z} \otimes U_{x, p_x}$, U_{x, p_x} (resp. U_{z, p_z}) contains the first p_x (resp. p_z) columns of U_x (resp. U_z), and Σ_{x, p_x} (resp. Σ_{z, p_z}) contains the rows and columns associated with the p_x (resp. p_z) smallest eigenvalues of Σ_x (resp. Σ_z).

In the two-dimensional case, there are several acceptable strategies for choosing the pair (p_x, p_z) . One simple solution is to set an upper bound p_{\max} on the number of parameters to be retained and choose p_x and p_z proportionally to n_x and n_z . Let $\kappa = \sqrt{p_{\max}/n_x n_z}$, and define $p_x = \lfloor \min(\kappa, 1)n_x \rfloor$ and $p_z = \lfloor \min(\kappa, 1)n_z \rfloor$. By construction, the pair $(p_x, p_z) \in \mathbb{N}^2$ defined this way satisfies the condition $p_x \times p_z \leq \kappa^2 n_x n_z = p_{\max}$. Other strategies may prove more efficient. Indeed, if the underlying law is known to have greater smoothness in one dimension compared to the other, retaining a smaller proportion of components associated with that dimension is advisable.

The results presented in this Section can be directly extended to *generalized Whittaker-Henderson smoothing* by replacing \mathbf{y} with \mathbf{z}_k and W with W_k in Equations 19 and 20.

6.5 Impact of rank reduction method

Here, we aim to assess the impact of using the rank-reduced estimators defined in Section 6.4 on the smoothing results and computation time. To do so, we start from the 100 replicate datasets (one-dimensional and two-dimensional) used in Section 5.6. The approach from Section 6.4 may be used for the selection of the smoothing parameter as well as for the final smoothing step. In what follows however, once the optimal parameter is chosen, we use the estimator $\hat{\mathbf{y}}$

(or $\hat{\theta}$) associated with the full-rank problem in Equation 3 or Algorithm 1 in order to make it easier to compare it to the full-rank method. We may therefore use the error measure defined by Equation 14 to quantify the impact of the rank reduction method on the smoothing results. The quantity $\ell_{\text{ML}}^m(\lambda)$ involved in this formula will be computed based on the values of λ obtained from the *performance iteration* approach for different values of the number of retained components p . In addition, the quantities $\ell_{\text{ML}}^m(\lambda_{\text{outer}})$ and $\ell_{\text{ML}}^m(\infty)$ appearing in the expression of the criterion will always be calculated using the *outer iteration* method and retaining all components in order to capture potential cross-effects between the retained approach and the use of rank reduction.




Figure 9: Distribution of the discrepancies induced by the proposed rank reduction method, depending on the number of retained eigenvectors during the selection of the smoothing parameter, in the one-dimensional and two-dimensional cases.

Figure 9 represents the impact of rank reduction on the smoothing results as a function of the number of retained components. In the one-dimensional case, reducing the number of retained eigenvectors from 74 to 10 only marginally impacts the results obtained. However, as shown in Figure 10, it reduces the computation time by a factor of 3. In the two-dimensional case, reducing the number of components from 627 to 160 reduces the computation time by a factor of close to 10, with imperceptible differences in the obtained results. The proposed rank reduction method, in combination with the *performance iteration* approach, thus allows, in the two-dimensional case, by considering only 160 components, to reduce the computation time by a factor of 40 at the cost of an average error of 0.03%, which exceeds 0.5% in only 1 of the 100 simulated datasets considered. This solution therefore offers significant operational advantages, especially for applications mentioned in Section 6.1 where computation times are far higher than in the presented analysis.




Figure 10: Distribution of computation time for the proposed rank reduction method, depending on the number of retained eigenvectors during the selection of the smoothing parameter, in the one-dimensional and two-dimensional cases.

7 How to extrapolate the smoothing?

Non-parametric methods such as P-splines or Whittaker-Henderson smoothing may be naturally extrapolated beyond the range of the original data, in the same way as parametric models. Nonetheless, this requires to take some precautions especially in the two-dimensional case as we show in this Section.

7.1 Defining the (unconstrained) extrapolation of the smoothing

Let $\hat{\mathbf{y}}$ be the result of Whittaker-Henderson smoothing for an observation vector \mathbf{y} , which corresponds in the one-dimensional case to an explanatory variable vector \mathbf{x} or in the two-dimensional case to a combination of explanatory variables \mathbf{x} and \mathbf{z} . Suppose we have an observation vector \mathbf{x}_+ or combinations of vectors \mathbf{x}_+ and \mathbf{z}_+ for which we want to make predictions using the model. Since Whittaker-Henderson smoothing applies only to evenly spaced observations, without loss of generality, let \mathbf{x}_+ and \mathbf{z}_+ be sequences of consecutive integers such that $\mathbf{x} \subset \mathbf{x}_+$ and $\mathbf{z} \subset \mathbf{z}_+$. In the one-dimensional case, let n_+ be the length of \mathbf{x}_+ , and in the two-dimensional case, let n_{x_+} and n_{z_+} be the respective lengths of \mathbf{x}_+ and \mathbf{z}_+ , with $n_+ = n_{x_+} \times n_{z_+}$. Let C_x (resp. C_z) be a matrix of size $n_x \times n_{x_+}$ (resp. $n_z \times n_{z_+}$) defined as $C_j = (0_{n_j, \min j - \min j_+} | I_{n_j, n_j} | 0_{n_j, \max j_+ - \max j})$ for $j \in \{x, z\}$, and let us note:

$$C = \begin{cases} C_x & \text{in the one-dimensional case} \\ C_z \otimes C_x & \text{in the two-dimensional case} \end{cases}$$

The matrix C of size $n \times n_+$ defined in this way has the following properties:

- For any vector \mathbf{y}_+ of size n_+ , $C\mathbf{y}_+$ is a vector of size n containing only the n values of \mathbf{y}_+ matching the positions of the observations in the initial smoothing,
- $C^T\mathbf{y}$ is a vector of size n_+ containing the values of \mathbf{y} at the positions corresponding to the initial observations and zeros elsewhere,
- $CC^T = I_n$.

The extrapolation of Whittaker-Henderson smoothing can now be defined as finding the solution $\hat{\mathbf{y}}_+$ to the extended optimization problem:

$$\hat{\mathbf{y}}_+ = \underset{\boldsymbol{\theta}_+}{\operatorname{argmin}} \left\{ (\mathbf{y}_+ - \boldsymbol{\theta}_+)^T W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+) + \boldsymbol{\theta}_+^T P_+ \boldsymbol{\theta}_+ \right\} \quad (21)$$

where $\mathbf{y}_+ = C^T\mathbf{y}$, $W_+ = C^TWC$, and

$$P_+ = \begin{cases} \lambda D_{n_+,q}^T D_{n_+,q} & \text{in the one-dimensional case} \\ \lambda_x I_{z_+} \otimes D_{n_{x+,q_x}}^T D_{n_{x+,q_x}} + \lambda_z D_{n_{z+,q_z}}^T D_{n_{z+,q_z}} \otimes I_{x_+} & \text{in the two-dimensional case.} \end{cases}$$

In this last expression, λ , λ_x and λ_z are the parameters chosen for the smoothing. Similar to the initial smoothing, Equation 21 involves a fidelity criterion and a smoothness criterion. It should be noted that:

$$(\mathbf{y}_+ - \boldsymbol{\theta}_+)^T W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+) = (C^T\mathbf{y} - \boldsymbol{\theta}_+)^T C^T W C (C^T\mathbf{y} - \boldsymbol{\theta}_+) = (\mathbf{y} - \boldsymbol{\theta})^T W (\mathbf{y} - \boldsymbol{\theta})$$

hence the fidelity criterion remains unchanged compared to the initial estimation problem. This is consistent with the fact that extrapolation does not involve any additional data. The smoothness criterion on the other hand applies to all elements of the vector $\boldsymbol{\theta}_+$, whether initially present or not.

7.2 Unconstrained solution for the one-dimensional case

The solution to the extended optimization problem of Equation 21 is readily obtained by taking the derivatives in $\boldsymbol{\theta}_+$ and setting them to 0, as in Section 1.1.3. This yields the solution:

$$\hat{\mathbf{y}}_+ = (W_+ + P_+)^{-1} W_+ \mathbf{y}_+ \quad \text{where} \quad \mathbf{y}_+ = C^T\mathbf{y} \quad \text{and} \quad W_+ = C^T W C \quad (22)$$

Let us further assume that $\mathbf{y}_+ | \boldsymbol{\theta}_+ \sim \mathcal{N}(\boldsymbol{\theta}_+, W_+^-)$ where $\boldsymbol{\theta}_+ \sim \mathcal{N}(0, P_+^-)$ and proceed as in Section 2 to obtain the following credible intervals:

$$\mathbb{E}(\mathbf{y}_+) | \mathbf{y}_+ \in \left[(W_+ + P_+)^{-1} W_+ \mathbf{y}_+ \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\operatorname{diag} \{ (W_+ + P_+)^{-1} \}} \right] \quad (23)$$

To get a better understanding about how the variance-covariance matrix $\Psi_+ = (W_+ + P_+)^{-1}$ for the unconstrained extrapolation problem of Equation 21 is related to the variance-covariance matrix $\Psi = (W + P_\lambda)^{-1}$ of the original smoothing problem, let us introduce:

$$\bar{C}_j = \begin{bmatrix} I_{\min j - \min j_+, \min j - \min j_+} & 0_{\min j - \min j_+, n_j} & 0_{\min j - \min j_+, \max j_+ - \max j} \\ 0_{\max j_+ - \max j, \min j - \min j_+} & 0_{\max j_+ - \max j, n_j} & I_{\max j_+ - \max j, \max j_+ - \max j} \end{bmatrix} \quad \text{for } j \in \{x, z\},$$

$$\bar{C} = \begin{cases} \bar{C}_x & \text{in the one-dimensional case} \\ \bar{C}_z \otimes \bar{C}_x & \text{in the two-dimensional case} \end{cases} \quad \text{and } Q = \begin{bmatrix} C \\ \bar{C} \end{bmatrix}.$$

With this definition, Q is a permutation matrix of size $n_+ \times n_+$ which selects the rows whose indices correspond to those of the initial observation positions in the extended observation vector and move them to the first positions. It also verifies $Q^{-1} = Q^T = Q$.

In the unidimensional case, the extended difference matrix $D_{n_+,q}$ takes the form:

$$D_{n_+,q} = \begin{bmatrix} D_{2-} & D_{1-} & 0 \\ 0 & D_{n,q} & 0 \\ 0 & D_{1+} & D_{2+} \end{bmatrix} = Q \begin{bmatrix} D_{n,q} & 0 & 0 \\ D_{1-} & D_{2-} & 0 \\ D_{1+} & 0 & D_{2+} \end{bmatrix} Q = Q \begin{bmatrix} D_{n,q} & 0 \\ D_1 & D_2 \end{bmatrix} Q$$

$$\text{where } D_1 = \begin{bmatrix} D_{1-} \\ D_{1+} \end{bmatrix} \quad \text{and} \quad D_2 = \begin{bmatrix} D_{2-} & 0 \\ 0 & D_{2+} \end{bmatrix}$$

Detailed expressions of matrices D_{1+} and D_{2+} for the most common values of q may be found in Carballo et al. (2021) where they are simply noted D_1 and D_2 . To cover extrapolation on both sides of the initial observation vector, matrices D_{1-} and D_{2-} were introduced. Those may be simply obtained by taking the transpose of adequate size D_1 and D_2 matrices found in Carballo et al. (2021). The extended weight and penalization matrices thus take the form:

$$W_+ = Q \begin{bmatrix} W & 0 \\ 0 & 0 \end{bmatrix} Q \quad \text{and} \quad P_+ = D_{n_+,q}^T D_{n_+,q} = \lambda Q \begin{bmatrix} P_\lambda + P_+^{11} & P_+^{12} \\ P_+^{21} & P_+^{22} \end{bmatrix} Q \quad (24)$$

where $P_+^{11} = \lambda D_1^T D_1$, $P_+^{12} = \lambda D_1^T D_2$, $P_+^{21} = \lambda D_2^T D_1$ and $P_+^{22} = \lambda D_2^T D_2$.

Applying the formulas for the inversion of a symmetric matrix partitioned with 2×2 blocks to $Q(W_+ + P_+)Q$, we obtain a more detailed expression for Ψ_+ :

$$\Psi_+ = Q \begin{bmatrix} \Psi_+^{11} & \Psi_+^{12} \\ \Psi_+^{21} & \Psi_+^{22} \end{bmatrix} Q = Q \begin{bmatrix} \Psi_+^{11} & -\Psi_+^{11} P_+^{12} (P_+^{22})^{-1} \\ -(P_+^{22})^{-1} P_+^{21} \Psi_+^{11} & (P_+^{22})^{-1} + (P_+^{22})^{-1} P_+^{21} \Psi_+^{11} P_+^{12} (P_+^{22})^{-1} \end{bmatrix} Q. \quad (25)$$

where $\Psi_+^{11} = [W + P_\lambda + P_+^{11} - P_+^{12} (P_+^{22})^{-1} P_+^{21}]^{-1}$.

Let us denote by *initial positions coefficients* the subvector $C\hat{y}_+$ that corresponds to the coefficients at the position of the initial observations in \hat{y}_+ and by *new positions coefficients* the remaining coefficients $\bar{C}\hat{y}_+$. The *initial positions coefficients* may be recovered as $C\hat{y}_+ = C\Psi_+ C^T W_+ y_+ = \Psi_+^{11} W y$. This does not simplify to $\Psi W y = \hat{y}$ for all weighted observation vectors $W y$ unless $\Psi_+^{11} = \Psi$. Indeed, the *initial positions coefficients* are chosen to optimize

the overall smoothness of the extrapolated coefficient vectors and not only the smoothness associated with the subvector $C\hat{\mathbf{y}}_+$. Besides, the expression of Ψ_+^{22} contains two terms: an *innovation error* $(P_+^{22})^{-1}$ associated with the prior on the *new positions coefficients* while $(P_+^{22})^{-1}P_+^{21}\Psi_+^{11}P_+^{12}(P_+^{22})^{-1}$ represents an additional uncertainty on the *new positions coefficients* caused by the uncertainty on the *initial positions coefficients*.

In the one-dimensional case, D_2 is a block-diagonal matrix of triangular matrices with non-zero diagonal elements and is thus non-singular. Hence $P_+^{11} - P_+^{12}(P_+^{22})^{-1}P_+^{21} = D_1^T D_1 - D_1^T D_2 (D_2^T D_2)^{-1} D_2^T D_1 = 0$ which means that $\Psi_+^{11} = (W + P_\lambda)^{-1} = \Psi$ and $C\hat{\mathbf{y}}_+ = \hat{\mathbf{y}}$. Therefore, in the one-dimensional case, the solution given by Equation 21 preserves the values from the original fit. As shown by Carballo et al. (2021), it is indeed always possible in the case of penalizations based on difference matrices to pick the *new positions coefficients* so that the smoothness criterion does not increase or in other words to find a perfectly smooth extrapolation for the fit. This is illustrated by Figure 11 which shows the extrapolation associated with the Whittaker-Henderson smoothing applied to the data used to produce Figure 1.




Figure 11: Extrapolation of Whittaker-Henderson smoothing in the one-dimensional case. The smoother extrapolation is achieved, on both sides of the smoothing, using a polynomial of degree $q - 1$ (in this case a straight line as $q = 2$).

7.3 Constrained solution for the two-dimensional case

In the two-dimensional case, while the extended penalization matrix P_+ still takes the form of Equation 24, expressions of P_+^{11} , P_+^{12} , P_+^{21} and P_+^{22} are more complex. In particular, $P_+^{11} - P_+^{12}(P_+^{22})^{-1}P_+^{21} \neq 0$ which implies that $\Psi_+^{11} \neq \Psi$ and $C\hat{\mathbf{y}}_+ \neq \hat{\mathbf{y}}$. Solving Equation 22 thus leads to a change in the value of the *initial positions coefficients* compared to the coefficients

obtained during the initial smoothing as shown by Carballo, Durban, and Lee (2021). Indeed, the smoothness criterion includes penalizations on both rows and columns and it is no longer possible, as in the one-dimensional case, to extrapolate the fit without increasing that criterion. As the smoothness criterion carries more weight in the extrapolation problem compared to the original problem, the optimal solution to the extended optimization problem will compromise on the fidelity to the initial observations in order to improve the overall smoothness.

To obtain an estimator $\hat{\mathbf{y}}_+^*$ that minimizes the penalized regression problem under the constraint of preserving the initial coefficients, *i.e.* $C\hat{\mathbf{y}}_+^* = \hat{\mathbf{y}}$, we follow the approach proposed by Carballo, Durban, and Lee (2021) and introduce the Lagrange multiplier $\boldsymbol{\omega}$. The associated constrained extended optimization problem is now written as:

$$(\hat{\mathbf{y}}_+^*, \hat{\boldsymbol{\omega}}) = \underset{\boldsymbol{\theta}_+^*, \boldsymbol{\omega}}{\operatorname{argmin}} \left\{ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*)^T W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*) + \boldsymbol{\theta}_+^{*T} P_+ \boldsymbol{\theta}_+^* + 2\boldsymbol{\omega}^T (C\boldsymbol{\theta}_+^* - \hat{\mathbf{y}}) \right\}. \quad (26)$$

Taking the partial derivatives of Equation 26 with respect to $\boldsymbol{\theta}_+^*$ and $\boldsymbol{\omega}$ gives:

$$\begin{aligned} \frac{\partial}{\partial \boldsymbol{\theta}_+^*} \left\{ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*)^T W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*) + \boldsymbol{\theta}_+^{*T} P_+ \boldsymbol{\theta}_+^* + 2\boldsymbol{\omega}^T (C\boldsymbol{\theta}_+^* - \hat{\mathbf{y}}) \right\} &= -2W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*) + 2P_+ \boldsymbol{\theta}_+^* + 2\boldsymbol{\omega}^T C \\ \frac{\partial}{\partial \boldsymbol{\omega}} \left\{ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*)^T W_+ (\mathbf{y}_+ - \boldsymbol{\theta}_+^*) + \boldsymbol{\theta}_+^{*T} P_+ \boldsymbol{\theta}_+^* + 2\boldsymbol{\omega}^T (C\boldsymbol{\theta}_+^* - \hat{\mathbf{y}}) \right\} &= 2(C\boldsymbol{\theta}_+^* - \hat{\mathbf{y}}) \end{aligned}$$

Setting these derivatives to zero yields the linear system:

$$\begin{bmatrix} W_+ + P_+ & C^T \\ C & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{y}}_+^* \\ \hat{\boldsymbol{\omega}} \end{bmatrix} = \begin{bmatrix} W_+ \mathbf{y}_+ \\ \hat{\mathbf{y}} \end{bmatrix}$$

The solution for $\hat{\mathbf{y}}_+^*$ can be derived using formulas for the inversion of a symmetric partitioned matrix with 2×2 blocks:

$$\begin{aligned} \hat{\mathbf{y}}_+^* &= (W_+ + P_+)^{-1} \left\{ I - C^T [C(W_+ + P_+)^{-1} C^T]^{-1} C (W_+ + P_+)^{-1} \right\} W_+ \mathbf{y}_+ \\ &\quad + (W_+ + P_+)^{-1} C^T [C(W_+ + P_+)^{-1} C^T]^{-1} \hat{\mathbf{y}} \end{aligned}$$

Since $W_+ = C^T W C$, the first term is actually zero, and this expression simplifies to:

$$\hat{\mathbf{y}}_+^* = (W_+ + P_+)^{-1} C^T [C(W_+ + P_+)^{-1} C^T]^{-1} \hat{\mathbf{y}} = Q \begin{bmatrix} I \\ -(P_+^{22})^{-1} P_+^{21} \end{bmatrix} \hat{\mathbf{y}} \quad (27)$$

which is a linear transformation of $\hat{\mathbf{y}}$. Defining $A_+^* = Q \begin{bmatrix} I \\ -(P_+^{22})^{-1} P_+^{21} \end{bmatrix}$, a natural candidate for the variance-covariance of $\mathbf{y}_+^* | \boldsymbol{\theta}_+$ is given by:

$$A_+^* \Psi A_+^{*T} = Q \begin{bmatrix} \Psi & -\Psi P_+^{12} (P_+^{22})^{-1} \\ -(P_+^{22})^{-1} P_+^{21} \Psi & (P_+^{22})^{-1} P_+^{21} \Psi P_+^{12} (P_+^{22})^{-1} \end{bmatrix} Q \quad (28)$$

Equation 28 is very similar to Equation 25 with however two differences. First, every occurrence of Ψ_+^{11} is replaced by Ψ . This is consistent with the constraint that the *initial positions coefficients* are forced to take the value of the coefficients used during the initial smoothing. Second, as the solution to the constrained extended optimization problem of Equation 26 was expressed as a linear transformation of $\hat{\mathbf{y}}$, Equation 28 is missing the *innovation error* term $(P_+^{22})^{-1}$ associated with the prior on the *new positions coefficients*. Not including this term would be tantamount to considering that $\boldsymbol{\theta}_+$ has some degree of wigglyness in the region of the initial data but is perfectly smooth everywhere else. Adding the *innovation error* back, we obtain the following variance-covariance matrix for the constrained optimization problem:

$$\Psi_+^* = Q \begin{bmatrix} \Psi & -\Psi P_+^{12} (P_+^{22})^{-1} \\ -(P_+^{22})^{-1} P_+^{21} \Psi & (P_+^{22})^{-1} + (P_+^{22})^{-1} P_+^{21} \Psi P_+^{12} (P_+^{22})^{-1} \end{bmatrix} Q \quad (29)$$

which still verifies

$$\Psi_+^* W_+ \mathbf{y}_+ = \Psi_+^* C W \mathbf{y} = Q \begin{bmatrix} I \\ -(P_+^{22})^{-1} P_+^{21} \end{bmatrix} \Psi W \mathbf{y} = \hat{\mathbf{y}}_+^*.$$

The associated $100(1 - \alpha)\%$ credible intervals are readily obtained as:

$$\mathbb{E}(\mathbf{y}_+) | \mathbf{y}_+ \in \left[\hat{\mathbf{y}}_+^* \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\mathbf{diag}(\Psi_+^*)} \right].$$

Figure 12 represents the results of the constrained extrapolation presented in this Section, with the associated standard deviation which accounts for the *innovation error*. It is based on the data used to generate Figure 2. Without the dotted lines marking the boundaries of the initial smoothing region, it would not be possible to tell from those plots where the extrapolation starts, which is precisely the goal of this procedure. Figure 13 represents the ratio between the mortality rates obtained from the unconstrained solution of Equation 22 and the constrained solution of Equation 27. The unconstrained solution shows significant discrepancies both in the initial smoothing region and the extrapolated region compared to the constrained solution. Finally, Figure 14 represents ratios between standard deviation derived using all of the three presented extrapolation methods. The denominator is the constrained extrapolation method which accounts for the *innovation error*, which is used as a reference, while the numerator shows the unconstrained method as well as the constrained method which ignores the *innovation error*. The constrained extrapolation method which includes the *innovation error* has a higher standard deviation compared to the unconstrained method as it always result in estimates that are less smooth.




Figure 12: Constrained extrapolation of Whittaker-Henderson smoothing in the two-dimensional case. The contour lines of mortality rates and the associated standard deviation are depicted. The dotted lines delimitate the boundaries of the initial smoothing region.




Figure 13: Ratio of mortality rates resulting from the extrapolation of the two-dimensional Whittaker-Henderson smoothing. The numerator corresponds to the unconstrained extrapolation and the denominator is the constrained extrapolation presented in Figure 12.




Figure 14: Ratio of standard deviation of log-mortality rates resulting from the extrapolation of the two-dimensional Whittaker-Henderson smoothing. The numerator on the left plot corresponds to the unconstrained extrapolation and the one on the right plot to the constrained extrapolation that does not account for the *innovation error*. The denominator of both ratios is the constrained extrapolation presented in Figure 12.

8 Discussion

In this paper, we first showed that Whittaker-Henderson smoothing, a method introduced in 1922, can be naturally interpreted in a Bayesian framework, which allows for credible intervals to be built, provided that the observations are independent and their variances are known and used as weights. By linking it to the framework of duration models, we have proved that in the construction of survival laws from experience, smoothing can be applied to the vector of crude exit rates, using the number of observed terminations as vector of weights. This is justified by the asymptotic normality of the maximum likelihood estimator of the raw exit rates. We have then established that the use of this asymptotic property comes at the cost of significant bias when the number of observations is limited, and that a more precise iterative version of the smoothing may instead be used by iteratively solving penalized likelihood equations.

Furthermore, we introduced an empirical Bayesian approach for selecting the smoothing parameter, based on maximizing a marginal likelihood function. The parameter selection introduces a second iterative process that is added to the one generated by solving the penalized likelihood equations for a fixed smoothing parameter. In this case, there are two main possible approaches for parameter selection, depending on the order in which the iterations are nested: the *outer iteration* and the *performance iteration* approaches. The latter has significantly lower computation time compared to the former, at the expense of lacking theoretical convergence guarantees, which however does not seem an issue in practice. We also introduced a rank

reduction method based on eigendecomposition of one-dimensional penalty matrices, which greatly accelerates the selection of the smoothing parameter - our study on simulated data show a 40-fold reduction in computation time - with negligible loss of precision.

Finally, we addressed the issue of extrapolation of smoothing and shown that it requires solving a new optimization problem. Extrapolation is straightforward in the one-dimensional case. However, in the two-dimensional case constraints needs to be imposed in order to preserve the values of the coefficients obtained during the initial smoothing step and the variance-covariance matrix needs to be corrected to include the innovation error associated with the extrapolated data.

Whittaker-Henderson smoothing has been used for nearly a century by actuaries, with little change in the approach. In parallel, statistical theory on smoothing methods has undergone numerous developments, particularly in the past 30 years with the emergence of generalized additive models, covered notably by Hastie and Tibshirani (1990) and Wood (2006). Our goal in this paper was to bridge the gap between those two perspectives. We also created an R package named `WH`, which implements all the steps mentioned in the paper and should be straightforward to use. Most results derived in this paper are directly applicable to other types of smoothing, such as the widely used P-splines smoothing method introduced by Eilers and Marx (1996). Compared to Whittaker-Henderson smoothing, P-splines are low-rank smoothers which naturally consider fewer parameters than observations, removing the need for the method introduced in Section 6.

While Whittaker-Henderson smoothing is applicable to both the one-dimensional and two-dimensional case, it does have some limitations. Firstly, it requires regularly spaced observations. This aligns well with the format of life insurance pricing and reserving assumptions, which traditionally include age and duration spent in certain states which trigger claim payment. However, this can be a limit when the terminations are not evenly distributed. For example, in the case of disability and long-term care risks, most terminations occur in the first few months following entry into the state. In such situations, using a spline basis that is arranged to prioritize areas with more observations, as proposed by Wood (2017), could yield better results. Secondly, Whittaker-Henderson smoothing in its original form does not allow for the incorporation of additional explanatory variables in the experience table, starting with gender, which plays a major role in most biometric risks. Fortunately, it is possible to introduce these additional variables as *random effects* in the model by adopting a *Smoothing Splines ANOVA approach*, as described in Lee and Durban (2011) and Gu (2013), yielding results that are both precise and robust compared to a unisex or stratified model.

Appendix A : Algorithms

Algorithm 1 provides an implementation of *generalized WH smoothing* as defined in Section 4 while Algorithms 2 and 3 also perform the estimation of the smoothing parameters respectively in the *outer iteration* and *performance iteration* approaches introduced in Section 5.

Algorithm 1: Iterative solution of generalized Whittaker-Henderson smoothing

inputs : \mathbf{d} and \mathbf{e}_c

outputs : $\hat{\boldsymbol{\theta}}$ and $\text{Var}(\hat{\boldsymbol{\theta}})$

parameters: $\lambda, q, \epsilon_\ell = 10^{-8}$

begin

Construct the penalty matrix P_λ based on the difference matrices of order q .

$k \leftarrow 0$

$\boldsymbol{\theta}_0 \leftarrow \ln(\mathbf{d}/\mathbf{e}_c)$

$l_0 \leftarrow -\infty, \text{cond}_\ell \leftarrow \text{true}$

while cond_ℓ **do**

$W_k \leftarrow \text{Diag}(\exp(\boldsymbol{\theta}_k) \odot \mathbf{e}_c)$

$\mathbf{z}_k \leftarrow \boldsymbol{\theta}_k + W_k^{-1}[\mathbf{d} - \exp(\boldsymbol{\theta}_k) \odot \mathbf{e}_c]$

$\Psi_{k+1} \leftarrow (W_k + P_\lambda)^{-1}$

$\boldsymbol{\theta}_{k+1} \leftarrow \Psi_{k+1} W_k \mathbf{z}_k$

$l_{k+1} \leftarrow \ell_P(\boldsymbol{\theta}_{k+1}), \text{cond}_\ell \leftarrow l_{k+1} - l_k \geq \epsilon_\ell \times \text{sum}(d)$

$k \leftarrow k + 1$

$\hat{\boldsymbol{\theta}} \leftarrow \boldsymbol{\theta}_k, \text{Var}(\hat{\boldsymbol{\theta}}) \leftarrow \Psi_k$

Algorithm 2: Parameter selection for generalized Whittaker-Henderson smoothing - outer iteration approach.

inputs : \mathbf{d} and \mathbf{e}_c
outputs : $\hat{\lambda}$, $\hat{\boldsymbol{\theta}}_{\hat{\lambda}}$ and $\text{Var}(\hat{\boldsymbol{\theta}}_{\hat{\lambda}})$
parameters: q , $\epsilon_\ell = 10^{-8}$, $\epsilon_{\text{ml}} = 10^{-8}$
begin
 $k \leftarrow 0$
 $\text{ml}_0 \leftarrow -\infty$, $\text{cond}_{\text{ml}} \leftarrow \text{true}$
 while cond_{ml} **do**
 If $k = 0$, choose an arbitrary value λ_0 for the smoothing parameter(s); otherwise, choose the next value λ_k using the selected heuristic
 Use Algorithm 1 to determine the vector $\boldsymbol{\theta}_{\lambda_k}$ associated with the choice of λ_k with an accuracy greater than ϵ_ℓ
 Calculate the marginal likelihood $\ell_{\text{ML}}^m(\lambda_k)$ associated with the choice of λ_k using the intermediate quantities calculated during the estimation of $\boldsymbol{\theta}_{\lambda_k}$
 $\text{ml}_{k+1} \leftarrow \ell_{\text{ML}}^m(\lambda_k)$, $\text{cond}_{\text{ml}} \leftarrow \text{ml}_{k+1} - \text{ml}_k \geq \epsilon_{\text{ml}} \times \text{sum}(d)$
 $k \leftarrow k + 1$
 $\hat{\boldsymbol{\theta}} \leftarrow \hat{\boldsymbol{\theta}}_{\hat{\lambda}}$, $\text{Var}(\hat{\boldsymbol{\theta}}) \leftarrow \Psi_k$

Algorithm 3: Parameter selection for generalized Whittaker-Henderson smoothing - performance iteration approach

inputs : \mathbf{d} and \mathbf{e}_c
outputs : $\hat{\lambda}$, $\hat{\boldsymbol{\theta}}_{\hat{\lambda}}$ and $\text{Var}(\hat{\boldsymbol{\theta}}_{\hat{\lambda}})$
parameters: q , $\epsilon_\ell = 10^{-8}$, $\epsilon_{\text{ml}} = 10^{-8}$
begin
 $k \leftarrow 0$
 $\boldsymbol{\theta}_0 \leftarrow \ln(\mathbf{d}/\mathbf{e}_c)$
 $l_0 \leftarrow -\infty$, $\text{cond}_\ell \leftarrow \text{true}$
 while cond_ℓ **do**
 $W_k \leftarrow \text{Diag}(\exp(\boldsymbol{\theta}_k) \odot \mathbf{e}_c)$
 $\mathbf{z}_k \leftarrow \boldsymbol{\theta}_k + W_k^{-1}[\mathbf{d} - \exp(\boldsymbol{\theta}_k) \odot \mathbf{e}_c]$
 Estimate the parameter λ_k maximizing the regularized marginal likelihood ℓ_{norm}^m associated with the observation vector \mathbf{z}_k and the weight matrix W_k , using the selected heuristic, with an accuracy greater than $\epsilon_{\text{ml}} \times \text{sum}(d)$
 $\Psi_{k+1} \leftarrow (W_k + P_{\lambda_k})^{-1}$
 $\boldsymbol{\theta}_{k+1} \leftarrow \Psi_{k+1} W_k \mathbf{z}_k$
 $l_{k+1} \leftarrow \ell_P(\boldsymbol{\theta}_{k+1} | \lambda_k)$, $\text{cond}_\ell \leftarrow l_{k+1} - l_k \geq \epsilon_\ell \times \text{sum}(d)$
 $k \leftarrow k + 1$
 $\hat{\lambda} \leftarrow \lambda_k$; Use Algorithm 1 to obtain the values of $\hat{\boldsymbol{\theta}}_{\hat{\lambda}}$ and $\text{Var}(\hat{\boldsymbol{\theta}}_{\hat{\lambda}})$ corresponding to the parameter $\hat{\lambda}$

References

- Akaike, Hirotugu. 1973. "Information Theory and an Extension of the Maximum Likelihood Principle." In *2nd International Symposium on Information Theory, 1973*.
- Anderssen, RS, and Peter Bloomfield. 1974. "A Time Series Approach to Numerical Differentiation." *Technometrics* 16 (1): 69–75.
- Brent, Richard P. 1973. "Algorithms for Minimization Without Derivatives, Chap. 4." Prentice-Hall, Englewood Cliffs, NJ.
- Brooks, RJ, M Stone, FY Chan, and LK Chan. 1988. "Cross-Validatory Graduation." *Insurance: Mathematics and Economics* 7 (1): 59–66.
- Carballo, Alba, Maria Durban, Göran Kauermann, and Dae-Jin Lee. 2021. "A General Framework for Prediction in Penalized Regression." *Statistical Modelling* 21 (4): 293–312.
- Carballo, Alba, Maria Durban, and Dae-Jin Lee. 2021. "Out-of-Sample Prediction in Multidimensional p-Spline Models." *Mathematics* 9 (15): 1761.
- Currie, Iain D, Maria Durban, and Paul HC Eilers. 2004. "Smoothing and Forecasting Mortality Rates." *Statistical Modelling* 4 (4): 279–98.
- Eilers, Paul H. C., and Brian D. Marx. 1996. "Flexible Smoothing with B-Splines and Penalties." *Statistical Science* 11 (2): 89–102.
- Giesecke, Lee, and Defense Manpower Data Center. 1981. "Use of the Chi-Square Statistic to Set Whittaker-Henderson Smoothing Coefficients." Smoothing.
- Gilks, Walter R, Sylvia Richardson, and David Spiegelhalter. 1995. *Markov Chain Monte Carlo in Practice*. CRC press.
- Gschlössl, Susanne, Pascal Schoenmaekers, and Michel Denuit. 2011. "Risk Classification in Life Insurance: Methodology and Case Study." *European Actuarial Journal* 1: 23–41.
- Gu, Chong. 1992. "Cross-Validating Non-Gaussian Data." *Journal of Computational and Graphical Statistics* 1 (2): 169–79.
- . 2013. *Smoothing Spline ANOVA Models*. Vol. 297. Springer.
- Hastie, Trevor J, and Robert J Tibshirani. 1990. *Generalized Additive Models*. Vol. 43. CRC press.
- Henderson, Robert. 1924. "A New Method of Graduation." *Transactions of the Actuarial Society of America* 25: 29–40.
- Hoem, Jan M. 1971. "Point Estimation of Forces of Transition in Demographic Models." *Journal of the Royal Statistical Society: Series B (Methodological)* 33 (2): 275–89.
- Kauermann, Göran. 2005. "A Note on Smoothing Parameter Selection for Penalized Spline Smoothing." *Journal of Statistical Planning and Inference* 127 (1-2): 53–69.
- Kimeldorf, GS, and Donald A Jones. 1967. "Bayesian Graduation." *Transactions of the Society of Actuaries* 19 (54 part 1): 66–112.
- Lee, Dae-Jin, and Maria Durban. 2011. "P-Spline ANOVA-Type Interaction Models for Spatio-Temporal Smoothing." *Statistical Modelling* 11 (1): 49–69.
- Marra, Giampiero, and Simon N Wood. 2012. "Coverage Properties of Confidence Intervals for Generalized Additive Model Components." *Scandinavian Journal of Statistics* 39 (1): 53–74.

- Nelder, John A, and Roger Mead. 1965. "A Simplex Method for Function Minimization." *The Computer Journal* 7 (4): 308–13.
- Nelder, John Ashworth, and Robert WM Wedderburn. 1972. "Generalized Linear Models." *Journal of the Royal Statistical Society: Series A (General)* 135 (3): 370–84.
- Patterson H. D., Thompson R. 1971. "Recovery of Inter-Block Information When Block Sizes Are Unequal." *Biometrika* 58: 545–54.
- Reinsch, Christian H. 1967. "Smoothing by Spline Functions." *Numerische Mathematik* 10 (3): 177–83.
- Reiss, Philip T, and R Todd Ogden. 2009. "Smoothing Parameter Selection for a Class of Semiparametric Linear Models." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71 (2): 505–23.
- Taylor, Greg. 1992. "A Bayesian Interpretation of Whittaker—Henderson Graduation." *Insurance: Mathematics and Economics* 11 (1): 7–16.
- Wahba, Grace. 1980. *Spline Bases, Regularization, and Generalized Cross Validation for Solving Approximation Problems with Large Quantities of Noisy Data*. University of Wisconsin.
- . 1985. "A Comparison of GCV and GML for Choosing the Smoothing Parameter in the Generalized Spline Smoothing Problem." *The Annals of Statistics*, 1378–1402.
- Whittaker, Edmund Taylor. 1922. "On a New Method of Graduation." *Proceedings of the Edinburgh Mathematical Society* 41: 63–75.
- Whittaker, Edmund Taylor, and George Robinson. 1924. *The Calculus of Observations: A Treatise on Numerical Mathematics*. Blackie; Son limited.
- Wood, Simon N. 2006. *Generalized Additive Models: An Introduction with r*. Chapman; hall/CRC.
- . 2011. "Fast Stable Restricted Maximum Likelihood and Marginal Likelihood Estimation of Semiparametric Generalized Linear Models." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (1): 3–36.
- . 2017. "P-Splines with Derivative Based Penalties and Tensor Product Smoothing of Unevenly Distributed Data." *Statistics and Computing* 27: 985–89.
- . 2020. "Inference and Computation with Generalized Additive Models and Their Extensions." *Test* 29 (2): 307–39.