


RESEARCH ARTICLE

Marked Cox models for IBNR claims count: continuous and discretized approaches with Dirichlet-driven reporting delays

Hassan Abdelrahman , Andrei L. Badescu, Radu V. Craiu and X. Sheldon Lin

Department of Statistical Sciences, University of Toronto, Toronto, ON M5G 1Z5, Canada

Corresponding author: Hassan Abdelrahman; Email: hassan.abdelrahman@mail.utoronto.ca

Received: 16 September 2024; **Revised:** 29 May 2025; **Accepted:** 2 June 2025; **First published online:** 11 July 2025

Keywords: Dirichlet distribution; EM algorithm; hidden Markov model; IBNR reserving; marked Cox model

Abstract

We propose a novel micro-level Cox model for incurred but not reported (IBNR) claims count based on hidden Markov models. Initially formulated as a continuous-time model, it addresses the complexity of incorporating temporal dependencies and policyholder risk attributes. However, the continuous-time model faces significant challenges in maximizing the likelihood and fitting right-truncated reporting delays. To overcome these issues, we introduce two discrete-time versions: one incorporating unsystematic randomness in reporting delays through a Dirichlet distribution and one without. We provide the EM algorithm for parameter estimation for all three models and apply them to an auto-insurance dataset to estimate IBNR claim counts. Our results show that while all models perform well, the discrete-time versions demonstrate superior performance by jointly modeling delay and frequency, with the Dirichlet-based model capturing additional variability in reporting delays. This approach enhances the accuracy and reliability of IBNR reserving, offering a flexible framework adaptable to different levels of granularity within an insurance portfolio.

1. Introduction

In property and casualty (P&C) insurance, accurate loss reserving is vital for financial stability, regulatory compliance, and effective risk management (Bjarnason and Sjögren, 2014). Reserves typically comprise two parts: Reported but not settled (RBNS) and incurred but not reported (IBNR) claims. This paper focuses on the latter, where the number of unreported claims at the valuation date is unknown. Estimating IBNR requires advanced modeling to predict both the frequency and severity of claims, making it a complex yet crucial aspect of the reserving process.

Most of the literature on modeling insurance claims and reserving is based on aggregated claims data, such as the chain ladder method (Mack, 1993, 1999) and the Bornhuetter-Ferguson method (Bornhuetter and Ferguson, 1972). These “macro-level” models summarize claim development by aggregating payments across occurrence and development years. Subsequently, stochastic models were introduced to account for the variability in reserves derived from these models; see Taylor (2012), England and Verrall (2002), and Wüthrich and Merz (2008) for a comprehensive overview of stochastic reserving. Despite their simplicity, these models come with several disadvantages, including the loss of valuable information on the characteristics of individual claims, the high uncertainty of parameter estimates, and potential bias under certain conditions (Charpentier and Pigeon, 2016). To resolve these shortcomings, “micro-level” models have been proposed with the aim to use individual-level data to depict the development of individual claims.

Arjas (1989) and Norberg (1993), Norberg (1999) are the first to propose a stochastic loss reserving framework at the individual claim level. They proposed a marked non-homogeneous Poisson process (NHPP) model and a general mathematical framework for the development of individual claims. Haastrup and Arjas (1996) then implemented the NHPP model using non-parametric Bayesian statistics. The subject of micro-level reserving remained unopened before emerging back in recent years. Some early contributions include Larsen (2007), who revisited the framework of Arjas (1989) and Norberg (1993), Norberg (1999), and Rosenlund (2012), who proposed a deterministic reserving approach that allows conditioning on individual claim characteristics. Antonio and Plat (2014) presented a case study in which they applied Norberg's model to an auto-insurance dataset. Data-driven comparisons between macro-level and micro-level reserving have shown the superiority of the micro-level models; see Charpentier and Pigeon (2016) and Huang *et al.* (2015), Huang *et al.* (2016).

Since then, various research streams have emerged in the micro-level estimation of IBNR reserves. One stream employs machine learning and neural network techniques for IBNR claim estimation. As described in Bücher and Rosenstock (2023), this approach utilizes Chain Ladder (CL)-based methods to estimate IBNR reserves over discrete time intervals, as demonstrated by studies such as Wüthrich (2018a), Wüthrich (2018b), Baudry and Robert (2019), De Felice and Moriconi (2019), Delong *et al.* (2022) and Bücher and Rosenstock (2023). Another stream focuses on modeling the claim arrival process, where the NHPP model has traditionally played a prevalent role. However, NHPPs assume temporal independence in claims count and overlook portfolio-wide environmental effects, as noted by Badescu *et al.* (2016) and Grandell (2012), p. 54. This limits their ability to capture dependencies and explain the calendar year effect observed in run-off triangles (Holmberg, 1994, Shi *et al.*, 2012).

Addressing these issues involves incorporating a temporal dependence structure into the claim arrival process model, which can be achieved by transitioning to a Cox process. Notable research in this direction includes Avanzi *et al.* (2016), who utilized a shot noise Cox process, and Badescu *et al.* (2016), who employed a hidden Markov model (HMM) with an Erlang state-dependent distribution to model the claim arrival process. It is crucial to note that while both studies incorporated the Cox model, they focused on modeling the claim arrival process for the entire portfolio, overlooking policyholder information that is likely to enhance the accuracy of the different components within the model.

To this end, we propose a novel micro-level Cox model that flexibly accommodates various levels of granularity, including the policyholder level, lines of business, or the entire portfolio. Building on Badescu *et al.* (2016), our model incorporates a shared HMM that governs claim arrivals across granular units. At the policyholder level, the HMM captures environmental dynamics that influence all policyholders simultaneously. Conditional on the HMM state, each policyholder's claim arrival intensity is constant and depends on their risk attributes through a regression function. Thus, the intensity reflects both shared external influences and individual risk profiles. This framework captures temporal dependence and environmental variation often overlooked in traditional reserving models. We show that, within our framework, the discretely observed claim arrival, reported claim, and IBNR claim processes for each policyholder belong to the family of Poisson-HMMs, with state-dependent distributions shaped by the policyholder's risk attributes.

We initially propose our model in continuous time. In general, fitting continuous-time marked count process models to fully granular reported claims data – that is, data including exact accident times and reporting delays – presents significant challenges. One of the main challenges is maximizing the likelihood of observed data. Maximizing the likelihood of the observed data entails addressing two distinct components: the likelihood for the truncated reporting delay and the likelihood for the claim frequency model of reported claims, which in turn depends on the reporting delay. Thus, maximizing the likelihood poses a formidable task due to the interdependence of the two components. Traditionally, researchers have adopted a two-step approach to maximize the likelihood of the observed data: first fitting the right-truncated reporting delay distribution, and then estimating the frequency model by treating the estimated delay distribution as fixed. However, this method overlooks the inherent interaction between the frequency and reporting delay components; an increase in reported claims could result from a higher claim frequency, faster reporting, or a combination of both.

We address such a challenge within our marked Cox model framework by integrating reporting delay and claim frequency modeling cohesively, transitioning from a continuous-time framework to a discrete-time framework. Specifically, we aggregate continuous-time events to a coarser time scale, which is more common in practice. This coarser time scale (or period) can be days, weeks, months, etc. We assume a hierarchical structure: if $N_{i,t}$ denotes the number of claims that occurred in period t for the policyholder i , and $Z_{i,t,d}$ denotes the number of these claims reported d periods later, then we assume that $Z_{i,t,d} | N_{i,t}$ follows a multinomial distribution. We simultaneously model both the delay and the occurrence components, capturing the intricate interplay between claim occurrence and reporting behavior. This type of discrete-time hierarchical structure has been explored in the literature. For instance, Verbelen *et al.* (2022) considered a model in which the underlying claim counts follow a non-homogeneous Poisson process, and employed the EM algorithm to maximize the joint likelihood of the claim occurrence and reporting delay components. Similarly, Avanzi *et al.* (2016) modeled claim arrivals via a shot noise Cox process, using reversible jump MCMC for joint inference within a discrete-time framework. Notably, both of these approaches do not incorporate policyholder-specific covariates. Crevecoeur *et al.* (2023) extended the structure used in Verbelen *et al.* (2022) by incorporating covariates into the model, though assuming a homogeneous Poisson process for the underlying claim counts and applying a similar EM-based estimation procedure.

Nevertheless, the existing literature often assumes static or time-dependent covariates to model reporting delay variations. However, practical scenarios frequently involve unexplainable changes in delay structures that would not be captured through conventional covariates. As previously demonstrated by Swamy (1970) and Sentürk and Müller (2005), accounting for such unexplainable changes requires us to assume that the delay probabilities are random. To address this, we extend the discrete-time model to account for the likely occurrence of unobserved confounding covariates by assuming that the probabilities of the multinomial distribution follow a Dirichlet distribution.

We outline the EM algorithm needed to obtain the maximum likelihood estimates for our proposed models. In particular, for the discrete-time model with the Dirichlet assumption, we employ the Monte Carlo EM (MCEM) algorithm, which uses Monte Carlo simulations to approximate the E-step and handles the complexities introduced by the Dirichlet distribution. We also fit our models to an auto-insurance dataset from a major European insurance company. We show that despite losing information by switching to the discrete-time framework, the discrete-time models perform better than the continuous-time model because they account for the joint modeling of occurrence and reporting. Moreover, we demonstrate that adding the Dirichlet assumption provides more realistic interval estimates for the IBNR claim count.

This work contributes to the literature by:

- Extending the HMM-based Cox model of Badescu *et al.* (2016); Badescu *et al.* (2019) to various levels of granularity (e.g., policyholder or business line), which is essential when portfolio composition shifts and claim behaviors vary across units.
- Extending the EM-based estimation framework of Verbelen *et al.* (2022) by modeling the underlying occurrence process as a Poisson-HMM rather than a Poisson process and conducting the analysis at the micro (policyholder-level) scale.
- Unsystematic variation in reporting delays by modeling the delay probabilities via a Dirichlet distribution – introducing randomness that was not accounted for in previous approaches such as Verbelen *et al.* (2022) and Avanzi *et al.* (2016) – and outlining the corresponding MCEM algorithm for model fitting.

The paper is structured as follows. In Section 2, we outline our modeling framework and discuss its interpretability. The properties of the model and the associated individual-level reported claim process and IBNR process are presented in Section 3. In Section 4, we outline the discrete-time version of the model and show the likelihood for the three different models under consideration. In Section 5, we provide the EM algorithm needed to estimate the parameters of the three different models. Section 6

details the application of the model in predicting the IBNR claims count and reserve, while we apply the models to real-life data in Section 7 to assess their performance in estimating the IBNR claim count. We conclude in Section 8.

2. Modeling framework

We present our proposed model for modeling the claim arrival process at the policyholder level. In line with the notation established by Badescu *et al.* (2016), we describe the arrival of the k -th claim for the i th policy ($i = 1, \dots, m$) through the random variables (T_{ik}, U_{ik}) , where T_{ik} denotes the occurrence time of the claim and U_{ik} denotes its reporting delay. Chronologically ordered, $\{(T_{ik}, U_{ik}), k = 1, 2, \dots\}$ constitutes the claim arrival process for policy i , whose risk attributes are denoted by $\mathbf{x}_i \in \mathcal{X}$, where \mathcal{X} is the covariate space. Additionally, the total claim count process for the i th policy is defined as $N_i(t) = \sum_{k=1}^{\infty} \mathbb{1}\{T_{ik} < t\}$, where $\mathbb{1}\{\cdot\}$ is the indicator function.

In our modeling approach, we represent $\{N_i(t), t \geq 0\}$ as a marked Cox process with two components:

1. The marks $\{U_{i1}, U_{i2}, \dots\}$, and
2. The stochastic intensity function $\Lambda_i(t)$.

The assumption is made that the marks are independent, with a density function given by $f_{U|t, \mathbf{x}_i}(u)$. Note that the distribution of the reporting delay is conditioned on the time of claim occurrence and the risk attributes of the policyholder. Additionally, the stochastic intensity function $\Lambda_i(t)$ is modeled as a piecewise stochastic process, where $\Lambda_i(t) = e_{i,l} \Lambda_l(t; \mathbf{x}_i)$ for $d_{l-1} \leq t < d_l$, with $l = 1, 2, \dots$ and $d_0 = 0$. The time points $d_l, l = 1, 2, \dots$, are predetermined, and $t = d_0 = 0$ marks the beginning of the observation window for the entire portfolio. Moreover, $e_{i,l} \in [0, 1]$ denotes the exposure of the i th policy in the interval $d_{l-1} \leq t < d_l$, defined as the proportion of days in that interval for which the contract is in force. Notably, $e_{i,l}$ can be zero, indicating the case when the i th policy's contract is not active during $d_{l-1} \leq t < d_l$.

To describe $\{\Lambda_1(t; \mathbf{x}_i), \Lambda_2(t; \mathbf{x}_i), \dots\}$, we propose a structure involving two components:

1. A hidden parameter process $\{C_1, C_2, \dots\}$, which constitutes a time-homogeneous Markov chain with a finite state space $\{1, \dots, g\}$. Denoting its initial distribution and transition probability matrix as $\boldsymbol{\pi}_1$ and $\boldsymbol{\Gamma} = \{\gamma_{jk}\}$, respectively, where $\gamma_{jk} = P(C_l = k | C_{l-1} = j)$, we assume the Markov chain is irreducible, aperiodic, and has all states being positive recurrent, leading to a unique limiting distribution.
2. A state-dependent process $\{\Lambda_1(t; \mathbf{x}_i), \Lambda_2(t; \mathbf{x}_i), \dots\}$, which depends on the current state C_l . Given $C_l = j$, we assume that $\Lambda_l(t; \mathbf{x}_i)$ is constant such that

$$\Lambda_l(t; \mathbf{x}_i) | (C_l = j) = \lambda^{(j)}(\mathbf{x}_i),$$

where $\lambda^{(j)} : \mathcal{X} \rightarrow \mathbb{R}^+$ is a regression function.

Specifically, each policyholder's risk attributes \mathbf{x}_i are incorporated into the model through the regression function $\lambda^{(j)}(\mathbf{x}_i)$. These risk attributes could include factors such as age, location, coverage type, and any other relevant information about the policyholder that may affect claim occurrence rates. By including these covariates, we tailor the intensity function $\Lambda_l(t; \mathbf{x}_i)$ to reflect the unique risk profile of each policyholder.

In summary, the intensity for each policy i remains constant during each period l , contingent on the realization of hidden parameter C_l , and is determined by the regression function based on policyholder's risk attributes. The hidden states represent environmental variation (e.g. weather or seasonal effect), affecting all policies in the portfolio, with policies assumed independent given the environmental variation. It is important to note that the state of the l th period is identical for all policies i .

Consequently, we have that, for each policy i ($i = 1, \dots, m$),

$$\Lambda_i(t; \mathbf{x}_i) = \begin{cases} \lambda^{(1)}(\mathbf{x}_i), & \text{with prob.} = \pi_{i1} \\ \vdots & \vdots \\ \lambda^{(g)}(\mathbf{x}_i), & \text{with prob.} = \pi_{ig} \end{cases}, \quad t \in [d_{l-1}, d_l),$$

where $\pi_{ij} = P(C_l = j)$ is the j th element of the row vector $\boldsymbol{\pi}_i = \boldsymbol{\pi}_1 \Gamma^{l-1}$.

Remarks

- Our Cox model generalizes the framework of Badescu *et al.* (2016) by modeling claim arrivals at the individual policyholder level, using covariates \mathbf{x}_i to capture heterogeneity across risk profiles. This allows the claim intensity to reflect both individual characteristics and broader environmental variation such as seasonality. Moreover, in contrast to their assumption of an Erlang-distributed intensity over $[d_i, d_{i+1})$, we model intensity as a covariate-dependent constant within each period. This choice simplifies estimation and leads to a Poisson mixture distribution, which is appropriate for overdispersed claim count data. While the Erlang assumption yields a Pascal mixture, it adds complexity without a clear improvement in model performance. Moreover, with appropriate segmentation of the timeline, the constant-intensity assumption remains a reasonable approximation.
- A special case of our framework arises when modeling claim arrivals across lines of business, with index i denoting each line. If the regression function $\lambda^{(j)}(\cdot)$ is constant, the model reduces to the contemporaneously conditionally independent multivariate HMM of MacDonald and Zucchini (2016). This HMM structure would capture dependence between lines of business through shared hidden states – an important feature given the emphasis in the literature on modeling such dependencies in reserving. Alternatively, i can represent claim types within a single line of business, allowing the model to account for heterogeneity in mark distributions across types.
- While our model assumes predetermined time intervals $[d_{l-1}, d_l)$, it would be ideal to select these intervals in a data-driven manner to better capture changes in intensity over time. Nonetheless, using fixed intervals – such as weekly, monthly, or quarterly – offers a practical and sufficiently flexible solution in most applications.

3. Properties of the model

In this section, we outline the immediate properties of our model, focusing on its application from the policy-holder level. At a given reserve valuation date τ , the complete claim arrival process is not fully observed; only reported claims are observable. Denoting the reported claim process for the i th policy with respect to τ as $\{N_i^r(t); 0 \leq t \leq \tau\}$, it is defined as $N_i^r(t) = \sum_{k=1}^{\infty} \mathbb{1}\{T_{ik} < t; T_{ik} + U_{ik} \leq \tau\}$, where T_{ik} denotes the claim occurrence time and U_{ik} denotes the reporting delay for the k th claim of the i th policyholder. Claims not reported by τ (i.e., $T_{ik} + U_{ik} > \tau$) contribute to the IBNR claim process, denoted as $\{N_i^{\text{ibnr}}(t); 0 \leq t \leq \tau\}$, and defined as $N_i^{\text{ibnr}}(t) = \sum_{j=1}^{\infty} \mathbb{1}\{T_{ik} < t; T_{ik} + U_{ik} > \tau\}$. Although both processes should be indexed by τ , we omit it for simplicity. Note that $N_i(t) = N_i^r(t) + N_i^{\text{ibnr}}(t)$.

It's notable that both the reported claim process and the IBNR claim process remain marked Cox processes, retaining easily convertible stochastic intensity functions and mark densities. We revisit a proposition proven in Theorem 3.1 of Badescu *et al.* (2016) to demonstrate this.

Proposition 1. Assume that the claim arrival process for the i th policyholder, $\{N_i(t); t \geq 0\}$, is a marked Cox process with stochastic intensity function $\Lambda_i(t)$ and independent marks $\{U_{i1}, U_{i2}, \dots\}$ with common density function $f_{U|r, \mathbf{x}_i}(u)$. Then for a given valuation date τ , its associated reported claim process

$\{N_i^r(t); 0 \leq t \leq \tau\}$ and IBNR claim process $\{N_i^{ibnr}(t); 0 \leq t \leq \tau\}$ are also marked Cox processes. Their adjusted stochastic intensity functions are

$$\Lambda_i^r(t) = \Lambda_i(t)F_{U|t,\mathbf{x}_i}(\tau - t)\mathbb{1}\{0 \leq t \leq \tau\} \quad \text{and} \quad \Lambda_i^{ibnr}(t) = \Lambda_i(t)(1 - F_{U|t,\mathbf{x}_i}(\tau - t))\mathbb{1}\{0 \leq t \leq \tau\},$$

respectively, and their independent marks follow adjusted position-dependent mark distributions

$$\frac{f_{U|t,\mathbf{x}_i}(u)\mathbb{1}\{u > \tau - t\}}{F_{U|t,\mathbf{x}_i}(\tau - t)} \quad \text{and} \quad \frac{f_{U|t,\mathbf{x}_i}(u)\mathbb{1}\{u \leq \tau - t\}}{1 - F_{U|t,\mathbf{x}_i}(\tau - t)},$$

respectively, where $F_{U|t,\mathbf{x}_i}$ is the distribution function for the reporting delay U given a claim occurrence time t and risk attributes \mathbf{x}_i .

Recall that we model the stochastic intensity function $\Lambda_i(t)$ of the claim arrival process $\{N_i(t; \mathbf{x}_i); t \geq 0\}$ as a piecewise stochastic process, where $\Lambda_i(t) = e_{i,l}\Lambda_i(t; \mathbf{x}_i)$ for $d_{l-1} \leq t < d_l$, with $d_0 = 0$, and the time points $d_l(l = 1, 2, \dots)$ are predetermined. These predetermined time points constitute the time units which make up the time series on which the HMM is defined. We define $N_{i,l} := N_i(d_l) - N_i(d_{l-1})$ as the number of claims occurring during $[d_{l-1}, d_l)$ for the i th policyholder ($i = 1, \dots, m$). Hence, $\{N_{i,1}, N_{i,2}, \dots\}$ represents the discrete observations of the claim arrival process at these time points d_l ($l = 1, 2, \dots$). The corresponding discrete observations of the reported claim process and the IBNR claim process with respect to a valuation date τ are denoted by $\{N_{i,1}^r, N_{i,2}^r, \dots\}$ and $\{N_{i,1}^{ibnr}, N_{i,2}^{ibnr}, \dots\}$, respectively. Notably, only $\{N_{i,1}^r, N_{i,2}^r, \dots\}$ is observed at time τ . It is straight-forward to see that these three discretely observed processes follow Poisson-HMM.

Proposition 2. For the proposed Cox model for $N_i(t; \mathbf{x}_i)$, the discretely observed claim arrival process, $\{N_{i,1}, N_{i,2}, \dots\}$, the discretely observed reported claim process, $\{N_{i,1}^r, \dots, N_{i,k}^r\}$, and the discretely observed IBNR claim process, $\{N_{i,1}^{ibnr}, \dots, N_{i,k}^{ibnr}\}$, all fall under the class of Poisson-HMMs, sharing the same hidden parameter process $\{C_1, C_2, \dots\}$ with $N_i(t)$. That is, for each $n \in \mathbb{N}_0$,

$$P(N_{i,l}^* = n \mid C_l = j) = p(n; \theta_l^{(j)}(\mathbf{x}_i)) := \frac{\theta_l^{(j)}(\mathbf{x}_i)^n e^{-\theta_l^{(j)}(\mathbf{x}_i)}}{n!},$$

where $N_{i,l}^*$ and $\theta_l^{(j)}$ denote any of the pairs $\{N_{i,l}, \tilde{\lambda}_l^{(j)}\}$, $\{N_{i,l}^r, \mu_l^{(j)}\}$, $\{N_{i,l}^{ibnr}, \nu_l^{(j)}\}$, with

$$\tilde{\lambda}_l^{(j)}(\mathbf{x}_i) = e_{i,l}\lambda^{(j)}(\mathbf{x}_i)(d_l - d_{l-1}), \quad \mu_l^{(j)}(\mathbf{x}_i) = e_{i,l}\lambda^{(j)}(\mathbf{x}_i) \int_{d_{l-1}}^{d_l} F_{U|t,\mathbf{x}_i}(\tau - t) dt,$$

$$\nu_l^{(j)}(\mathbf{x}_i) = e_{i,l}\lambda^{(j)}(\mathbf{x}_i) \int_{d_{l-1}}^{d_l} (1 - F_{U|t,\mathbf{x}_i}(\tau - t)) dt.$$

Consequently, for each $l = 1, 2, \dots$, the marginals $N_{i,l}$, $N_{i,l}^r$, and $N_{i,l}^{ibnr}$ follow mixed Poisson distributions with probability functions

$$P(N_{i,l}^* = n) = \sum_{j=1}^g \pi_{lj} p(n; \theta_l^{(j)}(\mathbf{x}_i)),$$

where $N_{i,l}^*$ and $\theta_l^{(j)}$ denote any of the pairs above, and $p(n; \lambda) = \frac{\lambda^n e^{-\lambda}}{n!}$.

The above proposition is straightforward given our assumption of constant intensity functions given a realization of the hidden state. When $C_l = j$, the intensity functions for the claim arrival process, reported claim process, and IBNR claim process become constant over the period $[d_{l-1}, d_l)$, resulting in Poisson processes for the claim arrival process and time-varying Poisson processes for the reported and IBNR claim processes. It is also evident that the three processes share the same hidden parameter process. This observation leads directly to the conclusion that all three observed processes are Poisson-HMMs.

While Poisson-HMMs have traditionally been employed to model claim count processes for entire portfolios, our application extends this methodology to the individual level (or other granular unit).

Properties similar to those in Badescu et al. (2016) can be easily derived, but we refrain from showing them in the manuscript, out of considerations related to the manuscript’s length.

4. Likelihood of the observed data

In this section, we propose the discrete-time version of our model, both with and without the Dirichlet assumption for the reporting probability vector. We begin by outlining the likelihood of observed data in the continuous-time framework and briefly discussing the associated complexities. The motivation for switching to a discrete-time framework arises from these complexities, particularly the difficulties in maximizing the likelihood of the observed data. This will serve as a motivation for our discrete-time alternative.

4.1. Continuous-time model

To estimate the IBNR reserve, we seek to maximize the likelihood of the observed data. Let $n_{i,l}^r$ denotes the number of claims incurred by policyholder (or any granular unit) i during period $[d_{l-1}, d_l]$ and are reported by $d_r = \tau$, where T represents the total number of observed periods. These claims, ordered by their arrival time, are described through the pairs $\{t_{ijl}, u_{ijl} : i = 1, \dots, m; j = 1, \dots, n_{i,l}^r\}$, where t_{ijl} is the claim arrival time of the j -th ordered reported claim for policyholder i in period $[d_{l-1}, d_l]$, and u_{ijl} is its reporting delay. We consider the set of observed data given by

$$\mathcal{O}_C = \{n_{i,l}^r, u_{ijl} : i = 1, \dots, m; j = 1, \dots, n_{i,l}^r; l = 1, \dots, T\}.$$

Note that since the Poisson intensity is assumed constant within each interval given the realization of the hidden state, the exact claims occurrence times do not influence the inference for the model parameters and are, therefore, omitted from the analysis. Let $N_i^{(r,s:t)} = (N_{i,s}^r, \dots, N_{i,t}^r)$ be the discretely observed reported claim process from period s to period t for the policyholder i , and let $\mathbf{n}_i^{(r,s:t)} = (n_{i,s}^r, \dots, n_{i,t}^r)$ be its realization, it is easy to show that the likelihood of the observed data is equal to (see, e.g., Antonio and Plat, 2014 and Fung et al., 2021):

$$\mathcal{L}^{(1)}(\Phi_1 | \mathcal{O}_C) = \underbrace{P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)})}_{\text{First Term (Claim Frequency)}} \times \underbrace{\prod_{i=1}^m \prod_{l=1}^T \prod_{j=1}^{n_{i,l}^r} \frac{f_{U|t_{ijl}, x_i}(u_{ijl})}{F_{U|t_{ijl}, x_i}(\tau - t_{ijl})}}_{\text{Second Term (Reporting Delay)}} \tag{4.1}$$

where Φ_1 represents the vector of the model’s parameters, including: the initial distribution π_1 and the transition matrix Γ of the HMM, the regression coefficients $\theta_j \in \mathbb{R}^p$ for the $\lambda^{(j)}$ ($j = 1, \dots, g$), where p is the number of regression parameters in $\lambda^{(j)}$, and the parameters associated with modeling the reporting delay.

The “Claim Frequency” component captures the probability of observing the reported claim counts across all policyholders and time periods. It characterizes the underlying claim frequency patterns in the data. On the other hand, the “Reporting Delay” component accounts for the right-truncated reporting delay for claims that are reported before the valuation date τ . It is modeled as a regression that depends on the characteristics of the policy and the time of occurrence.

As discussed in the introduction, maximizing the above likelihood is challenging. Proposition 2 establishes that the distribution of the discretely observed reported claim process depends on the reporting delay distribution, making the maximization of $\mathcal{L}^{(1)}$ particularly complex. A common approach in the micro-level reserving literature is to perform a two-step maximization: first fitting the “Reporting Delay” component, then using the fitted parameters to estimate the “Claim Frequency” component. This approach has been employed in all studies we are aware of that adopt a similar framework (see, e.g., Antonio and Plat, 2014; Fung et al., 2021, etc.), particularly those following the framework of Norberg (1993), with the exception of Wahl (2019), where the likelihood is maximized jointly under the assumption that claim arrivals follow a non-homogeneous Poisson process with piecewise constant

intensity and the delay distribution belongs to the exponential family. However, this iterative two-step strategy may lead to suboptimal parameter estimates, as it fails to account for the interdependence between the two components of the model.

4.2. Discrete-time models

4.2.1. The multinomial model

With the aim of joint modeling of frequency and delay, we deviate from the continuous-time framework to a discrete-time framework by aggregating the events towards a coarser time scale. To this end, we assume that claims can be reported with a maximum delay of D periods. This means that if a claim occurs in period t , it will be reported in periods $t + d$ ($d = 0, 1, \dots, D$). Note that a period t corresponds to the time interval $[d_{t-1}, d_t)$. We also, for simplicity, assume that the predetermined time points d_t ($t = 0, \dots, T$) are equidistant, so that $d_0 = 0, d_{t+1} - d_t = 1$ and T represent the last observed period before the valuation date τ . That is, we assume that we consider periods of the same length, where each period represents one unit of time.

We denote the probability that a policyholder’s claim with risk attributes \mathbf{x}_i that occurred in period t is reported in period $t + d$ by $p_t(d; \mathbf{x}_i)$. It is easy to show that the discretely observed claim arrival process, $\{N_{i,1}, N_{i,2}, \dots, N_{i,T}\}$, the discretely observed reported claim process, $\{N_{i,1}^r, \dots, N_{i,T}^r\}$, and the discretely observed IBNR claim process, $\{N_{i,1}^{ibnr}, \dots, N_{i,T}^{ibnr}\}$, fall under the class of Poisson-HMMs, with the same hidden parameter process $\{C_1, C_2, \dots\}$ as $N(t; \mathbf{x}_i)$, and with state-dependent Poisson intensities given by

$$\tilde{\lambda}_t^{(j)}(\mathbf{x}_i) = e_{i,t} \lambda^{(j)}(\mathbf{x}_i), \quad \mu_t^{(j)}(\mathbf{x}_i) = e_{i,t} \lambda^{(j)}(\mathbf{x}_i) p_{i,t}^r, \quad \nu_t^{(j)}(\mathbf{x}_i) = e_{i,t} \lambda^{(j)}(\mathbf{x}_i) (1 - p_{i,t}^r),$$

respectively, where $p_{i,t}^r = \sum_{d=0}^{\min\{D, T-t\}} p_t(d; \mathbf{x}_i)$ is the probability that a claim from policyholder i that occurred in period t is reported before the reserve valuation date.

Let $Z_{i,t,d} = \sum_{k=1}^{\infty} 1 \{T_{i,k} \in [t - 1, t), T_{i,k} + U_{i,k} \in [t + d - 1, t + d)\}$ denote the number of claims from policyholder i that occurred in period t and are reported in period $t + d$. Note that $N_{i,t}^r = \sum_{d=0}^{\min\{D, T-t\}} Z_{i,t,d}$. Our observed data are now given by

$$\mathcal{O}_D = \{n_{i,t}^r, z_{i,t,d} \mid i = 1, \dots, m; t = 1, \dots, T; d = 0, \dots, D; t + d \leq T\},$$

where $n_{i,t}^r$ and $z_{i,t,d}$ are the realizations of $N_{i,t}^r$ and $Z_{i,t,d}$, respectively.

Similar to Verbelen *et al.* (2022), we make the assumption that

$$(Z_{i,t,0}, \dots, Z_{i,t,\min\{D, T-t\}} \mid N_{i,t}^r) \sim \text{Multinomial} \left(N_{i,t}^r, \left(\frac{p_t(d; \mathbf{x}_i)}{p_{i,t}^r} \right)_{d=0}^{\min\{D, T-t\}} \right).$$

Note that for $t \leq T - D$, all claims occurred in period t is fully observed, that is, $N_{i,t}^r = N_{i,t}$, and the counts $(Z_{i,t,0}, \dots, Z_{i,t,D}) \mid N_{i,t}^r \sim \text{Multinomial}(N_{i,t}^r, \mathbf{p}_t(\mathbf{x}_i))$, where $\mathbf{p}_t(\mathbf{x}_i) = (p_t(0; \mathbf{x}_i), \dots, p_t(D; \mathbf{x}_i))$.

The likelihood of the observed data is then given by

$$\begin{aligned} \mathcal{L}^{(2)}(\Phi_2 \mid \mathcal{O}_D) &\propto P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \frac{n_{i,t}^r!}{\prod_{d=0}^{\min\{D, T-t\}} z_{i,t,d}!} \prod_{d=0}^{\min\{D, T-t\}} \left(\frac{p_t(d; \mathbf{x}_i)}{p_{i,t}^r} \right)^{z_{i,t,d}}, \end{aligned} \tag{4.2}$$

where Φ_2 represents the vector of the model’s parameters, including the initial distribution $\boldsymbol{\pi}_1$ and the transition matrix $\boldsymbol{\Gamma}$ of the HMM, the regression coefficients $\boldsymbol{\theta}_j$ for the $\lambda^{(j)}$ ’s, and the parameters related to the modeling of $\mathbf{p}_t(\mathbf{x}_i)$, denoted by $\boldsymbol{\delta}$. An equivalent way to write the multinomial component of the likelihood above is to decompose it into a series of conditional binomial likelihoods. That is, we can

write $\mathcal{L}^{(2)}$ as:

$$\begin{aligned} \mathcal{L}^{(2)}(\Phi_2 \mid \mathcal{O}_D) &\propto P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \prod_{d=1}^{\min\{D, T-t\}} \binom{\sum_{j=0}^d z_{i,t,j}}{z_{i,t,d}} q_t(d; \mathbf{x}_i)^{z_{i,t,d}} (1 - q_t(d; \mathbf{x}_i))^{\sum_{j=0}^{d-1} z_{i,t,j}}, \end{aligned} \tag{4.3}$$

where

$$q_t(d; \mathbf{x}_i) = \frac{p_t(d; \mathbf{x}_i)}{\sum_{j=0}^d p_t(j; \mathbf{x}_i)}$$

represents the conditional probability that a claim incurred at period t by policyholder with risk attributes \mathbf{x}_i is reported in period $t + d$, given that the claim is reported by period $t + d$. Thus, we break down the multinomial likelihood by considering the binomial likelihood of observing each count given the cumulative counts and adjusting the success probabilities accordingly. Note that $p_t(d; \mathbf{x}_i)$ can be obtained iteratively from $q_t(d; \mathbf{x}_i)$ by:

$$\begin{aligned} p_t(D; \mathbf{x}_i) &= q_t(D; \mathbf{x}_i), \quad p_t(d; \mathbf{x}_i) = q_t(d; \mathbf{x}_i) \times \left(1 - \sum_{j=d+1}^D p_t(j; \mathbf{x}_i) \right) \quad (d = D - 1, \dots, 1), \\ p_t(0; \mathbf{x}_i) &= 1 - \sum_{j=1}^D p_t(j; \mathbf{x}_i). \end{aligned}$$

Remark

Note that t used in a discrete-time setting differs from t in $N_i(t) = \sum_{j=1}^{\infty} \mathbb{1}\{T_{ij} < t\}$. In the continuous-time setting, t represents calendar time and $N_i(t)$ counts the number of claims of the policyholder i that occurred before t . In the discrete-time setting, t represents calendar time discretized into periods (e.g., months), where $t = 1$ corresponds to the first period, $t = 2$ to the second, and so on. For simplicity, we use the same notation t in both settings, but its meaning is clear from the context.

4.2.2. The Dirichlet-multinomial model

The probabilities of reporting delay often exhibit temporal variability, which can be characterized as either explainable or unexplainable. Explainable variations in reporting delay are systematic variations that arise from differences in the period of occurrence of claims and/or the risk attributes associated with policyholders. For example, when working with daily data, claims occurring on weekends may exhibit different reporting patterns compared to those occurring on weekdays. Similarly, the month in which a claim occurs can also influence reporting delay probabilities. In contrast, unexplainable variability refers to unsystematic fluctuations in reporting delay that are not captured by our covariates. In cases where data exhibit high random variability in reporting delay, both continuous-time and multinomial models may produce interval estimates for IBNR claim counts that are significantly off from the actual IBNR claim count, as we will observe in our data analysis. Although an exact cause for the high fluctuations cannot be precisely identified without additional data collection, one possible explanation is the presence of unobserved covariates that have direct and interaction effects on the delay probabilities.

To this end, we add randomness to the reporting delay probabilities $p_t(d; \mathbf{x}_i)$'s by assuming that $\mathbf{p}_t(\mathbf{x}_i) = (p_t(0; \mathbf{x}_i), \dots, p_t(D; \mathbf{x}_i))$ follows a Dirichlet distribution with parameters $\boldsymbol{\eta}_{i,t} = (\eta_{i,t,0}(\mathbf{x}_i), \dots, \eta_{i,t,D}(\mathbf{x}_i))$. This choice is motivated by the fact that the Dirichlet distribution serves as a conjugate prior for the multinomial distribution, which significantly simplifies the mathematical treatment and estimation of our model.

Parameters $\boldsymbol{\eta}_{i,t} = (\eta_{i,t,0}(\mathbf{x}_i), \dots, \eta_{i,t,D}(\mathbf{x}_i))$ can be specified in various ways. One approach is to model $\boldsymbol{\eta}_{i,t}$ using a regression with fixed effects, where the parameters are expressed as functions of observed covariates such as policyholder characteristics or temporal factors. Alternatively, the Dirichlet assumption can

be incorporated as a random effect by assuming that, for each period t , the delay probability vector \mathbf{p}_t follows a Dirichlet distribution with global parameters $\boldsymbol{\eta}_t$. In this paper, we adopt the fixed-effects regression structure for $\boldsymbol{\eta}_{i,t}$.

We make use of the following proposition which can also be found in Lawless (1994):

Proposition 3. *If $\mathbf{p}_i(\mathbf{x}_i) = (p_i(0; \mathbf{x}_i), \dots, p_i(D; \mathbf{x}_i))$ follows a Dirichlet distribution with parameters $\boldsymbol{\eta}_{i,t} = (\eta_{i,0}(\mathbf{x}_i), \dots, \eta_{i,D}(\mathbf{x}_i))$, then the conditional probabilities $q_i(d; \mathbf{x}_i) = \frac{p_i(d; \mathbf{x}_i)}{\sum_{j=0}^d p_i(j; \mathbf{x}_i)}$ ($d = 1, \dots, D$) are independently distributed beta random variables with*

$$q_i(d; \mathbf{x}_i) \sim \text{Beta}(\eta_{i,d}(\mathbf{x}_i), \sum_{j=0}^{d-1} \eta_{i,j}(\mathbf{x}_i)).$$

Thus, the binomial likelihoods components in $\mathcal{L}^{(2)}$ would be replaced with beta-binomial likelihoods components, and the likelihood for our model with the Dirichlet assumption becomes:

$$\begin{aligned} \mathcal{L}^{(3)}(\boldsymbol{\Phi}_3 \mid \mathcal{O}_D) &\propto P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \prod_{d=1}^{\min(D,T-t)} \binom{\sum_{j=0}^d z_{i,t,j}}{z_{i,t,d}} \frac{B(z_{i,t,d} + \eta_{i,d}(\mathbf{x}_i), \sum_{j=0}^{d-1} (z_{i,t,d} + \eta_{i,j}(\mathbf{x}_i)))}{B(\eta_{i,d}(\mathbf{x}_i), \sum_{j=0}^{d-1} \eta_{i,j}(\mathbf{x}_i))}, \end{aligned} \tag{4.4}$$

where B is the beta function, and $\boldsymbol{\Phi}_3$ represents the vector of the model’s parameters, including: the initial distribution $\boldsymbol{\pi}_1$ and the transition matrix $\boldsymbol{\Gamma}$ of the HMM, the regression coefficients $\boldsymbol{\theta}_j$ for the $\lambda^{(j)}$ ’s, and the parameters related to the modeling of the $\boldsymbol{\eta}$ ’s, denoted by $\boldsymbol{\delta}$.

Remark

With $g = 1$ (single state), we have that $N_{i,t}^r$ follows a Poisson distribution, and so, this Dirichlet assumption can be seen as an extension of the model of Verbelen *et al.* (2022). We provide the EM algorithm for this model in the next section.

5. Parameter estimation

In this section, we present the estimation methodology employed to fit the discrete-time models from the previous section. To achieve this, we utilize the EM algorithm, providing an iterative framework for updating the model parameters to maximize the likelihood of the observed reported claims while considering the unobserved hidden states of the HMM. The EM algorithm for fitting the continuous-time model using the two-step maximization can be found in Appendix A. Additionally, to assess the ability of our algorithm to recover model parameters, we conduct a simulation study based on the proposed models. The simulation details are provided in the Online Supplementary Material, and the corresponding code is available at: https://github.com/HassanAbdelrahman/HMM_Simulations.

The EM algorithm for the multinomial model extends the algorithm outlined in Verbelen *et al.* (2022), where the reported claims process now follows a Poisson-HMM, rather than a Poisson process. On the other hand, the EM algorithm for the Dirichlet-multinomial model requires special attention. This is because the Dirichlet assumption introduces additional complexity in estimating the probability vector, making the E-step analytically intractable. Therefore, we will need to switch from the EM algorithm to the MCEM algorithm, which uses Monte Carlo simulations to approximate the E-step and handle the complexities introduced by the Dirichlet distribution.

5.1. The multinomial model

Recall that the likelihood of the observed data for the discrete-time multinomial model, $\mathcal{L}^{(2)}$, is given by

$$\mathcal{L}^{(2)}(\Phi_2 | \mathcal{O}_D) \propto P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) \times \prod_{i=1}^m \prod_{t=1}^T \frac{n_{i,t}^r!}{\prod_{d=0}^{\min\{D,T-t\}} z_{i,t,d}!} \prod_{d=0}^{\min\{D,T-t\}} \left(\frac{p_t(d; \mathbf{x}_i)}{p_{i,t}^r} \right)^{z_{i,t,d}},$$

where $N_i^{(r,1:T)} = (N_{i,1}^r, \dots, N_{i,T}^r)$ denote the discretely observed reported claim process from period 1 to T for policy i , and $\mathbf{n}_i^{(r,1:T)} = (n_{i,1}^r, \dots, n_{i,T}^r)$ is its realization. As we established in Proposition 2, our model dictates that the discretely observed reported claims process for policy i , $\{N_{i,1}^r, \dots, N_{i,k}^r\}$, originates from the class of Poisson-HMM. Specifically, we demonstrated that:

$$N_{i,t}^r | C_t = j \sim \text{Poisson}(\mu_t^{(j)}(\mathbf{x}_i) = e_{i,t} \lambda^{(j)}(\mathbf{x}_i) p_{i,t}^r).$$

The likelihood for the ‘‘Claim Frequency’’ component is thus given by

$$P(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) = \boldsymbol{\pi}_1 \mathbf{P}_1(\mathbf{n}_1^r) \boldsymbol{\Gamma} \mathbf{P}_2(\mathbf{n}_2^r) \dots \boldsymbol{\Gamma} \mathbf{P}_T(\mathbf{n}_T^r) \mathbf{1}^T =: L_T, \tag{5.1}$$

where

$$\mathbf{n}_t^r = (n_{1,t}^r, \dots, n_{m,t}^r), \quad \mathbf{P}_t(\mathbf{n}_t^r) = \text{diag} \left\{ \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r | C_t = 1), \dots, \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r | C_t = g) \right\}$$

is a diagonal matrix whose diagonal entries correspond to the probabilities associated with each state, and $\mathbf{1}$ is a row vector of ones with dimension g , where g is the number of states.

The Forward-Backward Algorithm

Before sketching the EM algorithm, we first define the forward and backward probabilities as follows:

$$\alpha_{ij} := P(N_1^{(r,1:t)} = \mathbf{n}_1^{(r,1:t)}, \dots, N_m^{(r,1:t)} = \mathbf{n}_m^{(r,1:t)}, C_t = j),$$

$$\beta_{ij} := P(N_1^{(r,t+1:T)} = \mathbf{n}_1^{(r,t+1:T)}, \dots, N_m^{(r,t+1:T)} = \mathbf{n}_m^{(r,t+1:T)} | C_t = j),$$

respectively. These probabilities are used to iteratively update the hidden state probabilities during the EM algorithm, and they have the following properties (see MacDonald and Zucchini, 2016 for details):

$$P(C_t = j | N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) = \frac{\alpha_{ij} \beta_{ij}}{L_T},$$

$$P(C_{t-1} = j, C_t = k | N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) = \frac{\alpha_{t-1,j} \gamma_{jk} \mathbb{P}_k(\mathbf{n}_t^r) \beta_{tk}}{L_T},$$

where $\mathbb{P}_k(\mathbf{n}_t^r) = \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r | C_t = k)$.

They are computed using the following recursive relations:

$$\boldsymbol{\alpha}_1 = \boldsymbol{\pi}_1 \mathbf{P}_1(\mathbf{n}_1^r), \quad \boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} \boldsymbol{\Gamma} \mathbf{P}_t(\mathbf{n}_t^r), \quad t = 2, \dots, T,$$

$$\boldsymbol{\beta}_T = \mathbf{1}, \quad \boldsymbol{\beta}_t = \boldsymbol{\Gamma} \mathbf{P}_{t+1}(\mathbf{n}_{t+1}^r) \boldsymbol{\beta}_{t+1}, \quad t = T - 1, T - 2, \dots, 1.$$

We now proceed with the details of the EM algorithm.

Complete Data Likelihood

The EM algorithm iteratively estimates the parameters Φ_2 , where the number of hidden states, g , is treated as preset. The algorithm consists of two main steps: the E-step and the M-step, which will be demonstrated in the following subsections. We first introduce the following notations:

- $\mathbf{c}^{(T)} = (c_1, \dots, c_T)$: the unobserved states of the HMM,
- $u_{ij} := \mathbb{1}\{C_t = j\}$, and
- $v_{ijk} := \mathbb{1}\{C_{t-1} = j, C_t = k\}$.

We let our complete data consist of all claims that occurred before our reserve valuation date (both reported and unreported), as well as the realizations of the states of the HMM. Thus, the complete data is given by:

$$\mathcal{C} = \{n_{i,t}, z_{i,t,d}, c_t \mid i = 1, \dots, m; t = 1, \dots, T; d = 0, \dots, D\}.$$

Note that the complete data assumes that we know the total number of claims that occurred in each period for each policyholder, as well as the periods in which these claims are reported. The likelihood of the complete data is then given by

$$\begin{aligned} \mathcal{L}_C^{(2)}(\Phi_2 \mid \mathcal{C}) &= \pi_{1,c_1} \prod_{t=2}^T \gamma_{c_{t-1},c_t} \prod_{t=1}^T \prod_{i=1}^m \mathbb{P}_{c_t}(n_{i,t}) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \prod_{d=1}^D \left(\sum_{j=0}^d z_{i,t,j} \right) q_t(d; \mathbf{x}_i)^{z_{i,t,d}} (1 - q_t(d; \mathbf{x}_i))^{\sum_{j=0}^{d-1} z_{i,t,d}}, \end{aligned}$$

where $\mathbb{P}_{c_t}(n_{i,t}) = P(N_{i,t} = n_{i,t} \mid C_t = c_t)$.

It is straightforward to see that we can write the log-likelihood of the complete data as

$$\begin{aligned} l_C^{(2)}(\Phi_2 \mid \mathcal{C}) &= \sum_{j=1}^g u_{1j} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T v_{ijk} \right) \log \gamma_{jk} + \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m u_{ij} \log \mathbb{P}_j(n_{i,t}) \\ &+ \sum_{i=1}^m \sum_{t=1}^T \sum_{d=1}^D \left[\log \left(\sum_{j=0}^d z_{i,t,j} \right) + z_{i,t,d} \log q_t(d; \mathbf{x}_i) + \left(\sum_{j=0}^{d-1} z_{i,t,d} \right) \log (1 - q_t(d; \mathbf{x}_i)) \right], \end{aligned}$$

where $l_C^{(2)}(\Phi_2 \mid \mathcal{C}) := \log(\mathcal{L}_C^{(2)}(\Phi_2 \mid \mathcal{C}))$. We now outline the E-step and M-step of the EM algorithm that are applied to maximize the complete-data log-likelihood.

E-step

In the k -th E-step, we compute the conditional expectation of the complete-data log-likelihood given the observed data and the current estimator $\Phi_2^{(k-1)}$ of the model parameters Φ_2 . Note that the unobserved components of our complete data consist of the hidden states, $c^{(T)}$, as well as the IBNR claims, $\{n_{i,t}^{\text{ibnr}}, z_{i,t,d} \mid i = 1, \dots, m; t = 1, \dots, T; t + d > T\}$. The expectation of these components given the observed data and $\Phi_2^{(k-1)}$ should thus be derived. The conditional expectation is given by

$$\begin{aligned} \mathbb{E}[l_C^{(2)}(\Phi_2 \mid \mathcal{C}) \mid \Phi_2^{(k-1)}, \mathcal{O}_D] &= \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T \hat{v}_{ijk}^{(k)} \right) \log \gamma_{jk} \\ &+ \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{ij}^{(k)} \left[\hat{n}_{i,t,j}^{(k)} \log(\lambda^{(j)}(\mathbf{x}_i)) - \lambda^{(j)}(\mathbf{x}_i) \right] \\ &+ \sum_{i=1}^m \sum_{t=1}^T \sum_{d=1}^D \left[\hat{z}_{i,t,d}^{(k)} \log q_t(d; \mathbf{x}_i) + \left(\sum_{j=0}^{d-1} \hat{z}_{i,t,d}^{(k)} \right) \log (1 - q_t(d; \mathbf{x}_i)) \right] \\ &+ \text{constant}, \end{aligned} \tag{5.2}$$

where the expressions for $\hat{u}_{ij}^{(k)}$ and $\hat{v}_{ijl}^{(k)}$ are given by

$$\hat{u}_{ij}^{(k)} = \mathbb{E} (u_{ij} | \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)}) = \mathbb{P} (C_t = j | \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)}) = \frac{\alpha_{ij}^{(k-1)} \beta_{ij}^{(k-1)}}{L_T^{(k-1)}}, \tag{5.3}$$

and

$$\begin{aligned} \hat{v}_{ijl}^{(k)} &= \mathbb{E} (v_{ijl} | \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)}) = \mathbb{P} (C_{t-1} = j, C_t = l | \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)}) \\ &= \frac{\alpha_{t-1,j}^{(k-1)} \gamma_{jl}^{(k-1)} \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r | C_t = l, \Phi_2^{(k-1)}) \beta_{il}^{(k-1)}}{L_T^{(k-1)}} \end{aligned} \tag{5.4}$$

In these expressions, the forward probabilities, backward probabilities, and observed likelihood are obtained using the model parameters $\Phi_2 = \Phi_2^{(k-1)}$ from the $(k - 1)$ -th iteration. The terms $\hat{u}_{ij}^{(k)}$ and $\hat{v}_{ijl}^{(k)}$ represent the updated probabilities of being in state j at time t and the transition probability from state j to state l at time t , respectively, based on the current model parameter estimates $\Phi_2^{(k-1)}$ and the observed data.

The expressions for $\hat{n}_{i,t,j}^{(k)}$ and $\hat{z}_{i,t,d}^{(k)}$ are given by

$$\begin{aligned} \hat{n}_{i,t,j}^{(k)} &:= \mathbb{E}(N_{i,t} | C_t = j, \mathcal{O}_D, \Phi_2^{(k-1)}) = n_{i,t}^r + \mathbb{1}\{t > T - D\} e_{i,t} \lambda_j^{(k-1)}(\mathbf{x}_i) \sum_{d=T-t+1}^D p_t^{(k-1)}(d; \mathbf{x}_i), \\ \hat{z}_{i,t,d}^{(k)} &:= \mathbb{E}(Z_{i,t,d} | \mathcal{O}_D, \Phi_2^{(k-1)}) = \begin{cases} z_{i,t,d}, & t \leq T - D \\ (\sum_{j=1}^g \pi_{ij}^{(k-1)} e_{i,t} \lambda_j^{(k-1)}(\mathbf{x}_i)) p_t^{(k-1)}(d; \mathbf{x}_i), & t > T - D \end{cases} \end{aligned}$$

where $p_t^{(k-1)}(d; \mathbf{x}_i)$ and $\lambda_j^{(k-1)}(\mathbf{x}_i)$ are the estimates of $p_t(d; \mathbf{x}_i)$ and $\lambda_j^{(i)}(\mathbf{x}_i)$, respectively, computed using the parameter estimates obtained from the M-step in the $(k - 1)$ -th iteration. The term $\hat{n}_{i,t,j}^{(k)}$ represents the updated expectation of the total number of claims that occurred in period t for policyholder i , given that the state of the HMM is j , while the term $\hat{z}_{i,t,d}^{(k)}$ represents the updated expectation of the portion of these claims that are reported in period $t + d$. Note that all claims that occurred at $t \leq T - D$ are reported and are part of the observed data \mathcal{O}_D , and thus, the expected values of $N_{i,t}$ and $Z_{i,t,d}$ ($t \leq T - D; d = 0, \dots, D$) given the observed data are equal to $n_{i,t}$ and $z_{i,t,d}$, respectively.

By maximizing this expected complete-data log-likelihood, we can iteratively update the model parameters in the M-step of the algorithm, leading to improved parameter estimates that better fit the observed data.

M-step

The M-step in the EM algorithm involves maximizing the expected complete-data log-likelihood (Equation (5.2)) with respect to the model parameters Φ_2 , subject to the constraints

$$\sum_{j=1}^g \pi_{1j} = 1, \text{ and } \sum_{l=1}^g \gamma_{jl} = 1, \quad j = 1, \dots, g.$$

For this step, we can obtain the updated estimates for π_1 and Γ separately by maximizing the first two terms of Equation (5.2), respectively. The optimal values for $\pi_{1j}^{(k)}$ and $\gamma_{jl}^{(k)}$ at the k th iteration are given by

$$\pi_{1j}^{(k)} = \hat{u}_{1j}^{(k)}, \quad j = 1, \dots, g, \tag{5.5}$$

$$\gamma_{jl}^{(k)} = \frac{\sum_{t=2}^T \hat{v}_{ijl}^{(k)}}{\sum_{l=1}^g \sum_{t=2}^T \hat{v}_{ijl}^{(k)}}, \quad j, l = 1, \dots, g. \tag{5.6}$$

As for the third term:

$$\sum_{t=1}^T \sum_{i=1}^m \hat{w}_{ij}^{(k)} [\hat{n}_{i,t,j}^{(k)} \log(\lambda^{(j)}(\mathbf{x}_i)) - \lambda^{(j)}(\mathbf{x}_i)],$$

it is easy to see that for each j , this is a weighted Poisson log-likelihood, with the exception that $\hat{n}_{i,t,j}^{(k)}$ can be non-integer. Yet, we can still obtain the updated regression parameters for each $\lambda^{(j)}$ using quasi-likelihood methods available through standard packages in R. Finally, for the term:

$$\sum_{i=1}^m \sum_{t=1}^T \left[\hat{z}_{i,t,d}^{(k)} \log q_t(d; \mathbf{x}_i) + \left(\sum_{j=0}^{d-1} \hat{z}_{i,t,d}^{(k)} \right) \log(1 - q_t(d; \mathbf{x}_i)) \right],$$

it is easy to see that for each d , this is a binomial log-likelihood, again with the exception that $\hat{z}_{i,t,d}^{(k)}$ can be non-integer. For each d , we model $q_t(d; \mathbf{x}_i)$ as a regression function of the policy-holder risk attributes, \mathbf{x}_i , and time-dependent covariates (e.g. the month of the period t). The binomial log-likelihood can then be optimized using quasi-likelihood methods available through standard packages in R.

The model parameters are then given by $\Phi_2 = \{\pi_1, \Gamma, \theta_1, \dots, \theta_g, \delta_1, \dots, \delta_D\}$, where θ_j is the vector of regression parameters for $\lambda^{(j)}(\mathbf{x})$, and δ_d is the vector of regression parameters for $q_t(d; \mathbf{x})$.

Convergence

We continue iterating the E-step and M-step until the relative distance between the estimates from consecutive iterations falls below a predetermined threshold. The relative distance is defined as

$$d(\Phi_2^{(k-1)}, \Phi_2^{(k)}) = \sum_{j=1}^g \left| \frac{\pi_{1j}^{(k-1)} - \pi_{1j}^{(k)}}{\pi_{1j}^{(k-1)}} \right| + \sum_{j=1}^g \sum_{l=1}^g \left| \frac{\gamma_{jl}^{(k-1)} - \gamma_{jl}^{(k)}}{\gamma_{jl}^{(k-1)}} \right| + \sum_{j=1}^g \sum_{s=1}^{\|\theta_j\|_0} \left| \frac{\theta_{js}^{(k-1)} - \theta_{js}^{(k)}}{\theta_{js}^{(k-1)}} \right| + \sum_{d=1}^D \sum_{s=1}^{\|\delta_d\|_0} \left| \frac{\delta_{ds}^{(k-1)} - \delta_{ds}^{(k)}}{\delta_{ds}^{(k-1)}} \right|.$$

The EM algorithm converges when the relative distance is below the threshold.

Initialization of Model Parameters

To effectively apply the EM algorithm, appropriate initialization of model parameters is crucial. Note that we have complete data (all claims are reported) for periods $t \leq T - D$, so we can fit the model for such data to obtain parameter estimates as initialization

Model Selection

Finally, we need to decide on the number of states g . Typically, a larger g would provide a better fit, but it could also lead to overfitting. We usually select the model with the lowest chosen information criterion (e.g., AIC or BIC). For fast and efficient computation, we can employ a backward fitting strategy similar to Badescu *et al.* (2019). We start by fitting the model with a large g , and then iteratively delete a state until the chosen information criterion stops decreasing. In this strategy, we use the parameter estimates obtained from fitting with g states to initialize the algorithm for $g - 1$ states, with the initial values for π_1 and Γ being normalized estimates excluding the deleted state. This approach should make the estimation process very efficient.

Under-/Overflow Problems

As noted by MacDonald and Zucchini (2016), the computation of forward and backward probabilities, along with the subsequent calculations in the E- and M- steps, is susceptible to under- or overflow issues.

MacDonald and Zucchini (2016) recommend scaling the forward and backward probability computations by taking the logarithms of these quantities and provides code solutions for this purpose. However, for our model, additional considerations are needed.

The computation of forward and backward probabilities involves matrix multiplication of the matrices $\mathbf{P}_t(\mathbf{n}_t^r)$. Unlike normal HMMs fitted to single time series, where these matrices are diagonal matrices containing the probability of the count process obtaining its observed value given the different states, our model’s probabilities are the product of the observed count process for each individual. For our model, this matrix is defined as

$$\mathbf{P}_t(\mathbf{n}_t^r) = \text{diag} \left\{ \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r \mid C_t = 1), \dots, \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r \mid C_t = g) \right\}.$$

It becomes evident that the elements of the diagonal matrix are extremely small, potentially causing underflow problems in computing forward and backward probabilities, as well as subsequent calculations, even after scaling. To address this issue, we implement the LogSumExp algorithm (e.g., see Blanchard et al., 2021), which is effective in mitigating numerical instability arising from the small probabilities.

5.2. The Dirichlet-multinomial model

Complete Data Likelihood

We now aim to find the parameter estimates that maximize $\mathcal{L}^{(3)}$ using the EM algorithm. As we illustrated, the Dirichlet-multinomial model differs from the multinomial model by the assumption that the probability vector $\mathbf{p}_t(\mathbf{x}_t)$ is random. To this end, we let our complete data be given by

$$\tilde{\mathcal{C}} = \{n_{i,t}, z_{i,t,d}, p_{i,t,d}, c_t \mid t = 1, \dots, T; d = 0, \dots, D\},$$

where $p_{i,t,d}$ is the realization of $p_d(t; \mathbf{x}_i)$. That is, we now assume we know the actual reporting probability vector $\mathbf{p}_t(\mathbf{x}_i)$. The likelihood of the complete data is then given by:

$$\begin{aligned} \mathcal{L}_C^{(3)}(\Phi_3 \mid \tilde{\mathcal{C}}) &= \pi_{1,c_1} \prod_{t=2}^T \gamma_{c_{t-1},c_t} \prod_{t=1}^T \prod_{i=1}^m \mathbb{P}_{c_t}(n_{i,t}) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \left(\underbrace{\frac{n_{i,t}!}{\prod_{d=0}^D z_{i,t,d}!} \prod_{d=1}^D p_{i,t,d}^{z_{i,t,d}}}_{\text{Multinomial Distribution}} \times \underbrace{\frac{1}{B(\boldsymbol{\eta}_{i,t})} \prod_{d=0}^D p_{i,t,d}^{\eta_{i,t}(\mathbf{x}_i)-1}}_{\text{Dirichlet Distribution}} \right), \end{aligned}$$

where $\mathbb{P}_{c_t}(n_{i,t}) = P(N_{i,t} = n_{i,t} \mid C_t = c_t)$. Note that the ‘‘Multinomial Component’’ in the complete data likelihood is constant since it does not depend on the model’s parameters. The log-likelihood of the complete data is then given by

$$\begin{aligned} l_C^{(3)}(\Phi_3 \mid \tilde{\mathcal{C}}) &= \sum_{j=1}^g u_{1j} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T v_{jk} \right) \log \gamma_{jk} + \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m u_{ij} \log \mathbb{P}_j(n_{i,t}) \\ &+ \sum_{i=1}^m \sum_{t=1}^T - \log B(\boldsymbol{\eta}_{i,t}) \sum_{d=1}^D (\eta_{i,t}(\mathbf{x}_i) - 1) \log p_{i,t,d} + \text{constant}, \end{aligned}$$

where $l_C^{(3)}(\Phi_2 \mid \mathcal{C}) := \log(\mathcal{L}_C^{(3)}(\Phi_2 \mid \mathcal{C}))$.

E-Step

In the k -th iteration of the E-step, we take the expectation of $l_C^{(3)}(\Phi_3 | \tilde{C})$ with respect to $\Phi_3^{(k-1)}$ and \mathcal{O}_D . We have:

$$\begin{aligned} \mathbb{E} \left[l_C^{(3)}(\Phi_3 | \tilde{C}) | \Phi_3^{(k-1)}, \mathcal{O}_D \right] &= \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{l=1}^g \left(\sum_{t=2}^T \hat{v}_{tjl}^{(k)} \right) \log \gamma_{jl} \\ &+ \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{tj}^{(k)} \left[\hat{n}_{i,t,j}^{(k)} \log (\lambda^{(j)}(\mathbf{x}_i)) - \lambda^{(j)}(\mathbf{x}_i) \right] \\ &+ \sum_{i=1}^m \sum_{t=1}^T -\log B(\eta_{i,t}) \sum_{d=1}^D (\eta_{t,d}(\mathbf{x}_i) - 1) \widehat{\log p_{i,t,d}}^{(k)} \\ &+ \text{constant}, \end{aligned} \tag{5.7}$$

where $\hat{u}_{ij}^{(k)}$ and $\hat{v}_{tjl}^{(k)}$ are given by Equations (5.3) and (5.4), respectively, and

$$\begin{aligned} \hat{n}_{i,t,j}^{(k)} &:= \mathbb{E} \left[N_{i,t} | C_t = j, \mathcal{O}_D, \Phi_3^{(k-1)} \right] \\ &= n_{i,t}^r + \mathbb{1}\{t > T - D\} e_{i,t} \lambda_j^{(k-1)}(\mathbf{x}_i) \mathbb{E} \left[\sum_{d=T-t+1}^D p_t(d; \mathbf{x}_i) | C_t = j, \mathcal{O}_D, \Phi_3^{(k-1)} \right], \\ \widehat{\log p_{i,t,d}}^{(k)} &:= \mathbb{E} \left[\log p_t(d; \mathbf{x}_i) | \mathcal{O}_D, \Phi_3^{(k-1)} \right], \end{aligned}$$

where $\lambda_j^{(k-1)}(\mathbf{x}_i)$ is the estimate of $\lambda^{(j)}(\mathbf{x}_i)$ computed using the parameter estimates obtained from the M-step in the $(k - 1)$ -th iteration. We do not have closed-form analytical expressions for $\hat{n}_{i,t,j}^{(k)}$ ($t > T - D$) and $\widehat{\log p_{i,t,d}}^{(k)}$, and thus, numerical methods are needed to compute these expectations. We can obtain Monte Carlo estimates of these expectations using samples from the conditional distribution $\mathbf{p}_t(\mathbf{x}_i) = (p_t(0; \mathbf{x}_i), \dots, p_t(D; \mathbf{x}_i))$, using the parameters Φ_3 . This, however, adds an extra layer of challenge because there is no guarantee that the likelihood will increase with each iteration, unlike in the case of the normal EM algorithm, as we are replacing the actual expectation with a numerical approximation.

Remark

The computation of the forward probabilities, backward probabilities, the likelihood, and consequently, \hat{u}_{ij} and \hat{v}_{tjk} in the E-step differs slightly for the two models considered. The differences arise when computing the probability $P(N_{i,t}^r = n_{i,t}^r | \mathcal{O}_D, \Phi_3^{(k-1)})$. As previously discussed, this probability depends on the reporting delay. For the discrete-time Multinomial model, the estimation of the probability vector $\mathbf{p}_t(\mathbf{x}_i)$ changes after each iteration of the algorithm, and so we use $\mathbf{p}_t(\mathbf{x}_i)$ estimated from the $k - 1$ -th iteration in our computation in the k -th iteration. As for the Dirichlet-multinomial model, $\mathbf{p}_t(\mathbf{x}_i)$ is a Dirichlet random vector, and thus we compute

$$\begin{aligned} P(N_{i,t}^r = n_{i,t}^r | \mathcal{O}_D, \Phi_3^{(k-1)}) &= \\ \int_{\mathbf{p}} P(N_{i,t}^r = n_{i,t}^r | \mathbf{p}_t(\mathbf{x}_i) = \mathbf{p}, \mathcal{O}_D, \Phi_3^{(k-1)}) &\times P(\mathbf{p}_t(\mathbf{x}_i) = \mathbf{p} | \mathcal{O}_D, \Phi_3^{(k-1)}) d\mathbf{p}. \end{aligned}$$

Distribution of $\mathbf{p}_t | \mathcal{O}_D$

In the case where $t \leq T - D$, we have:

$$\mathbf{p}_t(\mathbf{x}_i) | \mathcal{O}_D \sim \text{Dirichlet}(\eta_{t,0}(\mathbf{x}_i) + z_{i,t,0}, \dots, \eta_{t,D}(\mathbf{x}_i) + z_{i,t,D}),$$

since the Dirichlet distribution is a conjugate prior of the multinomial distribution. Challenges arise in the case where $t > T - D$, as we do not observe $z_{i,t,d}$ for $d > T - t$, and thus we lose the nice conjugate property. Therefore, we need to derive the conditional distribution of \mathbf{p}_t for $t > T - D$.

The probability of $\mathbf{p}_i(\mathbf{x}_i)$ equal $\mathbf{p} = (p_0, \dots, p_D)$ given the observed data is given by

$$P(\mathbf{p}_i(\mathbf{x}_i) = \mathbf{p} \mid \mathcal{O}_D) \propto P(\mathcal{O}_D \mid \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \times P(\mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}),$$

where

$$P(\mathcal{O}_D \mid \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \propto P(N_{i,t}^r = n_{i,t}^r \mid \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \times P(Z_{i,t,0} = z_{i,t,0}, \dots, Z_{i,t,T-t} = z_{i,t,T-t} \mid N_{i,t}^r = n_{i,t}^r, \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}).$$

Thus,

$$\begin{aligned} P(\mathbf{p}_i(\mathbf{x}_i) = \mathbf{p} \mid \mathcal{O}_D) &\propto P(\mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \times P(N_{i,t}^r = n_{i,t}^r \mid \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \times \\ &P(Z_{i,t,0} = z_{i,t,0}, \dots, Z_{i,t,T-t} = z_{i,t,T-t} \mid N_{i,t}^r = n_{i,t}^r, \mathbf{p}_i(\mathbf{x}_i) = \mathbf{p}) \\ &\propto \prod_{d=0}^D (p_d)^{\eta_{i,d}(\mathbf{x}_i)-1} \times \left(\sum_{j=1}^g \pi_{ij} \frac{e^{-\lambda^{(j)}(\mathbf{x}_i)} p^r (\lambda^{(j)}(\mathbf{x}_i) p^r)^{n_{i,t}^r}}{n_{i,t}^r!} \right) \times \frac{n_{i,t}^r!}{\prod_{d=0}^{T-t} z_{i,t,d}!} \prod_{d=0}^{T-t} \left(\frac{p_d}{p^r} \right)^{z_{i,t,d}} \\ &\propto \underbrace{\left(\prod_{d=0}^{T-t} (p_d)^{z_{i,t,d} + \eta_{i,d}(\mathbf{x}_i) - 1} \right)}_h \times \underbrace{\left(\prod_{d=T-t+1}^D (p_d)^{\eta_{i,d}(\mathbf{x}_i) - 1} \right)}_g \times \underbrace{\sum_{j=1}^g \pi_{ij} (\lambda^{(j)}(\mathbf{x}_i))^{n_{i,t}^r} e^{-\lambda^{(j)}(\mathbf{x}_i) p^r}}_g, \end{aligned}$$

where $p^r = \sum_{d=0}^{T-t} p_d$. Note that the conditional distribution in this case is not a standard distribution, and thus, we cannot sample from it directly.

Rejection Sampling

Samples from the conditional distribution of $\mathbf{p}_i(\mathbf{x}_i)$ can be selected by multivariate rejection sampling. We can write $P(\mathbf{p}_i(\mathbf{x}_i) = \mathbf{p} \mid \mathcal{O}_D) = ah(\mathbf{p})g(\mathbf{p})$. It is straightforward to see that h is the density function of Dirichlet distribution with parameters $(z_{i,t,0} + \eta_{i,0}(\mathbf{x}_i), \dots, z_{i,t,T-t} + \eta_{i,T-t}(\mathbf{x}_i), \eta_{i,T-t+1}(\mathbf{x}_i), \dots, \eta_{i,D}(\mathbf{x}_i))$, and thus, we can easily sample from it. We perform the rejection sampling as follows:

- Step 1: sample \mathbf{p} from h , and independently, sample u from a Uniform(0,1) distribution.
- Step 2: If $u \leq \frac{g(\mathbf{p})}{\sup_{\mathbf{p}}\{g(\mathbf{p})\}}$, then accept \mathbf{p} . If not, go to Step 1.

Note that $\sup_{\mathbf{p}}\{g(\mathbf{p})\} \leq \sum_{j=1}^g \pi_{ij} (\lambda^{(j)}(\mathbf{x}_i))^{n_{i,t}^r}$. By obtaining a sample for $\mathbf{p}_i(\mathbf{x}_i)$ given \mathcal{O}_D and $\Phi_3^{(k-1)}$, we can estimate $\hat{n}_{i,t,j}^{(k)}$ and $\widehat{\log p_{i,t,d}^{(k)}}$ using Monte Carlo methods.

Remark: The formulas presented above in the derivation of the conditional distribution of \mathbf{p}_i implicitly assume an exposure of $e_{i,t} = 1$. To account for exposure, simply multiply the Poisson intensity by the exposure factor.

M-Step

We now maximize Equation (5.7) from the E-step after replacing $\hat{n}_{i,t,j}^{(k)}$ ($t > T - D$) and $\widehat{\log p_{i,t,d}^{(k)}}$ with their Monte Carlo estimates. Again, the updated estimates of $\boldsymbol{\pi}_1$ and $\boldsymbol{\Gamma}$ are given by Equations (5.5) and (5.6), respectively. Moreover, similar to the Multinomial model case, given each j , it is straightforward to see that the third term is a weighted Poisson log-likelihood, and thus can be fit using standard packages in R. Finally, for the fourth and final term, it is straightforward to see that this is a Dirichlet log-likelihood. As previously illustrated, there are several flexible ways to model $\boldsymbol{\eta}$. For our data analysis, we model $\boldsymbol{\eta}$ as a regression function of \mathbf{x}_i and time-dependent covariates. We obtain the parameter estimates for this Dirichlet regression using the `DirichletReg` package in R (Maier, 2014), which rescales the probability estimates from the E-step to ensure that the constraints are satisfied.

Convergence Criterion, Model Initialization, and Model Selection

We deviate from the standard EM algorithm and use the MCEM algorithm, as we compute the expectations in the E-step using Monte Carlo methods. In the MCEM algorithm, it is not guaranteed that the likelihood will increase after each iteration, unlike the standard EM algorithm (see Booth and Hobert, 1999). Thus, special care should be taken. Many methods have been proposed for approximating the EM algorithm (see Ruth, 2024 for a review). From a practical point of view, we use the same convergence criterion as we have for the previous model (the relative distance criterion), as this would have a negligible effect on the parameter estimates even though the likelihood can fluctuate.

For model initialization, we fit the model to our complete data up to period $T - D$ (as we did in the Multinomial case) and use the parameter estimates as the initialization for the algorithm. Finally, we choose the number of states g using a backward strategy as explained before.

6. IBNR reserve prediction

The estimation of the IBNR claims count varies slightly across the three models. Below, we outline the steps for the continuous-time model:

IBNR Claim Reserve Estimation Steps

1. *Fit Right-Truncated Reporting Delay Regression:* Estimate the parameters of the right-truncated reporting delay regression model to obtain the fitted distribution of the reporting delay $F_{U|t,x}(u)$. The choice of model depends on the data at hand. Options include flexible mixture-based models, as proposed in Fung *et al.* (2022) and Bücher and Rosenstock (2023). Alternatively, the *flexsurv* package in R (Jackson, 2016) provides tools for fitting parametric proportional hazard models for data with random right truncation.
2. *Compute Reporting Delay Probability:* Use the fitted reporting delay regression to compute the probability $\int_{d_{l-1}}^{d_l} F_{U|t,x_i}(\tau - t) dt$ for each policyholder $i = 1, \dots, m$ and each period $l = 1, \dots, T$. This probability is used in maximizing the likelihood of the discretely observed reported claims processes and simulating the IBNR claim counts.
3. *Fit the Discretely Observed Reported Claim Process:* Maximize the likelihood of the observed reported claim counts using the EM algorithm (see Appendix A). This involves estimating the initial state probabilities and the transition matrix for the HMM, as well as the claim frequency regression parameters.
4. *Viterbi's Global Decoding:* Apply Viterbi's algorithm (see Chapter 5 in MacDonald and Zucchini, 2016 for details) to find the most likely sequence of hidden states c^* given the observed reported claim counts $\mathbf{n}^{(r,1:T)}$ and the fitted model parameters from step 3.
5. *Simulate IBNR Claim Counts:* For each policyholder $i = 1, \dots, m$ and each period $l = 1, \dots, T$ where the exposure is positive, simulate the IBNR claim count $N_{i,l}^{IBNR}$ using the Poisson distribution with mean $e_{i,l} \lambda_l^{(c_l^*)}(\mathbf{x}_i) \times (1 - \int_{d_{l-1}}^{d_l} F_{U|t,x_i}(\tau - t) dt)$ for the corresponding hidden state from the previous step. The total IBNR claim count at time τ is then given by $\sum_{i=1}^m \sum_{l=1}^T \hat{N}_{i,l}^{IBNR}$, where $\hat{N}_{i,l}^{IBNR}$ is the simulated IBNR claim count for policyholder i and period l .

Note that *Step 5* is repeated 1000 times to obtain a distribution for the IBNR claim counts. The point estimate for the IBNR claim count is the mean of the 1000 simulated values.

The estimation process for the discrete-time models (Multinomial and Dirichlet-Multinomial) begins at *Step 3*, where the relevant likelihood is maximized, allowing the delay and frequency components to be simultaneously fitted.

In the multinomial model, maximizing the likelihood $\mathcal{L}^{(2)}$ provides direct estimates of the reporting probability vector $\mathbf{p}_t(\mathbf{x}_i)$. In *Step 5*, these probabilities are used to simulate the IBNR claim counts. Specifically, for each policyholder i and each period t , the IBNR claim count $N_{i,t}^{IBNR}$ is simulated by

drawing from a Poisson distribution with mean

$$e_{i,t} \lambda_i^{(c_t)}(\mathbf{x}_i) \times \left(1 - \sum_{d=0}^{\min\{D,T-t\}} p_i(d; \mathbf{x}_i) \right).$$

In contrast, the Dirichlet-Multinomial model involves additional steps to account for the variability in reporting probabilities. Instead of directly estimating $\mathbf{p}_i(\mathbf{x}_i)$, *Step 3* yields the Dirichlet parameters $\boldsymbol{\eta}_{i,t}$, which characterize the distribution of reporting probabilities. These parameters are used in *Step 5* to simulate $\mathbf{p}_i(\mathbf{x}_i)$ by drawing from a Dirichlet distribution. The simulated $\mathbf{p}_i(\mathbf{x}_i)$ is then used in the same way as in the Multinomial model to estimate the IBNR claim counts. The additional simulation step in the Dirichlet-Multinomial model introduces variability, which is captured by repeating the process multiple times to better account for uncertainty in the reporting probabilities.

7. Case study: real-life data

In this section, we apply the proposed models to estimate the IBNR claim count for a major European auto insurance company. We then compare these estimates with those obtained from the traditional Chain Ladder method. This analysis aims to evaluate the performance of our models in a practical insurance context. Additionally, we investigate how the joint modeling of occurrence and reporting in the discrete-time models performs compared to the continuous-time model, where we maximize the likelihood of the observed data using a two-step maximization process. This is particularly interesting given the loss of information when moving to discrete-time models.

7.1. Data analysis

Data Description

Our dataset contains information on Physical Damage policies issued by the company between January 2009 and December 2017. During these years, a total of 293,709 one-year policies were issued, resulting in 106,187 reported claims. For each policy, the dataset includes the policy start and end dates, as well as various vehicle and policyholder attributes. These attributes include the age of the driver (18–90), the age of the car (0–26), the class of the car (A, B, or C), the fuel type (Gasoline or Diesel), whether the contract is a renewal or issued to a new policyholder, and the region (five regions). For each claim, the dataset records the day on which the claim occurred, the day on which the claim was reported, and the progression of claim payments. In this study, we focus on estimating the IBNR claim count, which requires us to concentrate on the time of occurrence and reporting delay of each claim. Therefore, we will not consider the payments in our analysis.

Reporting Delay

From the 106,197 claims reported between January 2009 and December 2017, only two were reported more than two years after the occurrence date. This suggests that nearly all claims occurring before 2016 have been reported. Table 1 displays the distribution of reporting delays in months for these claims. A delay of zero months indicates a claim was reported in the same month it occurred, while a delay of one month means it was reported the following month, and so on. The majority of claims (81.71%) were reported in the same month they occurred. Over 99% of claims were reported within five months, and only 0.1% were reported after ten months. This indicates that most claims are reported promptly after occurrence, with only a small fraction experiencing delays beyond ten months.

Table 1. Reporting delay distribution for claims that occurred before 2016.

Delay (months)	0	1	2	3	4	5	6	7	8	9	10+
No. of claims	66,475	11,617	1580	642	342	192	141	105	80	52	130
% of total	81.71	14.28	1.94	0.79	0.42	0.24	0.17	0.13	0.10	0.06	0.16
Cumulative %	81.71	95.99	97.93	98.72	99.14	99.38	99.55	99.68	99.78	99.84	100.0




Figure 1. Monthly exposure (left) and claim count per exposure (right) between January 2009 and December 2015.

Frequency

Figure 1 illustrates the trends in exposure (left) and claim count per exposure (right) from January 2009 to December 2015. The exposure, which represents the number of policies at risk, exhibited a steady increase from January 2009 until December 2012. Following this period, it decreased consistently until January 2015, after which it began to rise again. In contrast, the claim count per exposure displayed significant fluctuations over the same period. It reached a minimum of 0.023 claims per exposure in December 2014 and peaked at 0.045 claims per exposure in May 2013. This variation in claim frequency highlights the need for a sophisticated model capable of capturing such dynamic changes. The HMM employed in our model is well-suited for this purpose, as it can account for the underlying state changes that drive fluctuations in claim frequency over time.

7.2. IBNR claim count prediction

In this subsection, we outline the four different models used to estimate the IBNR claim count for our data, along with their specifications. The models considered are the continuous-time model, the multinomial model, the Dirichlet-multinomial model, and the traditional Chain Ladder method.

For the discrete-time models, we set the maximum reporting delay, D , to 9 months. This choice is based on the observation that nearly 99.9% of claims are reported within this period, ensuring that the majority of claims are captured within the model.

Specifications for the Fitted Models

- *The Continuous-time Model (CM):* As previously explained, our approach begins by fitting the distribution for the reporting delay, which is then utilized to fit our model, and consequently




Figure 2. Cox-Snell residual plot for assessing the fit of the reporting delay model

estimate the IBNR claim count. A notable challenge in this process is the right truncation of the reporting delay, which complicates the fitting of the distribution. We provide the continuous-time model with an advantage over the discrete-time models as we fit the reporting delay using all claims that occurred before the reserve valuation date τ , including those that have not yet been reported. This overcomes the issue of right truncation, allowing for a more accurate fit.

While this method is not feasible in practice – since unreported claims cannot be known beforehand – it serves to highlight the strengths of the discrete-time models, which we will demonstrate in the results section. In a practical scenario, one could use a “back censoring” technique, where a time threshold within the training period is identified. Claims occurring before this threshold would have been reported by the valuation date, allowing the model to be fit without concern for right truncation.

In our analysis, the log-logistic proportional hazard regression performed well in fitting the reporting delay data, as evidenced by the Cox-Snell residuals plot in Figure 2. The covariates used in the regression include the vehicle and policyholder risk attributes introduced in the data description (Section 7.1), the month in which the claim occurred and the day of occurrence.

- *The Multinomial Model (MM)*: In Section 5.1, we outlined the EM algorithm required to fit the multinomial model, where we model the binomial probabilities $q_i(d; \mathbf{x}_i)$ for each $d \in \{1, \dots, D\}$ as a regression function of the vehicle and policyholder’s risk attributes \mathbf{x}_i and time-dependent covariates. The vehicle and policyholder’s covariates are the vehicle and policyholder risk attributes introduced in the data description (Section 7.1), while the time-dependent covariates include the month of occurrence and the weekday on which the period ends, both of which were found to be significant in our regression analysis. Given that $q_i(d; \mathbf{x}_i)$ is small for $d > 1$, we use a complementary log-log link function to model it. For $d = 1$, we use a standard log link function, as is typical for binomial GLMs. We fit the binomial component of the likelihood using the standard `glm` function in R.
- *The Dirichlet-Multinomial Model (DM)*: The fitting process for the Dirichlet-Multinomial model is thoroughly detailed in Section 5.2. We model the Dirichlet parameters $\boldsymbol{\eta}$ as a regression function of the same covariates used for $q_i(d; \mathbf{x}_i)$ in the Multinomial Model.
- *The Chain Ladder (CL)*: We estimate the IBNR claim count using the CL method on data aggregated monthly, serving as the baseline for comparing our models. As with the discrete-time models, we assume that claims are reported no later than 9 months after the month of occurrence.

Remarks

- For the continuous-time model, we do not assume a maximum reporting delay, so by following the steps in Section 6, we can estimate the full number of IBNR claims, unlike the discrete-time models. However, for a fair comparison with the discrete-time models, we use the estimated parameters after fitting the continuous-time model to estimate only the IBNR claims with a maximum delay period of 9 months after the month of claim occurrence. This is achieved by adjusting Step 5 so that the probability factor $\left(1 - \int_{d_{t-1}}^{d_t} F_{U|t,x_i}(\tau - t) dt\right)$ is replaced with:

$$\int_{d_{t-1}}^{d_t} P(U + t \in [\tau, t_{\max}] | t, x_i) dt,$$

where t is the day of claim occurrence, U is the reporting delay random variable, and t_{\max} is the date by which the claim should be reported if we assume a maximum delay of 9 months after the month of claim occurrence. For example, if the claim occurrence date is in January, then t_{\max} will be the end of October for the same year. This adjustment ensures that only the first 9 months of reporting delay are considered in the estimation, aligning the continuous-time model’s output with the assumptions made in the discrete-time models.

- The covariates used for the state-dependent intensities, $\lambda^{(j)}(\cdot)$, are the vehicle and policyholder risk attributes introduced in the data description (Section 7.1). More specifically, these attributes include the age of the driver, the age of the car, the car class, whether the contract is a renewal or issued to a new policyholder, and the policyholder’s region.

7.3. Results

We present the results of estimating the IBNR claim count using the four models specified in the previous section, focusing on claims reported within 9 months after the month of occurrence. Estimations were performed at the end of each month from January 2014 to December 2016, resulting in 36 valuation points. For each valuation date, the training set consists of all claims reported up to that valuation date. The models are fitted using this reported claims data to estimate the parameters, which are then used to predict the IBNR claim count as outlined in Section 6. Since our dataset extends to December 2017, we have the actual IBNR claim counts to compare against our model estimates. This process results in 36 valuation points, providing a robust assessment of the models’ predictive performance.

To evaluate the performance of each model, we calculated the absolute percentage error (APE) of the IBNR claim count point estimates at each of the 36 valuation points. The APE is defined as

$$APE = \left| \frac{\text{Estimated IBNR} - \text{Actual IBNR}}{\text{Actual IBNR}} \right|$$

The results are summarized in boxplots that display the distribution of APE for each model across the valuation period, as well as in a table that reports the mean and standard deviation of these errors (see Figure 3 and Table 2). We considered models with 2, 3, and 4 states for our HMMs. While more than 4 states could potentially capture additional granularity – especially for later reserve valuation dates with longer time series – such complexity may be excessive for the relatively short monthly time series available at earlier dates. Moreover, fitting models with more than 4 states for early dates often resulted in spurious states. Even for later valuation dates where models with 5 states did not produce spurious states, the Bayesian information criterion (BIC) generally favored the 4-state model, indicating no clear benefit from increased complexity. To ensure comparability, we therefore adopt a unified analysis using up to 4 states across all valuation dates.

Table 2. Mean and standard deviation of APE for IBNR estimates by model and HMM states (g): $g = 1$ corresponds to a Poisson process.

g	Metric	CL	CM	MM	DM
1	Mean APE	0.1689	0.1180	0.1246	0.0911
	SD APE	0.1171	0.0851	0.0979	0.0757
2	Mean APE	0.1689	0.1127	0.1021	0.0784
	SD APE	0.1171	0.0901	0.0619	0.0537
3	Mean APE	0.1689	0.1121	0.0931	0.0774
	SD APE	0.1171	0.0875	0.0631	0.0630
4	Mean APE	0.1689	0.1166	0.0875	0.0760
	SD APE	0.1171	0.0914	0.0597	0.0560




Figure 3. Boxplots of absolute percentage errors for IBNR claim count estimates by model and number of HMM states across 36 valuation dates.

Performance Comparison of the Chain Ladder Method and HMM-Based Models

The classical CL method performs worse than our proposed HMM-based models. The CL method shows a mean and median APE of around 16%, with high variability in the error, as indicated by both the boxplot and the standard deviation values presented in Table 2. This highlights the limitations of the CL method compared to more advanced approaches.

Performance of the Proposed Models

Our proposed models – the continuous-time model (CM), the multinomial model (MM), and the Dirichlet-mMultinomial model (DM) – perform much better than the CL method, demonstrating their effectiveness in estimating IBNR claim counts. The discrete-time models (MM and DM) generally outperform the CM, as evidenced by their lower median and mean APE and lower variability (see Figure 3) across all values of g (2, 3, and 4). The CM model shows a median absolute error close to that of the

Table 3. BIC values for three fitted HMM-based models with $g = 2, 3, 4$ at the first reserve valuation date.

Model	$g = 2$	$g = 3$	$g = 4$
Continuous-time	169,372.9	169,098.1	169,064.5
Multinomial	169,265.0	168,993.6	168,962.2
Dirichlet-Multinomial	169,240.4	169,021.5	168,996.8

MM and DM models, but the mean absolute error is significantly larger due to its high variability, as reflected in the boxplot.

In general, the CM model tends to produce good estimates of IBNR claim counts when the integral

$$\int_{d_{t-1}}^{d_t} F_{U|t,x_i}(\tau - t) dt,$$

representing the probability that a claim occurring within the period $(d_{t-1}, d_t]$ is reported before the valuation date τ , aligns closely with the actual percentage of such claims. However, when this integral deviates significantly from the actual percentage, the CM model’s estimates deteriorate. For example, at the end of January 2016, the average value of the integral

$$\int_{d_{T-1}}^{d_T} F_{U|t,x_i}(\tau - t) dt,$$

where $d_T = \tau$ (which represents the probability that a claim that occurred in January 2016 was reported before the end of the month), across policyholders was 0.81095, while the actual percentage was 0.7465753 (i.e., 75% of such claims were reported by the end of the month) – a difference that diverged from the usual pattern. This discrepancy, which was not captured by our model for reporting delay despite using all available data, significantly impacted the IBNR claim count estimate; the underestimated value of

$$1 - \int_{d_{T-1}}^{d_T} F_{U|t,x_i}(\tau - t) dt$$

led to an underestimated IBNR claim count.

Effect of Increasing States on Model Performance

As the number of states g increases, the performance of the discrete-time models (MM and DM) improves, suggesting that these models benefit from more states in capturing the underlying dynamics of the data. However, this trend does not hold for the continuous-time model (CM). The CM model’s performance remains inconsistent due to its sensitivity to the fitting of reporting delays and the two-step maximization of the likelihood process, as explained in the previous analysis. Therefore, while adding states improves the accuracy of discrete-time models, it does not necessarily translate into better performance for the continuous-time model.

BIC for the Fitted Models

We report the BIC for the fitted models at the first reserve valuation date. The BIC values for the Continuous-time, Multinomial, and Dirichlet-Multinomial models with $g = 2, 3, 4$ are shown in Table 3.

The results indicate that the BIC is lowest for $g = 4$ across all three models, suggesting that $g = 4$ provides the best fit. However, the BIC values for $g = 3$ and $g = 4$ are very close, indicating that the improvement in fit from $g = 3$ to $g = 4$ is relatively small.

Effect of Incorporating the HMM

Table 2 also shows summary statistics for the APE in the case $g = 1$, which translates to the Poisson process. It is clear that the performance of both the discrete-time models improves by incorporating the HMM. For both models, the mean APE decreases significantly as the number of states increases. This indicates that incorporating the HMM structure, with its ability to account for different state transitions, provides better accuracy in the IBNR claim count estimates. The improvement is seen not only in the mean APE but also in the reduction of the standard deviation of the APE, reflecting the model's increased stability and precision with more states.

In contrast, the continuous-time CM model shows minimal change when incorporating the HMM. The APE remains largely consistent, indicating that the continuous-time model is less affected by the number of states. The same reasoning we outlined in the previous subsection, when talking about the effect of increasing states on model performance, applies here.

Effect of Including Covariates

To evaluate the effect of including the vehicle and policyholder's risk characteristics, we fitted our models without incorporating any covariates in the framework. The results indicate that, across all numbers of hidden states g , the APE ranged between 16 and 17% for all three models (the errors are all very similar so we do not include a table). This is comparable to the results obtained using the CL method. These findings highlight the significant role of covariates in improving model performance, as their inclusion enables the models to account for variations in claim counts arising from individual-level characteristics, which cannot be captured in a purely aggregate framework.

Comparison with the IPW Method (Vanegas et al., 2024)

To provide additional context for the performance of our proposed models, we compared them to the inverse probability weighting (IPW) method proposed in Vanegas et al. (2024). This method was implemented using the same reporting delay regression structure as in our continuous-time model. The mean APE for the IPW method was 15%, only 2% less than the APE obtained by the CL.

It is important to note that our proposed models fundamentally differ from the IPW method in their approach. While the IPW focuses on estimating development factors through reporting delay regressions, our models incorporate the temporal dynamics of the claim arrival process. This temporal component enables our models to capture additional patterns and dependencies in the data, resulting in superior performance relative to the IPW method.

Comparison with Chain Ladder under Yearly Aggregation

Figure 4 reveals a clear seasonal pattern in the IBNR claim counts. Although Tee and Käärik (2022) argue that CL estimation tends to be more accurate under finer (e.g., monthly) aggregation, it is important to also consider coarser aggregations when seasonal effects in reporting structure are present. In practice, practitioners may prefer yearly aggregation to better capture full seasonal cycles in reporting structures.

To assess the impact of aggregation choice on model performance, we compare our models against the CL method applied to yearly aggregated data. Since the CL method with yearly aggregation does not allow restriction to claims reported within a maximum delay of 9 months, we estimate the total IBNR claim count for all models to ensure a fair comparison. For our discrete-time models, which assume a maximum delay of $D = 9$, we estimate the long-delay component (i.e., for delays $d > 9$) using the empirical average:

$$\hat{z}_{i,d} = \frac{1}{|\mathcal{L}_d|} \sum_{l \in \mathcal{L}_d} z_{l,d}, \quad \text{where } \mathcal{L}_d = \{l \mid l + d \leq T\},$$

Table 4. Mean APE for IBNR claim count estimates across models.

Model	CL-monthly	CL-yearly	CM ($g = 4$)	MM ($g = 4$)	DM ($g = 4$)
Mean APE	0.1555	0.1269	0.1142	0.0871	0.0810

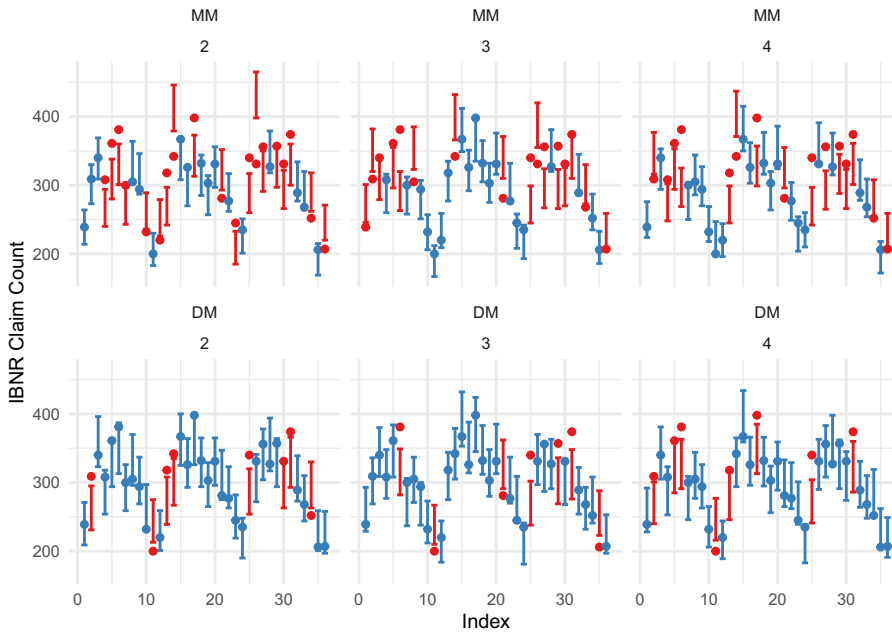


Figure 4. Error bar plots showing the 95% confidence intervals for simulated IBNR claim count estimates using the DM and MM models with $g = 2, 3, 4$. The red points indicate when the actual IBNR claim count falls outside the interval, and the blue points indicate when it falls within the interval

and $z_{l,d}$ is the number of claims occurring at time l and reported with delay d . This correction is applied to all censored entries with $d > 9$ and $t + d > T$, up to the maximum observed reporting delay in the data, allowing for a complete IBNR claim count estimate.

We compute IBNR claims counts at 24 monthly valuation dates spanning January 2014 to December 2015. Given that the maximum reporting delay in the dataset is two years and data are available up to December 2017, we are able to compare these estimates with the actual observed IBNR claim counts. The mean APE for IBNR claim counts across the 24 valuation dates - calculated for the CL method with monthly aggregation, the CL method with yearly aggregation, and the three proposed HMM-based models with 4 states - is shown in Table 4.

The results confirm that yearly aggregation improves CL performance compared to monthly aggregation, likely due to better alignment with seasonal patterns in reporting delays. Notably, the performance of CL-Yearly is comparable to our continuous-time model. Nevertheless, both of our discrete-time HMM-based models outperform the CL benchmark under both aggregation schemes.

Analysis of Confidence Intervals for DM and MM Models

The error bar plots in Figure 4 illustrate the 95% confidence intervals for the simulated IBNR claim count estimates using the DM and MM models with $g = 2, 3, 4$, while the points represent the actual IBNR claim counts. The red points indicate instances where the actual IBNR claim count falls outside the interval, and the blue points indicate instances where it falls within the interval. It is evident that the DM model’s confidence intervals are generally wider and more frequently contain the actual

IBNR claim counts compared to the MM model. Specifically, for the DM model, the confidence intervals include the actual IBNR claim count in 28, 29, and 28 out of the 36 valuation dates for $g = 2, 3, 4$, respectively. In contrast, the MM model's intervals contain the actual IBNR claim counts only 17, 20, and 21 times for $g = 2, 3, 4$, respectively. The increased width of the confidence intervals in the DM model reflects the introduction of variability in the reporting probability vector through the Dirichlet assumption. This added variability leads to a more appropriate distribution of the IBNR claim counts, capturing the inherent uncertainty and providing a more realistic estimate. This highlights the importance of considering such variability for better risk management, as underestimating or overestimating reserves can have significant financial implications.

8. Conclusion

In this paper, we proposed a novel micro-level Cox model designed to enhance the estimation of IBNR claim counts in P&C insurance. Our model incorporates a HMM to capture the temporal dependencies in the claim arrival process, allowing for the integration of policyholder-level data and environmental factors. This framework supports analysis at varying levels of granularity, providing flexibility for different applications. We initially presented the model in a continuous-time framework, then extended it to a discrete-time framework to enable simultaneous modeling of claim occurrence and reporting delays. Additionally, we introduced a Dirichlet distribution assumption for the reporting delay probabilities, addressing nonsystematic or unexplainable variations in delay structures.

Our empirical analysis demonstrated that the discrete-time models, particularly the one incorporating the Dirichlet assumption, outperformed the continuous-time model, providing more accurate estimates of the IBNR claim count, with both frameworks outperforming the classical CL method. This finding highlights the importance of jointly modeling claim occurrence and reporting behavior, as well as the value of accounting for randomness in reporting delays.

For future research, we plan to conduct a deeper analysis of the applicability of continuous-time models in the literature. Given the superior performance of our discrete-time approach, it is essential to reassess the contexts in which continuous-time models are most appropriate. Additionally, the inclusion of the Dirichlet assumption in our model suggests a promising direction for the application of Bayesian methods in IBNR reserving. By transitioning to a Bayesian framework, we aim to further improve the robustness and flexibility of our reserve estimates.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/asb.2025.10056>

Funding statement. The authors acknowledge the financial support provided by the Natural Sciences and Engineering Research Council of Canada (Andrei L. Badescu: Grant No. RGPIN-2017-06684; Radu Craiu: Grant No. RGPIN-2024-04506; X. Sheldon Lin: Grant No. RGPIN-2023-04326).

Competing interests. The authors declare none.

References

- 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., Wong, B. and Yang, X. (2016) A micro-level claim count model with overdispersion and reporting delays. *Insurance: Mathematics and Economics*, **71**, 1–14.
- Badescu, A.L., Chen, T., Lin, X.S. and Tang, D. (2019) A marked cox model for the number of ibnr claims: Estimation and application. *ASTIN Bulletin: The Journal of the IAA*, **49**(3), 709–739.
- Badescu, A.L., Lin, X.S. and Tang, D. (2016) A marked cox model for the number of ibnr claims: Theory. *Insurance: Mathematics and Economics*, **69**, 29–37.

- Baudry, M. and Robert, C.Y. (2019) A machine learning approach for individual claims reserving in insurance. *Applied Stochastic Models in Business and Industry*, **35**(5), 1127–1155.
- Bjarnason, T. and Sjögren, M. (2014) Insurance loss reserving.
- Blanchard, P., Higham, D.J. and Higham, N.J. (2021) Accurately computing the log-sum-exp and softmax functions. *IMA Journal of Numerical Analysis*, **41**(4), 2311–2330.
- Booth, J.G. and Hobert, J.P. (1999) Maximizing generalized linear mixed model likelihoods with an automated Monte Carlo em algorithm. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **61**(1), 265–285.
- Bornhuetter, R.L. and Ferguson, R.E. (1972) The actuary and ibnr. *Proceedings of the Casualty Actuarial society*, Vol. 59, pp. 181–195.
- Bücher, A. and Rosenstock, A. (2023) Micro-level prediction of outstanding claim counts based on novel mixture models and neural networks. *European Actuarial Journal*, **13**(1), 55–90.
- Charpentier, A. and Pigeon, M. (2016) Macro vs. micro methods in non-life claims reserving (an econometric perspective). *Risks*, **4**(2), 12.
- 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.
- De Felice, M. and Moriconi, F. (2019) Claim watching and individual claims reserving using classification and regression trees. *Risks*, **7**(4), 102.
- Delong, L., Lindholm, M. and Wüthrich, M.V. (2022) Collective reserving using individual claims data. *Scandinavian Actuarial Journal*, 2022(1), 1–28.
- England, P.D. and Verrall, R.J. (2002) Stochastic claims reserving in general insurance. *British Actuarial Journal*, **8**(3), 443–518.
- Fung, T.C., Badescu, A.L. and Lin, X.S. (2021) A new class of severity regression models with an application to ibnr prediction. *North American Actuarial Journal*, **25**(2), 206–231.
- Fung, T.C., Badescu, A.L. and Lin, X.S. (2022) Fitting censored and truncated regression data using the mixture of experts models. *North American Actuarial Journal*, **26**(4), 496–520.
- Grandell, J. (2012) *Aspects of Risk Theory*. Springer Science & Business Media.
- Haastrup, 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.
- Holmberg, R.D. (1994) Correlation and the measurement of loss reserve variability. *Casualty Actuarial Society Forum*, Vol. 1, pp. 247–278. Citeseer.
- Huang, J., Qiu, C. and Wu, X. (2015) Stochastic loss reserving in discrete time: Individual vs. aggregate data models. *Communications in Statistics-Theory and Methods*, **44**(10), 2180–2206.
- Huang, J., Wu, X. and Zhou, X. (2016) Asymptotic behaviors of stochastic reserving: Aggregate versus individual models. *European Journal of Operational Research*, **249**(2), 657–666.
- Jackson, C. (2016) flexsurv: A platform for parametric survival modeling in R. *Journal of Statistical Software*, **70**, 1–33.
- Larsen, C.R. (2007) An individual claims reserving model. *Astin Bulletin*, **37**(1), 113–132.
- Lawless, J. (1994) Adjustments for reporting delays and the prediction of occurred but not reported events. *Canadian Journal of Statistics*, **22**(1), 15–31.
- MacDonald, I.L. and Zucchini, W. (2016) Hidden Markov models for discrete-valued time series. In *Handbook of Discrete-Valued Time Series*, pp. 267–286.
- 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. (1999) The standard error of chain ladder reserve estimates: Recursive calculation and inclusion of a tail factor. *ASTIN Bulletin: The Journal of the IAA*, **29**(2), 361–366.
- Maier, M. (2014) Dirichletreg: Dirichlet regression for compositional data in R.
- Norberg, R. (1993) Prediction of outstanding liabilities in non-life insurance I. *ASTIN Bulletin: The Journal of the IAA*, **23**(1), 95–115.
- Norberg, R. (1999) Prediction of outstanding liabilities II. Model variations and extensions. *ASTIN Bulletin: The Journal of the IAA*, **29**(1), 5–25.
- Rosenlund, S. (2012) Bootstrapping individual claim histories. *ASTIN Bulletin: The Journal of the IAA*, **42**(1), 291–324.
- Ruth, W. (2024) A review of Monte Carlo-based versions of the em algorithm. arXiv preprint arXiv:2401.00945.
- Sentürk, D. and Müller, H.-g. (2005) Covariate adjusted correlation analysis via varying coefficient models. *Scandinavian Journal of Statistics*, **32**(3), 365–383.
- Shi, P., Basu, S. and Meyers, G.G. (2012) A bayesian log-normal model for multivariate loss reserving. *North American Actuarial Journal*, **16**(1), 29–51.
- Swamy, P.A. (1970) Efficient inference in a random coefficient regression model. *Econometrica: Journal of the Econometric Society*, 311–323.
- Taylor, G. (2012) *Loss Reserving: An Actuarial Perspective*, Vol. 21. Springer Science & Business Media.
- Tee, L. and Käärrik, M. (2022) Estimating ibnr claim counts using different levels of data aggregation. *Communications in Statistics-Simulation and Computation*, **51**(6), 3013–3036.
- Vanegas, S.C., Badescu, A.L. and Lin, X.S. (2024) Claim reserving via inverse probability weighting: A micro-level chain-ladder method. *European Actuarial Journal*, 1–47.
- Verbelen, R., Antonio, K., Claeskens, G. and Crevecoeur, J. (2022) Modeling the occurrence of events subject to a reporting delay via an em algorithm. *Statistical Science*, **37**(3), 394–410.

Wahl, F. (2019) Explicit moments for a class of micro-models in non-life insurance. *Insurance: Mathematics and Economics*, **89**, 140–156.
 Wüthrich, M.V. (2018a) Machine learning in individual claims reserving. *Scandinavian Actuarial Journal*, **2018**(6), 465–480.
 Wüthrich, M.V. (2018b) 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. EM Algorithm for the Continuous-time Model

We maximize the likelihood for the continuous-time model (Equation 4.1) using the two-step procedure commonly adopted in the micro-level reserving literature: first, we fit the right-truncated reporting delay distribution, and then estimate the frequency model while treating the delay distribution as fixed. This appendix focuses on the estimation of the “Claim Frequency” component under this framework.

Recall from Proposition 2 that $N_{i,t}^r | C_t = j \sim \text{Poisson}(\lambda^{(j)}(\mathbf{x}_t) a_t(\mathbf{x}_t))$, where $a_t(\mathbf{x}_t) = e_{i,l} \left(\int_{d_{t-1}}^{d_t} P_{U|\mathbf{x}_t}(\tau - t) dt \right)$. We now describe the EM algorithm needed to fit the “Claim Frequency” component (Equation 5.1), assuming that $a_t(\mathbf{x}_t)$ is a known constant for all $l = 1, \dots, T$ and $i = 1, \dots, m$.

Complete Log-Likelihood

The complete data are given by the observed data (the discretely observed reported claims) and the unobserved data (the states of the HMM). It is straightforward to show that the log-likelihood of the complete data is given by

$$\begin{aligned}
 l_C^{(1)}(\Phi_1 | N_i^{(r,1:T)} = \mathbf{n}_i^{(r,1:T)}, C_t = c_t, i = 1, \dots, m, t = 1, \dots, T) \\
 = \sum_{j=1}^g u_{1j} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T v_{ijk} \right) \log \gamma_{jk} + \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m u_{ij} \log \mathbb{P}_j(n_{i,t}^r).
 \end{aligned}$$

where $u_{ij} = \mathbb{1}\{C_t = j\}$, $v_{ijk} = \mathbb{1}\{C_{t-1} = j, C_t = k\}$, and $\mathbb{P}_j(n_{i,t}^r) = P(N_{i,t}^r = n_{i,t}^r | C_t = j)$.

E-step

In the E-step, we compute the conditional expectation of the complete log-likelihood given the observed data and the current estimator $\Phi_1^{(k-1)}$ of the model parameters Φ_1 . This conditional expectation is given by: A

$$\begin{aligned}
 \mathbb{E} \left(l^{(1)}(\Phi_1 | \mathbf{n}^{(r,1:T)}, \mathbf{c}^{(T)} | \mathbf{n}^{(r,1:T)}, \Phi_1^{(k-1)}) \right) &= \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{l=1}^g \left(\sum_{t=2}^T \hat{v}_{jl}^{(k)} \right) \log \gamma_{jl} \\
 &+ \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{ij}^{(k)} \log \mathbb{P}_j(n_{i,t}^r), \tag{A1}
 \end{aligned}$$

where $\hat{u}_{ij}^{(k)}$ and $\hat{v}_{jl}^{(k)}$ are given by Equations (5.3) and (5.4), respectively.

M-step

The M-step in the EM algorithm involves maximizing Equation (A1) with respect to the model parameters $\Phi_1 = \{\pi_1, \Gamma, \theta_1, \dots, \theta_g\}$, subject to the necessary constraints. The optimal values for $\pi_{1j}^{(k)}$ and $\gamma_{jl}^{(k)}$ at the k th iteration are given by Equations (5.5) and (5.6), respectively. Finally, we aim to obtain the regression parameter $\theta_j^{(k)}$ that maximizes the third term of Equation (A1), which is the likelihood for a weighted Poisson regression. The maximization of this likelihood is readily available in many software packages in R.

Convergence Criterion, Model Initialization and Model Selection

We use similar procedures as those for the discrete-time models.