

# PHASE-TYPE DISTRIBUTIONS FOR CLAIM SEVERITY REGRESSION MODELING

BY

MARTIN BLADT

## ABSTRACT

This paper addresses the task of modeling severity losses using segmentation when the data distribution does not fall into the usual regression frameworks. This situation is not uncommon in lines of business such as third-party liability insurance, where heavy-tails and multimodality often hamper a direct statistical analysis. We propose to use regression models based on phase-type distributions, regressing on their underlying inhomogeneous Markov intensity and using an extension of the expectation–maximization algorithm. These models are interpretable and tractable in terms of multistate processes and generalize the proportional hazards specification when the dimension of the state space is larger than 1. We show that the combination of matrix parameters, inhomogeneity transforms, and covariate information provides flexible regression models that effectively capture the entire distribution of loss severities.

## KEYWORDS

Phase-type distributions, claim severity modeling, heavy-tailed regression.

**JEL codes:** 60J27, 91B30, 62J99

## 1. INTRODUCTION

The task of modeling claim severities is a well known and difficult challenge in actuarial science, and their correct description is of large interest for the practicing actuary or risk manager. In some lines of business, the size of claims may manifest features which are difficult to capture correctly by simple distributional assumptions, and *ad hoc* solutions are very often used to overcome this practical drawback. In particular, heavy-tailed and multimodal data may

*Astin Bulletin* 52(2), 417–448. doi:[10.1017/asb.2021.40](https://doi.org/10.1017/asb.2021.40) © The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

be present, which together with a high count of small claim sizes, produces an odd-shaped distribution.

In this paper, we introduce phase-type (PH) distributions as a powerful vehicle to produce probabilistic models for the effective description of claim severities. PH distributions are roughly defined as the time it takes for a pure-jump Markov processes on a finite state space to reach an absorbing state. Their interpretation in terms of traversing several states before finalizing is particularly appealing actuarial science, since one may think of claim sizes as the finalized “time” after having traversed some unobserved states to reach such magnitude, for instance, legal cases, disabilities, unexpected reparation costs, etc. Many distributions such as the exponential, Erlang, Coxian, and any finite mixture between them fall into PH domain, and they possess a universality property in that they are dense (in the sense of weak convergence, cf. Asmussen, 2008) on the set of probability measures with positive support.

Despite many particular instances of PH being classical distributions, a unified and systematic approach was only first studied in Neuts (1975, 1981), which laid out the modern foundations of matrix analytic methods (see Bladt and Nielsen, 2017 for a recent comprehensive treatment). Since then, they have gained popularity in applied probability and statistics due to their versatility and closed-form formulas in terms of matrix exponentials. A key development for their use in applied statistics was the development of a maximum likelihood estimation (MLE) procedure via the expectation–maximization (EM) algorithm (cf. Asmussen *et al.*, 1996), which considers the full likelihood arising from the path representation of a PH distribution. Ahn *et al.* (2012) explore the case of exponentially transformed PH distributions to generate heavy-tailed distributions. Furthermore, transforming PH distributions parametrically has been considered in Albrecher and Bladt (2019) and Albrecher *et al.* (2020b), resulting in inhomogeneous phase-type (IPH) distributions, which leads to non-exponential tail behaviors by allowing for the underlying process to be time-inhomogeneous (see also Bladt and Rojas-Nandayapa, 2017; Albrecher *et al.*, 2020a for alternative heavy-tailed PH specifications). The latter development allows to model heavy- or light-tailed data with some straightforward adaptations to the usual PH statistical methods.

Incorporating rating factors, or covariates, is particularly relevant when segmentation of claims is sought for, for example, for insurance pricing, and is known to be problematic when the underlying model does not belong to the exponential dispersion family, i.e. when a Generalized Linear Model (GLM) is not sufficiently flexible. Regression models for PH distributions have been considered in the context of survival analysis, cf. McGrory *et al.* (2009) and references therein for the Coxian PH case, and more recently (Albrecher *et al.*, 2021) for IPH distributions with applications to mortality modeling. The current work uses the specification of the model within the latter contribution but applied to claim severity modeling for insurance. The methodology is simpler, since we treat the fully observed case, allowing for some closed-form formulas for, for example, mean estimation or inferential tools.

Several other statistical models for insurance data can be found in the literature, with the common theme being constructing a global distribution consisting of smaller simpler components, either by mixing, splicing (also referred to as composite models), or both. In Lee and Lin (2010), a mixture of Erlang (ME) distributions with common scale parameter was proposed and subsequently extended to more general mixture components in Tzougas *et al.* (2014), Miljkovic and Grün (2016), and Fung *et al.* (2020). Concerning composite models, many different tail and body distribution combinations have been considered, and we refer to Grün and Miljkovic (2019) for a comparison of some of them, where a substantial body of the relevant literature can be found as well. Global models is dealt with a combination of mixtures and splicing, with Reynkens *et al.* (2017) being the main reference (see also Fung *et al.*, 2021 for a feature selection approach).

Our present PH regression model is not only mathematically but also conceptually different to all previous approaches. With the use of matrix calculus, it allows to specify multimodal and heavy-tailed distributions without the need of threshold selection or having to specify the number of mixture components. Instead, an inhomogeneity function and the size of the underlying state space have to be specified. The former is usually straightforward according to the tail behavior of the data, and the latter can be chosen either in a data-driven way or in a more dogmatic way if there is a certain belief of a specific number of unobserved states driving the data-generating process.

The structure of the remainder of the paper is as follows. In Section 2, we provide the necessary preliminaries on IPH distributions and formulate our PH regression model including their parametrizations, their estimation procedure, statistical inference, and dimension and structure selection. We then perform a simulation study in Section 3 to illustrate its effectiveness in a synthetic heterogeneous dataset. Subsequently, we perform the estimation on a real-life French insurance dataset in Section 4. Finally, Section 5 concludes.

## 2. MODEL SPECIFICATIONS

In this section, we lay down some preliminaries on IPH distributions which are needed in order to introduce the relevant key concepts to be used in our regression and subsequently proceed to introduce the specifications of the claim severity models. We skip the proofs, and the interested reader can find a comprehensive treatment of the marginal models in the following references: Bladt and Nielsen (2017) and Albrecher and Bladt (2019), the latter taking a more abstract approach. We will, however, tailor and emphasize the most relevant features of the models with respect to claim severity modeling for insurance.

**2.1. Mathematical formulation of IPH distributions**

Let  $(J_t)_{t \geq 0}$  be a time-inhomogeneous Markov pure-jump process on the finite state space  $\{1, \dots, p, p + 1\}$ , where states  $1, \dots, p$  are transient and  $p + 1$  is absorbing. In our setting, such a process can be regarded as an insurance claim evolving through different states during its lifetime until the absorbing state is reached, which determines the total size of the claim. In its most general form, the transition probabilities of the jump process:

$$p_{ij}(s, t) = \mathbb{P}(J_t = j | J_s = i), \quad i, j \in \{1, \dots, p + 1\},$$

may be written in succinct matrix form in terms of the product integral as follows:

$$P(s, t) = \prod_s^t (I + \Lambda(u)du) := I + \sum_{k=1}^\infty \int_s^t \int_s^{u_k} \dots \int_s^{u_2} \Lambda(u_1) \dots \Lambda(u_k) du_1 \dots du_k,$$

for  $s < t$ , where  $\Lambda(t)$  is a matrix with negative diagonal elements, nonnegative off-diagonal elements, and rows sums equal to 0, commonly referred to as an intensity matrix. Since the process  $(J_t)_{t \geq 0}$  has a decomposition in terms of transient and absorbing states, we may write

$$\Lambda(t) = \begin{pmatrix} T(t) & \mathbf{t}(t) \\ \mathbf{0} & 0 \end{pmatrix} \in \mathbb{R}^{(p+1) \times (p+1)}, \quad t \geq 0,$$

where  $T(t)$  is a  $p \times p$  sub-intensity matrix and  $\mathbf{t}(t)$  is a  $p$ -dimensional column vector providing the exit rates from each state directly to  $p + 1$ . The intensity matrix has this form, since there are no positive rates from the absorbing state to the transient ones, and since the rows should sum to 0, we have that for  $t \geq 0$ ,  $\mathbf{t}(t) = -T(t) \mathbf{e}$ , where  $\mathbf{e}$  is a  $p$ -dimensional column vector of ones. Thus,  $T(t)$  is sufficient to describe the dynamics of process after time zero. We write  $\mathbf{e}_k$  in the sequel for the  $k$ -th canonical basis vector of  $\mathbb{R}^p$ , so in particular  $\mathbf{e} = \sum_{k=1}^p \mathbf{e}_k$ .

If the matrices  $T(s)$  and  $T(t)$  commute for any  $s < t$ , we may write

$$P(s, t) = \begin{pmatrix} \exp\left(\int_s^t T(u)du\right) & \mathbf{e} - \exp\left(\int_s^t T(u)du\right) \mathbf{e} \\ \mathbf{0} & 1 \end{pmatrix}, \quad s < t.$$

The matrices  $(T(t))_{t \geq 0}$  together with the initial probabilities of the Markov process,  $\pi_k = \mathbb{P}(J_0 = k)$ ,  $k = 1, \dots, p$ , or in vector notation:

$$\boldsymbol{\pi} = (\pi_1, \dots, \pi_p)^\top,$$

fully parametrize IPH distributions, and with the additional assumption that  $\mathbb{P}(J_0 = p + 1) = 0$  we can guarantee that absorption happens almost surely at a positive time. We thus define the positive and finite random variable given as

the absorption time as follows:

$$Y = \inf\{t > 0 : J_t = p + 1\},$$

and say that it follows an IPH distribution with representation  $(\boldsymbol{\pi}, (\mathbf{T}(t))_{t \geq 0})$ . For practical applications, however, we focus on a narrower class which allows for effective statistical analysis. We make the following assumption regarding the inhomogeneity of the sub-intensity matrix:

$$\mathbf{T}(t) = \lambda(t) \mathbf{T},$$

with  $\lambda(t)$  some parametric and positive function such that the map:

$$y \mapsto \int_0^y \lambda(s) ds \in (0, \infty), \quad \forall y > 0,$$

converges to infinity as  $y \rightarrow \infty$ , and  $\mathbf{T}$  is a sub-intensity matrix that does not depend on time  $t$ . When dealing with particular entries of the sub-intensity matrix, we introduce the notation:

$$\mathbf{T} = (t_{kl})_{k,l=1,\dots,p}, \quad \mathbf{t} = -\mathbf{T}\mathbf{e} = (t_1, \dots, t_p)^\top, \quad \alpha_i = -t_{ii} > 0, \quad i = 1, \dots, p.$$

We then write

$$Y \sim \text{IPH}(\boldsymbol{\pi}, \mathbf{T}, \lambda),$$

where IPH stands for inhomogeneous phase type. Similarly, we write PH in place of phase type in the sequel, that is, for the case when  $\lambda$  is a constant.

IPH distributions are the building blocks to model claim severity for two reasons. Firstly, they have the interpretation of being absorption times of an underlying unobserved multistate process evolving through time, which can be very natural if the losses are linked with an evolving process, such as a lifetime or a legal case; and secondly, PH distributions are the archetype generalization of the exponential distribution in that many closed-form formulas are often available in terms of matrix functions, and thus, when armed with matrix calculus and a good linear algebra software, a suite of statistical tools can be developed explicitly.

It is worth mentioning that PH distributions possess a universality property, since they are known to be dense on the set of positive-valued random variables in the sense of weak convergence. Consequently, they possess great flexibility for atypical histogram shapes that the usual probabilistic models struggle to capture. Do note, however, that the tail can be grossly misspecified if no inhomogeneity function is applied, since weak convergence does not guarantee tail equivalence between the approximating sequence and the limit.

Another feature of IPH distributions is that they can have a different tail behavior than their homogeneous counterparts, allowing for non-exponential tails. To see this, first note that if  $Y \sim \text{IPH}(\boldsymbol{\pi}, \mathbf{T}, \lambda)$ , then we may write

$$Y \sim g(Z), \tag{2.1}$$

where  $Z \sim \text{PH}(\boldsymbol{\pi}, \mathbf{T})$  and  $g$  is defined in terms of  $\lambda$  by:

$$g^{-1}(y) = \int_0^y \lambda(s)ds, \quad y \geq 0,$$

which is well defined by the positivity of  $\lambda$ , or equivalently (see for instance Theorem 2.9 in Albrecher and Bladt, 2019):

$$\lambda(s) = \frac{d}{ds}g^{-1}(s).$$

Since the asymptotic behavior of PH distributions can be deduced from its Jordan decomposition, the following result in the inhomogeneous case is available. It is a generalization of the theory developed in Albrecher *et al.* (2021), and the proof is hence omitted.

**Proposition 2.1.** *Let  $Y \sim \text{IPH}(\boldsymbol{\pi}, \mathbf{T}, \lambda)$ . Then the survival function  $S_Y = 1 - F_Y(y)$ , density  $f_Y$ , hazard function  $h_Y$ , and cumulative hazard function  $H_Y$  of  $Y$  satisfy, respectively, as  $t \rightarrow \infty$ ,*

$$\begin{aligned} S_Y(y) &= \boldsymbol{\pi}^T \exp\left(\int_0^y \lambda(s)ds \mathbf{T}\right) \mathbf{e} \sim c_1 [g^{-1}(y)]^{n-1} e^{-\chi [g^{-1}(y)]}, \\ f_Y(y) &= \lambda(y) \boldsymbol{\pi}^T \exp\left(\int_0^y \lambda(s)ds \mathbf{T}\right) \mathbf{t} \sim c_2 [g^{-1}(y)]^{n-1} e^{-\chi [g^{-1}(y)]} \lambda(y), \\ h_Y(y) &\sim c \lambda(y), \\ H_Y(y) &\sim k g^{-1}(y), \end{aligned}$$

where  $c_1, c_2, c$ , and  $k$  are positive constants,  $-\chi$  is the largest real eigenvalue of  $\mathbf{T}$ , and  $n$  is the dimension of the Jordan block associated with  $\chi$ .

In particular, we observe that the tail behavior is determined almost entirely by function  $g$ . We also note that any model based on the intensity of the process (as will be the case for our regression model below) is asymptotically a model on the hazard of the distribution.

If  $g$  is the identity, we get Erlang tails, but for other choices, we move away from the exponential domain. The choice of  $g$  is typically determined according to the desired tail behavior required for applications. For heavy-tailed claim severities in the Fréchet max-domain of attraction, the exponential transformation  $g(z) = \eta (\exp(z) - 1)$ ,  $\eta > 0$  leading to Pareto tails is the simplest choice, although the log-logistic can also have such a behavior. When data are moderately heavy-tailed, as understood by being in the Gumbel max-domain of attraction, heavy-tailed Weibull  $g(z) = z^{1/\eta}$ ,  $\eta > 0$ , or Lognormal  $g(z) = \exp(z^{1/\gamma}) - 1$ ,  $\gamma > 1$  transforms can provide useful representations. In life insurance applications, the Matrix-Gompertz distribution with  $g(z) = \log(\eta z + 1)/\eta$ ,  $\eta > 0$  can be particularly useful to model human lifetimes.

One tool which is particularly useful to deal with expressions involving IPH distributions is matrix calculus. Here, the main tool is Cauchy’s formula for matrices, which is as follows. If  $u$  is an analytic function and  $A$  is a matrix, we define

$$u(A) = \frac{1}{2\pi i} \oint_{\Gamma} u(w)(wI - A)^{-1} dw, \quad A \in \mathbb{R}^{p \times p}.$$

where  $\Gamma$  is a simple path enclosing the eigenvalues of  $A$ , and  $I$  is the  $p$ -dimensional identity matrix. This allows for functions of matrices to be well defined, which in turn gives explicit formulas for IPH distributions, simplifying calculations and numerical implementation alike.

**Example 2.2.** Consider the following underlying Markov process parameters:

$$\boldsymbol{\pi} = (1, 0, 0)^T, \quad \mathbf{T} = \begin{pmatrix} -100 & 50 & 0 \\ 0 & -1 & 1/2 \\ 0 & 0 & -1/100 \end{pmatrix},$$

and a Weibull inhomogeneity parameter  $\eta = 8$ , which was chosen significantly light-tailed for visualization purposes. Compared to a conventional Weibull distribution, with  $\boldsymbol{\pi} = 1$ ,  $\mathbf{T} = -1/10$ , and  $\eta = 8$ , we see in Figure 1 that just a few additional parameters in the augmented matrix version of a common distribution can result in very different behavior in the body of the distribution (the tail behavior is guaranteed to be the same). For this structure, the case  $p = 3$  shows that the Markov process can traverse three states before absorption, resulting in three different modes of the density. For other matrix structures or parameters, there is in general no one-to-one correspondence between  $p$  and the number of modes of the resulting density.

### 2.2. The PH regression model

We now introduce a regression model based on IPH distributions and use it to provide a segmentation model for claim severities. Classically, for pure premium calculation, the mean is the main focus. Presently, we do not aim to improve on other data-driven procedures for pure mean estimation but instead focus on the entire distribution of the loss severities. That said, a consequence of a well-fitting global model is a reasonable mean estimate in the small sample case. However, when considering large sample behavior, any scoring rule in terms of means is not a proper scoring rule for target distributions that are not fully characterized by their first moment (see Gneiting and Raftery, 2007), and thus model choice is less important in this case.

Let  $\mathbf{X} = (X_1, \dots, X_d)^T$  be a  $d$ -dimensional vector of rating factors associated with the loss severity random variable  $Y$ . We keep the convention that  $\mathbf{X}$  does *not* contain the entry of 1 associated with an intercept in order to obtain an identifiable regression specification (the intercept is included in the

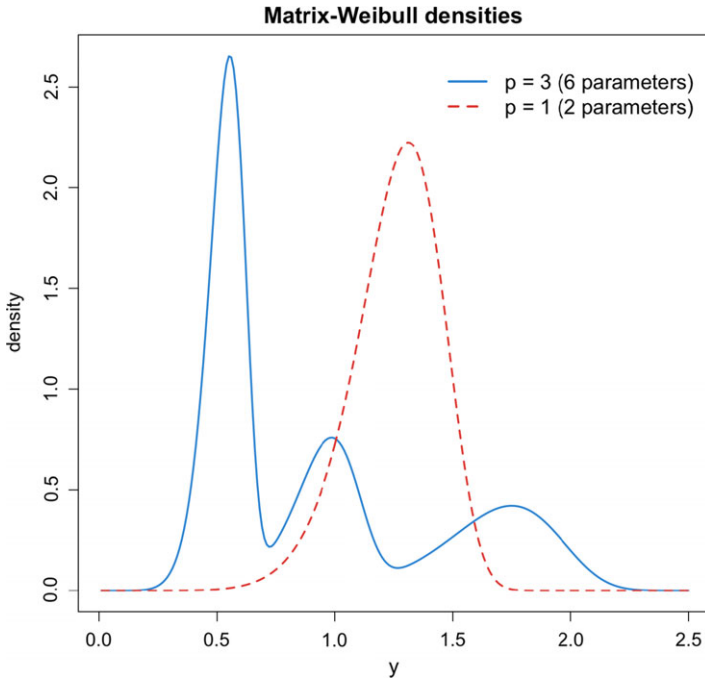


FIGURE 1: Densities corresponding to a scalar Weibull distribution and a Matrix-Weibull distribution.

matrix  $T$ ). Let  $\beta$  be a  $d$ -dimensional vector of regression coefficients so that  $X^T \beta$  is a scalar. Then, we define the conditional distribution of  $Y|X$  as having density function, see Proposition 2.1,

$$f(y) = m(X^T \beta) \lambda(y; \theta) \pi^T \exp \left( m(X^T \beta) \int_0^y \lambda(s; \theta) ds T \right) \mathbf{t}, \quad (2.2)$$

where  $m : \mathbb{R} \rightarrow \mathbb{R}_+$  is any positive function (monotonicity is not required). We call this the *phase-type regression model*. In particular, we recognize an IPH distribution  $\text{IPH}(\boldsymbol{\pi}, T, \lambda(\cdot | X, \beta, \theta))$ , where

$$\lambda(t | X, \beta, \theta) := \lambda(t; \theta) m(X^T \beta). \quad (2.3)$$

One possible interpretation of this model is that covariates act multiplicatively on the underlying Markov intensity, thus creating proportional intensities among different policies. In fact, since the intensity of an IPH distribution is asymptotically equivalent to its hazard function, we have that covariates satisfy a generalized proportional hazards specification in the far-right tail.

The conditional mean can be written in the form:

$$\mu(Y|X) = \int_0^\infty \pi^T \exp \left( m(X^T \beta) \int_0^y \lambda(s; \theta) ds T \right) \mathbf{e} dy, \quad (2.4)$$



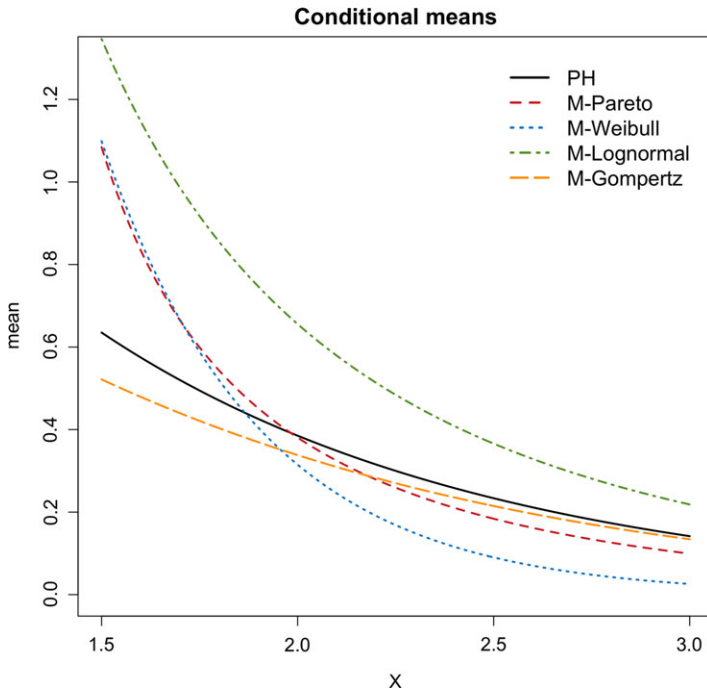


FIGURE 2: Mean functions of the PH regression model, as a function of a univariate regressor and for different specifications of the inhomogeneity function  $\lambda$ .

which is simply obtained by integrating the survival function over  $\mathbb{R}_+$ , see Proposition 2.1. A simple special case is obtained by the following choices, giving a Gamma GLM with canonical link: take  $T = -1$ , and  $\lambda \equiv 1$  to receive

$$\mu(Y|X) = \int_0^\infty \exp(-m(X^T \beta)y) dy = \frac{1}{m(X^T \beta)}.$$

Another slightly more complex special case is that of regression for Matrix-Weibull distributions, which contains the pure PH specification (when  $\lambda \equiv 1$ ). In this setting, it is not hard to see that

$$\mu(Y|X) = \int_0^\infty \pi^T \exp(m(X^T \beta)Ty^\theta) e dy = \frac{\Gamma(1 + \theta^{-1})\pi^T T^{-\theta^{-1}} e}{m(X^T \beta)^{\theta^{-1}}}. \tag{2.5}$$

Figure 2 shows (2.4) as a function of a unidimensional  $X$  for various inhomogeneity functions  $\lambda$ , and fixing the underlying Markov parameters at  $\pi = (1/2, 1/4, 1/4)^T$  and

$$T = \begin{pmatrix} -1 & 1/2 & 1/4 \\ 1/4 & -1 & 1/4 \\ 1/4 & 1/4 & -1 \end{pmatrix}.$$

### 2.3. Estimation

Claim severity fitting can be done by MLE in much the same way whether there are rating factors or not, and hence we exclusively present the more general case. For this purpose, we adopt a generalization of the EM algorithm, which we outline below. Such an approach is much faster to converge than naive multivariate optimization, provided a fast matrix exponential evaluation is available (see Moler and Van Loan, 1978).

A full account of the EM algorithm for ordinary PH distributions can be found in Asmussen *et al.* (1996) which led to the subsequent developments and extensions to censored and parameter-dependent transformations in Olsson (1996) and Albrecher *et al.* (2020b). Further details are provided in Appendix A, since the formulas therein are the building blocks of the EM algorithm for the PH regression model.

By defining  $g(\cdot; \theta)$  and  $g(\cdot | \mathbf{X}, \boldsymbol{\beta}, \theta)$  as follows in terms of their inverse functions  $g^{-1}(y; \theta) = \int_0^y \lambda(s; \theta) ds$  and

$$g^{-1}(y | \mathbf{X}, \boldsymbol{\beta}, \theta) = \int_0^y \lambda(s | \mathbf{X}, \theta) ds = g^{-1}(y; \theta) \exp(\mathbf{X}^T \boldsymbol{\beta}),$$

we note the simple identity:

$$g(y | \mathbf{X}, \boldsymbol{\beta}, \theta) = g(y \exp(-\mathbf{X}^T \boldsymbol{\beta}); \theta), \quad (2.6)$$

which yields that  $g^{-1}(Y | \mathbf{X}, \boldsymbol{\beta}, \theta) \sim \text{PH}(\boldsymbol{\pi}, \mathbf{T})$ . In other words, there is a parametric transformation, depending both on the covariates and regression coefficients, which brings the PH regression model into a conventional PH distribution. The generalized EM algorithm hinges on this fact and optimizes the transformed PH distribution and the transformation itself in alternating steps.

A standard fact of the above procedure, which is not hard to verify directly, is the following result.

**Proposition 2.3.** *The likelihood function is increasing at each iteration of Algorithm 1. Thus, for fixed  $p$ , since in that case the likelihood is also bounded, we obtain convergence to a (possibly local) maximum.*

**Remark 2.1 (Computational remarks).** *Algorithm 1 is simple to comprehend and very often will converge not only to a maximum but to a global maximum. However, its implementation can quickly turn prohibitively slow, if one fails to implement fast matrix analytic methods. The main difficulty is the computation of matrix exponentials, which appears in the E-step of step 1, and then repeatedly in the optimization of step 2. This problem is by no means new, with Moler and Van Loan (1978) already providing several options on how to calculate the exponential of a matrix.*

**Algorithm 1** Generalized EM algorithm for phase-type regression

**Input:** positive data points  $\mathbf{y} = (y_1, y_2, \dots, y_N)^T$ , covariates  $\mathbf{x}_1, \dots, \mathbf{x}_N$ , and initial parameters  $(\boldsymbol{\pi}, \mathbf{T}, \theta)$

- (1) *Transformation:* transform the data into  $z_i = g^{-1}(y_i; \theta) \exp(\mathbf{x}_i^T \boldsymbol{\beta})$ ,  $i = 1, \dots, N$ .
- (2) *E-step:* compute the statistics (see Appendix A for precise definitions of the random variables  $B_k, Z_k, N_{ks}$ , and  $N_k$ ):

$$\mathbb{E}(B_k | \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N \frac{\pi_k \mathbf{e}_k^T \exp(\mathbf{T}z_i) \mathbf{t}}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(Z_k | \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N \frac{\int_0^{z_i} \mathbf{e}_k^T \exp(\mathbf{T}(z_i - u)) \mathbf{t} \boldsymbol{\pi}^T \exp(\mathbf{T}u) \mathbf{e}_k du}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(N_{ks} | \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N t_{ks} \frac{\int_0^{z_i} \mathbf{e}_s^T \exp(\mathbf{T}(z_i - u)) \mathbf{t} \boldsymbol{\pi}^T \exp(\mathbf{T}u) \mathbf{e}_k du}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(N_k | \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N t_k \frac{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{e} y_k}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

- (3) *M-step:* let

$$\hat{\pi}_k = \frac{\mathbb{E}(B_k | \mathbf{Z} = \mathbf{z})}{N}, \quad \hat{t}_{ks} = \frac{\mathbb{E}(N_{ks} | \mathbf{Z} = \mathbf{z})}{\mathbb{E}(Z_k | \mathbf{Z} = \mathbf{z})}$$

$$\hat{t}_k = \frac{\mathbb{E}(N_k | \mathbf{Z} = \mathbf{z})}{\mathbb{E}(Z_k | \mathbf{Z} = \mathbf{z})}, \quad \hat{t}_{kk} = - \sum_{s \neq k} \hat{t}_{ks} - \hat{t}_k.$$

- (4) Maximize

$$\begin{aligned} (\hat{\theta}, \hat{\boldsymbol{\beta}}) &= \arg \max_{(\theta, \boldsymbol{\beta})} \sum_{i=1}^N \log(f_Y(y_i; \hat{\boldsymbol{\pi}}, \hat{\mathbf{T}}, \theta, \boldsymbol{\beta})) \\ &= \arg \max_{(\theta, \boldsymbol{\beta})} \sum_{i=1}^N \log \left( m(\mathbf{x}_i \boldsymbol{\beta}) \lambda(y; \theta) \boldsymbol{\pi}^T \exp \left( m(\mathbf{x}_i \boldsymbol{\beta}) \int_0^y \lambda(s; \theta) ds \mathbf{T} \right) \mathbf{t} \right) \end{aligned}$$

- (5) Update the current parameters to  $(\boldsymbol{\pi}, \mathbf{T}, \theta, \boldsymbol{\beta}) = (\hat{\boldsymbol{\pi}}, \hat{\mathbf{T}}, \hat{\theta}, \hat{\boldsymbol{\beta}})$ . Return to step 1 unless a stopping rule is satisfied.

**Output:** fitted representation  $(\boldsymbol{\pi}, \mathbf{T}, \theta, \boldsymbol{\beta})$ .

In Asmussen et al. (1996), matrix exponentiation was done for PH estimation by converting the problem into a system of ordinary differential equations (ODEs) and then subsequently solved by the Runge–Kutta method of order 4. The C implementation, called EMpht, is still available online, cf. Olsson (1998). Using the same ODE approach, Albrecher et al. (2021) implemented a C++ based algorithm available from the `matrixdist` package in R (cf. Bladt and Yslas, 2021). However, when extended to IPH distributions, this approach requires reorderings and lacks the necessary speed to estimate datasets of the magnitude of those found in insurance.

For this reason, we presently choose to make use of the uniformization method, which expresses the dynamics of a continuous-time Markov jump process in terms of a discrete-time chain where jumps occur at a Poisson intensity. More precisely, taking  $\phi = \max_{k=1,\dots,p} (-t_{kk})$  and  $\mathbf{Q} = \phi^{-1} (\phi \mathbf{I} + \mathbf{T})$ , which is a transition matrix, we have that

$$\exp(\mathbf{T}y) = \sum_{n=0}^{\infty} \frac{\mathbf{Q}^n (\phi y)^n}{n!} e^{-\phi y}.$$

and a truncated approximation of this series has error smaller than

$$\sum_{n=M+1}^{\infty} \frac{(\phi y)^n}{n!} e^{-\phi y} = \mathbb{P}(N_{\phi y} > M),$$

with  $N_{\phi y} \sim \text{Poisson}(\phi y)$ . In other words, with the introduction of the regularization parameter  $\phi$ , we are able to tame the error incurred in the truncation approximation by considering  $M$  large enough according to a Poisson law. Further improvements can be achieved by artifacts as the following one: since trivially  $e^{\mathbf{T}y} = (e^{\mathbf{T}y/2^m})^{2^m}$ , then the Poisson law with mean  $\phi y/2^m < 1$  implies good approximations with small  $M$  for  $e^{\mathbf{T}y/2^m}$  and then  $e^{\mathbf{T}y}$  is obtained by simple sequential squaring.

Integrals of matrix exponentials can be computed using an ingenious result of Van Loan (1978), which states that

$$\exp\left(\begin{pmatrix} \mathbf{T} & \mathbf{t} \boldsymbol{\pi} \\ \mathbf{0} & \mathbf{T} \end{pmatrix} y\right) = \begin{pmatrix} e^{\mathbf{T}y} & \int_0^y e^{\mathbf{T}(y-u)} \mathbf{t} \boldsymbol{\pi} e^{\mathbf{T}u} du \\ \mathbf{0} & e^{\mathbf{T}y} \end{pmatrix}.$$

Hence, integrals of the form  $\int_0^y e^{\mathbf{T}(y-u)} \mathbf{t} \boldsymbol{\pi} e^{\mathbf{T}u} du$  can easily be computed at a stroke by one single matrix exponential evaluation of the left-hand-side expression and extracting the corresponding upper-right sub-matrix.

Bundled together with a new `ReppArmadillo` implementation (which is now also part of the latest Github version of `matrixdist`), the increase in speed was of about two to three orders of magnitude. Consequently, larger datasets with more covariates can be estimated in reasonable time. Naturally, the methods are still slower than those for simple regression models such as GLMs.

**2.4. Inference and goodness of fit**

Inference and goodness of fit can always be done via parametric bootstrap methods. However, re-fitting a PH regression can be too costly. Consequently, we will take a more classical approach and derive certain quantities of interest which will allow us to perform variable selection and assess the quality of a fitted model.

*2.4.1. Inference for phase-type regression models*

For the following calculations, we focus only in the identifiable parameters of the regression, which are the relevant ones in terms of segmentation. Furthermore, we now make the assumption that  $m$  is differentiable and write  $\ell_y(\boldsymbol{\beta}, \theta)$  for the joint log-likelihood function of the observed severities  $\mathbf{y} = (y_1, \dots, y_N)$  with corresponding rating factors  $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ .

From (2.2), we obtain for  $j = 1, \dots, d$  and  $h(y; \theta) := \int_0^y \lambda(s; \theta) ds$ ,

$$\frac{d\ell_y(\boldsymbol{\beta}, \theta)}{d\beta_j} = \sum_{i=1}^N G_1(i, j | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}), \tag{2.7}$$

$$\frac{d\ell_y(\boldsymbol{\beta}, \theta)}{d\theta} = \sum_{i=1}^N G_2(i | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}), \tag{2.8}$$

where

$$G_1(i, j | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) = x_{ij} m'(\mathbf{x}_i^T \boldsymbol{\beta}) \left( \frac{1}{m(\mathbf{x}_i^T \boldsymbol{\beta})} + \frac{\boldsymbol{\pi}^T \exp(m(\mathbf{x}_i^T \boldsymbol{\beta}) h(y_i; \theta) \mathbf{T}) h(y_i; \theta) \mathbf{T} \mathbf{t}}{\boldsymbol{\pi}^T \exp(m(\mathbf{x}_i^T \boldsymbol{\beta}) h(y_i; \theta) \mathbf{T}) \mathbf{t}} \right),$$

$$G_2(i | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) = \frac{\frac{d}{d\theta} \lambda(y_i; \theta)}{\lambda(y_i; \theta)} + \frac{\boldsymbol{\pi}^T \exp(m(\mathbf{x}_i^T \boldsymbol{\beta}) h(y_i; \theta) \mathbf{T}) m(\mathbf{x}_i^T \boldsymbol{\beta}) \frac{d}{d\theta} h(y_i; \theta) \mathbf{T} \mathbf{t}}{\boldsymbol{\pi}^T \exp(m(\mathbf{x}_i^T \boldsymbol{\beta}) h(y_i; \theta) \mathbf{T}) \mathbf{t}}$$

Such expressions should be equal to 0 whenever convergence is achieved in step 2 of Algorithm 1.

In general, the score functions (2.7) and (2.8) will not have an explicit solution when equated to 0. Concerning the  $(d + 1) \times (d + 1)$ -dimensional Fisher information matrix, we may write it as:

$$[\mathcal{I}]_{jk} = \begin{cases} \sum_{i=1}^N G_1(i, j | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) G_1(i, k | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) & 1 \leq j, k \leq d, \\ \sum_{i=1}^N G_1(i, j | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) G_2(i | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) & 1 \leq j \leq d, k = d + 1, \\ \sum_{i=1}^N G_2^2(i | \boldsymbol{\pi}, \mathbf{T}, \boldsymbol{\beta}, \theta, \mathbf{y}, \bar{\mathbf{x}}) & j = k = d + 1. \end{cases} \tag{2.9}$$

Then, we may use the asymptotic MLE approximation, as  $N \rightarrow \infty$ ,

$$(\widehat{\boldsymbol{\beta}}, \widehat{\theta}) \stackrel{d}{\approx} \mathcal{N}((\boldsymbol{\beta}, \theta), \mathcal{I}^{-1}),$$

which can be used to compute standard errors and Neyman-type confidence intervals. For instance,

$$(\widehat{\beta}_1 - 1.96\sqrt{[\mathcal{I}^{-1}]_{11}}, \widehat{\beta}_1 + 1.96\sqrt{[\mathcal{I}^{-1}]_{11}}),$$

constitutes an approximate 95% confidence interval for the parameter  $\beta_1$ . We obtain  $p$ -values in much the same manner.

Observe that we have deliberately omitted inference for the PH parameters  $\boldsymbol{\pi}$  and  $\mathbf{T}$ . The reason being, this is a particularly difficult task due to the *non-identifiability* issue of PH distributions, namely that several representations can result in the same model. Papers such as Bladt *et al.* (2011) and Zhang *et al.*, 2021 provide methods of recovering the information matrix for these parameters, but the validity of the conclusions drawn from such approach are not fully understood. Consequently, and similarly to other regression models, we will only perform estimation on the regression coefficients themselves and use the distributional parameters merely as a vehicle to obtain  $\boldsymbol{\beta}$ .

**Remark 2.2.** *We performed estimation using numerical optimization, with and without using the gradient of the log-likelihood for step 2 in Algorithm 1. Not using the gradient was always faster, since when dealing with matrix exponentials the evaluation of the derivatives is equally costly as the objective function itself, and thus best avoided. In practice, the Fisher information matrix may be near-singular when large numbers of covariates are considered. The negative of the the Hessian matrix obtained by numerical methods such as the Broyden-Fletcher-Goldfarb-Shanno algorithm is a more reliable alternative in this case, whose inverse may also be used to approximate the asymptotic covariance matrix.*

2.4.2. *Goodness of fit for phase-type regression models*

Below is a brief description of a goodness-of-fit diagnostic tool for the above models, which, when ordered and plotted against theoretical uniform order statistics, constitute a PP plot. The idea is to transform the covariate-specific data into a parameter-free scale using the probability integral transformation (PIT).

Indeed, by applying the PIT transform to the data:

$$r_i = \left( \boldsymbol{\pi}^T \exp \left( m \left( \mathbf{x}_i^T \boldsymbol{\beta} \right) \int_0^{y_i} \lambda(s; \theta) ds \mathbf{T} \right) \mathbf{e} \right), \quad i = 1, \dots, N. \tag{2.10}$$

The null hypothesis under the PH regression model is

$$\mathcal{H}_0 : Y_i \sim \text{IPH} \left( \boldsymbol{\pi}, \mathbf{T}, m \left( \mathbf{x}_i^T \boldsymbol{\beta} \right) \lambda \right), \quad i = 1, \dots, N,$$

and it standard see that the dataset:

$$\mathcal{D} = \{r_1, r_2, \dots, r_N\}$$

constitutes a sample of uniform variables.

Thus, it only remains to assess the goodness of fit of  $\mathcal{D}$  against a standard uniform distribution, for which a suite of statistical tests and visual tools can be applied.

### 2.5. On the choice of matrix dimension and structure

Often a general sub-intensity matrix structure is too general to be either practical or parsimonious, and special structures can do the job almost as well with fewer parameters (see, e.g., Bladt and Nielsen, 2017, Section 3.1.5, and also Section 8.3.2). This is because IPH distributions are not identifiable: two matrices of different dimension and/or structure may yield exactly the same distribution.

A zero in the off-diagonal of the sub-intensity matrix  $T$  indicates absence of jumps between two states, whereas a zero in the diagonal is not possible. Below we describe the most commonly used special structures for  $T$  (and its corresponding  $\pi$ ). Note that we borrow the names of the structures from the homogeneous case, but the resulting distribution will in general not be linked to such name. For instance, a Matrix-Pareto of exponential structure will yield a Pareto distribution, and not an exponential one.

*Exponential structure.* This is the simplest structure one can think of and applies only for  $p = 1$ , that is, the scalar case. We have  $\pi^T = 1$ ,  $T = -\alpha < 0$ , and so  $t = \alpha$ . The corresponding stochastic process is schematically depicted in the top-left panel of Figure 3.

*Erlang structure.* This structure traverses  $p$  identical states consecutively from beginning to end, or in matrix notation:

$$\pi^T = (1, 0, \dots, 0),$$

$$T = \begin{pmatrix} -\alpha & \alpha & & 0 \\ & \ddots & \ddots & \vdots \\ & & -\alpha & \alpha \\ 0 & & & -\alpha \end{pmatrix}, \quad t = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \alpha \end{pmatrix}, \quad \alpha > 0.$$

The corresponding stochastic process is schematically depicted in the top-right panel of Figure 3.

*Hyperexponential PH distribution.* This structure will always give rise to a mixture of scalar distributions, since it starts in one of the  $p$  states but has no

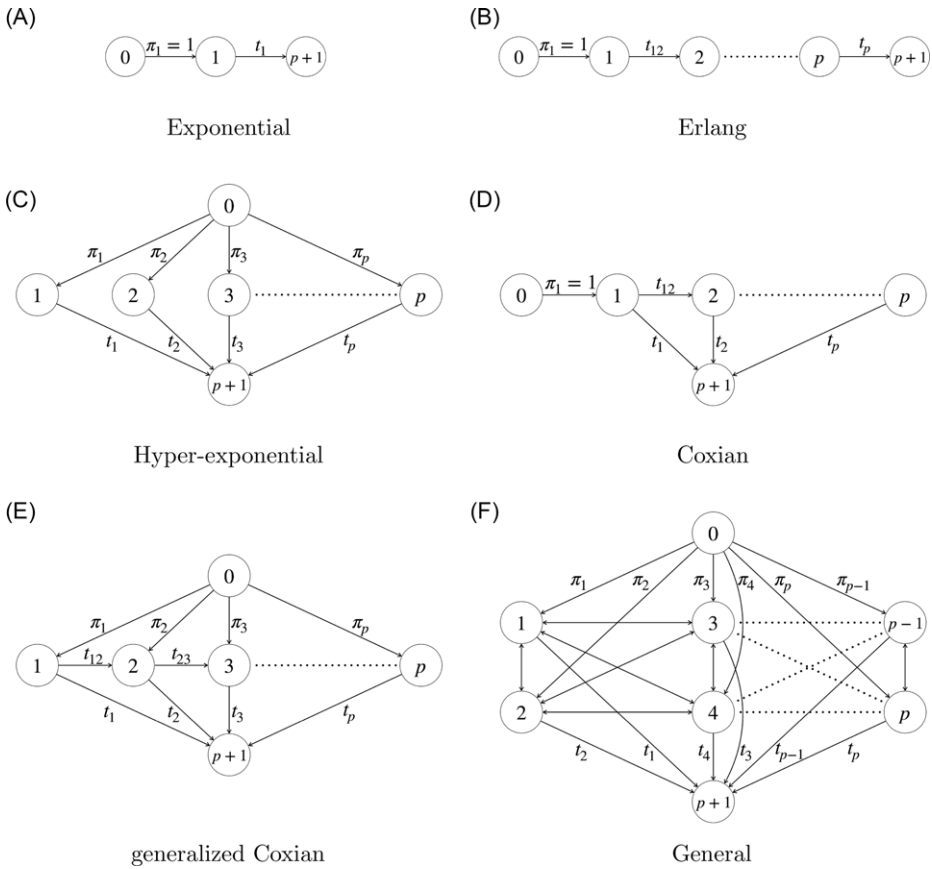


FIGURE 3: Underlying Markov structures. Names are borrowed from the corresponding PH representations but apply to our inhomogeneous setup as well. The state 0 is added for schematic reasons but is not part of the actual state space of the chain. The (F) general case has the intensities  $t_{ij}$  and  $t_{ji}$  between each pair of states  $i, j \in \{0, \dots, p\}$  omitted for display purposes.

transitions between them thereafter. The matrix representation is the following:

$$\boldsymbol{\pi}^T = (\pi_1, \pi_2, \dots, \pi_p),$$

$$\mathbf{T} = \begin{pmatrix} -\alpha_1 & & 0 \\ & \ddots & \\ 0 & & -\alpha_p \end{pmatrix}, \quad \mathbf{t} = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_p \end{pmatrix}, \quad \alpha_i > 0, \quad i = 1, \dots, p.$$

The corresponding stochastic process is schematically depicted in the center-left panel of Figure 3.

*Coxian structure.* This structure is a generalization of the Erlang one, where different parameters are allowed for the diagonal entries, and also absorption



can happen spontaneously from any intermediate state. The parametrization is as follows:

$$\boldsymbol{\pi}^T = (1, 0, \dots, 0),$$

$$\mathbf{T} = \begin{pmatrix} -\alpha_1 & q_1\alpha_1 & & 0 \\ & -\alpha_2 & q_2\alpha_2 & \\ & & \ddots & \ddots \\ & & & -\alpha_{p-1} & q_{p-1}\alpha_{p-1} \\ 0 & & & & -\alpha_p \end{pmatrix}, \quad \mathbf{t} = \begin{pmatrix} (1 - q_1)\alpha_1 \\ (1 - q_2)\alpha_2 \\ \vdots \\ (1 - q_{p-1})\alpha_{p-1} \\ \alpha_p \end{pmatrix},$$

with  $\alpha_i > 0$ ,  $q_i \in [0, 1]$ ,  $i = 1, \dots, p$ . The corresponding stochastic process is schematically depicted in the center-right panel of Figure 3.

*Generalized Coxian structure.* This structure has the same  $\mathbf{T}$  and  $\mathbf{t}$  parameters as in the Coxian case, the only difference being that the initial vector is now distributed among all  $p$  states, that is,

$$\boldsymbol{\pi}^T = (\pi_1, \pi_2, \dots, \pi_p).$$

The corresponding stochastic process is schematically depicted in the bottom-left panel of Figure 3.

**Remark 2.3.** *One convenient feature of having zeros in a sub-intensity matrix is that the EM algorithm will forever keep those zeros in place, effectively searching only within the subclass of IPH distributions of a given special structure. A prevalent observation which has nearly become a rule of thumb is that Coxian structures perform nearly as well as general structures and are certainly much faster to estimate.*

**Remark 2.4.** *When dealing with matrices whose entries are the parameters of the model, the usual systematic measures such as Akaike’s Information Criteria (AIC) or Bayesian Information Criteria (BIC) are overly conservative (yielding too simple models), and the appropriate amount of penalization is a famous unsolved problem in the PH community, and out of the scope of the current work. The selection of the number of phases and the special substructure of such matrix is commonly done by trial and error, much in the same manner as is done for the tuning of hyperparameters of a machine learning method, or the selection of the correct combination of covariates in a linear regression. Such approach is also taken presently. Despite this disadvantage, there is one advantage from a purely statistical standpoint, which is that standard errors (and thus significance of regression coefficients) arising from asymptotic theory are more truthful if no automatic selection procedure is used.*

*Hence, when selecting between fitted dimensions and structures, it is advised to keep track of the negative log-likelihood (rather than AIC or BIC) and its decrease when increasing the number of phases or changing the structure. Theoretically, larger matrices will always yield better likelihood, but in practice,*

it is often the case that the likelihood stops increasing (effectively getting stuck in a good local optimum), or increases are not so significant. Usually, EM algorithms for PH distributions for marginal distributions estimate effectively and in reasonable time up to about 20–30 phases, while for a PH regression the number is closer to 5–10. The additional inhomogeneity function of IPH distributions, however, makes models more parsimonious than the PH counterparts, and usually less than 10 phases are needed.

Once the dimension and structure is chosen, one can perform model selection with respect to the regression coefficients. These fall into the usual framework, and the AIC or BIC criteria can then be used to compare and select between various models.

### 3. A SIMULATION STUDY

Before applying the above models to a real-life insurance dataset, we would first like to confirm that PH regression is an effective tool for estimating data exhibiting features which are common in insurance. Namely, claim severity can exhibit multimodality and different behavior in the body and tails of the distribution, the latter being heavy-tailed in some lines of business. In order to create synthetic data with these features, we take the following approach.

We simulate  $N = 1000$  times from a bivariate Gaussian copula  $\mathbf{X} = (X_1, X_2)^T$  with correlation coefficient 0.7, which will play the role of sampling from a bivariate feature vector with uniform marginals. Then, we create three mean specifications based on the first entry of such vector:

$$\mu_1 = \exp(X_1), \quad \mu_2 = \exp(3 + X_1), \quad \mu_3 = \exp(X_1 - 1),$$

in a sense individually specifying a log-link function. We finally sample with probabilities 0.4, 0.4, and 0.2 from three different regression specifications:

$$\begin{aligned} &\text{Gamma}(\mu = \mu_1, \phi = 1), \quad \text{Gamma}(\mu = \mu_2, \phi = 1), \\ &\exp(\text{Gamma}(\mu = \mu_3, \phi = 1)). \end{aligned}$$

The data are depicted in Figure 4.

Observe that the third component implies Pareto-type tails, although this component has the smallest probability of occurrence. We also remark that the true driver of the regression is  $X_1$ , and thus we would ideally like to obtain the nonsignificance of  $X_2$  from the inference.

Next, consider four model specifications:

- (1) a Gamma GLM with log-link function and considering  $X_1$  as covariate.
- (2) a Gamma GLM with log-link function and considering  $X_1$  and  $X_2$  as covariates.
- (3) a Matrix-Pareto PH regression of Coxian type and dimension 3, considering  $X_1$  as covariate.

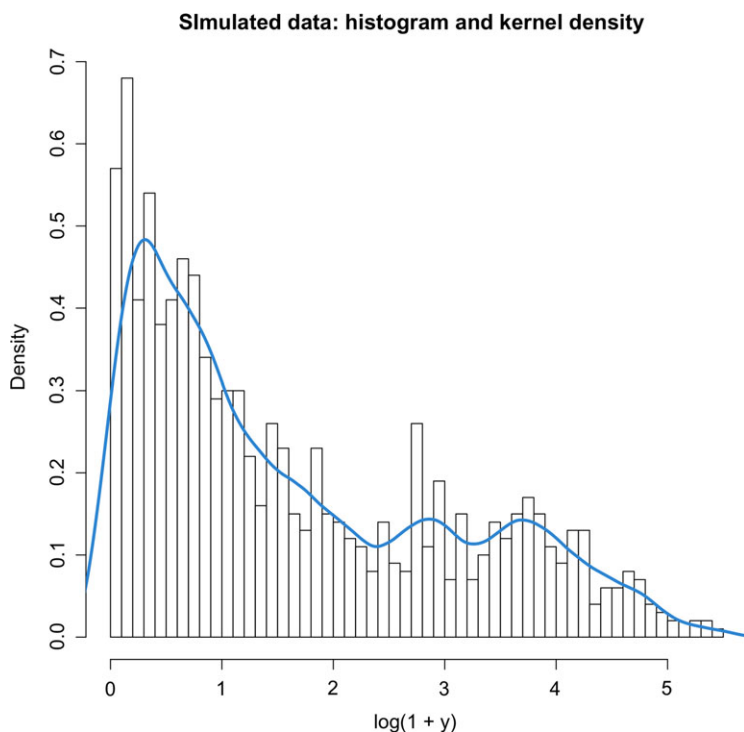


FIGURE 4: Histogram and kernel density estimate of the log-transformed simulated data.

- (4) a Matrix-Pareto PH regression of Coxian type and dimension 3, considering  $X_1$  and  $X_2$  as covariates.

The analysis was carried using the `matrixdist` package in R, Bladt and Yslas (2021). The results are given in Table 1 (observe that the IPH models do not have intercept since it is included in their sub-intensity matrices), showing good segmentation for the matrix-based methods. Uninformative covariates can play the role of mixing in the absence of a correctly specified probabilistic model, and thus we see that  $X_2$  is significant for the GLM model, whereas this is no longer the case for the PH regression. The AIC and BIC select the overall best model to be the Matrix-Pareto with only  $X_1$  as rating factor. In terms of goodness of fit, we visually confirm from Figure 5 that the distributional features of the data are captured substantially better by considering an underlying Markov structure.

#### 4. APPLICATION TO A THE FRENCH MOTOR PERSONAL LINE DATASET

In this section, we consider a real-life insurance dataset: the publicly available French Motor Personal Line dataset. We apply the IPH and PH regression

TABLE 1  
GLMs AND PH REGRESSION MODELS.

	Gamma GLM	Gamma GLM	M-Pareto(3)	M-Pareto(3)
Intercept	1.936*** (0.127)	2.054*** (0.121)		
$X_1$	0.762*** (0.280)	1.138*** (0.2105)	-0.813*** (0.187)	-1.039*** (0.147)
$X_2$	-0.586* (0.273)		-0.266 (0.181)	
AIC	6407	6411	6100	<b>6098</b>
BIC	6426	6426	6139	<b>6133</b>
Log-likelihood	-3199	-3202	-3042	-3042
Degrees of freedom	4	3	8	7
Num. obs.	1000	1000	1000	1000

The smallest AIC and BIC values are given in boldface.

\*\*\* $p < 0.001$ ; \*\* $p < 0.01$ ; \* $p < 0.05$




FIGURE 5: Ordered PITs from equation (2.10) versus uniform order statistics for the simulated dataset. KS refers to the Kolmogorv–Smirnov statistic for testing uniformity.

estimation procedures to the distribution of claim severities, illustrating how the theory and algorithms from the previous sections can be effectively applied to model claim sizes.

We consider jointly the four datasets `fremMPL1`, `fremMPL2`, `fremMPL3`, and `fremMPL4`, from the `CASdatasets` package in R. These data describe claim frequencies and severities pertaining to four different coverages in a French motor personal line insurance for about 30,000 policies in the year 2004. The data have 18 covariates, which are described in the Supplementary Material. The 7008 observations with positive claim severity possess a multimodal density, arising from mixing different types of claims.

### 4.1. Marginal estimation

In the first step, we estimate the marginal behavior of the data, without rating factors. The model we choose to estimate is an IPH distribution with the exponential transform, namely, a Matrix-Pareto distribution. This choice is motivated by a preliminary analysis of the data by which a Hill plot indicated the presence of regularly varying tails, indicating heavy-tailedness. Since the data on the log-scale have a pronounced hump in the middle, a dimension larger than 1 or 2 is needed, so we chose 5. Finally, since there are no particular specificities in the right or left tail, a Coxian sub-matrix structure can do an equally good job as a general structure, and thus we select the former in the name of parsimony.

We also estimate a 20-dimensional Matrix-Weibull distribution on the data to illustrate the denseness property of IPH distributions. The Matrix-Weibull distribution generates modes more easily than the Matrix-Pareto, and we choose such a high dimension in order to illustrate the denseness of IPH distributions. However, this model is over-parametrized and thus mainly of academic interest.

As the first reference, we also fit a two-parameter Gamma distribution to the data. We also consider a splicing (or composite) model, cf. Albrecher *et al.* (2017) and Reynkens *et al.* (2017) (see also Miljkovic and Grün, 2016), defined as follows. Let  $f_1$  and  $f_2$  be two densities supported on the positive real line. A spliced density based on the the latter components is then given by:

$$f_S(x) = v \frac{f_1(x)}{F_1(t)} 1_{(0,t]}(x) + (1 - v) \frac{f_2(x)}{1 - F_2(t)} 1_{(t,\infty)}(x), \quad t > 0, \quad v \in [0, 1],$$

where  $t$  is the splicing location. We consider a popular implementation found in the `ReIns` package in R, Reynkens and Verbelen (2020), where

$$f_1(x) = \sum_{j=1}^M \omega_j \frac{x^{r_j-1} \exp(-x/\zeta)}{\zeta^{r_j} (r_j - 1)!}, \quad x > 0, \quad f_2(x) = \frac{1}{vt} \left(\frac{x}{t}\right)^{-\frac{1}{v}-1}, \quad x > t,$$

are, respectively, a ME densities with common scale parameter and a Pareto density, which we call the ME-Pareto specification.

TABLE 2

SUMMARY OF FITTED MARGINAL MODELS TO SEVERITIES FROM THE FREMPL DATASET.

	Gamma	Spliced ME-Pareto	Matrix-Pareto	Matrix-Weibull
Log-likelihood	-60,653	-59,611	-59,605	-59,167
Degrees of freedom	2	12	10	404
AIC	121,311	119,247	119,231	<b>119,142</b>
BIC	121,325	119,330	<b>119,299</b>	121,912
Num. obs.	7008	7008	7008	7008

The smallest AIC and BIC values are given in boldface.

The parameter estimates are the following for the Gamma distribution, found by MLE,

$$\text{Gamma shape} = 0.75, \quad \text{Gamma scale} = 2925.67,$$

while for the splicing model we obtain, selecting  $\nu$  through EVT techniques (Hill estimator) and the remaining parameters through MLE (using an EM algorithm), and choosing  $M$  from 1 to 10 according to the best AIC:

$$M = 5, \quad \zeta = 363.7, \quad r = (1, 2, 4, 11, 21),$$

$$\omega = (0.26, 0.07, 0.53, 0.09, 0.02), \quad \nu = 0.56, \quad \nu = 0.96,$$

and  $t \approx 10,000$  was selected by visual inspection according to when the Hill plot stabilizes for the sample. Finally, for the IPH Matrix-Pareto distribution, we obtain the following estimates by applying Algorithm 1 without covariates (which is implemented in the `matrixdist` package in R, Bladt and Yslas, 2021):

$$\pi = e_1, \quad T = \begin{pmatrix} -12.61 & 12.48 & 0 & 0 & 0 \\ 0 & -12.61 & 10.33 & 0 & 0 \\ 0 & 0 & -1.99 & 1.99 & 0 \\ 0 & 0 & 0 & -7.34 & 7.34 \\ 0 & 0 & 0 & 0 & -7.34 \end{pmatrix}, \quad \eta = 1149.57,$$

where  $e_1 = (1, 0, 0, 0, 0)^T$ . The 404 fitted parameters of the Matrix-Weibull, most of the zero, are omitted, since they do not provide further intuition on IPH fitting.

Table 2 provides numerical results on the three fitted models, and in particular it shows that an IPH Matrix-Pareto distribution is well aligned with the state-of-the-art models for severity modeling in insurance such as composite methods. The AIC and BIC are not good measures for IPH distributions, since they can over-penalize them when the underlying sub-intensity matrix is not of minimal order. However, we see that even with this hinderance, the performance in terms of information criteria is positive.

The quality of the fit is also assessed in Figure 6. In particular, we observe that the data have a big spike at the log-severity around the value 7.3, which

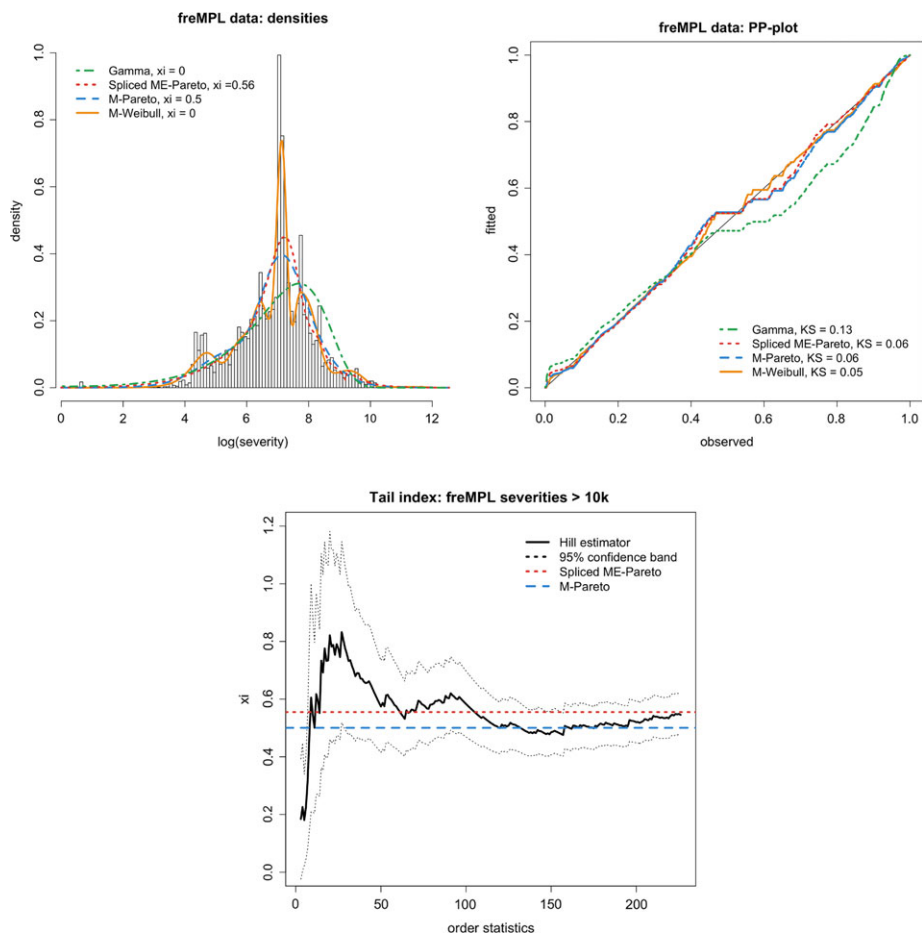


FIGURE 6: Fitted distributions to claim severities from the MPL dataset. The log-transform is used exclusively for visual purposes, the estimation having been carried out in the usual scale. KS refers to the Kolmogorv–Smirnov statistic for testing uniformity.

poses problems to the first three distributions. This disturbance manifests itself in the PP plot as a slight “S” shape. Not surprisingly, the visual quality of the fit of the splicing and IPH models is, as expected, far superior to its Gamma counterpart, while the difference between the two Pareto-tailed models is not very noticeable. Their implied tail indices agree, as given respectively by the formulas  $\hat{\xi} = \nu = 0.56$ , and

$$\hat{\xi} = -1 / \max\{\Re\text{Eigen}(\hat{T})\} = 0.50,$$

where  $\Re\text{Eigen}(A)$  denotes the set of real eigenvalues of a matrix  $A$ . For reference, recall that in terms of the Pareto tail parameter, the following relationship holds  $\alpha = 1/\xi$ . As a sanity check, we see that both estimates for the tail index

TABLE 3  
SUMMARY STATISTICS OF THE FRENCH MPL DATASET: CLAIM SEVERITY.

Variable	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
ClaimAmount	2176.9	5807.2	0.4	500.4	1204	2163.2	163,427

TABLE 4  
SUMMARY FOR GLM AND PH REGRESSION MODELS FOR THE FREMPL DATASET.

	Gamma GLM	Pareto PH regression	Weibull PH regression
Log-likelihood	-60,368	-59,464	-59,446
Degrees of freedom	26	34	34
AIC	120,788	118,996	<b>118,961</b>
BIC	120,966	119,229	<b>119,194</b>
Num. obs.	7008	7008	7008
Loss ratio (pure)	101.03%	105.18%	101.13%

The smallest AIC and BIC values are given in boldface.

fall perfectly within the Hill estimator's bounds. In particular, the estimated models have finite mean (their tail index is less than 1; alternatively, their Pareto index is larger than 1). We would like to remark that fitting a conventional Pareto distribution as a global model will result in a particularly bad fit, such that in the IPH case we have managed use matrix parameters to our advantage, in order to automatically engineer an all-inclusive model, featuring no threshold selection.

## 4.2. Incorporating rating factors

Having already analyzed the marginal behavior of the French MPL data, we proceed to incorporate rating factors through the PH regression model that we have presented above. In Table 3, we show summarizing statistics for claim severity, while two further tables of summarizing statistics for continuous rating factors and categorical rating factors, respectively, can be found in Supplementary Material.

To reduce the large number of categories and variables, we perform a pre-processing step where we individually, for each categorical variable, prune a regression tree according to the best predictive performance on the mean of the log-severities (since the variables exhibit heavy-tails). We then merge categories together according to the resulting optimal tree. For all the continuous variables, we perform a simple shift-and-scale transformation.

The Matrix-Pareto is the only distribution which according to the Hill estimator will have the correct tail behavior, and thus we consider it as a good target model for PH regression. However, we also consider a model which does not have the correct tail behavior, namely the Matrix-Weibull distribution. The



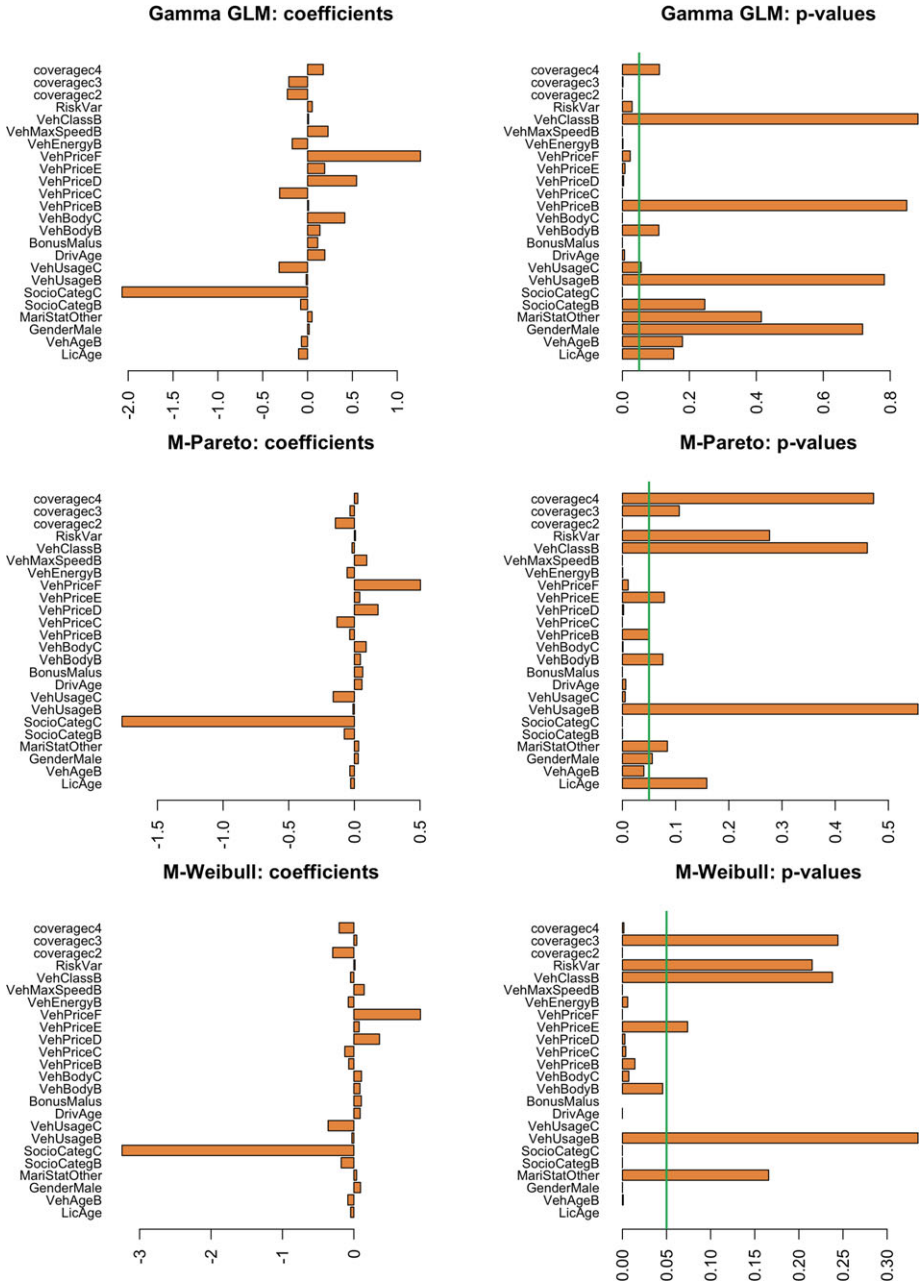


FIGURE 7: Coefficients and p-values of IPH and GLM regression. For display: IPH coefficients multiplied by  $-1$  and intercept of GLM omitted.




FIGURE 8: Left panels: aggregate observed losses versus aggregate implied model premia (expected value), normalized to sum to 1, for the Gamma GLM, and Pareto and Weibull PH regressions; right panels: number of claims within each category (right).

latter is convenient since it has a closed-form formula for mean prediction as per Equation (2.5).

The result of the fit using all the processed rating factors is given in Table 4 and Figure 7, with all the coefficients of the five-dimensional IPH models multiplied by  $-1$ , to be comparable with the log-link Gamma GLM. The significance is obtained from a normal approximation using the implied Fisher matrix based on the Hessian matrix from the numerical optimization, as detailed in the goodness-of-fit subsection above. When considering less covariates, Equation (2.9) can be equivalently used, yielding very close results. The AIC and BIC are well below what the corresponding GLM can achieve, even if this is not a generous metric for IPH models.

We observe that the loss ratio (where we understand as losses the claim severities, i.e., disregarding frequencies) is kept much better for the Weibull variant than for the Pareto, despite the latter being a better model for large claims. In practice, this inconvenience can be amended by regressing the losses with respect to the premium. Since estimation is done via MLE, it is not uncommon for a model with misspecified tail behavior to globally perform better. The fact that the Weibull PH regression performs almost as well as the

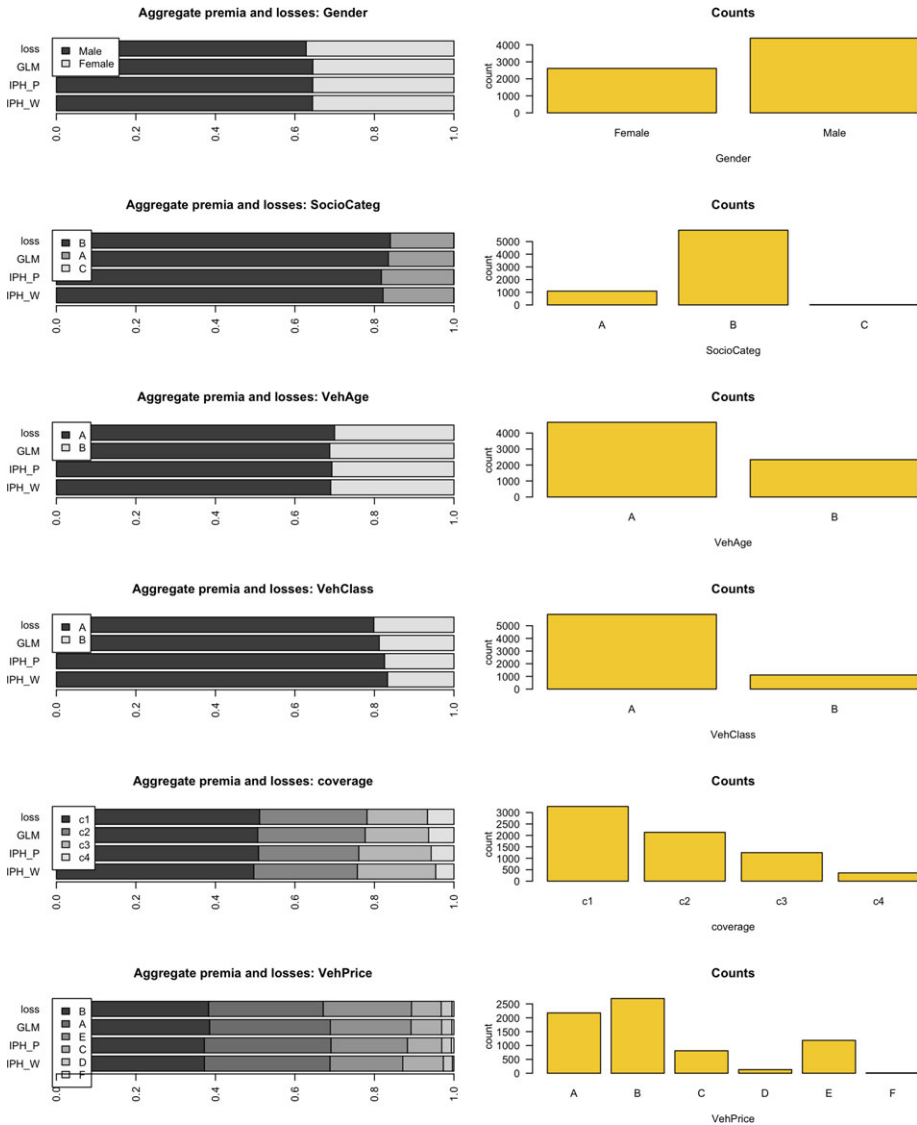


FIGURE 9: Empirical quantiles by coverage category versus mean (across all other covariates) quantiles implied by the Gamma GLM, and Pareto and Weibull PH regressions.

GLM when predicting the mean is remarkable, since the former model does not specifically target averages, and the latter does.

In Figure 8, we illustrate the behavior of the aggregation of the implied mean predictions across all policies. We consider prediction from all three models, for selected rating factors and their categories, as computed by the general formula (2.4) in the Pareto case, and the explicit formula (2.5) in the

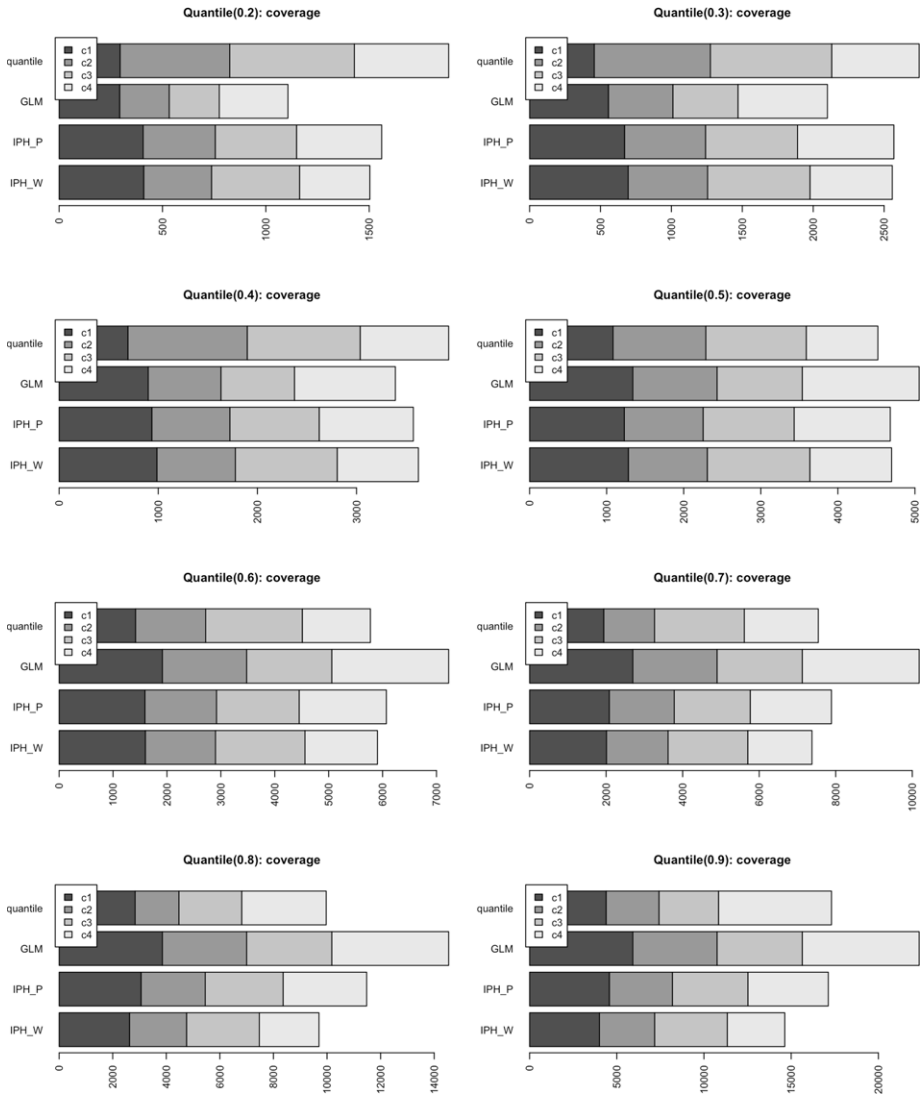


FIGURE 10: Ordered PITs from Equation (2.10) versus uniform order statistics for the French MPL dataset. KS refers to the Kolmogorv–Smirnov statistic for testing uniformity.

Weibull case. Each of the aggregate predictions are normalized to sum to 1 for display purposes. We again see that despite the Weibull PH regression not being specifically designed for such task, it comes very close to the performance of the GLM.

In Figure 9, we illustrate the prediction of quantiles, as implied by the fitted conditional distribution functions of each model. We average the quantile estimates across policies corresponding to each of the rating factors and their

categories. We observe that matrix distributions perform much better in this respect, and the correct specification of the Pareto tail can be appreciated in the bottom right panel, corresponding to the quantile at level 0.9. Figure 10 further supports the global estimation improvement in terms of the PP plot goodness-of-fit diagnostic.

In terms of out-of-sample predictive performance, the mean square error (MSE) when holding out the first 20% of the observations and training on the remaining 80% is given by<sup>1</sup> 3.66, 3.69, and 3.70, for the Gamma GLM, Pareto, and Weibull PH regression models, respectively. In general, if the risk manager is solely interested in mean prediction, our experiments suggest that other data-driven methods may outperform PH regression models, although the advantage is usually small. However, the tables turn when considering the entire distribution of loss severities. Thus, a key takeaway here is that PH regression models should be used when the aim is to capture the entire distributional properties of claim severities correctly. In contrast, models specifically designed for mean prediction may be more appropriate if only the average is of interest.

## 5. CONCLUSION

We have presented a novel claim severities model based on PH distributions, which implicitly assumes an underlying and unobservable multistate Markov structure. The inhomogeneity function is the key ingredient for translating the exponential tails of PH distributions to other tail domains, which is particularly useful for modeling of insurance data. We have shown that the use of a generalized EM algorithm performs effective estimation on the marginal distribution of severities, both with and without rating factors, the former being backed up by a simulation study. The flexibility of the PH regression model is particularly advantageous when data are multimodal, heavy- or light-tailed, and generally more heterogeneous than what the classical regression methods require. The practical implementation of PH, IPH, and PH regression can be carried out using the `matrixdist` (Bladt and Yslas, 2021) package in R.

Several questions remain open for further research, for instance, the incorporation of multi-parameter PH regression models, or using IPH models to describe data exhibiting incurred but not reported claims. A faster implementation of the generalized EM algorithms is needed, mainly in the case when numerous categorical rating factors are present, for widespread and systematic use of the PH regression model. Alternative estimation methods have not been explored and could prove competitive to MLE.

## ACKNOWLEDGMENT

MB would like to acknowledge financial support from the Swiss National Science Foundation Project 200021\_191984.

## CONFLICTS OF INTEREST

MB declares no conflict of interest related to the current manuscript.

## SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit <https://doi.org/10.1017/asb.2021.40>.

## REFERENCES

- AHN, S., KIM, J.H. and RAMASWAMI, V. (2012) A new class of models for heavy tailed distributions in finance and insurance risk. *Insurance: Mathematics and Economics*, **51**(1), 43–52.
- ALBRECHER, H., BEIRLANT, J. and TEUGELS, J.L. (2017) *Reinsurance: Actuarial and Statistical Aspects*, Chichester: John Wiley & Sons.
- ALBRECHER, H. and BLADT, M. (2019) Inhomogeneous Phase-type distributions and heavy tails. *Journal of Applied Probability*, **56**(4), 1044–1064.
- ALBRECHER, H., BLADT, M. and BLADT, M. (2020a) Matrix Mittag–Leffler distributions and modeling heavy-tailed risks. *Extremes*, **23**, 425–450.
- ALBRECHER, H., BLADT, M., BLADT, M. and YSLAS, J. (2021) Mortality modeling and regression with matrix distributions. arXiv:2011.03219.
- ALBRECHER, H., BLADT, M. and YSLAS, J. (2020b) Fitting inhomogeneous phase-type distributions to data: The univariate and the multivariate case. *Scandinavian Journal of Statistics*, 1–34.
- ASMUSSEN, S. (2008) *Applied Probability and Queues*, Vol. 51. New York: Springer Science & Business Media, New York.
- ASMUSSEN, S., NERMAN, O. and OLSSON, M. (1996) Fitting phase-type distributions via the EM algorithm. *Scandinavian Journal of Statistics*, **23**, 419–441.
- BLADT, M., ESPARZA, L.J.R. and NIELSEN, B.F. (2011) Fisher information and statistical inference for Phase-type distributions. *Journal of Applied Probability*, **48**(A), 277–293.
- BLADT, M. and NIELSEN, B.F. (2017) *Matrix-Exponential Distributions in Applied Probability*. New York: Springer.
- BLADT, M. and ROJAS-NANDAYAPA, L. (2017) Fitting Phase-type scale mixtures to heavy-tailed data and distributions. *Extremes*, **21**(2), 285–313.
- BLADT, M. and YSLAS, J. (2021) *matrixdist: Statistics for Matrix Distributions*. R package version 1.1.2.
- FUNG, T.C., BADESCU, A.L. and LIN, X.S. (2020) A new class of severity regression models with an application to IBNR prediction. *North American Actuarial Journal*, **25**(2), 206–231.
- FUNG, T.C., TZOUGAS, G. and WUTHRICH, M. (2021) Mixture composite regression models with multi-type feature selection. arXiv preprint arXiv:2103.07200.
- GNEITING, T. and RAFTERY, A.E. (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, **102**(477), 359–378.
- GRÜN, B. and MILJKOVIC, T. (2019) Extending composite loss models using a general framework of advanced computational tools. *Scandinavian Actuarial Journal*, **2019**(8), 642–660.
- LEE, S.C. and LIN, X.S. (2010) Modeling and evaluating insurance losses via mixtures of Erlang distributions. *North American Actuarial Journal*, **14**(1), 107–130.
- MCGRORY, C.A., PETTITT, A.N. and FADDY, M.J. (2009) A fully bayesian approach to inference for Coxian phase-type distributions with covariate dependent mean. *Computational Statistics & Data Analysis*, **53**(12), 4311–4321.

MILJKOVIC, T. and GRÜN, B. (2016) Modeling loss data using mixtures of distributions. *Insurance: Mathematics and Economics*, **70**, 387–396.

MOLER, C. and VAN LOAN, C. (1978) Nineteen dubious ways to compute the exponential of a matrix. *SIAM Review*, **20**(4), 801–836.

NEUTS, M.F. (1975) Probability distributions of Phase type. In *Liber Amicorum Professor Emeritus H. Florin*, pp. 173–206. Department of Mathematics, University of Louvain, Belgium.

NEUTS, M.F. (1981) *Matrix-Geometric Solutions in Stochastic Models*, p. 352, Vol. 2. Baltimore: Johns Hopkins Series in the Mathematical Sciences, The Johns Hopkins University Press, Baltimore.

OLSSON, M. (1996) Estimation of Phase-type distributions from censored data. *Scandinavian Journal of Statistics*, **23**, 443–460.

OLSSON, M. (1998) The EMpht programme. Manual. Chalmers University of Technology and Göteborg University.

REYNKENS, T. and VERBELEN, R. (2020) *ReIns: Functions from “Reinsurance: Actuarial and Statistical Aspects”*. R package version 1.0.10.

REYNKENS, T., VERBELEN, R., BEIRLANT, J. and ANTONIO, K. (2017) Modelling censored losses using splicing: A global fit strategy with mixed Erlang and extreme value distributions. *Insurance: Mathematics and Economics*, **77**, 65–77.

TZOUGAS, G., VRONTOS, S. and FRANGOS, N. (2014). Optimal bonus-malus systems using finite mixture models. *Astin Bulletin*, **44**(2), 417–444.

VAN LOAN, C. (1978) Computing integrals involving the matrix exponential. *IEEE Transactions on Automatic Control*, **23**(3), 395–404.

ZHANG, J., ZHENG, J., OKAMURA, H. and DOHI, T. (2021) An efficient algorithm for computation of information matrix in phase-type fitting. *International Journal for Computational Methods in Engineering Science and Mechanics*, **22**(3), 193–199.

MARTIN BLADT (Corresponding author)

*Department of Actuarial Science*  
*Faculty of Business and Economics*  
*University of Lausanne*  
*UNIL-Dorigny*  
*1015 Lausanne, Switzerland*  
*E-Mail: martin.bladt@unil.ch*

## APPENDIX A. DETAILS ON THE EM ALGORITHM FOR PH DISTRIBUTIONS

Here, we assume that  $Z \sim \text{PH}(\boldsymbol{\pi}, T)$ . Let  $B_k$  be the number of times that the process  $\{J_t\}_{t \geq 0}$  starts in state  $k$ ,  $N_{ks}$  is the total number of jumps from state  $k$  to  $s$ ,  $N_k$  is the number of times that we reach the absorbing state  $p + 1$  from state  $k$ , and let  $Z_k$  be the total time that the underlying Markov jump process spends in state  $k$  prior to absorption. These statistics are not recoverable from  $Z$ . Given a sample of absorption times  $\mathbf{z}$ , the completely observed likelihood can be written in terms of these sufficient statistics as follows:

$$\mathcal{L}_c(\boldsymbol{\pi}, T; \mathbf{z}) = \prod_{k=1}^p \pi_k^{B_k} \prod_{k=1}^p \prod_{s \neq k} t_{ks}^{N_{ks}} e^{-t_{ks} Z_k} \prod_{k=1}^p t_k^{N_k} e^{-t_k Z_k}, \tag{A1}$$

which is seen to conveniently fall into the exponential family of distributions, and thus has explicit maximum likelihood estimators.

However, the full data are not observed, and hence we employ the EM algorithm as an iterative way to obtain the maximum likelihood estimators. At each iteration, the conditional expectations of the sufficient statistics  $B_k, N_{ks}, N_k,$  and  $Z_k$  given the absorption times  $\mathbf{z}$  are computed, commonly referred to as the E-step. Subsequently,  $\mathcal{L}_c(\boldsymbol{\pi}, \mathbf{T}, \mathbf{z})$  is maximized using the estimates of the sufficient statistics from the previous step, in this way obtaining  $(\boldsymbol{\pi}, \mathbf{T})$ , commonly known as the M-step. The latter maximization is simple to perform because of the closed-form maximum likelihood estimators of exponential families.

Below are the explicit formulas needed for the E- and M-steps for a sample of size  $N$ . We denote by  $\mathbf{e}_k$  the column vector with all elements equal to 0 besides the  $k$ th entry which is equal to 1, that is, the  $k$ th element of the canonical basis of  $\mathbb{R}^d$ .

(1) *E-step, conditional expectations:*

$$\mathbb{E}(B_k \mid \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N \frac{\pi_k \mathbf{e}_k^T \exp(\mathbf{T}z_i) \mathbf{t}}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(Z_k \mid \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N \frac{\int_0^{x_i} \mathbf{e}_k^T \exp(\mathbf{T}(z_i - u)) \mathbf{t} \boldsymbol{\pi}^T \exp(\mathbf{T}u) \mathbf{e}_k du}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(N_{ks} \mid \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N t_{ks} \frac{\int_0^{z_i} \mathbf{e}_s^T \exp(\mathbf{T}(z_i - u)) \mathbf{t} \boldsymbol{\pi}^T \exp(\mathbf{T}u) \mathbf{e}_k du}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}$$

$$\mathbb{E}(N_k \mid \mathbf{Z} = \mathbf{z}) = \sum_{i=1}^N t_k \frac{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{e}_k}{\boldsymbol{\pi}^T \exp(\mathbf{T}z_i) \mathbf{t}}.$$

(2) *M-step, explicit maximum likelihood estimators:*

$$\hat{\pi}_k = \frac{\mathbb{E}(B_k \mid \mathbf{Z} = \mathbf{z})}{N}, \quad \hat{t}_{ks} = \frac{\mathbb{E}(N_{ks} \mid \mathbf{Z} = \mathbf{z})}{\mathbb{E}(Z_k \mid \mathbf{Z} = \mathbf{z})}$$

$$\hat{t}_k = \frac{\mathbb{E}(N_k \mid \mathbf{Z} = \mathbf{z})}{\mathbb{E}(Z_k \mid \mathbf{Z} = \mathbf{z})}, \quad \hat{t}_{kk} = - \sum_{s \neq k} \hat{t}_{ks} - \hat{t}_k.$$

We set

$$\hat{\boldsymbol{\pi}} = (\hat{\pi}_1, \dots, \hat{\pi}_p)^T, \quad \hat{\mathbf{T}} = \{\hat{t}_{ks}\}_{k,s=1,2,\dots,p}, \quad \hat{\mathbf{t}} = (\hat{t}_1, \dots, \hat{t}_p)^T.$$

If we repeat the above two steps, it can be shown that the likelihood increases at each iteration, and thus convergence to a possibly local maximum is guaranteed.