

Marked Cox Models for IBNR Claims Count: Continuous and Discretized Approaches with Dirichlet-Driven Reporting Delays

Hassan Abdelrahman^{*1}, Andrei Badescu¹, Radu Craiu¹, Sheldon Lin¹

¹Department of Statistical Sciences, University of Toronto

^{*}hassan.abdelrahman@mail.utoronto.ca

Abstract

Accurate loss reserving is crucial in Property and Casualty (P&C) insurance for financial stability, regulatory compliance, and effective risk management. We propose a novel micro-level Cox model based on hidden Markov models (HMMs). 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.

Keywords: Dirichlet distribution, EM algorithm, hidden Markov model, IBNR claims, marked Cox model, micro-level reserving

1 Introduction

In Property and Casualty (P&C) insurance, accurate loss reserving is essential for insurers to guarantee financial stability, regulatory compliance, and effective risk management (Bjarnason and Sjögren (2014)). Loss reserving involves estimating the funds needed to cover future payments for incurred claims. Two key components of loss reserves are the Reported But Not Settled (RBNS) reserve, accounting for claims reported but not yet settled, and the Incurred But Not Reported (IBNR) reserve, accounting for incurred claims that are yet to be reported. While both components are essential, this paper focuses on IBNR reserving. Unlike RBNS claims, the number of IBNR claims is unknown to the insurer at the time of reserve valuation, necessitating sophisticated modeling techniques to accurately estimate this reserve. The complexity of IBNR reserving lies in forecasting the frequency and severity

of unreported claims, making it a challenging yet crucial aspect of the insurance reserving process.

The majority 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 accumulate historical data on the development of claims over time in a two dimensional table, referred to as run-off triangle, by aggregating payments by occurrence and development year. Stochastic models were subsequently 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 insights into the characteristics of individual claims, high parameter uncertainty due to the limited number of observations in aggregated data leading to reduced predictive power, and potential bias under certain conditions (Crevecoeur et al. (2022)). 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, 1999) are the first to propose a stochastic loss reserving framework at 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. 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, 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 outlined by Bücher and Rosenstock (2023), this approach utilizes Frequency-Severity or Chain Ladder-based methods to estimate IBNR reserves over discrete time intervals, as demonstrated by studies such as Wüthrich (2018a,b), 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, considerations for enhancing this aspect of loss reserving modeling have been highlighted, as discussed by Badescu et al. (2016). NHPP, assuming independence among claim numbers from different periods, fails to capture dependence among individual claim arrivals due to environmental variations affecting the entire portfolio (Grandell (2012)) and contradicts the calendar year effect seen in the run-off triangle (Holmberg (1994) and Shi et al. (2012)).

Addressing these issues involves incorporating a temporal dependence structure into the claim arrival process model, achievable by transitioning to a Cox process. Notable research in this direction includes Avanzi et al. (2016), who utilized the shot noise Cox process to model claim arrival, employing a reversible jump Markov Chain Monte Carlo method for estimation. Conversely, Badescu et al. (2016) employed a hidden Markov model (HMM)

with an Erlang state-dependent distribution to model the claim arrival process for the entire portfolio. Their work showcased that the discretely observed claim arrival process, along with associated reported claims and IBNR processes, follows a Pascal-HMM. The authors detailed the estimation algorithm for their model and its application in estimating IBNR claim count in Badescu et al. (2019). 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 introduce a novel micro-level Cox model that offers a level of flexibility allowing actuaries to customize the analysis to various levels of granularity. Specifically, our model accommodates modeling at the policyholder level, lines of business level, or the entirety of the portfolio, providing a comprehensive framework for analyzing claim arrival patterns. In the same spirit as Badescu et al. (2016), we adopt an HMM as the foundation of our Cox model. In our framework, we assume a shared hidden Markov process governing the claim arrival processes across different granular units within the portfolio. From a policyholder-level modeling perspective, this HMM serves as an environmental dynamic that simultaneously influences all policyholders. Given a state of the HMM, our model assumes that the claim arrival intensity for a policyholder is a constant, which depends on the risk attributes of the policyholder via a regression function. Thus, the claim arrival intensity for a policyholder is influenced by both external factors affecting all policyholders and their own risk attributes. The temporal dependence and environmental variation, often not incorporated in most reserving models, are captured through the common hidden Markov process. We show that within our framework, the discretely observed claim arrival process, the discretely observed reported claim process, and the discretely observed IBNR claim process for each policyholder belong to the family of Poisson-HMMs, where the state-dependent distribution depends on the policyholder’s risk attributes.

We initially propose our model in continuous time. In general, fitting continuous marked count process models to reported claims data (whether they are marked Poisson or marked Cox processes) 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. In recent micro-level reserving literature, only Wahl (2019) have attempted to maximize the likelihood at once for a specific class of micro-level reserving models. Traditionally, researchers have adopted a two-step maximization approach to maximize the likelihood of the observed data, initially fitting the right-truncated reporting delay and subsequently estimating the frequency model using the reporting-delay components in the frequency model as constants (which are estimated from the first step). However, this approach rests on the assumption of independence between claim frequency and reporting delay, which is often invalid; an increase in reported claims may indicate either a surge in actual claims occurrence, accelerated reporting, or a combination of both, challenging the independence assumption.

We address such 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 policyholder i , and $Z_{i,t,d}$ denotes the number of these claims reported d periods later, then we assume that $Z_{i,t,d} \mid N_{i,t}$ follows a multinomial distribution. We simultaneously model both the delay and occurrence components, capturing the intricate interplay between claim occurrence and reporting behavior.

Moreover, existing literature often assumes either 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 Şentürk 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 the three different models: the continuous-time model assuming two-step maximization of the observed data likelihood, and the discrete-time models, both with and without the Dirichlet assumption for the reporting delay probability vector. We also fit the three models to an auto-insurance dataset from a major European insurance company, which exhibits significant randomness in the occurrence and delay structure. 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 given the randomness exhibited in our data, adding the Dirichlet assumption provides more realistic interval estimates for the IBNR claim count.

This work contributes to the literature by:

- Extending the Cox model of Badescu et al. (2016, 2019) to various levels of granularity (e.g., policyholder level, lines of business level). This is specifically important if the distribution of granular units in our portfolio changes over time and these granular units possess different characteristics in terms of claim occurrence, reporting delay, and claim severity. Ignoring these differences would lead to biased reserve estimates.
- Building upon the work of Verbelen et al. (2022) on modeling event occurrence subject to reporting delay via an EM algorithm, where they assume a Poisson process for claim occurrence. We extend their analysis by allowing the occurrence process to follow a Poisson-HMM, as well as conducting our analysis from a micro-level perspective.
- Incorporating unsystematic randomness in the reporting delay probability vector by assuming that it follows a Dirichlet distribution, and outlining the EM algorithm needed to fit the model. We show that this is also an extension to the work of Verbelen et al. (2022) who do not assume randomness of the reporting delay.

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 development of the j th claim for the i th policy ($i = 1, \dots, m$) through three random variables $(T_{ij}, U_{ij}, \mathbf{Z}_{ij})$, where T_{ij} denotes the occurrence time of the claim, U_{ij} represents its reporting delay, and \mathbf{Z}_{ij} characterizes the development process following claim reporting. Chronologically ordered, $\{(T_{ij}, U_{ij}, \mathbf{Z}_{ij}), j = 1, 2, \dots\}$ constitutes the claim history process for policy i , whose risk attributes are denoted by \mathbf{x}_i . Additionally, the total claim count process for the i th policy is defined as $N_i(t) = \sum_{j=1}^{\infty} \mathbb{1}\{T_{ij} < t\}$, 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 $\{(U_{i1}, \mathbf{Z}_{i1}), (U_{i2}, \mathbf{Z}_{i2}), \dots\}$, and
2. The stochastic intensity function $\Lambda_i(t)$.

The assumption is made that the marks are independent, with a common density function given by $f_{U, \mathbf{Z} | t, \mathbf{x}_i}(u, \mathbf{z}) = f_{U | t, \mathbf{x}_i}(u) f_{\mathbf{Z} | u, t, \mathbf{x}_i}(\mathbf{z})$. 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 $d_{l-1} \leq t < d_l$, given by $e_{i,l} = \frac{\# \text{ days for which the contract is in force in } d_{l-1} \leq t < d_l}{\# \text{ days in } d_{l-1} \leq t < d_l}$. 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 $\gamma_{jk} = P(C_l = k \mid C_{l-1} = j)$, 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

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

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

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 ($i = 1, \dots, m$),

$$\Lambda_l(t; \mathbf{x}_i) = \begin{cases} \lambda^{(1)}(\mathbf{x}_i), & \text{with prob.} = \pi_{l1} \\ \vdots & \vdots \\ \lambda^{(g)}(\mathbf{x}_i), & \text{with prob.} = \pi_{lg} \end{cases}, \quad t \in [d_{l-1}, d_l).$$

Remarks

- It is important to note that our proposed Cox model framework generalizes the work of Badescu et al. (2016), albeit with different modeling assumptions. While Badescu et al. (2016) models the claim arrival process for the entire portfolio, our approach takes a more granular stance by considering the claim arrival process for each individual policyholder (or other granular units), incorporating their respective risk covariates \mathbf{x}_i . Thus, our model accounts not only for seasonality or environmental variations but also explicitly considers the risk characteristics of individuals, allowing the number of claims to depend on both external factors and the unique risk profiles within the portfolio.

Moreover, in the work of Badescu et al. (2016), the intensity function for a time period $[d_l, d_{l+1})$ is assumed to follow an Erlang distribution. While this assumption might be appropriate for a macro-level perspective, it is less suitable for our micro-level approach, where we employ a regression function to model the intensity. Assuming the intensity is constant and dependent on covariates simplifies our setup and aligns more closely with existing literature. In the next section, we show that the distribution of the number of claims in a given period under our model will be a mixture of Poisson distributions, which a good fit for the over-dispersed claim frequency data. Assuming an Erlang distribution for the intensity would result in a mixture of Pascal distributions for the number of claims, which is not expected to significantly enhance model performance. Additionally, with appropriate choices of predetermined time points, the intensity can be safely assumed to be constant.

- A special case within the above framework arises when the practitioner aims to model the claim arrival process across different lines of business. In this scenario, the index i corresponds to the i th line of business. If the regression function $\lambda^{(j)}(\cdot)$ is a constant-regression function, our model aligns with the contemporaneously conditionally independent multivariate HMM; see MacDonald and Zucchini (2016). Importantly, our framework permits the inclusion of portfolio-level information through covariates in the regression model (e.g., the average age of policyholders in the line of business). Many papers argue for the necessity of accounting for the dependence between different lines of business in reserving practices; the HMM captures this dependence via hidden states that influence the claim arrival process across all lines of business. Conversely, if the practitioner models the claim arrival process for a single line of business, they may choose i to correspond to the i th claim type, thereby modeling claim arrivals for different claim types. This distinction is crucial because even within the same line of business, the distribution of the marks may vary across different claim types.
- An essential assumption in our model is that, given a realization of the hidden state during $[d_{l-1}, d_l)$, the intensity function is considered constant - with respect to an individual's risk attributes - within that period. While data-driven choices for d_l could ideally capture variations in intensity over time, such an approach might overcomplicate the model and lead to overfitting issues. Setting d_l to weekly, biweekly, or monthly intervals provides a reasonable balance between granularity and simplicity.
- In our proposed framework, risk attributes \mathbf{x}_i are not explicitly time-indexed, aligning with the common practice where policyholders' risk characteristics are assumed to remain constant over time. It's crucial to recognize that, in practice, insurance risk attributes can dynamically change. For example, in auto insurance, age is often a risk attribute that changes upon policy renewal. However, this standard assumption should have no effect on the performance of our model.
- Similar to the model in Badescu et al. (2016), it is evident that both the micro-level mixed Poisson process and the Ammeter process (Ammeter (1948)) are special cases of the proposed model.

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 $N_i^r(t) = \sum_{j=1}^{\infty} \mathbb{1}\{T_{ij} < t; T_{ij} + U_{ij} \leq \tau\}$, where T_{ij} represents the claim occurrence time and U_{ij} denotes the reporting delay for the j th claim of the i th policyholder. Claims not reported by τ (i.e., $T_{ij} + U_{ij} > \tau$) contribute to the IBNR claim process, denoted as $\{N_i^{\text{ibnr}}(t); 0 \leq t \leq \tau\}$, and defined as $N_i^{\text{ibnr}}(t) = \sum_{j=1}^{\infty} \mathbb{1}\{T_{ij} < t; T_{ij} + U_{ij} > \tau\}$. Although both processes should be indexed by τ , we omit it for simplicity.

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 $\{(U_{i1}, \mathbf{Z}_{i1}), (U_{i2}, \mathbf{Z}_{i2}), \dots\}$ with common density function $f_{U, \mathbf{Z}|t, \mathbf{x}_i}(u, \mathbf{z}) = f_{U|t, \mathbf{x}_i}(u) f_{\mathbf{Z}|U, t, \mathbf{x}_i}(\mathbf{z})$. 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 $\{N_i^{\text{ibnr}}(t); 0 \leq t \leq \tau\}$ are also marked Cox processes. Their adjusted stochastic intensity functions are $\Lambda_i^r(t) = \Lambda_i(t) F_U(\tau - t) \mathbb{1}\{0 \leq t \leq \tau\}$ and $\Lambda_i^{\text{ibnr}}(t) = \Lambda_i(t) (1 - F_U(\tau - t)) \mathbb{1}\{0 \leq t \leq \tau\}$, respectively, and their independent marks follow adjusted position-dependent mark distributions $f_{U, \mathbf{Z}|t, \mathbf{x}_i}^r(u, \mathbf{z}) = \frac{f_{U|t, \mathbf{x}_i}(u) \mathbb{1}\{u > \tau - t\}}{F_{U|t, \mathbf{x}_i}(\tau - t)} f_{\mathbf{Z}|U, t, \mathbf{x}_i}(\mathbf{z})$ and $f_{U, \mathbf{Z}|t, \mathbf{x}_i}^{\text{ibnr}}(u, \mathbf{z}) = \frac{f_{U|t, \mathbf{x}_i}(u) \mathbb{1}\{u \leq \tau - t\}}{1 - F_{U|t, \mathbf{x}_i}(\tau - t)} f_{\mathbf{Z}|U, t, \mathbf{x}_i}(\mathbf{z})$, 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 . \square*

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 $\Lambda_i(t) = e_{i,l} \Lambda_l(t; \mathbf{x}_i)$ for $d_{l-1} \leq t < d_l$, with $l = 1, 2, \dots$ and $d_0 = 0$, 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 $\{N_{i,1}^{\text{ibnr}}, N_{i,2}^{\text{ibnr}}, \dots\}$, 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 the claim arrival process of policyholder i , $N_i(t; \mathbf{x}_i)$, 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, $\{N_{i,1}^{\text{ibnr}}, \dots, N_{i,k}^{\text{ibnr}}\}$, all fall under the class of Poisson-HMMs. They share the same hidden parameter process $\{C_1, C_2, \dots\}$ with $N(t; \mathbf{x}_i)$. Furthermore, their state-dependent distributions are all Poisson with the following probability functions, respectively,*

$$\begin{aligned} P(N_{i,l} = n \mid C_l = j) &= p(n; \tilde{\lambda}_l^{(j)}(\mathbf{x}_i)), \\ P(N_{i,l}^r = n \mid C_l = j) &= p(n; \mu_l^{(j)}(\mathbf{x}_i)), \\ P(N_{i,l}^{\text{ibnr}} = n \mid C_l = j) &= p(n; \nu_l^{(j)}(\mathbf{x}_i)), \end{aligned}$$

where

$$\begin{aligned}\tilde{\lambda}_l^{(j)}(\mathbf{x}_i) &= e_{i,l} \lambda^{(j)}(\mathbf{x}_i) (d_l - d_{l-1}) \\ \mu_l^{(j)}(\mathbf{x}_i) &= e_{i,l} \left(\int_{d_{l-1}}^{d_l} F_{U|t,\mathbf{x}_i}(\tau - t) dt \right) \lambda^{(j)}(\mathbf{x}_i), \\ \nu_l^{(j)}(\mathbf{x}_i) &= e_{i,l} \left(\int_{d_{l-1}}^{d_l} (1 - F_{U|t,\mathbf{x}_i}(\tau - t)) dt \right) \lambda^{(j)}(\mathbf{x}_i),\end{aligned}$$

and

$$p(n; \lambda) = \frac{\lambda^n e^{-\lambda}}{n!}.$$

As a result, for $l = 1, 2, \dots$, $N_{i,l}$, $N_{i,l}^r$, and $N_{i,l}^{ibnr}$ follow mixed Poisson distributions with the probability functions, respectively,

$$\begin{aligned}P(N_{i,l} = n) &= \sum_{j=1}^g \pi_{lj} p(n; \tilde{\lambda}_l^{(j)}(\mathbf{x}_i)), \\ P(N_{i,l}^r = n) &= \sum_{j=1}^g \pi_{lj} p(n; \mu_l^{(j)}(\mathbf{x}_i)), \\ P(N_{i,l}^{ibnr} = n) &= \sum_{j=1}^g \pi_{lj} p(n; \nu_l^{(j)}(\mathbf{x}_i)).\end{aligned}$$

□

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. The motivation for switching to a discrete-time framework arises from the complexities associated with the continuous-time framework, particularly the difficulties in maximizing the likelihood of the observed data.

We begin by outlining the likelihood of the observed data in the continuous-time framework and briefly discussing these complexities. 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 $d_T = \tau$. These claims, ordered by their arrival time, are described through the three random variables $\{t_{ijl}, u_{ijl}, \mathbf{z}_{ijl} : i = 1, \dots, m; j = 1, \dots, n_{i,l}^r\}$, where t_{ijl} is the claim arrival time of the j -th ordered claim for policyholder i in period $[d_{l-1}, d_l]$, u_{ijl} is its reporting delay, and \mathbf{z}_{ijl} is its development information. Our set of observed data is thus given by:

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

Let $\mathbf{N}_i^{(r,s:t)} = (N_{i,s}^r, \dots, N_{i,t}^r)$ be the discretely observed reported claim process from period s to period t for policyholder i , and $\mathbf{n}_i^{(r,s:t)} = (n_{i,s}^r, \dots, n_{i,t}^r)$ be its realization, it is easy to show that the likelihood of the observed data is composed of three terms (e.g., see Antonio and Plat (2014) and Fung et al. (2021)):

$$\begin{aligned} \mathcal{L}^{(1)}(\Phi_1 \mid \mathcal{O}_C) &\propto P \left(\underbrace{\mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}}_{\text{First Term (Claim Frequency)}} \right) \times \underbrace{\prod_{i=1}^m \prod_{l=1}^T \prod_{j=1}^{n_{i,l}^r} \frac{f_{U|t_{ijl}, \mathbf{x}_i}(u_{ijl})}{F_{U|t_{ijl}, \mathbf{x}_i}(\tau - t_{ijl})}}_{\text{Second Term (Reporting Delay)}} \\ &\times \underbrace{\prod_{i=1}^m \prod_{l=1}^T \prod_{j=1}^{n_{i,l}^r} f_{\mathbf{Z}|u_{ijl}, t_{ijl}, \mathbf{x}_i}(\mathbf{z}_{ijl})}_{\text{Third Term (Claim Development)}}, \end{aligned} \tag{1}$$

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 θ_j for the $\lambda^{(j)}$'s, and the parameters associated with modeling the reporting delay and claim severity.

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.

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 dependent on the policy characteristics and the time of occurrence. Fitting the right-truncated reporting delay regression requires special attention and can be a challenging task because: random right-truncation is not as well studied in survival analysis as left-truncation, and

empirical data for reporting delay show the necessity of having a flexible mixture regression model, which is not straight-forward for right-truncated data.

The ‘‘Claim Development’’ component is concerned with the development process after a claim is reported and involves right-censored data. In case we need to estimate the IBNR reserve, we will need to model claim severity, and thus, this third term is represented as $p_{Y|U,t,\mathbf{x}}(y)$ where Y denotes the claim severity. This models the relationship between claim severity and the reporting delay U , considering the policy characteristics \mathbf{x} .

As discussed in the introduction, maximizing the above likelihood is a challenging task. We observe that the third term, representing the claim development \mathbf{z} , is independent of the parameters in the first two terms. Hence, it can be maximized separately. From Proposition 2, we know that the reported claim counts’ distribution depends on the reporting delay, making the maximization of the first two terms of $\mathcal{L}^{(1)}$ a formidable task. As previously mentioned, the common approach in the micro-level reserving literature is to employ a two-step maximization. This two-step maximization approach rests on the assumption of independence between claim frequency and reporting delay, which is not a realistic assumption.

In what follows, we drop the claim development component from our analysis as we focus on the interaction between frequency and delay.

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 happens at period t (i.e. the claim happens in $[d_t, d_{t+1})$), it can be reported in periods $t + d$ ($d = 0, 1, \dots, D$). We also, for simplicity, assume that the predetermined time points d_t ($t = 0, \dots, T$) are equidistant, so that $d_{t+1} - d_t = 1$ and $d_T = \tau$. We denote the probability that a claim from policyholder 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, $\{N_{i,1}^{\text{ibnr}}, \dots, N_{i,T}^{\text{ibnr}}\}$, 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:

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

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 $Z_{i,t,d}$ 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 is 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 given $N_{i,t}^r$, the counts $Z_{i,t,d}$ ($d = 0, \dots, \min\{D, T - t\}$) follow a multinomial distribution with probabilities $\frac{p_t(d; \mathbf{x}_i)}{p_{i,t}^r}$ ($d = 0, \dots, \min\{D, T - t\}$). Note that for $t \leq T - D$, all claims occurred in period t is fully observed, i.e., $N_{i,t}^r = N_{i,t}$, 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:

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

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:

$$\begin{aligned} \mathcal{L}^{(2)}(\Phi_2 \mid \mathcal{O}_D) &\propto P\left(\mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}\right) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \prod_{d=1}^{\min\{D, T-t\}} \binom{\sum_{j=0}^d z_{i,t,j}}{z_{i,t,d}} q_t(d; \mathbf{x}_i)^{z_{i,t,d}} (1 - q_t(d; \mathbf{x}_i))^{\sum_{j=0}^{d-1} z_{i,t,j}}, \end{aligned} \quad (3)$$

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:

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

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 the unsystematic fluctuations in reporting delay that are not captured by our covariates. In cases where data exhibits 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. While an exact cause for the high fluctuations cannot be precisely identified without additional data collection, we believe that there are unobserved covariates who 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 $\boldsymbol{\eta}_{i,t} = (\eta_{t,0}(\mathbf{x}_i), \dots, \eta_{t,D}(\mathbf{x}_i))$. This choice is motivated by the fact that the Dirichlet distribution serves as the conjugate prior for the multinomial distribution, which significantly simplifies the mathematical treatment and estimation of our model.

The parameters $\boldsymbol{\eta}_{i,t} = (\eta_{t,0}(\mathbf{x}_i), \dots, \eta_{t,D}(\mathbf{x}_i))$ can be specified in various ways. One approach is to model $\boldsymbol{\eta}_{i,t}$ using a regression with fixed effects, where these parameters are expressed as a function of observed covariates, such as policyholder characteristics or temporal factors. Alternatively, we can incorporate the Dirichlet assumption as a random effect. For instance, we could assume that for each period t , the reporting delay probability vector \mathbf{p}_t follows a Dirichlet distribution with parameters $\boldsymbol{\eta}_t$. This formulation represents a "global model," where the reporting delay probability vector at period t for each policyholder i is drawn from this global distribution. However, this model assumes that all policyholders have similar reporting patterns, which may not always hold in practice. To refine this approach further, we could cluster the policyholders into groups $\mathcal{I}_1, \dots, \mathcal{I}_{n'}$, and assume that for each period t and for each group \mathcal{I} , the reporting delay probability vector $\mathbf{p}_t^{\mathcal{I}}$ follows a Dirichlet distribution with parameters $\boldsymbol{\eta}_t^{\mathcal{I}}$. In this case, the reporting delay probability vector at period t for all policyholders $i \in \mathcal{I}$ would be drawn from this group-specific model, allowing for a more granular modeling of heterogeneity in reporting behaviors across different groups. Lastly, a random coefficients model could be employed, where both the intercept and the coefficients associated with the covariates in the regression model are allowed to vary randomly across different clusters or observations (in this case, different policies). This approach provides further flexibility, enabling us to account for substantial random fluctuations in the delay probabilities.

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

Proposition 3. *If $\mathbf{p}_t(\mathbf{x}_i) = (p_t(0; \mathbf{x}_i), \dots, p_t(D; \mathbf{x}_i))$ follows a Dirichlet distribution with parameters $\boldsymbol{\eta}_{i,t} = (\eta_{t,0}(\mathbf{x}_i), \dots, \eta_{t,D}(\mathbf{x}_i))$, then the conditional probabilities $q_t(d; \mathbf{x}_i) = \frac{p_t(d; \mathbf{x}_i)}{\sum_{j=0}^D p_t(j; \mathbf{x}_i)}$*

$(d = 1, \dots, D)$ are independently distributed beta random variables with

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

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:

$$\begin{aligned} \mathcal{L}^{(3)}(\Phi_3 | \mathcal{O}_D) &\propto P\left(\mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}\right) \\ &\times \prod_{i=1}^m \prod_{t=1}^T \prod_{d=1}^{\min\{D, T-t\}} \binom{\sum_{j=0}^d z_{i,t,j}}{z_{i,t,d}} \frac{B(z_{i,t,d} + \eta_{t,d}(\mathbf{x}_i), \sum_{j=0}^{d-1} (z_{i,t,j} + \eta_{t,j}(\mathbf{x}_i)))}{B(\eta_{t,d}(\mathbf{x}_i), \sum_{j=0}^{d-1} \eta_{t,j}(\mathbf{x}_i))}, \end{aligned} \quad (4)$$

where B is the beta function, and Φ_3 represents the vector of the model's parameters, including: the initial distribution $\boldsymbol{\pi}_1$ and the transition matrix $\mathbf{\Gamma}$ 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.

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 Monte Carlo EM (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:

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

where $\mathbf{N}_i^{(r,1:T)} = (N_{i,1}^r, \dots, N_{i,T}^r)$ 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:

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

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

$$P\left(\mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}\right) = \boldsymbol{\pi}_1 \mathbf{P}_1(\mathbf{n}_1^r) \boldsymbol{\Gamma} \mathbf{P}_2(\mathbf{n}_2^r) \dots \boldsymbol{\Gamma} \mathbf{P}_T(\mathbf{n}_T^r) \mathbf{1}^T =: L_T,$$

where

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

and $\mathbf{1}$ is a row vector of 1s.

The Forward-Backward Algorithm

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

$$\begin{aligned} \alpha_{tj} &:= P(\mathbf{N}_1^{(r,1:t)} = \mathbf{n}_1^{(r,1:t)}, \dots, \mathbf{N}_m^{(r,1:t)} = \mathbf{n}_m^{(r,1:t)}, C_t = j), \\ \beta_{tj} &:= P(\mathbf{N}_1^{(r,t+1:T)} = \mathbf{n}_1^{(r,t+1:T)}, \dots, \mathbf{N}_m^{(r,t+1:T)} = \mathbf{n}_m^{(r,t+1:T)} \mid C_t = j), \end{aligned}$$

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):

$$\begin{aligned} P(C_t = j \mid \mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) &= \frac{\alpha_{tj} \beta_{tj}}{L_T}, \\ P(C_{t-1} = j, C_t = k \mid \mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}) &= \frac{\alpha_{t-1,j} \gamma_{jk} \mathbb{P}_k(\mathbf{n}_t^r) \beta_{tk}}{L_T}, \end{aligned}$$

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

They are computed using the following recursive relations:

$$\begin{aligned}\boldsymbol{\alpha}_1 &= \boldsymbol{\pi}_1 \mathbf{P}_1(\mathbf{n}_1^r), \\ \boldsymbol{\alpha}_t &= \boldsymbol{\alpha}_{t-1} \boldsymbol{\Gamma} \mathbf{P}_t(\mathbf{n}_t^r), \quad t = 2, \dots, T, \\ \boldsymbol{\beta}_T &= \mathbf{1}, \\ \boldsymbol{\beta}_t &= \boldsymbol{\Gamma} \mathbf{P}_{t+1}(\mathbf{n}_{t+1}^r) \boldsymbol{\beta}_{t+1}, \quad t = T-1, T-2, \dots, 1.\end{aligned}$$

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,
- $u_{tj} := \mathbb{1}\{C_t = j\}$, and
- $v_{tjk} := \mathbb{1}\{C_{t-1} = j, C_t = k\}$.

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:

$$\begin{aligned}\mathcal{C} &= \mathcal{O}_D \cup \{n_{i,t}^{\text{ibnr}}, z_{i,t,d} \mid i = 1, \dots, m; t = 1, \dots, T; d = 1, \dots, D, t+d > T\} \cup \{c_1, \dots, c_T\} \\ &= \{n_{i,t}, z_{i,t,d}, c_t \mid i = 1, \dots, m; t = 1, \dots, T; d = 0, \dots, D\}.\end{aligned}$$

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:

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

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:

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

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:

$$\begin{aligned}
\mathbb{E}[l_C^{(2)}(\Phi_2 \mid \mathcal{C}) \mid \Phi_2^{(k-1)}, \mathcal{O}_D] &= \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T \hat{v}_{tjk}^{(k)} \right) \log \gamma_{jk} \\
&+ \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{tj}^{(k)} \left[\hat{n}_{i,t,j}^{(k)} \log(\lambda^{(j)}(\mathbf{x}_i)) - \lambda^{(j)}(\mathbf{x}_i) \right] \\
&+ \sum_{i=1}^m \sum_{t=1}^T \sum_{d=1}^D \left[\hat{z}_{i,t,d}^{(k)} \log q_t(d; \mathbf{x}_i) + \left(\sum_{j=0}^{d-1} \hat{z}_{i,t,d}^{(k)} \right) \log(1 - q_t(d; \mathbf{x}_i)) \right] \\
&+ \text{constant},
\end{aligned} \tag{5}$$

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

$$\begin{aligned}
\hat{u}_{tj}^{(k)} &= \mathbb{E} \left(u_{tj} \mid \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)} \right) \\
&= \mathbb{P} \left(C_t = j \mid \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)} \right) \\
&= \frac{\alpha_{tj}^{(k-1)} \beta_{tj}^{(k-1)}}{L_T^{(k-1)}},
\end{aligned} \tag{6}$$

and

$$\begin{aligned}
\hat{v}_{tjl}^{(k)} &= \mathbb{E} \left(v_{tjl} \mid \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)} \right) \\
&= \mathbb{P} \left(C_{t-1} = j, C_t = l \mid \mathbf{n}^{(r,1:T)}, \Phi_2^{(k-1)} \right) \\
&= \frac{\alpha_{t-1,j}^{(k-1)} \gamma_{jl}^{(k-1)} \prod_{i=1}^m P(N_{i,t}^r = n_{i,t}^r \mid C_t = l, \Phi_2^{(k-1)}) \beta_{tl}^{(k-1)}}{L_T^{(k-1)}}
\end{aligned} \tag{7}$$

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:

$$\hat{n}_{i,t,j}^{(k)} := \mathbb{E}(N_{i,t} \mid C_t = j, \mathcal{O}_D, \Phi_2^{(k-1)}) = n_{i,t}^r + \mathbb{1}\{t > T - D\} e_{i,t} \lambda_j^{(k-1)}(\mathbf{x}_i) \sum_{d=T-t+1}^D p_t^{(k-1)}(d; \mathbf{x}_i),$$

$$\hat{z}_{i,t,d}^{(k)} := \mathbb{E}(Z_{i,t,d} \mid \mathcal{O}_D, \Phi_2^{(k-1)}) = \begin{cases} z_{i,t,d}, & t \leq T - D \\ \left(\sum_{j=1}^g \pi_{tj}^{(k-1)} e_{i,t} \lambda_j^{(k-1)}(\mathbf{x}_i) \right) p_t^{(k-1)}(d; \mathbf{x}_i), & t > T - D \end{cases}$$

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 $\lambda_j(\mathbf{x}_i)$, 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)) 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 $\boldsymbol{\pi}_1$ and $\boldsymbol{\Gamma}$ separately by maximizing the first two terms of Equation (5), respectively. The optimal values for $\pi_{1j}^{(k)}$ and $\gamma_{jl}^{(k)}$ at the k th iteration are given by:

$$\pi_{1j}^{(k)} = \hat{u}_{1j}^{(k)}, \quad j = 1, \dots, g, \tag{8}$$

$$\gamma_{jl}^{(k)} = \frac{\sum_{t=2}^T \hat{v}_{tjl}^{(k)}}{\sum_{l=1}^g \sum_{t=2}^T \hat{v}_{tjl}^{(k)}}, \quad j, l = 1, \dots, g. \tag{9}$$

As for the third term:

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

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 package in R.

The model parameters are then given by $\Phi_2 = \{\boldsymbol{\pi}_1, \boldsymbol{\Gamma}, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_g, \boldsymbol{\delta}_1, \dots, \boldsymbol{\delta}_D\}$, where $\boldsymbol{\theta}_j$ is the vector of regression parameters for $\lambda^{(j)}(\mathbf{x})$, and $\boldsymbol{\delta}_d$ 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

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

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 $\boldsymbol{\pi}_1$ and $\boldsymbol{\Gamma}$ 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) recommends 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

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

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 $\mathbf{p}_t(\mathbf{x}_i)$ is random. To this end, we let our complete data be given by:

$$\begin{aligned} \tilde{\mathcal{C}} &= \mathcal{C} \cup \{p_{i,t,d} \mid i = 1, \dots, m; t = 1, \dots, T; d = 1, \dots, D\} \\ &= \{n_{i,t}, z_{i,t,d}, p_{i,t,d}, c_t \mid t = 1, \dots, T; d = 0, \dots, D\}, \end{aligned}$$

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:

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

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

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

E-Step

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

$$\begin{aligned} \mathbb{E} \left[l_C^{(3)}(\Phi_3 | \tilde{\mathcal{C}}) \mid \Phi_3^{(k-1)}, \mathcal{O}_D \right] &= \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{l=1}^g \left(\sum_{t=2}^T \hat{v}_{tjl}^{(k)} \right) \log \gamma_{jl} \\ &+ \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{tj}^{(k)} \left[\hat{n}_{i,t,j}^{(k)} \log(\lambda^{(j)}(\mathbf{x}_i)) - \lambda^{(j)}(\mathbf{x}_i) \right] \\ &+ \sum_{i=1}^m \sum_{t=1}^T -\log B(\boldsymbol{\eta}_{i,t}) \sum_{d=1}^D (\eta_{t,d}(\mathbf{x}_i) - 1) \widehat{\log p_{i,t,d}}^{(k)} \\ &+ \text{constant}, \end{aligned} \tag{10}$$

where $\hat{u}_{tj}^{(k)}$ and $\hat{v}_{tjl}^{(k)}$ are given by Equations (6) and (7), respectively, and

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

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}_{tj} and \hat{v}_{tjk} in the E-step differs slightly for the two models considered. The differences arise when computing the probability $P(N_{i,t}^r = n_{i,t}^r \mid \mathcal{O}_D, \Phi^{(k-1)})$. 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

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

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

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

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

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

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

where

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

Thus,

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

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 $\mathbf{p}_t(\mathbf{x}_i)$ can be selected by multivariate rejection sampling. We can write $P(\mathbf{p}_t(\mathbf{x}_i) = \mathbf{p} \mid \mathcal{O}_D) = ah(\mathbf{p})g(\mathbf{p})$. It is straightforward to see that h is the density function of Dirichlet distribution with parameters $(z_{i,t,0} + \eta_{t,0}(\mathbf{x}_i), \dots, z_{i,t,T-t} + \eta_{t,T-t}(\mathbf{x}_i), \eta_{t,T-t+1}(\mathbf{x}_i), \dots, \eta_{t,D}(\mathbf{x}_i))$, 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 $\sup_{\mathbf{p}}\{g(\mathbf{p})\} \leq \sum_{j=1}^g \pi_{tj} (\lambda^{(j)}(\mathbf{x}_i))^{n_{i,t}^r}$. By obtaining a sample for $\mathbf{p}_t(\mathbf{x}_i)$ 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.

M-Step

We now maximize Equation (10) 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 (8) and (9), 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, which rescales the probability estimates from the E-step to ensure that the constraints are satisfied.

Convergence Criterion, Model Initialization & 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 claim reserve 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 $F_{U|t,\mathbf{x}}(u)$.
2. *Compute Reporting Delay Probability:* Use the fitted reporting delay regression to compute the probability $\int_{d_{l-1}}^{d_l} F_{U|t,\mathbf{x}_i}(\tau - t) dt$ 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. 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 to find the most likely sequence of hidden states \mathbf{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 $N_{i,l}^{\text{IBNR}}$ using the Poisson distribution with mean $\lambda_l^{(\mathbf{c}_i^*)}(\mathbf{x}_i) \times (1 - \int_{d_{l-1}}^{d_l} F_{U|t,\mathbf{x}_i}(\tau - t) dt)$ for the corresponding hidden state from the previous step. The total IBNR claim count at time τ is then given by $\sum_{i=1}^m \sum_{t=1}^T \hat{N}_{i,l}^{\text{IBNR}}$, where $\hat{N}_{i,l}^{\text{IBNR}}$ is the simulated IBNR claim count for policyholder i and period l .

6. *Estimate the Reserve:* For the continuous-time model, simulate the occurrence time t for each IBNR claim from the previous step, assuming it is uniformly distributed across the period of occurrence. Then, simulate the reporting delay u from the fitted regression model for reporting delay, given that $u > \tau - t$. Finally, simulate the claim severity using a fitted regression model for claim severity, given the policyholder characteristics \mathbf{x}_i , the occurrence time t , and the simulated reporting delay u . The sum of these simulated amounts will constitute the reserve.

Note that *Steps 5 & 6* are repeated 1000 times to obtain a distribution for the IBNR claim counts and IBNR reserves. The point estimate for these amounts is the mean of the 1000 simulated values.

The process differs slightly for the discrete-time models. For the Multinomial model, the reporting delay component is fitted simultaneously with the frequency component using the EM algorithm. Thus, the estimation procedure starts from *Step 3*. The steps of estimating the IBNR claim counts (*Steps 4 & 5* from above) are essentially the same as in the continuous-time model, but we use the estimated reporting probability vector $\mathbf{p}_t(\mathbf{x}_i)$ in the simulation instead of the continuous distribution F_U . For the Dirichlet model, we obtain estimates for $\boldsymbol{\eta}_{i,t}$, the Dirichlet parameters of $\mathbf{p}_t(\mathbf{x}_i)$, in *Step 3*. This is used to first simulate $\mathbf{p}_t(\mathbf{x}_i)$, which is then used in *Step 5* to estimate the IBNR claim count as in the case of the Multinomial model. We can simulate the period of reporting d using the estimated reporting probability vector $\mathbf{p}_t(\mathbf{x}_i)$ and then use it to estimate the IBNR reserve by simulating from a fitted model for the average claim severity given the policyholder characteristics \mathbf{x}_i , the occurrence time t , and the simulated period of reporting d .

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, the age of the car, 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.

Delay (months)	0	1	2	3	4	5	6	7	8	9	10+
No. of claims	66475	11617	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

Table (1) Reporting Delay for Claims Occurring Before 2016

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




Figure (1) Exposure & Claim Count Per Exposure Between Jan 2009 and Dec 2015

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. Additionally, the choice of $D = 9$ is motivated by the limitations of the `DirichletReg` function in R, which encounters numerical instability with larger values due to the exceedingly small probabilities associated with longer delays.

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 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 policyholder-level covariates mentioned in the data description subsection and time-dependent covariates such as the month in which the claim occurred and the day of occurrence.

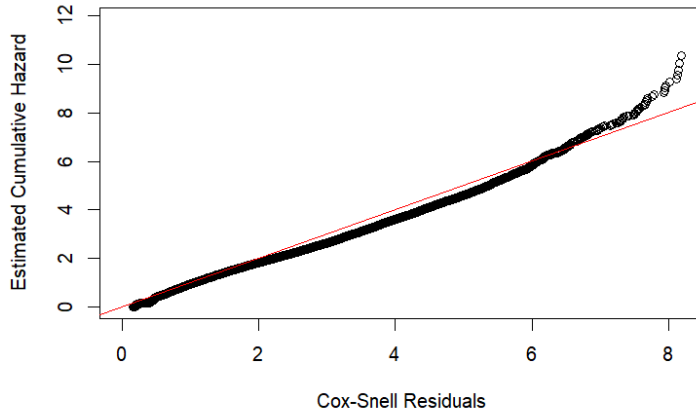


Figure (2) Cox-Snell Residual Plot for Assessing the Fit of the Reporting Delay Model

- *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 $q_t(d; \mathbf{x}_i) = \frac{p_t(d; \mathbf{x}_i)}{\sum_{j=0}^d p_t(j; \mathbf{x}_i)}$ for each $d \in \{1, \dots, D\}$ as a regression function of the policyholder’s risk attributes \mathbf{x}_i and time-dependent covariates. The policyholder-level covariates are those described in the data section, 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 $q_t(d; \mathbf{x}_i)$ 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 $q_t(d; \mathbf{x}_i)$ in the Multinomial Model.
- *The Chain Ladder (CL)*: We estimate the IBNR claim count using the Chain Ladder 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

- While setting $D = 9$ months simplifies the model, it may lead to underestimating the IBNR claim count because it ignores claims that are reported after this period. However, since nearly all claims are reported within this window, the number of missed claims should

be very small. In practice, practitioners can estimate this remaining number empirically using the formula:

$$\sum_{t=1}^T \sum_{d=T-t+1}^{\infty} \hat{z}_{t,d}$$

where T is such that $d_T = \tau$, $\hat{z}_{t,d}$ is the estimated number of claims that occurred in period t and were reported after a delay of d periods, and is calculated as:

$$\hat{z}_{t,d} = \frac{1}{T-1} \sum_{l=1}^{t-1} z_{l,d}$$

Here, $z_{l,d}$ represents the number of claims that occurred in period l and were reported with a delay of d periods. Essentially, $\hat{z}_{t,d}$ is the average number of claims reported after a delay of d periods, based on the observed claims data from previous periods.

- 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 $\left(1 - \int_{d_{l-1}}^{d_l} F_{U|t, \mathbf{x}_i}(\tau - t) dt\right)$ is replaced with:

$$\int_{d_{l-1}}^{d_l} P(U + t \in [\tau, t_{\max}] | t, \mathbf{x}_i) dt,$$

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.

7.3 Results

We present the results of estimating the IBNR claim count using the four different models, 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. Since our dataset extends to December 2017, and given that more than 99.99% of claims are reported within a year of occurrence, we have a reliable actual IBNR claim count to compare against our model estimates.

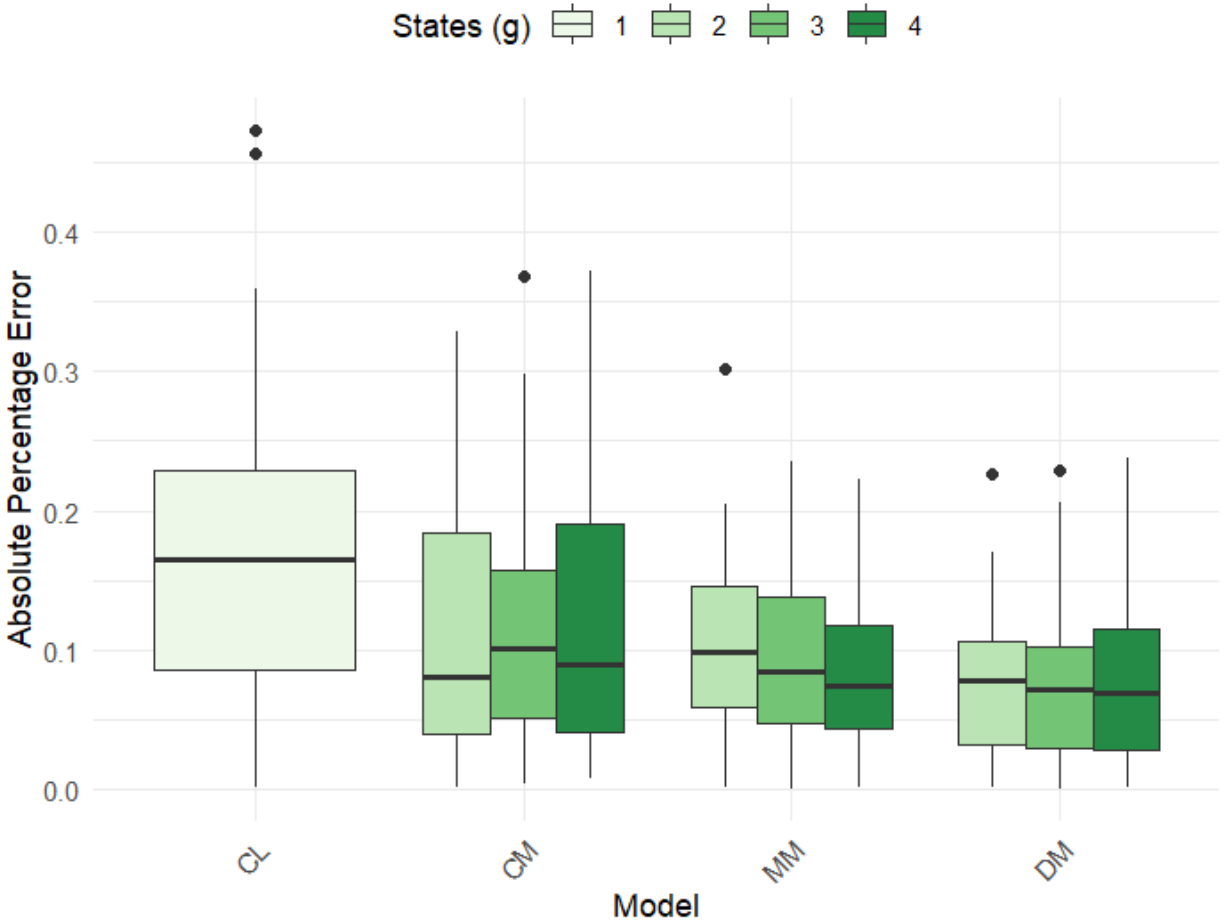


Figure (3) Boxplots of Absolute Percentage Errors for IBNR Claim Count Estimates by Model and Number of States

To evaluate the performance of each model, we calculated the absolute percentage error of the IBNR claim count point estimates at each of the 36 valuation points. The results are summarized in boxplots that display the distribution of absolute percentage errors for each model across the valuation period, as well as in a table that reports the mean, median, 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 offer additional granularity, it may be excessive for our relatively short monthly time series.

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

The classical Chain Ladder (CL) method performs worse than our proposed HMM-based models. The CL method shows a mean and median absolute percentage error 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.

	$g = 2$			
Models	CL	CM	MM	DM
Mean Absolute % Error	0.1689	0.1127	0.1021	0.0784
Median Absolute % Error	0.1648	0.0807	0.0978	0.0785
SD of Absolute % Error	0.1171	0.0901	0.0619	0.0537

	$g = 3$			
Models	CL	CM	MM	DM
Mean Absolute % Error	0.1689	0.1121	0.0931	0.0774
Median Absolute % Error	0.1648	0.1012	0.0848	0.0720
SD of Absolute % Error	0.1171	0.0875	0.0631	0.0630

	$g = 4$			
Models	CL	CM	MM	DM
Mean Absolute % Error	0.1689	0.1166	0.0875	0.0760
Median Absolute % Error	0.1648	0.0887	0.0747	0.0685
SD of Absolute % Error	0.1171	0.0914	0.0597	0.0560

Table (2) Summary of Absolute Percentage Errors for IBNR Claim Count Estimates Across Models

Performance of the Proposed Models

Our proposed models—the Continuous-time Model (CM), the Multinomial Model (MM), and the Dirichlet-Multinomial 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 continuous-time model (CM), as evidenced by their lower median and mean absolute percentage errors 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 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

$$\int_{d_{i-1}}^{d_i} F_{U|t, \mathbf{x}_i}(\tau - t) dt,$$

representing the probability that a claim occurring within the period $(d_{i-1}, d_i]$ 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

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

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

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

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.

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

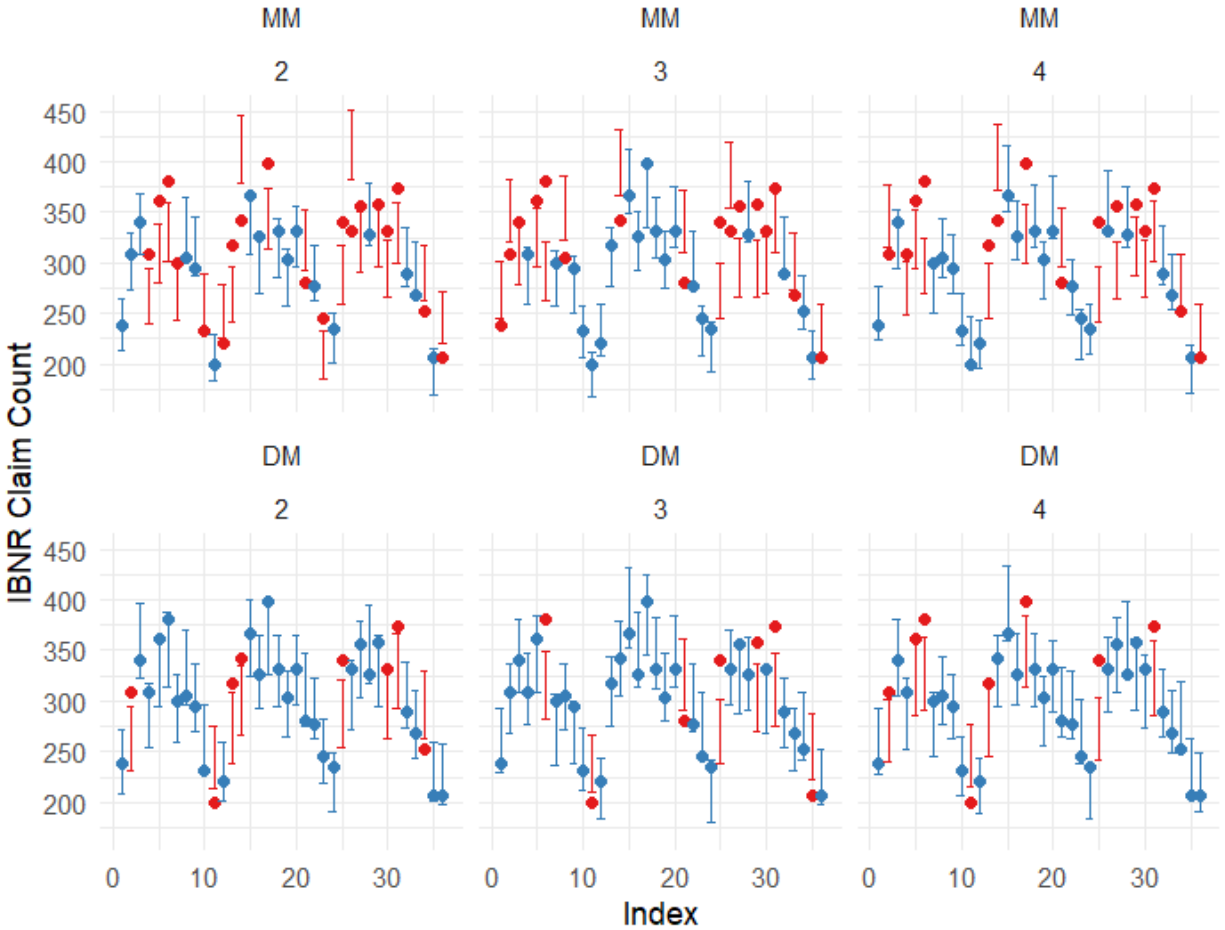


Figure (4) Error bar plots showing the 95% confidence intervals for simulated IBNR claim count estimates using the DM and MM models with $g = 2, 3, 4$. The red points indicate when the actual IBNR claim count falls outside the interval, and the blue points indicate when it falls within the interval.

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 non-systematic 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 Chain Ladder 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.

Competing interests: The authors declare none

References

- Ammeter, H. (1948). A generalization of the collective theory of risk in regard to fluctuating basic-probabilities. *Scandinavian Actuarial Journal*, 1948(1-2):171–198.
- 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. In *Proceedings of the casualty actuarial society*, volume 59, pages 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., Robben, J., and Antonio, K. (2022). A hierarchical reserving model for reported non-life insurance claims. *Insurance: Mathematics and Economics*, 104:158–184.
- 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.
- 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. In *Casualty Actuarial Society Forum*, volume 1, pages 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.
- 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. *Handbook of discrete-valued time series*, pages 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.
- Norberg, R. (1993). Prediction of outstanding liabilities in non-life insurance1. *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.
- Ruth, W. (2024). A review of monte carlo-based versions of the em algorithm. *arXiv preprint arXiv:2401.00945*.

- Şentürk, 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*, pages 311–323.
- Taylor, G. (2012). *Loss reserving: an actuarial perspective*, volume 21. Springer Science & Business Media.
- 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.

A EM Algorithm for the Continuous-time Model

Recall that the likelihood of the observed data for the continuous-time model, $\mathcal{L}^{(1)}$, is given by:

$$\mathcal{L}^{(1)}(\Phi_1 | \mathcal{O}_C) \propto P\left(N_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, N_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}\right) \times \prod_{i=1}^m \prod_{t=1}^T \prod_{j=1}^{n_{i,t}^r} \frac{f_{U|t_{itj}, \mathbf{x}_i}(u_{itj})}{F_{U|t_{itj}, \mathbf{x}_i}(\tau - t_{itj})},$$

where $\mathbf{N}_i^{(r,1:T)} = (N_{i,1}^r, \dots, N_{i,T}^r)$ 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. We maximize the likelihood by employing the two-step maximization widely used in the micro-level reserving literature. To this end, we describe the EM algorithm needed to fit the ‘Claim Frequency’ component assuming a known distribution for the reporting delay. Recall that the likelihood such component is given by

$$P\left(\mathbf{N}_1^{(r,1:T)} = \mathbf{n}_1^{(r,1:T)}, \dots, \mathbf{N}_m^{(r,1:T)} = \mathbf{n}_m^{(r,1:T)}\right) = \boldsymbol{\pi}_1 \mathbf{P}_1(\mathbf{n}_1^r) \boldsymbol{\Gamma} \mathbf{P}_2(\mathbf{n}_2^r) \dots \boldsymbol{\Gamma} \mathbf{P}_T(\mathbf{n}_T^r) \mathbf{1}^T =: L_T,$$

where

$$\mathbf{n}_t^r = (n_{1,t}^r, \dots, n_{m,t}^r),$$

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

and $\mathbf{1}$ is a row vector of 1s.

Complete Log-Likelihood

The complete data is given by the observed data (the discretely observed reported claims) and the unobserved data (the states of the HMM). The log-likelihood of the complete data is given by

$$\begin{aligned} l_C^{(1)}(\Phi_1 | \mathbf{N}_i^{(r,1:T)} = \mathbf{n}_i^{(r,1:T)}, C_t = c_t, i = 1, \dots, m, t = 1, \dots, T) \\ = \log \left(\pi_{1,c_1} \prod_{t=2}^T \gamma_{c_{t-1}, c_t} \prod_{t=1}^T \prod_{i=1}^m \mathbb{P}_{c_t}(n_{i,t}^r) \right) = \log \pi_{1,c_1} + \sum_{t=2}^T \log \gamma_{c_{t-1}, c_t} + \sum_{t=1}^T \sum_{i=1}^m \log \mathbb{P}_{c_t}(n_{i,t}^r) \\ = \sum_{j=1}^g u_{1j} \log \pi_{1j} + \sum_{j=1}^g \sum_{k=1}^g \left(\sum_{t=2}^T v_{tjk} \right) \log \gamma_{jk} + \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m u_{tj} \log \mathbb{P}_j(n_{i,t}^r). \end{aligned}$$

where $\mathbb{P}_j(n_{i,t}^r) = P(N_{i,t}^r = n_{i,t}^r | C_t = j)$.

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:

$$\begin{aligned} \mathbb{E} \left(l^{(1)}(\Phi_1 | \mathbf{n}^{(r,1:T)}, \mathbf{c}^{(T)}) | \mathbf{n}^{(r,1:T)}, \Phi_1^{(k-1)} \right) = \sum_{j=1}^g \hat{u}_{1j}^{(k)} \log \pi_{1j} + \sum_{j=1}^g \sum_{l=1}^g \left(\sum_{t=2}^T \hat{v}_{tjl}^{(k)} \right) \log \gamma_{jl} \\ + \sum_{j=1}^g \sum_{t=1}^T \sum_{i=1}^m \hat{u}_{tj}^{(k)} \log \mathbb{P}_j(n_{i,t}^r), \end{aligned} \quad (11)$$

where $\hat{u}_{ij}^{(k)}$ and $\hat{v}_{tjl}^{(k)}$ are given by Equations 6 and 7, respectively.

M-step

The M-step in the EM algorithm involves maximizing Equation 11 with respect to the model parameters $\Phi_1 = \{\boldsymbol{\pi}_1, \boldsymbol{\Gamma}, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_g\}$, 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 8 and 9, respectively. Finally, we aim to obtain the regression parameter $\boldsymbol{\theta}_j^{(k)}$ that maximizes the third term of Equation 11, given by:

$$\sum_{t=1}^T \sum_{i=1}^m \hat{u}_{tj}^{(k)} \log \mathbb{P}_j(n_{i,t}; \boldsymbol{\theta}_j) = \sum_{t,i} \hat{u}_{tj}^{(k)} \log \mathbb{P}_j(n_{i,t}; \boldsymbol{\theta}_j),$$

which is the likelihood for a weighted Poisson regression. The maximization of this likelihood is readily available in many software packages in R.

Remark: To expedite the maximization process, one can employ the binning technique for continuous covariates, treating them as nominal covariates. This approach significantly enhances the computational efficiency of the estimation. By creating categorical classes of policies based on binned covariates, we assume that policies within each class share similar risk characteristics. For an even faster maximization process, we can impose a restriction while binning continuous covariates, ensuring that each class has at least one claim occurrence for some $t \in 1, \dots, T$. This enables us to express the likelihood function as a weighted exponential likelihood regression. Consequently, the maximization step becomes notably faster and more manageable, allowing us to efficiently estimate the model parameters in just a few minutes.

Convergence Criterion, Model Initialization & Model Selection

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