


RESEARCH ARTICLE

Individual claims reserving using the Aalen–Johansen estimator

Martin Bladt¹ and Gabriele Pittarello² 

¹Department of Mathematical Sciences, University of Copenhagen, Copenhagen, Denmark

²Department of Statistical Sciences, Sapienza University of Rome, Rome, Italy

Corresponding author: Gabriele Pittarello; Email: gabriele.pittarello@uniroma1.it

Received: 14 November 2023; **Revised:** 3 June 2024; **Accepted:** 3 June 2024; **First published online:** 3 December 2024

Keywords Multi-state models; conditional Aalen–Johansen; reserving; non-life insurance

Abstract

We propose an individual claims reserving model based on the conditional Aalen–Johansen estimator, as developed in Bladt and Furrer ((2023a) arXiv:2303.02119.). In our approach, we formulate a multi-state problem, where the underlying variable is the individual claim size, rather than time. The states in this model represent development periods, and we estimate the cumulative density function of individual claim sizes using the conditional Aalen–Johansen method as transition probabilities to an absorbing state. Our methodology reinterprets the concept of multi-state models and offers a strategy for modeling the complete curve of individual claim sizes. To illustrate our approach, we apply our model to both simulated and real datasets. Having access to the entire dataset enables us to support the use of our approach by comparing the predicted total final cost with the actual amount, as well as evaluating it in terms of the continuously ranked probability score.

1. Introduction

This manuscript demonstrates that multi-state models provide a natural framework for modeling the cost of individual claims in an individual reserving application. Our approach builds upon the findings presented in a recent paper (Bladt and Furrer, 2023b), which introduces the conditional Aalen–Johansen estimator. This estimator is a versatile, non-parametric tool for estimating state occupation probabilities conditioned on specific features, and it also discusses its key properties. Note that Aalen–Johansen-type estimators have traditionally found applications in life insurance. However, in this work, we demonstrate how these estimators can also be applied to non-life insurance for reserving purposes.

As mentioned earlier, multi-state models have found widespread use in the domain of life insurance. In the works of Hoem (1969, 1972), they introduce a by-now standard approach, representing biometric states (active, deceased, etc.) as finite states within a spatial model. However, to the best of our knowledge, these models have seen less frequent application in non-life insurance literature. Notable exceptions can be found in the work of Hesselager (1994) and Maciak et al. (2022). In Hesselager (1994), the authors model outstanding claim amounts conditionally on the current state, such as incurred but not reported (IBNR), reported but not settled (RBNS), and settled. In this context, the time variable corresponds to calendar time, similar to that in life insurance. In contrast, our approach in this paper takes a different perspective, replacing time with payments as the fundamental evolution variable. In Maciak et al. (2022), the authors propose modeling the incremental paid amounts as outcomes of a homogeneous Markov chain on a finite state space, though only for aggregated data and in discrete time.

In this manuscript, we propose an individual claims reserving model based on a continuous-time non-explosive pure jump process denoted by J on a finite state space, $\mathcal{S} = \{1, \dots, k\}$, $k \in \mathbb{N}$, where states correspond to the development periods (DPs) within a development triangle, and the “time spent”

between state transitions corresponds to the claim size growth between DP's. The state k serves as an absorbing state, representing claim closure. It is worth noting that, in practice, DPs have a stochastic behavior that leads to variation in the state space dimension across individual claim histories. However, for simplicity, we address this by collapsing long claim developments into a single state, as the specific details are not of primary importance in this context.

In a broader context, the application of counting processes to individual reserving has been a subject of research since the seminal contributions (Arjas, 1989; Norberg, 1993, 1997), with subsequent refinements and practical considerations detailed in the work of Haastrup and Arjas (1996) and Antonio and Plat (2014). In the study of Antonio and Plat (2014), they assume a framework where claims are generated by a position-dependent marked Poisson process, and they incorporate a severity modelling approach within the mark.

The primary contribution of this manuscript lies in presenting a novel interpretation of state spaces, which enables us to approach the claims reserving problem from a fresh perspective, using a non-parametric estimator from survival analysis. We denote state process by J_z when the cumulative claim size of an individual reaches z , where $z \in [0, +\infty)$. Note that in this approach, the fundamental variable of interest is the claim size itself, rather than time. The transition probabilities leading to the absorbing state k represent the cumulative distribution of individual claim sizes. In other words, as individuals progress through the states, their observations accumulate claim size, reflecting the evolution of their claims. Upon reaching the absorbing state, a claim is considered settled, and no further payments are possible.

Recent literature on individual reserving has focused on decomposing individual reserving data into different modules that describe the micro-level structure of the portfolio, for example, payment delay, settlement delay and payment size. By doing so, Huang et al. (2015) and Wang et al. (2021) extended the discrete model for predicting IBNR claims in Norberg (1986) to a model for RBNS and IBNR claims that discusses, theoretically and with empirical studies, the beneficial effect of including individual information in reserving models. An interesting approach to modelling the joint distribution of reporting delays and claim sizes of RBNS claims under censoring can be found in Lopez (2019). The work in Crevecoeur et al. (2023) embeds Crevecoeur et al. (2022) in a more general context and illustrates how pricing and reserving can be modelled under a unified framework and severity is modelled conditionally on the previous modules. A similar approach has been proposed in Delong et al. (2022), wherein the authors incorporate a gamma-distributed severity component into a neural network-based methodology.

In contrast to the aforementioned works, which rely on likelihood-based estimation or parametric assumptions on the claim size distribution, we propose a fully non-parametric estimation of the full cumulative distribution function of claim sizes, building on the results of Bladt and Furrer (2023b). The latter reference proposes a conditional version of the classic Aalen–Johansen estimator, which allows for continuous or discrete covariates, and where minimal assumptions on the process J are imposed, that is no Markov assumption is required. In particular, such lax probabilistic framework is desirable when it is unclear whether there exist duration effects in each state, as is the case for claim development processes.

Finally, we also propose a solution for modelling the IBNR and RBNS reserves separately. Our model is compared with the chain ladder in two case studies, one using simulated data and the other using a recent dataset from a Danish non-life insurer. In both the simulation case study and the real data case application, we have access to the total future claims costs and can thus compare the predicted amount with the target amount. The predicted curve for the cumulative density function of claims size is compared with the empirical cumulative density function using the continuously ranked probability score (CRPS), first proposed in Gneiting and Raftery (2007). Interestingly, we show how the CRPS can be used to perform model selection by answering the natural question of how to select the number of states in the state space \mathcal{S} .

This manuscript is organised as follows. In Section 2, we introduce the conditional Aalen–Johansen estimator after a discussion on the probabilistic setup and we explain how to reinterpret a state space

\mathcal{S} to set up our reserving model. In Section 3, we connect our framework to development triangles and after a discussion on the current literature in reserving, we illustrate how to estimate of the final cost of the claims. Section 4 will demonstrate the effectiveness of our approach through a comprehensive case study using simulated data. Recent contributions to the literature have provided various algorithms for generating individual loss reserve datasets (Gabrielli and Wüthrich, 2018; Avanzi et al., 2021, 2023; Wang and Wüthrich, 2022). Once reporting delays become available, it is possible to simulate both incurred but not settled (IBNR) claims and reported but not settled (RBNS) claims. We have chosen to focus on the exploration of RBNS claims in the next section, deliberately avoiding additional assumptions about the claims generation process (see for example Avanzi et al., 2021). This approach aims to draw the attention of the reader to the applicability of our models in different scenarios. We showcase the application of our IBNR modeling strategy in the section Section 5, where we rigorously test our models using a recent real individual data portfolio from a Danish insurer. Finally, Section 6 concludes.

2. The conditional Aalen–Johansen estimator

In this section, we introduce our statistical framework, the quantities of interest, and the conditional occupation probabilities. We present the conditional Aalen–Johansen estimator (Bladt and Furrer, 2023b), a non-parametric estimator for the conditional occupation probabilities of an arbitrary jump process. Note that the conditional Aalen–Johansen estimator is closely linked with the literature streams on kernel-based estimation for survival and semi-Markov processes (McKeague and Utikal, 1990; Nielsen and Linton, 1995; Dabrowska, 1997), as well as landmarking (Van Houwelingen, 2007). We conclude the section with some remarks on the interpretation of multi-state models for reserving. As we explain in the next paragraphs, this is a fundamental step to construct our reserving model.

2.1. The probabilistic setup

Let $J = (J_t)_{t \geq 0}$ be a non-explosive pure jump process on a finite state space \mathcal{S} and take $(\Omega, \mathcal{F}, \mathbb{P})$ to be the underlying probability space. Here, we let $\mathcal{S} = \{1, \dots, k\}$, with $k \in \mathbb{N}$. Denote by Y the possibly infinite absorption time of J . Let X be a random variable with values in \mathbb{R}^d equipped with the Borel σ -algebra. We assume that the distribution of X admits a density g with respect to the Lebesgue measure λ . Absolute continuity in case of known atoms may be relaxed as discussed in Bladt and Furrer (2023b).

We introduce a multivariate counting process N with components $N_{jh} = (N_{jh}(z))_{z \geq 0}$ given by

$$N_{jh}(z) = \# \{s \in (0, z] : J_{s-} = j, J_s = h\}, \quad z \geq 0,$$

for $j, h \in \mathcal{S}, j \neq h$. Let us introduce the following assumption.

Assumption 1. For every $z \geq 0$ and $j, h \in \mathcal{S}, j \neq h$ there exists some $\delta > 0$ such that

$$\mathbb{E} [N_{jh}(z)^{1+\delta}] < \infty.$$

Taking the largest z and the minimum across the whole states provides a delta which is independent of j, h, z .

For any x in the support of X define the conditional occupation probabilities according to

$$p_j(z | x) = \mathbb{E} [\mathbb{1}_{\{J_z=j\}} | X = x].$$

We denote by $p(z | x)$ the row vector with elements $p_j(z | x)$. We now introduce the cumulative conditional transition rates and cast the conditional occupation probabilities as a product integral of these cumulative rates. To this end, let

$$p_{jh}(z | x) = \mathbb{E} [N_{jh}(z) | X = x],$$

and define the cumulative conditional transition rates via

$$\Lambda_{jh}(z | x) = \int_0^z \frac{1}{p_j(s- | x)} p_{jh}(\mathrm{d}s | x).$$

A cumulative transition hazard matrix is defined by having the j, j entry $\Lambda_{jj}(z | x) = \sum_{j \neq h} \Lambda_{jh}(z | x)$ for $j, h \in \mathcal{S}, j \neq h$.

Let $q(\cdot | x)$ be the interval function associated with the conditional transition probabilities:

$$q_{jh}(s, z | x) = \mathbb{1}_{\{p_j(s|x) > 0\}} \frac{\mathbb{E} [\mathbb{1}_{\{J_z=h\}} \mathbb{1}_{\{J_s=j\}} | X=x]}{p_j(s|x)} + \mathbb{1}_{\{p_j(s|x)=0\}} \mathbb{1}_{\{j=h\}}.$$

In Overgaard (2019), it is shown that this assumption ensures that the product integral

$$\prod_0^t (\mathrm{Id} + \Lambda(\mathrm{d}s | x))$$

is well defined. Therefore,

$$p(t | x) = p(0 | x) \prod_0^t (\mathrm{Id} + \Lambda(\mathrm{d}s | x)). \quad (2.1)$$

We are interested in the situation where the observation of J is right-censored. To this end, we introduce a strictly positive random variable W describing right censoring, so that we actually observe the triplet

$$(X, (J_z)_{0 \leq z \leq W}, Y \wedge W).$$

Whenever the Markov property fails for the jump process J , the following assumption is key for our estimation procedure.

Assumption 2. *Right censoring is conditionally entirely random:*

$$W \perp\!\!\!\perp J | X.$$

Other censoring schemes are possible if J is Markovian, though we do not make that assumption in the sequel.

Remark 1. From a reserving perspective, the above assumption amounts to stating that, conditional on covariates, closure of a claim is simply specified through an independent random mechanism. One should note that this assumption is not verifiable, despite its widespread use in survival analysis.

We introduce

$$p_j^c(z | x) = \mathbb{E} [\mathbb{1}_{\{J_z=j\}} \mathbb{1}_{\{z < W\}} | X=x],$$

and

$$p_{jh}^c(z | x) = \mathbb{E} [N_{jh}(z \wedge W) | X=x].$$

from which we may conclude that

$$\Lambda_{jh}(z | x) = \int_0^z \frac{1}{p_j^c(s- | x)} p_{jh}^c(\mathrm{d}s | x), \quad (2.2)$$

within the interior of the support of W .

2.1.1. Estimators

Consider the i.i.d. replicates $\left(X^i, (J_z^i)_{0 \leq z \leq W^i}, Y^i \wedge W^i\right)_{i=1}^n$. Define $\delta^i := \mathbb{1}_{\{Y^i \leq W^i\}}$, which equals 1 when the observation is absorbed (closed claim), and zero otherwise. In a reserving context, this distinction is usually reflected by specifying two types of claims: closed claims and RBNS (open) claims, and we write, respectively, $n = n^{\text{closed}} + n^{\text{RBNS}}$.

We now construct an estimator for the conditional occupation probabilities $p_j(z | x)$, but from Equation (2.1), it essentially suffices to estimate the cumulative conditional transition rates $\Lambda_{jh}(z | x)$, which we do using Equation (2.2). Let K_b be kernel functions of bounded variation and bounded support, with $b = 1, \dots, d$ and let (a_n) be a bandwidth sequence. We impose the following standard assumption, conferring with Stute (1986).

Assumption 3. *The band sequence satisfies:*

1. $a_n \rightarrow 0$ as $n \rightarrow \infty$.
2. $na_n^d \rightarrow \infty$ as $n \rightarrow \infty$.
3. $\sum_{n=1}^{\infty} c_n^r < \infty$ for some $r \in (1, 1 + \delta)$ where $c_n := (na_n^d)^{-1} \log n$.

For instance, one may take, for $(1 + \delta)^{-1} < \eta < 1$, any sequence satisfying

$$a_n^d \sim \frac{\log n}{n^{1-\eta}}.$$

We introduce the usual kernel estimator for the density g :

$$\mathfrak{g}^{(n)}(x) = \frac{1}{n} \sum_{i=1}^n \mathfrak{g}^{(n,i)}(x) = \frac{1}{n} \sum_{i=1}^n \prod_{b=1}^d \frac{1}{a_n} K_b \left(\frac{x_b - X_b^i}{a_n} \right).$$

It is well-known (cf. Ghosh, 2018, Ch. 1.4) that

$$\mathfrak{g}^{(n)}(x) \xrightarrow{\text{a.s.}} g(x), \quad n \rightarrow \infty,$$

and, for any random variable V with i.i.d. replicates (V^i) satisfying $\mathbb{E}[|V|^{1+\delta}] < \infty$ that

$$\frac{1}{n} \sum_{i=1}^n \frac{V^i \mathfrak{g}^{(n,i)}(x)}{\mathfrak{g}^{(n)}(x)} \xrightarrow{\text{a.s.}} \mathbb{E}[V | X = x], \quad n \rightarrow \infty. \quad (2.3)$$

Based on the kernel estimator for g , we form the following Nadaraya–Watson type kernel estimators:

$$\mathbb{N}_{jh}^{(n)}(z | x) = \frac{1}{n} \sum_{i=1}^n \mathbb{N}_{jh}^{(n,i)}(z | x) = \frac{1}{n} \sum_{i=1}^n \frac{N_{jh}^i(z \wedge W^i) \mathfrak{g}^{(n,i)}(x)}{\mathfrak{g}^{(n)}(x)} \quad (2.4)$$

$$\mathbb{I}_j^{(n)}(z | x) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_j^{(n,i)}(z | x) = \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{1}_{\{z < W^i\}} \mathbb{1}_{\{J_z^i = j\}} \mathfrak{g}^{(n,i)}(x)}{\mathfrak{g}^{(n)}(x)} \quad (2.5)$$

Remark 2. Due to Assumption 1, it follows from Equation (2.3) that

$$\mathbb{N}_{jh}^{(n)}(z | x) \xrightarrow{\text{a.s.}} p_{jh}^c(z | x), \quad n \rightarrow \infty.$$

Furthermore, it holds that

$$\mathbb{I}_j^{(n)}(z | x) \xrightarrow{\text{a.s.}} p_j^c(z | x), \quad n \rightarrow \infty.$$

Definition 1. *The conditional Nelson–Aalen estimator for the cumulative conditional transition rates is defined as*

$$\Lambda_{jh}^{(n)}(z | x) = \int_0^z \frac{1}{\mathbb{I}_j^{(n)}(s- | x)} \mathbb{N}_{jh}^{(n)}(\text{d}s | x).$$




Figure 1. Example of multi-state model for claims reserving, with $k = 5$. We interpret time spent in each state as increasing claims size clock, instead of calendar time. The states represent the development periods (DP's) of the individual claims.

Definition 2. The conditional Aalen–Johansen estimator for the conditional occupation probabilities is defined according Equation (2.1), but with the cumulative conditional transition rates replaced by the conditional Nelson–Aalen estimator, that is,

$$\mathbb{P}^{(n)}(z \mid x) = \mathbb{P}^{(n)}(0 \mid x) \prod_0^z \left(\text{Id} + \mathbb{A}^{(n)}(ds \mid x) \right),$$

where $\mathbb{P}_j^{(n)}(0 \mid x) = \mathbb{I}_j^{(n)}(0 \mid x)$. The estimator $\mathbb{P}_j^{(n)}(0 \mid x)$ for the occupation probabilities $p_j(z \mid x)$ denotes the j th entry of the row vector $\mathbb{P}^{(n)}(z \mid x)$, an estimator for $p(z \mid x)$.

Remark 3. We refer to Bladt and Furrer (2023b) for the proof that the estimators are pathwise strongly consistent and asymptotically Gaussian distributed.

When all the subjects start in state 1 at inception, we can use the estimator of the occupation probabilities of the absorbing state k , $\mathbb{P}_k^{(n)}$ to model the cumulative density function of the individual claim size.

Remark 4. The Volterra integral equation defines the product integral of the matrix function $\Lambda(ds \mid x)$ as the unique solution $p(\cdot \mid x)$ to the equation

$$p(t \mid x) = \text{Id} + \int_{(0,t]} p(s- \mid x) \Lambda(ds \mid x).$$

However, when a piecewise constant estimator of Λ is given by $\Lambda^{(n)}$, as is described in detail in the previous section, the computation of the product integral simplifies. More precisely, denote $0 = t_0 < t_1 < t_2 < \dots$ the sequence of time points describing all observed jump times in the sample. Then the recursion

$$\mathbb{P}_j^{(n)}(t_{\ell+1} \mid x) - \mathbb{P}_j^{(n)}(t_\ell \mid x) = \sum_{\substack{k \in \mathcal{Z} \\ k \neq j}} \mathbb{P}_k^{(n)}(t_\ell \mid x) \Delta \Lambda_{kj}^{(n)}(t_{\ell+1} \mid x) - \mathbb{P}_j^{(n)}(t_\ell \mid x) \sum_{\substack{k \in \mathcal{Z} \\ k \neq j}} \Delta \Lambda_{jk}^{(n)}(t_{\ell+1} \mid x)$$

and starting value $\mathbb{P}^{(n)}(0 \mid x) = \mathbb{I}^{(n)}(0 \mid x)$ completely characterize the product integral $\mathbb{P}_j^{(n)}(\cdot \mid x)$ at jump times (and constant otherwise).

2.2. Practical considerations for claims development

In this section, we offer practical insights into our modeling approach, with the aim of assisting the reader in understanding our multi-state setup for claims reserving. Refer to Figure 1, which conceptually illustrates our model featuring a state space \mathcal{S} with $k = 5$ states.

In our framework, we represent the active state as IBNR claims, while we treat RBNS claims as transitioning through intermediate states. Any claim that reaches state k is considered closed in our system due to the presence of an absorbing state. Consequently, we do not allow claims to reopen in our framework. Notably, unlike a traditional survival study, we do not explicitly model the time component, and the number of IBNR claims, along with their characteristics, remains unknown at the time of evaluation. In Section 3, we delve into how to model RBNS and IBNR claims, showcasing how our model can accommodate such scenarios. Furthermore, we have designed our state space so that state DP3+ at step $k - 1$ in Figure 1 is the final relevant DP that we explicitly model.

Finally, it is worth mentioning that we have not introduced development triangles yet, which are the standard aggregate datasets commonly used in claims reserving applications. The choice of parameter k aligns with the data that, in a conventional modelling approach, would be aggregated into development triangles of varying sizes. Since our proposed framework does not explicitly model the time component, we assume that payments close at the development year following the last payment. This assumption results in a development triangle with $k - 1$ development times. This aspect will be particularly relevant for our data application, where we compare our model to the benchmark chain ladder method outlined in Mack (1993), and we need to construct a development triangle for this purpose.

Remark 5. By changing the multi-state model in Figure 1, our approach can be modified to take into account additional domain-specific aspects of reserving data. For instance, if we allow transitions back from the state k (closed) to states $3, \dots, k - 2$ (DP 2, \dots), our framework can be modified to take into account reopenings.

3. Claims reserving

In this section, we will introduce the definitions of final cost of the claims and of the claims reserve and provide a notation for development triangles, an aggregated representation of the reserving data. We then connect the multi-state model that we introduced in the previous section to the actuarial literature in reserving, with a comparative literature review that distinguishes between the approaches that rely on aggregated data and those that rely on individual data, as the methodology outlined in this manuscript. We remark that aggregated data are a known data representation in survival studies when the individual data are not available, Andersen *et al.* (1993, p. 168) and Gill (1986). We conclude this section with a discussion on the numerical estimation of the total final cost of the claims using the conditional Aalen–Johansen estimator $\mathbb{P}_k^{(n)}$ of $p_k(z|x)$.

Notation. Together with our sample, which is composed of n i.i.d. observations generated from J and its covariates (refer to the previous section), we additionally observe the following quantities that are *not* related to J^1, \dots, J^n :

$$\{T^i, U^i\}_{i=1}^n.$$

In detail, U^i and T^i represent the accident period and the reporting delay of the i th claim, with $U^i, T^i \in \{1, \dots, k - 1\}$. Let us also define the development triangle of the reporting delays as the set

$$\mathcal{D}^{(k-1)} = \{d_{\ell,j} : \ell + j \leq k; \ell, j = 1, \dots, k - 1\},$$

with $d_{\ell,j} = \sum_{i=1}^n \mathbb{1}_{\{U^i=\ell \text{ and } T^i=j\}}$ denoting the total claims reported in accident period ℓ with delay j . In a similar fashion, we can define the development triangle of cumulative payments,

$$\mathcal{C}^{(k-1)} = \{C_{\ell,j} : \ell + j \leq k; \ell, j = 1, \dots, k - 1\},$$

where $C_{\ell,j} = \sum_{i=1}^n \mathbb{1}_{\{U^i=\ell \text{ and } J_{z_{i,j}}^i=j\}} z_{i,j}$, with $0 \leq z_{i,j} \leq W^i$ such that $J_{z_{i,j}}^i \in \{1, \dots, j - 1\}$, $J_{z_{i,j}}^i = j$ denoting the total cumulative amount paid in accident period ℓ and DP j .

We are presently interested forecasting the ultimate cost of our claims,

$$Y^{\text{Closed}} + Y^{\text{RBNS}} := \sum_{i=1}^n \delta^i Y^i + \sum_{i=1}^n (1 - \delta^i) Y^i = \sum_{\ell=1}^{k-1} C_{\ell, k-1},$$

The claims reserve is simply $R = \sum_{\ell=1}^{k-1} (C_{\ell, k-1} - C_{\ell, k-\ell})$.

3.1. A comparative literature review

Research papers on loss reserving can be broadly categorized into two main streams of research: models based on development triangles (aggregate loss reserving models) and models based on individual data (individual loss reserving models). The relationship between individual data and aggregate modelling has been extensively studied in Miranda et al. (2012), Hiabu (2017), and Bischofberger et al. (2020), utilizing different survival analysis tools.

The existing body of literature on aggregate reserving has already provided the building blocks for predicting the distribution of loss reserves, with a specific emphasis on models that seek to emulate the empirical chain-ladder algorithm. In this section, we briefly mention some of these models and defer to the comprehensive discussion in Hess and Schmidt (2002) for their detailed categorization. For a distribution-independent approach to computing the mean squared prediction error of the loss reserve, one can refer to Mack (1993), where the author decomposes the prediction error of the loss reserve into its elemental components, namely process variance and estimation variance, as expressed by $\mathbb{E}[(\hat{R} - R)^2] = \text{Var}(R) + \mathbb{E}[(\hat{R} - \mathbb{E}[\hat{R}])^2]$. The model has received several statistical critiques through the years but remains one of the best-performing reserving models in an actuary's toolkit.

In England and Verrall (1999), the authors replicate the point estimates of the chain-ladder method and ingeniously employ bootstrap and generalized linear models (GLM) to compute $\mathbb{E}[(\hat{R} - R)^2]$. For an extension of the model, allowing for negative payments as well as an exploration of the model assumptions underpinning it, one may consult Verrall (2000). An interesting comparison of the fundamental assumptions underpinning the aforementioned models is presented in Mack and Venter (2000), where the authors conclude that only the approach outlined in Mack (1993) is the true underlying model of the chain-ladder algorithm. In a more recent development, in Al-Mudafer et al. (2022), the approach in England and Verrall (1999) is integrated into the framework of mixture density networks (Bishop, 1994).

Shifting our focus to non-parametric models for individual reserving, the current discussion primarily centers on providing precise point forecasts for the loss reserve. In contrast, Delong and Wüthrich (2023) discuss the role of the variance function in regression problems for pricing.

In another recent development, Wüthrich (2018) incorporates neural networks into the chain-ladder model, introducing features to the existing methodology. Further related work can be found in Lopez et al. (2019) and Lopez and Milhaud (2021), where the authors employ a tree-based algorithm for modeling RBNS claims using censored data. In particular, in Lopez and Milhaud (2021), the use of bootstrapping for estimating the uncertainty in the claims reserve is discussed.

Our manuscript introduces a novel approach to the individual reserving problem, allowing us to explicitly model curves of individual claim sizes denoted as $p_k(z|x)$ within our framework.

3.2. Prediction of RBNS and IBNR claim costs

We can calculate the final cost of RBNS claims numerically using the conditional Aalen–Johansen estimate $\mathbb{P}_k^{(n)}(z|x)$ for $p_k(z|x)$. The predictor of the final total cost of RBNS claims is defined by

$$\hat{Y}^{\text{RBNS}} = \sum_i^n \mathbb{1}_{\{\delta^i=0\}} \left(W^i + \mathbb{E}[Y|Y > \widehat{W}^i, X=x] \right)$$

with

$$\mathbb{E}[Y|Y > \widehat{W^i}, X = x] = \frac{1}{1 - \mathbb{P}_k^{(n)}(W^i | x)} \int_{W^i}^{+\infty} (1 - \mathbb{P}_k^{(n)}(y | x)) dy. \quad (3.1)$$

We remark that the general formula for the m th moment is explicitly given by

$$\mathbb{E}[Y^m | Y > \widehat{W^i}, X = x] = \frac{1}{1 - \mathbb{P}_k^{(n)}(W^i | x)} \int_{W^i}^{+\infty} m y^{m-1} (1 - \mathbb{P}_k^{(n)}(y | x)) dy.$$

Remark 6. Relatable approaches to the predictor in Equation (3.1) can be found in the survival literature on censored regression. For instance, similar techniques were presented for the estimation of the linear model in Buckley and James (1979) and for the construction of the artificial data points in Heuchenne and Van Keilegom (2007).

Strictly speaking, from a reserving standpoint, we also need to propose a strategy for reserving for the (unknown) IBNR claims. We denote the total number of IBNR claims by n^{IBNR} , which is a random variable with values in \mathbb{N}_0 . The individual cost of the IBNR claims is described by a sequence of iid copies of the random variable \tilde{Y} , namely $\{\tilde{Y}^{i'}\}_{i' \in \mathbb{N}}$. Under the framework of Bladt and Furrer (2023b), \tilde{Y} has cumulative distribution function $p_k(z)$. As we will describe shortly, we specify a model that does not consider the features, since in this case we do not have access to future exposures by feature. We propose to correct the model for the random presence of IBNR using a compound distribution, under the assumptions in Klugman et al. (2012, p. 148). We describe the total cost of IBNR claims as follows:

$$Y^{\text{IBNR}} = \sum_{i'=1}^{n^{\text{IBNR}}} \tilde{Y}^{i'}.$$

The expression for Y^{IBNR} is well known in non-life mathematics and is often referred to as the collective risk model, Parodi (2014). The moments of Y^{IBNR} can be calculated in closed form. In particular, the expected cost of IBNR claims is predicted by

$$\hat{Y}^{\text{IBNR}} = \mathbb{E}[\widehat{Y^{\text{IBNR}}}] = \hat{n}^{\text{IBNR}} \mathbb{E}[\tilde{Y}],$$

with the m -th moment of \tilde{Y} predicted by

$$\mathbb{E}[\tilde{Y}^m] = \int_0^\infty m y^{m-1} (1 - \mathbb{P}_k^{(n)}(y)) dy,$$

fitting the unconditional Aalen–Johansen to the sample n to obtain $\mathbb{P}_k^{(n)}(z)$. To project the number of IBNR claims, we simply use the chain ladder model on the triangle of reporting delays $\mathcal{D}^{(k-1)}$ and obtain \hat{n}^{IBNR} .

Similarly, the variance of the IBNR claims is predicted by

$$\widehat{\text{Var}}(Y^{\text{IBNR}}) = \hat{n}^{\text{IBNR}} \left[\mathbb{E}[\tilde{Y}^2] - \mathbb{E}[\tilde{Y}]^2 \right] + \mathbb{E}[\tilde{Y}]^2 \widehat{\text{Var}}(n^{\text{IBNR}}).$$

We denote by $\widehat{\text{Var}}(n^{\text{IBNR}})$ a predictor for the variance of n^{IBNR} . In this manuscript, we consider the predictor for the process variance discussed in Mack (1993).

The total size of claims is $Y^{\text{TOT}} = Y^{\text{Closed}} + Y^{\text{RBNS}} + Y^{\text{IBNR}}$ and we estimate it with $\hat{Y}^{\text{TOT}} = Y^{\text{Closed}} + \hat{Y}^{\text{RBNS}} + \hat{Y}^{\text{IBNR}}$. We remark that in several lines of business the contribution of the true cost of IBNR claims (Y^{IBNR}) to Y^{TOT} is minor compared to the cost of RBNS claims (Y^{RBNS}). This is the case for those lines of business that have data with mostly single payments and a short time lag between reporting and settlement. See for instance the results presented in Friedland (2010, p. 80) on the US market data, comparing for different lines of insurance the size of the estimated cost of IBNR claims to paid and outstanding claims. As we will show in Section 5, our real dataset exhibits exactly this behaviour and the IBNR correction will have a minor impact on the estimated final cost.

4. Simulation of RBNS claims

In this section, we perform two simulation studies to provide numerical evidence of the applicability of our models to claims reserving. To demonstrate the potential benefits to actuaries of using our modeling approach, we first compare our results with the chain ladder model (Mack, 1993). A second comparison with the hierarchical reserving model in Crevecoeur et al. (2023) can be found in Appendix A. To facilitate the presentation of our results, we refer to our model as AJ, the chain ladder model in Mack (1993) as CL, and the hierarchical reserving model as hirem.

4.1. Implementation

It is possible to simulate observations from a jump-process model using the function `sim_path` from the R package `AalenJohansen` (Bladt and Furrer, 2023a), by introducing a matrix of transition intensities. Markov and semi-Markov models are allowed. To obtain a realistic set of examples for the simulation study, we used the Nelson–Aalen implementation of the R package `survival` (Therneau, 2023) to estimate the (unconditional) cumulative transition intensities $\Lambda_{jh}^{(n)}(z)$ for different choices of $k = 4, 5, 6, 7$ on the real dataset available to us for writing this manuscript. Then, similar choices of parameters were employed for the present section. See Table 3 for the description of the real data. The matrix of intensities can be obtained from the cumulative intensities $\Lambda_{jh}^{(n)}(z)$. Simulation by knowing the data generation process underlying the data is important to understand model performance compared to the true estimation target. For example, in Figure 2, we show an example of the fitted $\mathbb{p}_k^{(n)}(z | x)$ on datasets simulated directly with Λ without any feature effect. For each choice of k we simulated two datasets. The first dataset contains $120(k - 1)$ observations. The second dataset contains $1200(k - 1)$ observations. By comparing the $\mathbb{p}_k^{(n)}(z | x)$ fitted to the smallest sample (red) with the curve fitted to the largest sample (blue), we observe that increasing the sample size yields fits that are closer to the true $p_k(z | x)$, empirically validating the consistency of the Aalen–Johansen estimator with claim size as operational time.

In particular, we have presented a model that can take into account the individual feature information. In the case studies that we present in the next section, we treat the accident period as a feature and add an effect by simply specifying $q_{jh}(s, z | x) = q_{jh}(s, z)/f(x)$, where $f(x)$ is a function of some feature x and we can derive $q_{jh}(s, z)$ from the cumulative transition matrix obtained with the `survival` package. In the simulations we present, we assume that observations are generated from $f(x) = 10 + k - x$, where in place of x we use the accident period, namely U^i introduced in the previous sections. For the simulation of the reserving datasets, we decide arbitrarily the number of observations in each accident period U^i . Each claimant is then associated with a numeric variable representing the accident period U^i , taking values in $\{1, \dots, k - 1\}$, from which we can calculate the effect $f(U^i) = 10 + k - U^i$. We also add an effect on the volumes, and generate fewer observations in most recent accident period. Generating decreasing volumes resembles cycles in the underwriting, a case study of an insurer that underwrites less policies in the more recent accident periods, resulting in less claims. In particular, we generated 1200 observations in the first accident period $U^i = 1$ and decreased by 100 the number of observations every accident periods. For instance, we have 1100 observations with $U^i = 2$ and 1000 observations with $U^i = 3$.

4.2. Evaluating the models performance

Using the approach introduced in Section 4.1, we present two simulated scenarios: Alpha and Beta. In scenario Alpha, we simulate directly with unconditional transition intensities that we have fitted to the real data. In the Beta scenario, we introduce the effect of the accident period characteristics and the occupation intensities. In each scenario we simulate, for the different choices of $k = 4, 5, 6, 7$, a total number of 20 datasets to smooth possible fluctuations due to the random generator and report the average results that we obtain. Each dataset in scenario Alpha contains $1200(k - 1)$ observations. The dataset




Figure 2. Comparison of the true severity curve (dark grey dotted line) to the fitted severity curve for different data sizes in the 4 simulated scenarios. The red curve is fitted on $120(k - 1)$ RBNS claims, the blue curve is fitted on 1200 RBNS claims. Z is rounded by millions.

in scenario Beta contains $1200(k - 1) - 100(k - 2)$ observations. As motivated in Section 4.1, we let the volumes decrease in recent accident periods. We chose arbitrarily to decrease the volume by 100 in every accident period subsequent to the first one.

The choice of measure of model performance can be an arduous task when models are very different, as are AJ and CL. The first performance measure that we propose is the error incidence,

$$EI = \frac{\hat{Y}^{\text{TOT}}}{Y^{\text{TOT}}} - 1,$$

where \hat{Y}^{TOT} is the predicted final claim size and Y^{TOT} is the actual final claim size, which is available from the simulation. The EI is an interesting benchmark that provides insight into the relative accuracy of the AJ compared to the CL, though it does not highlight the fact that AJ provides a much more nuanced description of the claim size than point estimates. In Figure 3, we provide a box plot of the EI for the Alpha scenario (left side) and the Beta scenario (right side) for the different values of k . Compared to the CL, the AJ model performs favourably for each choice of k and in both scenarios.

While the EI is an interesting benchmark, to better understand the performance of the models compared to the data at hand, we would ideally like to use a proper scoring measure (Gneiting and Raftery, 2007). In principle, a scoring metric with the property of being proper is able to evaluate the better model by assigning better scores to better models. In this manuscript we consider the continuously

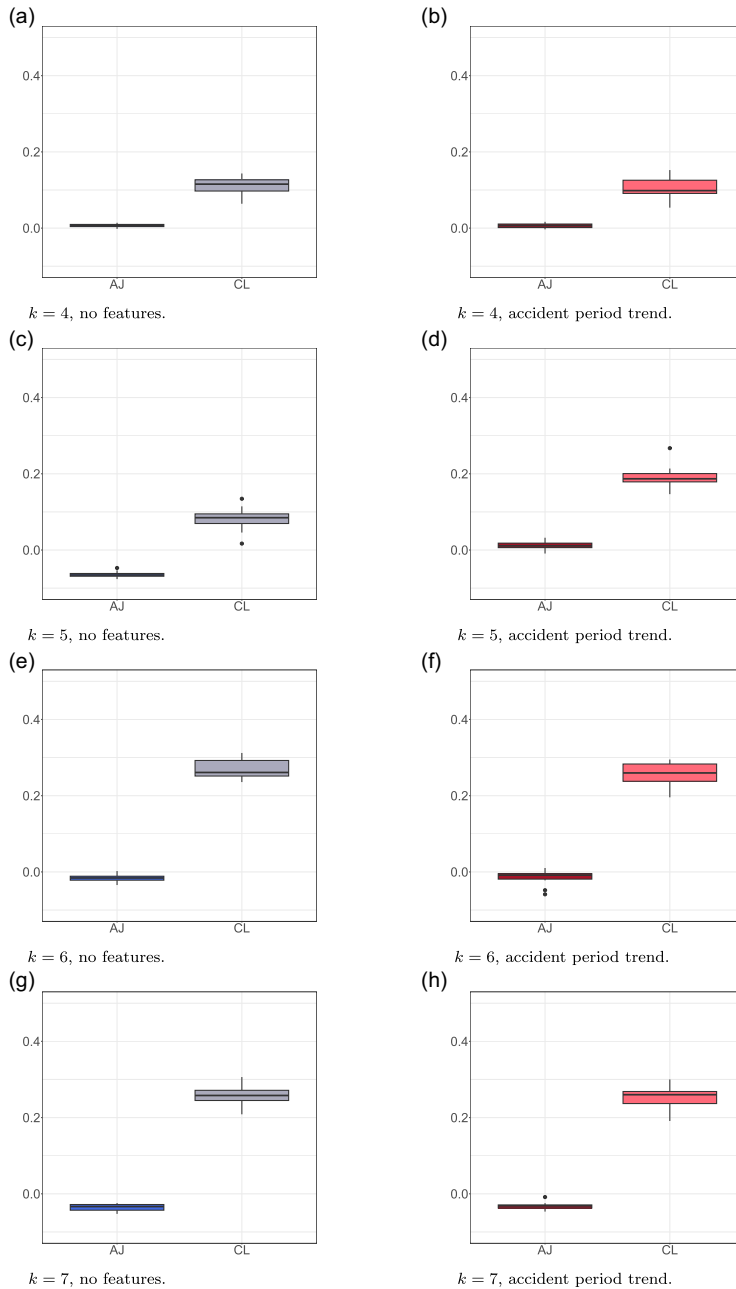


Figure 3. Box plots of EI for AJ and CL over the 20 simulations, for each value of k , in the Alpha scenario (left column) and the Beta scenario (right column).

ranked probability score (CRPS) as a proper scoring metric:

$$\begin{aligned} \text{CRPS}(\mathbb{P}_k^{(n)}(z|x), y) &= \int_0^{+\infty} (\mathbb{P}_k^{(n)}(z|x) - \mathbb{1}_{\{y \leq z\}})^2 dz \\ &= \int_0^y (\mathbb{P}_k^{(n)}(z|x))^2 dz + \int_y^{+\infty} (1 - \mathbb{P}_k^{(n)}(z|x))^2 dz, \end{aligned}$$

Table 1. Results for scenario Alpha. For each value of k (column one) we present the average results over the 20 simulations. Each row of the table corresponds to a different AJ model specification (column three). The table includes the (average) actual Y^{TOT} simulated total cost (column four) and the error incidence for the AJ and the CL (columns five and six). In columns seven and eight, we reported the coefficients of variation of Y^{TOT} . The results for the CRPS are reported in column nine.

k	Scenario	U	Actual Y^{TOT} (average)	EI (AJ)	EI (CL)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (AJ)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (CL)	CRPS (average)
4	Alpha	✗	598.483	0.006	0.111	0.008	0.009	0.145
5		✗	456.746	−0.085	0.082	0.007	0.008	0.095
6		✗	1952.683	−0.016	0.269	0.007	0.009	0.410
7		✗	2359.151	−0.035	0.257	0.007	0.007	0.557

Table 2. Results for scenario Beta. For each value of k (column one), we present the average results over the 20 simulations. Each row of the table corresponds to a different AJ model specification (column three). The table includes the (average) actual simulated total cost (column four) and the error incidence for the AJ and the CL (columns five and six). In columns seven and eight, we reported the coefficients of variation of Y^{TOT} . The results for the CRPS, relative to the AJ model with features, are reported in column nine.

k	Scenario	U	Actual Y^{TOT} (average)	EI (AJ)	EI (CL)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (AJ)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (CL)	CRPS (average, relative)
4	Beta	✓	42.279	0.006	0.103	0.008	0.008	1.000
		✗		0.005		0.008		1.008
✓		94.580	0.011	0.190	0.009	0.008	1.000	
✗			−0.006		0.007		1.064	
✓		108.423	−0.013	0.256	0.007	0.008	1.000	
✗			−0.016		0.007		1.039	
✓		116.032	−0.033	0.255	0.007	0.008	1.000	
✗			−0.031		0.008		1.260	

where $\mathbb{P}_k^{(n)}(z \mid x)$ is the predicted curve of the individual claims sizes and y_i is the (true) value of the final claim size that is available from the simulation. The interested reader can refer to the established statistical literature on proper scoring for a detailed discussion on this topic Gneiting and Raftery, 2007; Gneiting and Ranjan, 2011).

4.3. Empirical analysis

The average results of our 20 simulations for the Alpha scenario are reported in Table 1. Each column of the table represents one of the AJ models fitted for the different choices of k (column one). The target Y^{TOT} (average simulated total cost) assuming $Y^{\text{IBNR}} = 0$ is in column four. The predicted final cost is $\hat{Y}^{\text{TOT}} = Y^{\text{Closed}} + Y^{\text{RBNS}}$. The EI results show that, on average, our models outperform the CL for every choice of k except $k = 5$ (columns five and six). For all choices of k , we find that the relative variability of our model is less than that provided by the CL (columns seven and eight). We compute $\text{sd}(Y^{\text{TOT}}) = \sqrt{\sum_i^{n^{\text{RBNS}}} \text{Var}(Y^i)}$. In this manuscript, we will not discuss the estimation error, and the results we present with respect to the variability of the reserve refer only to the process variance. In column nine, we report the (average) CRPS.

Table 2 reports the average results for the 20 simulations of the Beta scenario. For each simulation, we fit two models, with and without the use of U as a feature (✓ and ✗ in the table). The model with U

is correctly specified assuming the Beta scenario. The aim of this experiment is to show that the CRPS is able to judge which is the better model in a context where we know the data generation process. As expected, we obtain a lower average CRPS for the models that include features. The CRPS results in Table 2 are relative to the AJ model with feature U . For example, for $k = 6$ we see that the AJ model without the feature U shows (on average) a CRPS 3.9% higher than the CRPS of the AJ model with feature U . For the purpose of an easier interpretation, the CRPS results will be shown relative to the AJ with feature model also in the following tables. The (average) target Y^{TOT} is reported in the fourth column. The results for EI indicate that the AJ model generally outperforms the CL model. The relative variability results show comparable results for AJ and CL.

5. A data application on an insurance portfolio

In this section, we propose a case study on our model using real data. For this data application, we have at our disposal a recent real dataset from a Danish non-life insurer. The dataset is not publicly available (Table 3).

Table 3. Description of the real dataset.

Covariates	Description
Claim_number	Policy identifier
claim_type $\in \{1, \dots, 20\}$	Type of claim
AM	Accident month
CM	Calendar month of report
DM	Development month
incPaid	Incremental paid amount
Delta	Indicator, 0 when the claim is open

We provide additional insight in Figure 4. The exploratory analysis of the data shows that most of the records have a single payment (Figure 4(b)) and a very fast settlement (Figure 4 (d)). The data we present in this manuscript come from a very stable line of business. Interestingly, we show that we are able to outperform the CL on data where the CL is expected to behave as well as it can. Most of the datasets belong to claim_type 1, see Figure 4 (a).

Remark 7. Notice that the covariate in our application is discrete, which with a uniform kernel effectively amounts to subsampling of data. However, the estimators of the paper provide consistent and asymptotically normal estimators even when covariates are continuous. Nonetheless, kernel-based methods are prone to the curse of dimesionality, and hence our method provides the best results when covariate dimensions are small.

To analyse the real dataset, we propose two strategies that an actuary could follow to calculate the loss reserve. In Section 5.1, we show an example where we censor our data by different calendar periods and fit an AJ model for the maximum depth of data available. For example, for a model with $k = 4$, we will have the individual data corresponding to a development triangle with 3 accident periods. The purpose of this application is to show the behaviour of our model compared to the CL on different datasets. In Section 5.2, we conclude with an application where we highlight the use of the CRPS measure to perform model selection, after having censored the data at an arbitrary calendar time, that is 5 calendar periods. Here we fit different models to the same datasets and select the best-performing model in terms of CRPS to calculate the loss reserve. The applications in this section also comprise the IBNR model that we have illustrated in the previous sections. In fact, we predict $\hat{Y}^{\text{TOT}} = Y^{\text{Closed}} + \hat{Y}^{\text{RBNS}} + \hat{Y}^{\text{IBNR}}$ and $\widehat{\text{Var}}(Y^{\text{TOT}}) = \widehat{\text{Var}}(Y^{\text{RBNS}}) + \widehat{\text{Var}}(Y^{\text{IBNR}})$.

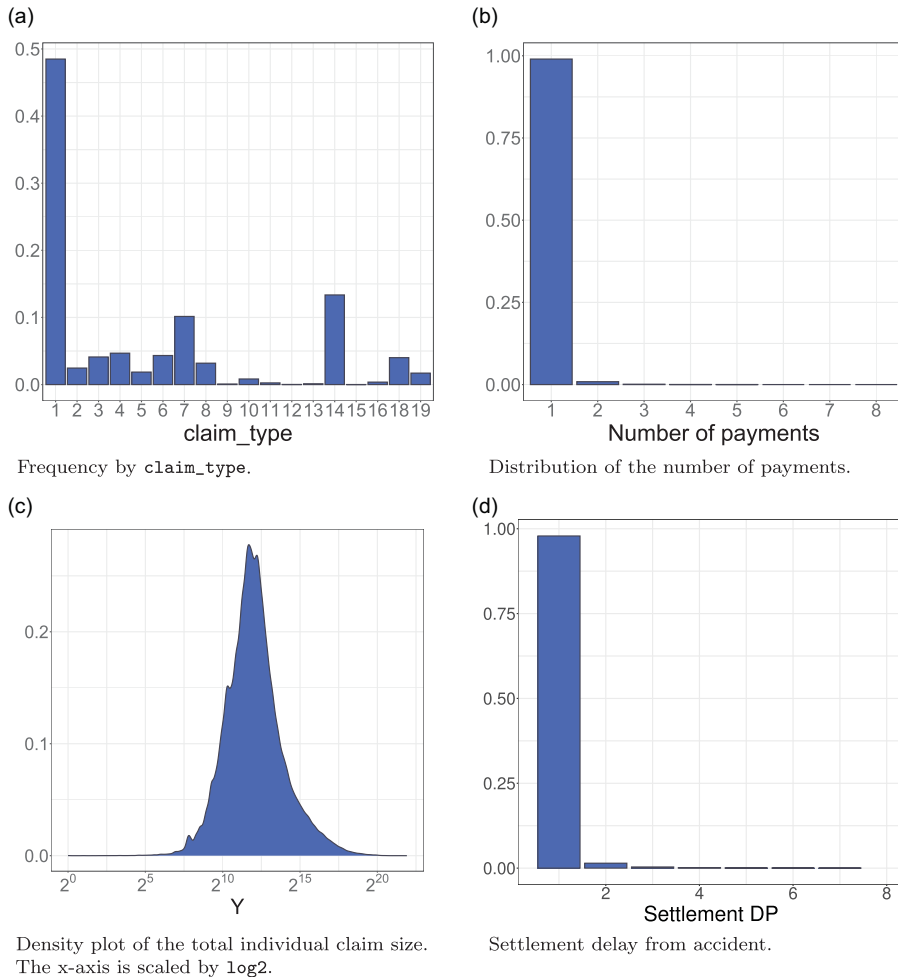


Figure 4. Exploratory data analysis on the real dataset that we use in this section. We show the relative frequencies by type of claim Figure (a), the distribution of the number of payments Figure (b), the density plot of the total individual claim size Figure (c) and the distribution of the settlement delay from accident Figure (d).

5.1. Model comparison on different datasets

Consider the first application, where we censor our data at different depths and fit an AJ model using the largest possible state space. Using our model, it is possible to explore and compare the individual claim size curves for the different values of `claim_type`, as shown in Figure 5.

The results are reported in Table 4, and they show that in general the AJ models outperform the CL in terms of EI, except for the scenario with $k = 7$. For each scenario, we are able to select a model with or without features using the CRPS measure. The rows with a \checkmark (column two) refer to models that include the `claim_type` feature in the fit. The target Y^{TOT} is shown in column three. While in other scenarios, the CRPS for the AJ model that includes the `claim_type` feature is close to or better than the CRPS of the model without the feature, in $k = 5$ the model without the feature has a much lower CRPS. The relative variability results show that AJ and CL have comparable results. In fact, only for $k = 4$ does CL have a lower relative variability.

For each choice of k , we can simply select the best-performing model in terms of CRPS and compute the claims reserve, the results are displayed in Table 5.

Table 4. For different choices of k (column one), we fit a model with and without `claim_type` (column two). The target Y^{TOT} is shown in column three. The EI is shown in columns four and five. The CV of Y^{TOT} are displayed in columns seven and eight for the AJ and CL respectively. The CRPS is shown in column nine.

k	<code>claim_type</code>	Y^{TOT}	EI (AJ)	EI (CL)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (AJ)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (CL)	CRPS (average, relative)
4	✓	616.1327	−0.0035	0.0157	0.0029	0.0023	1.0000
	✗		−0.0029		0.0029		1.1403
5	✓	822.5956	−0.0064	0.0209	0.0008	0.0024	1.0000
	✗		−0.0061		0.0007		0.4596
6	✓	999.6005	−0.0059	0.0173	0.0017	0.0017	1.0000
	✗		−0.0052		0.0017		0.9987
7	✓	1190.9112	−0.0146	0.0144	0.0011	0.0017	1.0000
	✗		−0.0142		0.0011		1.0022




Figure 5. For each different dimension $k = 4, 5, 6, 7, 8$, we provide the individual claim size curves for our observations by covariate value `claim_type`. The x-axis represents Z and we scale it by \log_2 to ease the plot visualization.

Table 5. We selected for each data k (column one), the best-performing model in terms of CRPS and present the reserve (column one) and the standard deviation of the reserve (column two).

k	R	$\text{sd}(Y^{\text{TOT}})$	claim_type
4	10.7832	1.8054	✓
5	7.4981	0.5720	✗
6	10.3268	1.6731	✗
7	5.6317	1.3381	✓

Table 6. For the dataset with depth 5 accident periods, we fit the AJ model both including and excluding the claim_type feature (column two). The target Y^{TOT} is displayed in column three. The EI for the AJ and the CL can be found in columns four and five. The coefficient of variation of Y^{TOT} is displayed in columns six and seven. The CRPS is displayed in column eight.

k	claim_type	Y^{TOT}	EI (AJ)	EI (CL)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (AJ)	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$ (CL)	CRPS (average, relative)
4	✗	999.6005	−0.0059	0.0173	0.0016	0.0017	1.0143
	✓		−0.0040		0.0020		1.0000
5	✗		−0.0069		0.0015		0.9916
	✓		−0.0052		0.0018		1.0000
6	✗		−0.0052		0.0017		0.9987
	✓		−0.0059		0.0017		1.0000

5.2. Model comparison on a single dataset

In the second application, we censor our data after 5 accident period and construct the AJ model for $k = 4, 5, 6$. For each choice of k , we fit a model both with and without the feature claim_type. Table 6 shows that in terms of CRPS, the model with $k = 4$ using the claim_type feature is the best model (minimum CRPS). This model is also the best in terms of EI.

Using the CRPS, we are able to select the model with $k = 4$ and that uses claim_type as a feature to compute the claims reserve. The model provides a reserve of 11.485 millions and a standard deviation of 2.0385 millions. The strategy presented in this last section is of particular interest to reserving actuaries. While the empirical study in Section 5.1 is interesting for understanding the behaviour of our model on different data, in practice the maximum depth of the data is selected using indicators such as claims settlement speed or using the expert judgement of an experienced actuary Wüthrich and Merz (2008, p. 12). The depth of the data is then known and the main interest is simply to select the best model, as shown here. Selecting the most appropriate model can be very difficult, see again Wüthrich and Merz (2008, p. 13), and the strategy presented in this section provides a mathematically sound approach to doing so.

6. Conclusion

The methodology described in this manuscript presents an improved approach for predicting loss reserves compared to commonly used aggregate loss reserving methods. We also introduce a comprehensive non-parametric estimator for the cumulative density function of individual claim sizes, distinguishing it from other individual loss reserving methods that mainly focus on point forecasts of claim reserves. While some questions remain open, particularly regarding the modeling of IBNR claims and their dependence on specific features, practical applications can address these issues by fitting our

model with the chain ladder method for future reports and projecting their future costs using the unconditional curve of individual claim sizes. Furthermore, this methodology can be extended in various ways. One possibility is to explore the use of reverse time models, as described in Hiabu (2017), to project the future exposure of IBNR claims. In a broader context, future research could investigate how to model not only the size but also the frequency of individual claims.

Supplementary material. The code at <https://github.com/gpitt71/conditional-aj-reserving> complements the results of this manuscript. The folder was registered with unique Zenodo DOI [10.5281/zenodo.10118895](https://doi.org/10.5281/zenodo.10118895).

Competing interests. The authors declare no competing interests of any kind.

References

- Al-Mudafer, M.T., Avanzi, B., Taylor, G. and Wong, B. (2022) Stochastic loss reserving with mixture density neural networks. *Insurance: Mathematics and Economics*, **105**, 144–174.
- Andersen, P.K., Borgan, O., Gill, R.D. and Keiding, N. (1993) *Statistical Models Based on Counting processes*. New York, NY: Springer-Verlag New York, Inc..
- Antonio, K. and Plat, R. (2014) Micro-level stochastic loss reserving for general insurance. *Scandinavian Actuarial Journal*, **2014**(7), 649–669.
- Arjas, E. (1989) The claims reserving problem in non-life insurance: Some structural ideas. *ASTIN Bulletin: The Journal of the IAA*, **19**(2), 139–152.
- Avanzi, B., Taylor, G. and Wang, M. (2023) SPLICE: A synthetic paid loss and incurred cost experience simulator. *Annals of Actuarial Science*, **17**(1), 7–35.
- Avanzi, B., Taylor, G., Wang, M. and Wong, B. (2021) SynthETIC: An individual insurance claim simulator with feature control. *Insurance: Mathematics and Economics*, **100**, 296–308.
- Bischofberger, S.M., Hiabu, M. and Isakson, A. (2020) Continuous chain-ladder with paid data. *Scandinavian Actuarial Journal*, **2020**(6), 477–502.
- Bishop, C.M. (1994) *Mixture Density Networks*. Birmingham: Aston University.
- Bladt, M. and Furrer, C. (2023a) *AalenJohansen: Conditional Aalen-Johansen Estimation*. R package version 1.0. url: <https://github.com/martinbladt/AalenJohansen>.
- Bladt, M. and Furrer, C. (2023b) Conditional Aalen–Johansen estimation. arXiv:2303.02119.
- Buckley, J. and James, I. (1979) Linear regression with censored data. *Biometrika*, **66**(3), 429–436.
- Crevecoeur, J., Antonio, K., Desmedt, S. and Masquelein, A. (2023) Bridging the gap between pricing and reserving with an occurrence and development model for non-life insurance claims. *ASTIN Bulletin: The Journal of the IAA*, **53**(2), 185–212.
- Crevecoeur, J. and Robben, J. (2024) *hirem: Hierarchical reserving models*. R package version 0.1.0.
- Crevecoeur, J., Robben, J. and Antonio, K. (2022) A hierarchical reserving model for reported non-life insurance claims. *Insurance: Mathematics and Economics*, **104**, 158–184.
- Dabrowska, D.M. (1997) Smoothed cox regression. *The Annals of Statistics*, **25**(4), 1510–1540.
- Delong, L., Lindholm, M. and Wüthrich, M.V. (2022) Collective reserving using individual claims data. *Scandinavian Actuarial Journal*, **2022**(1), 1–28.
- Delong, L. and Wüthrich, M.V. (2023) The Role of the Variance Function in Mean Estimation and Validation. *Available at SSRN* 4477677.
- England, P. and Verrall, R. (1999) Analytic and bootstrap estimates of prediction errors in claims reserving. *Insurance: Mathematics and Economics*, **25**(3), 281–293.
- Friedland, J. (2010) Estimating unpaid claims using basic techniques. *Casualty Actuarial Society*, Vol. **201**. 0.
- Gabrielli, A. and Wüthrich, M.V. (2018) An individual claims history simulation machine. *Risks*, **6**(2), 29.
- Ghosh, S. (2018) *Kernel Smoothing: Principles, methods and applications*, 1st ed. Hoboken, NJ: John Wiley & Sons, Inc.
- Gill, R.D. (1986) On estimating transition intensities of a Markov process with aggregate data of a certain type: “Occurrences but No Exposures”. *Scandinavian Journal of Statistics*, **13**(2), 113–134.
- Gneiting, T. and Raftery, A.E. (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, **102**(477), 359–378.
- Gneiting, T. and Ranjan, R. (2011) Comparing density forecasts using threshold-and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, **29**(3), 411–422.
- Haastруп, S. and Arjas, E. (1996) Claims reserving in continuous time; a nonparametric Bayesian approach. *ASTIN Bulletin: The Journal of the IAA*, **26**(2), 139–164.
- Hess, K.T. and Schmidt, K.D. (2002) A comparison of models for the chain-ladder method. *Insurance: Mathematics and Economics*, **31**(3), 351–364.
- Hesselager, O. (1994) A Markov model for loss reserving. *ASTIN Bulletin: The Journal of the IAA*, **24**(2), 183–193.
- Heuchenne, C. and Van Keilegom, I. (2007) Nonlinear regression with censored data. *Technometrics*, **49**(1), 34–44.

- Hiabu, M. (2017) On the relationship between classical chain ladder and granular reserving. *Scandinavian Actuarial Journal*, **2017**(8), 708–729.
- Hoem, J.M. (1969) Markov chain models in life insurance. *Blätter der DGVFM*, **9**(2), 91–107.
- Hoem, J.M. (1972) Inhomogeneous semi-Markov processes, select actuarial tables, and duration-dependence in demography. In *Population Dynamics*. Elsevier, pp. 251–296.
- Huang, J., Qiu, C., Wu, X. and Zhou, X. (2015) An individual loss reserving model with independent reporting and settlement. *Insurance: Mathematics and Economics*, **64**, 232–245.
- Klugman, S.A., Panjer, H.H. and Willmot, G.E. (2012) *Loss Models: From Data to Decisions*. Vol. **715**. Hoboken, NJ: John Wiley & Sons.
- Lopez, O. (2019) A censored copula model for micro-level claim reserving. *Insurance: Mathematics and Economics*, **87**, 1–14.
- Lopez, O. and Milhaud, X. (2021) Individual reserving and nonparametric estimation of claim amounts subject to large reporting delays. *Scandinavian Actuarial Journal*, **2021**(1), 34–53.
- Lopez, O., Milhaud, X. and Thérond, P.-E. (2019) A tree-based algorithm adapted to microlevel reserving and long development claims. *ASTIN Bulletin: The Journal of the IAA*, **49**(3), 919–919.
- Maciak, M., Mizera, I. and Pešta, M. (2022) Functional profile techniques for claims reserving. *ASTIN Bulletin: The Journal of the IAA*, **52**(2), 449–482.
- Mack, T. (1993) Distribution-free calculation of the standard error of chain ladder reserve estimates. *ASTIN Bulletin: The Journal of the IAA*, **23**(2), 213–225.
- Mack, T. and Venter, G. (2000) A comparison of stochastic models that reproduce chain ladder reserve estimates. *Insurance: Mathematics and Economics*, **26**(1), 101–107.
- McKeague, I.W. and Utikal, K.J. (1990) Inference for a nonlinear counting process regression model. *The Annals of Statistics*, **18**(3), 1172–1187.
- Miranda, M.D.M., Nielsen, J.P. and Verrall, R. (2012) Double chain ladder. *ASTIN Bulletin: The Journal of the IAA*, **42**(1), 59–76.
- Nielsen, J.P. and Linton, O.B. (1995) Kernel estimation in a nonparametric marker dependent hazard model. *The Annals of Statistics*, **23**(5), 1735–1748.
- Norberg, R. (1986) A contribution to modelling of IBNR claims. *Scandinavian Actuarial Journal*, **1986**(3-4), 155–203.
- Norberg, R. (1993) Prediction of outstanding liabilities in non-life insurance. *ASTIN Bulletin: The Journal of the IAA*, **23**(1), 95–115.
- Norberg, R. (1997) Prediction of outstanding liabilities II: Model variations and extensions. *Insurance Mathematics and Economics*, **3**(21), 251.
- Overgaard, M. (2019) State occupation probabilities in non-Markov models. *Mathematical Methods of Statistics*, **28**, 279–290.
- Parodi, P. (2014) *Pricing in General Insurance. first*. CRC Press. isbn: 9780429166471.
- Stute, W. (1986) On almost sure convergence of conditional empirical distribution functions. *The Annals of Probability*, **14**(3), 891–901.
- Therneau, T.M. (2023). *A Package for Survival Analysis in R*. R package version 3.5-5. url: <https://CRAN.R-project.org/package=%3Dsurvival>.
- Van Houwelingen, H.C. (2007) Dynamic prediction by landmarking in event history analysis. *Scandinavian Journal of Statistics*, **34**(1), 70–85.
- Verrall, R.J. (2000) An investigation into stochastic claims reserving models and the chain-ladder technique. *Insurance: Mathematics and Economics*, **26**(1), 91–99.
- Wang, M. and Wüthrich, M.V. (2022) Individual Claims Generator for Claims Reserving Studies: data simulation. R. UNSW Business School Research Paper Forthcoming. Available at SSRN.
- Wang, Z., Wu, X. and Qiu, C. (2021) The impacts of individual information on loss reserving. *ASTIN Bulletin: The Journal of the IAA*, **51**(1), 303–347.
- Wüthrich, M.V. (2018) Neural networks applied to chain–ladder reserving. *European Actuarial Journal*, **8**, 407–436.
- Wüthrich, M.V. and Merz, M. (2008) *Stochastic Claims Reserving Methods in Insurance*. John Wiley & Sons.

Appendix

A Comparison with hirem

We perform a first comparison with the `hirem` model on the scenario Beta that we introduced in Section 4.3. The results are reported in Table A1. For the different values of k (column one), we compare the AJ model (with feature U) to the CL and `hirem` in terms of EI and CRPS. Each row corresponds to a different model (column two). The (average) target Y^{TOT} is reported in column three. The EI (column four) shows that the AJ model is predicting the target Y^{TOT} more accurately than the CL and `hirem`. We also observe that the `hirem` model seems to be less accurate for higher values of k . Consistently with the data generating assumptions, the CRPS indicates that the AJ model better describes the curve of the claim size. The relative variation is reported from Table 2 the AJ model and the CL model (column five). We observe a similar relative variation of Y^{TOT} on our models compared to `hirem`.

Table A1. Results for scenario Beta. For each value of k (column one) we present the average results over the 20 simulations. Each row of the table corresponds to a different model (column two). The table includes the (average) actual simulated total cost (column three) and the error incidence for the AJ, the CL, and hirem (columns four). In columns five we show the average relative variation and in column six we reported the average CRPS relative to the AJ CRPS.

Scenario	Model	Y^{TOT} (average)	EI	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$	CRPS (average, relative)
Beta ($k = 4$)	AJ	42.279	0.006	0.008	1.000
	CL		0.103	0.008	–
	hirem		–0.028	0.009	1.582
Beta ($k = 5$)	AJ	94.580	0.011	0.009	1.000
	CL		0.190	0.008	–
	hirem		0.071	0.009	1.680
Beta ($k = 6$)	AJ	108.423	–0.013	0.007	1.000
	CL		0.256	0.008	–
	hirem		0.148	0.009	1.697
Beta ($k = 7$)	AJ	116.032	–0.033	0.007	1.000
	CL		0.255	0.008	–
	hirem		0.168	0.008	1.790

Table A2. Description of the hirem data.

Covariates	Description
claim.nr	Policy identifier
Type $\in \{T1, T2, T3\}$	Type of claim
Hidden $\in \{I, M, H\}$	Covariate unknown to the insurer (disregarded for modeling)
occ.year	Accident year
rep.year	Calendar year of reporting
dev.year	Development year
size	Incremental paid amount
settlement	Indicator, 1 in the development year of settlement

A second comparison is performed on synthetic data generated from the individual claims simulators available in the R package (Crevecoeur and Robben, 2024). The package includes the hirem models implementation and four individual claims simulators for reported claims for $k = 10$. While the four scenarios are briefly introduced in this section, we refer to the package documentation and the main manuscript for a detailed description. The generated features that are relevant for our application are reported in Table A2.

The data include one categorical feature, the type of claim (Type), which we use in our application and a hidden feature (Hidden) that is not known to the insurer at the evaluation. The package includes a baseline simulated scenario (Baseline) and three scenarios that include modifications compare to the Baseline. In the scenario Claim Mix, there is a change in the portfolio distribution with respect to the feature Type. In the scenario Extreme, a seasonality effect in calendar year 9 simulates an increase in the claim occurrences. Lastly, the Settlement scenario simulates a change of the claim settlement process to quicker settlement after calendar period 7.

The results of this second application are reported Table A3. We find that, both in terms of EI and CRPS and in the four simulated scenarios, our models always outperform the CL, while hirem performs better than the AJ (columns five and seven). The CRPS in column seven are relative to the CRPS of the

Table A3. The `hirem` package includes four data generators in four different scenarios with $k = 10$ (column one). We compare our models (AJ with and without features, column two) to the model in Crevecoeur et al. (2023) and the CL (column three). The actual Y^{TOT} target is reported in column four. We show the EI and the CRPS results in columns five and seven, respectively. The predicted relative variation of Y^{TOT} is shown in column six.

Scenario	Type	Model	Y^{TOT} (average)	EI	$\hat{\text{sd}}(Y^{\text{TOT}})/\hat{Y}^{\text{TOT}}$	CRPS (average, relative)
Baseline	✓	AJ	50.240	0.070	0.004	1.000
	✗			0.067	0.004	1.001
	✓	hirem		0.012	0.002	0.785
	✗	CL		0.232	0.010	
Claim Mix	✓	AJ	57.880	0.071	0.004	1.000
	✗			0.066	0.004	1.002
	✓	hirem		0.014	0.002	0.792
	✗	CL		0.240	0.010	
Extreme	✓	AJ	48.812	0.076	0.005	1.000
	✗			0.069	0.004	1.002
	✓	hirem		0.039	0.002	0.755
	✗	CL		0.231	0.010	
Settlement	✓	AJ	50.241	0.069	0.004	1.000
	✗			0.065	0.004	1.001
	✓	hirem		0.011	0.002	0.766
	✗	CL		0.231	0.001	

AJ model with features. These results are expected, as the data are generated using assumptions that are consistent with the `hirem` model, see Appendix A of Crevecoeur et al. (2023). Notwithstanding, the AJ model is performing comparatively well. The results are averaged over ten simulations.

We remark that the applications reported in Tables A1 and A3 indicate that, in general, for the scenarios that we inspected, individual models always outperform the chain ladder benchmark. However, as expected, the models perform best when the data generating process is consistent with the models assumptions.