Claim Reserving via Inverse Probability Weighting: A Micro-Level Chain-Ladder Method

Sebastián Calcetero Vanegas*, Andrei L. Badescu, X. Sheldon Lin

Department of Statistical Sciences University of Toronto Toronto, Ontario *sebastian.calcetero@mail.utoronto.ca

June 11, 2024

Abstract

Claim reserving primarily relies on macro-level models, with the Chain-Ladder method being the most widely adopted. These methods were heuristically developed without minimal statistical foundations, relying on oversimplified data assumptions and neglecting policyholder heterogeneity, often resulting in conservative reserve predictions. Micro-level reserving, utilizing stochastic modeling with granular information, can improve predictions but tends to involve less attractive and complex models for practitioners. This paper aims to strike a practical balance between aggregate and individual models by introducing a methodology that enables the Chain-Ladder method to incorporate individual information. We achieve this by proposing a novel framework, formulating the claim reserving problem within a population sampling context. We introduce a reserve estimator in a frequency and severity distribution-free manner that utilizes inverse probability weights (IPW) driven by individual information, akin to propensity scores. We demonstrate that the Chain-Ladder method emerges as a particular case of such an IPW estimator, thereby inheriting a statistically sound foundation based on population sampling theory that enables the use of granular information, and other extensions.

Keywords— Claim reserving, Population Sampling, Inverse Probability Weighting, Chain-Ladder, Survival modeling

1 Introduction

Claim reserving is a crucial aspect of insurance and risk management, and is vital for ensuring solvency, assessing risk, and setting appropriate premiums. The insurance industry employs several types of reserves but is primarily interested in the reserve for outstanding claims, which cover the estimated costs of unsettled and non-reported claims, representing the insurer's liability for future payments related to already occurred accidents. This reserve can be split into subcomponents depending on the source of the claim i.e., whether it is from a reported claim or not, and it is of interest to create this distinction for accounting purposes. These reserves play a vital role in maintaining financial stability and ensuring the availability of funds for future claim payments.

Reserving in general insurance is one of the most studied problems in actuarial research. See e.g., Schmidt (2017) for an extended list of most of the research covering this topic. Taha et al. (2021) and references therein provide a detailed overview of methods used in insurance reserving. Briefly, two primary approaches to reserving, namely micro-level and macro-level approaches, have been widely studied in the actuarial literature. The theoretical foundations of these methods can be found in the literature on stochastic claim reserving, e.g. Wüthrich and Merz (2008).

On the one hand, the macro-level approach to reserving focuses on estimating claim payments at an aggregate level. Among the macro-level approaches, Chain-Ladder-based techniques are widely employed in the insurance industry due to their ease of implementation, interpretation, and reliance on intuitive assumptions. See for e.g., Mack (1994), Quarg and Mack (2004), Martínez-Miranda et al. (2012). These methods avoid the use of very complex mathematical concepts, such as machine learning-based models or stochastic processes, instead relying on simple operations that can be implemented using spreadsheets. Additionally, they only require estimating development factors, which can be easily obtained from aggregate data without the need for specialized software. Consequently, the Chain-Ladder method and its variants are favored by insurance companies and regulators, with approximately more than 90% of insurers relying on them as their primary reserving methods, according to a survey made by Francis (2016).

However, these aggregate methods overlook the actual composition and heterogeneity of the insurance portfolio. The Chain-Ladder assumes homogeneity among claims within a given portfolio, disregarding valuable insights that can be gained by considering factors, such as attributes associated with the risk of each policyholder Wüthrich (2018b). The most recent literature on claim reserving (e.g., Crevecoeur et al. (2022) and literature therein) highlights the importance of using all the information available (i.e., the granular data) for the estimation of accurate reserves, and how ineffective is ignoring it. Consequently, the Chain-Ladder and similar macro-level models exhibit clear limitations and modest accuracy of estimation of the reserves when compared to models that do account for granular information i.e., micro-level models.

Moreover, the statistical foundation of aggregate models is most of the time questionable. The Chain-Ladder method was conceived as an ad-hoc method without reliance on a solid statistical basis that justifies the methodology's nature. Therefore, the Chain-Ladder has been the target of critics due to the lack of formalization with a well-founded statistical basis. Although research in stochastic claim reserving over the past decades has aimed at understanding the statistical framework behind the Chain-Ladder method, including the original work by Mack (1994) and subsequent studies through the decades e.g., England and Verrall (2002). Martínez-Miranda et al. (2012), Taylor and McGuire (2016) and Engler and Lindskog (2024); the answer is not yet clear and the literature is continuously evolving with new perspectives on the Chain-Ladder model. We refer to the introductions of Pittarello et al. (2023) and Engler and Lindskog (2024), and the literature therein, for a more comprehensive literature review on the understanding of the Chain-Ladder Method, as well as the most up-to-date developments on this matter.

On the other hand, the micro-level approach to reserving involves estimating individual claim payments by considering detailed characteristics, such as policyholder information, claim type, severity, and other relevant factors, e.g. Boumezoued and Devineau (2017). Micro-level reserving methods utilize probabilistic models that mimic the development of the claims, and are designed to directly capture the behavior of policyholders and their impact on reserves, resulting in relatively more accurate forecasts than the macro-level models. See for e.g., Taylor et al. (2008), Antonio and Plat (2014), Wüthrich (2018a), Fung et al. (2021) for various modeling examples.

While the theory behind micro-level models is relatively well understood given the nature of their construction based on statistical approaches, micro-level models pose challenges in terms of complexity, making them difficult to implement by insurance companies. These models incorporate both stochastic and predictive modeling, adding layers of complexity that may hinder their use in practice. Moreover, micro-level models often require assumptions about model components, such as distributions and simplifications of reality, which may be questionable. Consequently, these models are not widely adopted by actuarial practitioners due to the additional implementation effort required, and the lack of consensus on modeling practices from a regulatory standpoint. Indeed, according to the survey report by Francis (2016), micro-level reserving methods are virtually absent among insurance companies worldwide, with almost no one utilizing them, either as their primary method or for internal check-up purposes.

A main obstacle to the consideration of micro-level reserving by practitioners and regulators is the significant disparity in methodologies with respect to macro-level models, in addition to the associated effort required for their construction. Macro-level models, such as the Chain-Ladder, differ significantly from micro-level models in terms of how the reserve estimation is derived, as well as their theoretical foundations. Consequently, the transition from a macro-level to a micro-level model represents a substantial and challenging undertaking for any insurance company. Furthermore, regulators face difficulties in validating and accepting a micro-level model when its underlying principles deviate significantly from the familiar idea of Chain-Ladder and the construction of the reserves via development factors. Therefore, the substantial gap between these two key reserving methodologies hinders the adoption of micro-level modeling in the insurance industry.

In this paper, we focus on bridging the gap between macro-level and micro-level models by introducing a novel approach that translates the reserving problem in non-life insurance into a population sampling problem. By treating the reported claims as a sample from a larger population of claims, we develop a statistically sound approach based on an inverse probability weighting (IPW) method. We therefore develop a new methodology that accommodates for the introduction of individual claim information via a predictive model on the sampling probabilities, similar to how it is achieved to propensity scores in other fields. Our methodology enables the use of individual information in the Chain-Ladder method and therefore improves its performance while retaining most of its simplicity and interpretability.

One of the main features of the IPW estimator is that it exhibits a functional form reminiscent of the Chain-Ladder method and its development factors. However, it distinguishes itself by having claim-specific factors that depend on the attributes of the claims, and also does not require the specification of a certain granularity on the accident or development dates cells. As a result, our methodology can be viewed as a continuous "micro-level" version of the Chain-Ladder, where the development of each claim up to its ultimate value is performed at the individual level.

We observe that this similarity is not incidental and carries profound implications within the literature on claim reserving. We formally demonstrate that the Chain-Ladder method can be viewed as an empirical IPW estimator, operating under some specified homogeneity assumptions. Consequently, we assert that the theory on population sampling and IPW estimators establishes a robust and intuitively sound statistical framework for the Chain-Ladder method, distinct from the traditional approaches explored in the stochastic claim reserving literature that replicates, sometimes artificially, the Chain-Ladder estimate, as we mentioned earlier. In contrast, the population sampling framework addresses the reserving problem in a well-defined probabilistic setting, in which the Chain-Ladder method and its extensions naturally emerge without artificial constructions. Uncovering this connection introduces fresh perspectives on the Chain-Ladder method, allowing for extensions and a deeper comprehension of its theoretical properties inherited from its newly revealed nature as an IPW estimator.

The IPW approach signifies an enhancement over aggregate claim reserving models rooted in the traditional Chain-Ladder method, offering a cost-effective alternative to conventional microlevel reserving models. One of its primary strengths lies in its distribution-free methodology for estimating reserves, eliminating the need to specify parametric models for claim arrival processes (frequency) or claim amounts (severity). Notably, the IPW estimator only necessitates modeling the development of claims, encompassing reporting and payment delays, mirroring the approach of traditional aggregate models as the Chain-Ladder. Consequently, modeling efforts are streamlined to focus solely on estimating claim-specific inclusion probabilities based on observed delay distributions, simplifying the process compared to other individual reserving techniques.

The IPW approach integrates individual claims information in a simple, yet statistically justified manner. It maintains the practicality and interpretability characteristic of macro-level models, rendering it a more attractive choice for both practitioners and regulators. This approach may serve as an initial step to encourage practitioners, who typically rely on macro-level models, to explore the potential benefits and insights obtained from incorporating individual information in the reserving process. Ultimately, it paves the way for practitioners and regulators to consider tailored-made models based on micro-level techniques.

This paper is structured as follows: Section 2 introduces the reserving problem as a sampling problem, and shows the derivation of the IPW estimator for the reserve of all outstanding claims. Section 3 extends the methodology to consider other types of reserves, such as the incurred but not reported (IBNR) and the reported but not settled (RBNS) reserves, as well as their cumulative payments counterparts. Section 4 shows how the Chain-Ladder and some extensions can be viewed as IPW estimators, and discusses its implications. Section 5 discusses how to estimate the inclusion probabilities. Section 6 provides a numerical study on a real insurance dataset. Lastly, Section 7 provides the conclusion and future research directions.

2 Claim reserving via inverse probability weighting

In this section, we present the claim reserving problem and demonstrate how it can be effectively tackled using inverse probability weighting methods. Since there are various types of reserves in general insurance, in this section we provide the overall idea of the methodology for the total reserve of outstanding claims only. Section 3 will delve into the specific details of the methodology for the most prevalent and significant reserves in general insurance, namely RBNS and IBNR reserves.

2.1 The claim reserving problem

Suppose an insurance company is analyzing its total liabilities associated with claims whose accident times occur between t = 0 and $t = \tau$, where τ is the valuation time of analysis as defined by the actuary. In general insurance, accidents are often not immediately reported to the insurance company for various reasons, resulting in a significant delay between the occurrence of a claimable accident and the time the insurance company is notified. Therefore, at a given valuation time τ , the insurance company only has information on the claims reported by τ and is unaware of the unreported claims. Furthermore, the complexity of the problem increases due to another delay in the payment process. When a claim is reported, it is common for it to be paid in several subpayments over time rather than as a lump sum. This is because the impact of an accident can evolve, requiring additional payments until its associated claim is fully settled. Therefore, at a given valuation time τ , the insurance company is only aware of the claims that were reported on time, and for each one, it may have paid only a partial amount of the associated claim size.

Along those lines, the insurance company is interested in estimating the total claim amount of unreported claims, as well as the remaining payments of the reported claims, to construct the overall reserve of outstanding claims. This reserve is also known in the insurance jargon as the Incurred But Not Settled (IBNS) (Asmussen and Steffensen, 2020, p. 433), and is usually decomposed into further subcomponents depending on whether the payment is associated with a reported or not reported claim. For simplicity, we will use the terms IBNS reserve and the outstanding claims reserve interchangeably. Furthermore, in this Section, we consider the IBNS reserve without referring to the decomposition. Let us now describe the payment process as follows:

- Let $N(\tau)$ represent the total number of different payments associated with all the claims whose accident time is before the valuation time τ .
- Let Y_i , $i = 1, ..., N(\tau)$ denote the sequence of payment amounts. Note that some payments may belong to the same claim/accident, but we will not make any explicit distinction notationwise here. Similarly, we assume that the indexes *i* are sorted according to the payment time, for simplicity of exposition.
- Let T_i , $i = 1, ..., N(\tau)$ denote the sequence of accident times associated with the claim underlying each payment; let R_i , $i = 1, ..., N(\tau)$ denote the sequence of the associated reporting times; and let W_i , $i = 1, ..., N(\tau)$ denote the sequence of the associated times in which the payments take place. Clearly, $T_i < R_i < W_i$ and note that the values T_i, R_i would be the same for payments associated with the same claim, but the W_i would differ.
- Let $U_i = R_i T_i$, $i = 1, ..., N(\tau)$ be the sequence of the reporting delay times associated with the claim underlying each payment, $V_i = W_i - R_i$, $i = 1, ..., N(\tau)$ be the sequence of the associated payment delay time from reporting, and $Z_i = W_i - T_i = U_i + V_i$, $i = 1, ..., N(\tau)$ be the sequence of the associated payment delay time from the accident (i.e., total delay). Note that U_i is the same for all the payments associated with the same claim.
- Let X_i , $i = 1, ..., N(\tau)$ be the sequence of information/attributes of relevance, that is associated with the accident, the type of claim, the policyholder attributes, or the characteristics of the payment itself.
- Let $N^{P}(\tau)$ the number of payments made by valuation time τ out of the total $N(\tau)$, i.e., the number of payments made to the claims reported by τ .

Along those lines, the total liability of the insurance company associated with accidents occurring before the valuation time τ , which we will denote as $L(\tau)$, is given by

$$L(\tau) = \sum_{i=1}^{N(\tau)} Y_i.$$

Similarly, the portion of liability that is known to the insurance company (i.e., the so-called paid amount) by valuation time τ , which we will denote as $L^{P}(\tau)$, is

$$L^P(\tau) = \sum_{i=1}^{N^P(\tau)} Y_i.$$

Lastly, an actuary is interested in estimating the remaining liability associated with the outstanding claims i.e., the IBNS reserve. We will denote this quantity as $L^{IBNS}(\tau)$, and it is given by just the difference

$$L^{IBNS}(\tau) = L(\tau) - L^P(\tau).$$

This value is what the insurance company requires to set up the reserve for all outstanding claims, either non-reported, non-settled, or both, and is our goal for estimation. For further details of the claim reserving problem, we refer the reader to Wüthrich and Merz (2008).

2.2 A population sampling framework for claim reserving

Our proposal in this paper is based on a simple yet novel idea that allows us to frame the reserving problem in the context of *population sampling*, enabling us to leverage techniques from this field to our advantage. Population sampling is a statistical technique used to estimate population totals based on a smaller sample, especially in contexts where data collection from the entire population is impractical. The *sampling design* is the systematic process of selecting individuals or units from the population to be included in the sample. Different sampling methods are used depending on the research objectives and resources available. By using statistical techniques based on the sampling design, researchers can make reliable inferences about the population based on the sample.

Applying this concept to our reserving problem, we can consider all the Y_i , $i = 1, ..., N(\tau)$ payments as the population under study, while the current $N^P(\tau)$ payments made by the valuation date, i.e., the Y_i , $i = 1, ..., N^P(\tau)$, serve as the selected sample for understanding this population. It is important to note that the sampling design and the actual sampling process are not determined or performed by the investigator, but are purely driven by the randomness associated with whether a payment is made or not by the valuation date. Thus, the sample is given according to an unknown mechanism rather than being selected by the actuary.

The sampling mechanism based on the payment data can be conceptualized as a two-stage sampling process (Thompson, 2012, p. 171). In the first stage, a *Poisson* sampling without replacement is employed to sample the reported claims. This means, for each of the claims in the population, a Bernoulli experiment is conducted, where success is defined as the claim being reported by the valuation time, and failure occurs if it is not reported. Refer to (Särndal et al., 2003, p. 85) for more details on the Poisson sampling.

Moving to the second stage, we focus on the payments associated with each of the sampled claims from the previous stage (i.e., the reported claims). In this case, another sampling procedure is carried out to determine which payments of a claim are made before the valuation time and which are not. This is also achieved by Bernoulli-like experiments, however, do note that these are not independent because of the ordering of the payments e.g., a second payment of a claim can be sampled as long as the first payment is sampled.

As a result of the sampling, we can assign a dichotomic random variable $\mathbf{1}_i(\tau)$, $i = 1, \ldots, N(\tau)$ with success probability $\pi_i(\tau)$, to each of the payments Y_i in the population. Such a variable takes the value of 1 or 0, indicating whether the payment Y_i belongs to the sample of payments made or not, respectively, by a given valuation time τ . These variables are referred to as the *membership indicators* of the payments and are determined based on the delay in reporting (for the first stage of sampling) and the delay in payment (for the second stage) by the valuation time. Mathematically,

$$\mathbf{1}_{i}(\tau) = \mathbf{1}_{\{W_{i} \leq \tau\}} = \mathbf{1}_{\{T_{i}+U_{i}+V_{i} \leq \tau\}} = \mathbf{1}_{\{Z_{i} \leq \tau-T_{i}\}} = \mathbf{1}_{\{U_{i} \leq \tau-T_{i}\}} \mathbf{1}_{\{V_{i} \leq \tau-R_{i}\}}$$

where the indicators in the product on the right-hand side are the indicators of the first and second stages of sampling, respectively. The probabilities $\pi_i(\tau)$ are known as *inclusion probabilities* and can be interpreted as the likelihood of payment Y_i belonging to the sample or, equivalently, being paid by the valuation time τ . These probabilities are dependent on the valuation time and vary across payments due to the heterogeneity of the portfolio as reflected in the different accident times, different types of accidents and in general, different attributes associated with each payment. We also remark that it may be the case that these probabilities are also associated with the actual value of the payment sizes Y_i , which is the unit of measure in the sampling context. This potential dependence is known in the literature on population sampling as the *sampling being informative* as the actual values of the payments may be associated with the sampling design. Along those lines, the inclusion probabilities are given by

$$\pi_i(\tau) = P(U_i \le \tau - T_i | X_i, Y_i, T_i) \times P(V_i \le \tau - R_i | X_i, Y_i, T_i, R_i, U_i) = \pi_i^U(\tau) \times \pi_i^V(\tau)$$
(1)

where $\pi_i^U(\tau) = P(U_i \leq \tau - T_i | X_i, Y_i, T_i) = P(U_i \leq \tau - T_i | \mathcal{F}_i)$ and $\pi_i^V(\tau) = P(V_i \leq \tau - R_i | \mathcal{F}_i)$ $R_i | X_i, Y_i, T_i, R_i, U_i) = P(V_i \leq \tau - R_i | \mathcal{F}_i)$ are the inclusion probabilities of the first and second stage of sampling, respectively, which also depend on the characteristics of the payment. Here we use \mathcal{F}_i as the set of all characteristics of the *i*-th claim, i.e., the X_i, Y_i, T_i, R_i, U_i etc, to simplify the notation of these conditional probabilities. Note that the value of the second probability depends on the outcome of the first stage of sampling, and so $\pi_i^V(\tau)$ is in principle a conditional probability given the realization of U_i . Additionally, it is important to observe that the first probability is contingent solely upon the reporting of a claim, and thus remains consistent for all payments linked to the same claim. In contrast, the second probability is contingent upon the occurrence of the payment itself, hence it may fluctuate from one payment to another owing to the varying attributes/characteristics of each payment. We will write these probabilities as $\pi_i^U(\tau)$ and $\pi_i^V(\tau)$ to streamline the notation and emphasize that the indexation on *i* corresponds to the probabilities being specific for each payment as determined based on their attributes. It is important to note that these probabilities are not predefined and are therefore unknown to the investigator. We will delve into their estimation in Section 5, via the use of an auxiliary model.

Finally, note that the sample size in the design is not a fixed quantity. The sample size, which in our case is equivalent to the number of payments currently made $N^P(\tau)$, is a random variable defined as $N^P(\tau) = \sum_{i=1}^{N(\tau)} \mathbf{1}_i(\tau)$, which we can identify as the thinning of the counting process of the total number of payments.

2.3 Point estimation using the Horvitz-Thompson estimator

As motivated by the population sampling literature (e.g., (Särndal et al., 2003, p. 45)), a wellestablished estimator of the population total, which in our case is the total of payments (i.e., the ultimate $L(\tau)$) is provided by the *Horvitz-Thompson* (HT) estimator described as follows

$$\hat{L}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \frac{Y_{i}}{\pi_{i}(\tau)},$$
(2)

and therefore an estimator of the outstanding claims is the difference between the estimated total and the currently paid amount,

$$\hat{L}^{IBNS}(\tau) = \hat{L}(\tau) - L^{P}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \frac{Y_{i}}{\pi_{i}(\tau)} - \sum_{i=1}^{N^{P}(\tau)} Y_{i} = \sum_{i=1}^{N^{P}(\tau)} \frac{1 - \pi_{i}(\tau)}{\pi_{i}(\tau)} Y_{i}.$$
(3)

The intuition behind the HT estimator lies in the fact that only a portion of all payments Y_i is reported, proportionally to $\pi_i(\tau)$, and so each payment in the sample is "augmented" by a factor of $1/\pi_i(\tau)$ to approximate the actual total amount. The HT estimator has been demonstrated to be unbiased for the population total, regardless of the underlying sampling design (Thompson, 2012, p. 69), making it a highly practical choice. One of its most attractive features is that this property holds without assumptions on the probabilistic mechanism generating the data. Indeed, most of the inference results for the HT estimator are based on the so-called *design-based* approach (see for e.g. Arnab (2017)). This approach relies solely on the randomness associated with the sampling process, rather than the *model-based* approach, which relies on a probabilistic model for the population itself. From an actuarial viewpoint, this essentially implies that the HT estimator is distribution-free, meaning that it leads to an estimation of the reserve that does not require any assumptions on the underlying distribution of the number of claims (frequency) or the distribution of claim sizes (severity), which is a desirable property for reserving. Additionally, we emphasize the fact that even though the estimator is based on a population-level target (i.e., a macro-level scale), the inclusion probabilities are dependent on the individual attributes of policyholders, claims, and payments. Therefore the estimator incorporates granular information as part of the aggregate estimation.

The HT estimator is widely recognized as one of the most influential estimators in the statistics literature, having been extensively studied for over 70 years in population sampling (e.g., Arnab (2017)). Consequently, the HT estimator has a solid theoretical foundation and possesses numerous desirable properties that directly inherit to the claim-reserving problem, including consistency, unbiasedness, and sufficiency, among others. More recently, it has also been applied in inverse probability weighting (IPW) methods for estimation in causal inference, e.g. Yao et al. (2021), including applications in fairness in insurance. The terminology "IPW estimator" is more extended in and outside the statistics literature, and so we will mostly refer to the estimator of the reserve as the IPW estimator, and reserve the naming of the HT estimator when referring to the general concept.

In our specific context, it is essential to highlight that the estimator presented in Equation (3) can be evaluated once the inclusion probabilities have been estimated. Consequently, some of the properties of the HT estimator may be marginally influenced by this preliminary estimation. It is noteworthy that situations requiring such estimation are not uncommon in the context of IPW estimators and are indeed prevalent in practice. For example, in the field of causal inference, inclusion probabilities, often referred to as propensity scores, are routinely estimated before the application of IPW estimators (e.g. Salditt and Nestler (2023)). Other similar examples occur in the context of handling missing data, e.g., Section on Kim and Shao (2021). It is well-documented that, as long as these preliminary estimations are conducted reasonably, the resulting estimator maintains its reliability and properties. References such as Hirano et al. (2003) and Fattorini (2006) provide further insights into the robustness of estimators under such conditions.

A specific case of interest of the IPW estimator arises when we set $Y_i = 1$. In this scenario, all the sums above simplify to a count of the number of payments, allowing us to obtain an estimator for the number of the outstanding payments, $N^{IBNS}(\tau)$, to be made after the valuation time τ , as

$$\hat{N}^{IBNS}(\tau) := \hat{N}(\tau) - N^{P}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \frac{1}{\pi_{i}(\tau)} - \sum_{i=1}^{N^{P}(\tau)} 1 = \sum_{i=1}^{N^{P}(\tau)} \frac{1 - \pi_{i}(\tau)}{\pi_{i}(\tau)}.$$
(4)

It is noteworthy that this particular expression coincides, structure-wise, with the one employed by Hiabu et al. (2021) and Fung et al. (2022), specifically for estimating the number of incurred but not reported (IBNR) claims. In the former, the author proposed a density estimation of the distribution of claims in the run-off triangle and used an expression similar to the one in Equation (4) to estimate the number of IBNRs. However, this estimator does not account for the heterogeneity of the claims, nor include payment data. In the latter, the authors derived a similar expression to Equation (4) under the assumption that the number of unreported claims follows a geometric distribution, demonstrating its unbiasedness when the overall number of claims is driven by a Poisson process. It is therefore crucial to emphasize that within the framework of the HT estimator, this result generalizes immediately without the assumption of the geometric distribution, and it also accommodates the case of payment data.

A "Micro-level" Chain-Ladder method

From an actuarial standpoint, our proposed IPW estimator for the ultimate claims can be perceived as an individual-level adaptation of the Chain-Ladder method, by expressing the estimator in Equation (2) as

$$\hat{L}(\tau) = \sum_{i=1}^{N^P(\tau)} f_i(\tau) Y_i,$$

where, we can interpret $f_i(\tau) := 1/\pi_i(\tau)$ as an individual development factor assigned to each payment Y_i . These factors serve to project the payment to its ultimate value which aligns with the fundamental principle of the Chain-Ladder method. As the factors $f_i(\tau)$ are influenced by the policyholder's attributes, we can think of this methodology as a "micro-level" version of the Chain-Ladder method, as it implies the development factors on an individual level while retaining the essential characteristics of the Chain-Ladder. This analogy provides the IPW estimator with an intuitive and interpretable estimation of the reserve that is already well-established in the actuarial community and makes its use more appealing to practitioners. We will revisit this discussion with more in-depth details in Section 4.

2.4 Confidence interval of the estimation

Confidence intervals for the reserve can be constructed based on the sampling distribution of the HT estimator, as discussed by (Thompson, 2012, p. 70). In summary, under minimal regularity conditions, the HT estimator follows approximately a normal distribution under the two-stage sampling design for large populations Chauvet and Vallée (2020). Thus, an approximate $1 - \alpha$ confidence interval can be constructed using normal quantiles. However, as explained by (Thompson, 2012, p. 70), the accuracy of the normal approximation relies on the sample size and the distribution of Y_i , which tends to exhibit skewness and heavy tails. Consequently, the normal distribution might provide a suboptimal approximation for our reserving application.

Alternatively, one can construct a confidence interval by applying a log transformation to the liability. This approach utilizes the delta method to construct an interval for the logarithm of the liability, which tends to exhibit behavior closer to normality. Subsequently, the interval is transformed back to the original scale using the reverse transformation. This log-transformed confidence interval can be a more appropriate choice, considering the distribution of the data and its potential skewness and heavy-tailed characteristics.

Therefore, an approximate $1 - \alpha$ confidence interval for $L^{IBNS}(\tau)$ can be constructed as

$$\left(\exp\left(\log(\hat{L}^{IBNS}(\tau)) - Z_{\alpha/2}\frac{\sqrt{Var(\hat{L}^{IBNS}(\tau))}}{\hat{L}^{IBNS}(\tau)}\right), \exp\left(\log(\hat{L}^{IBNS}(\tau)) + Z_{\alpha/2}\frac{\sqrt{Var(\hat{L}^{IBNS}(\tau))}}{\hat{L}^{IBNS}(\tau)}\right)\right),$$
(5)

where $Z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal distribution. As noted by (Thompson, 2012, p. 71), the computation of the variance of the HT estimator can be quite laborious. To address this challenge, Berger (1998) propose the use of a simpler estimator given by:

$$\hat{Var}(\hat{L}^{IBNS}(\tau)) = \frac{\sum_{i=1}^{N^{P}(\tau)} \left(N^{P}(\tau) \frac{1 - \pi_{i}(\tau)}{\pi_{i}(\tau)} Y_{i} - \hat{L}^{IBNS}(\tau) \right)^{2}}{N^{P}(\tau)(N^{P}(\tau) - 1)}$$
(6)

This formulation is computationally simple to obtain and tends to provide a conservative estimate of the variance (Thompson, 2012, p. 70), which is desirable for the claim reserving problem. Once again, it is crucial to note that the previous procedure for the construction of confidence intervals is conditional upon the estimated inclusion probabilities and, therefore, does not account for parameter uncertainty associated with such estimation. Numerous studies have addressed the distribution of the HT estimation while considering the estimation of the inclusion probabilities, such as Hirano et al. (2003), Tan (2006), Fattorini (2009), and references therein. In brief, under certain regularity conditions, the HT estimator remains consistent and approximately normally distributed, with variance depending on the actual process used for estimation. Nevertheless, in general practice, it is not uncommon either to neglect the variance stemming from the uncertainty of the parameters for constructing prediction intervals or to adopt more practical approaches. For instance, one can construct confidence intervals employing the bootstrap procedure, as outlined by (Arnab, 2017, p. 636). This method allows for the incorporation of parameter uncertainty without introducing excessive complexity, and simultaneously adopting a data-driven approach without relying on the normal approximation.

2.5 Adjustments to the IPW estimate

end

A key issue that makes the claim reserving estimation process to be a very challenging one, is that the inclusion probabilities can vary significantly impacting the stability of the IPW estimator. An extreme case is when the inclusion probability of a claim is close to 0, which mostly occurs when a claim is recently reported, which represents the majority of the claims included in the IBNR reserve. Such circumstances can lead to instability in the estimator, potentially resulting in abnormally high values of the reserve when compared with the experience of previous reserving exercises. This behavior has been widely documented in the population sampling literature of the HT estimator in the context of heterogeneous inclusion probabilities. See for e.g., Hulliger (1995), Chen et al. (2017), Ma and Wang (2020) and references therein.

A standard approach to address the instabilities of the HT is by trimming the inclusion probabilities. In this approach, if the inclusion probability is too small, it is replaced with a larger value to get rid of the instability (see for e.g. Chen et al. (2017), Ma and Wang (2020)). As a result, the adjusted inclusion probabilities lead to an adjusted IPW estimator that is more robust towards these scenarios. Data-driven methods, such as Ma and Wang (2020) or Algorithm 1, as proposed by citezong2018improved and illustrated below, offer a systematic approach for this task. In the latter case, it can be shown that the mean square error of the adjusted IPW estimator is less or equal to its not adjusted counterpart for several sampling mechanisms, including the Poisson sampling.

 Algorithm 1 Trimming of inclusion probabilities for the IPW estimator

 Obtain the ordered inclusion probabilities $\{\pi_{(1)}, \pi_{(2)}, \dots, \pi_{(N^P(\tau))}\}$ from smallest to largest.

 for $j = 1, \dots, N^P(\tau)$ do

 if $\pi_{(j)} \leq \frac{1}{j+1}$ and $\pi_{(j+1)} > \frac{1}{j+2}$ then

 Modify inclusion probabilities as:

 $\{\underline{\pi_{(j)}, \dots, \pi_{(j)}}, \pi_{(j)}, \pi_{(j+1)}, \dots, \pi_{(N^P(\tau))}\}$

 end

While these adjustments are theoretically sound, increasing the inclusion probabilities may introduce a downward bias in the estimation of the reserve. Therefore, it is advisable to make adjustments only if the estimation demonstrates sensitivity to changes in the inclusion probabilities. Lastly, we note that there is a vast literature on improving the raw estimation of IPW estimators to make them more robust, as seen in, for example, Ma and Wang (2020). These adjustments are beyond the scope of this paper, yet it is important to acknowledge their existence and potential application in reserving. Nevertheless, it is worth emphasizing that the decision on whether or not to implement these adjustments should be left to the actuary's discretion, based on their expertise.

3 Calculation of RBNS, IBNR and incremental claims reserves

The reserve for outstanding claims, as discussed earlier, accounts for both unreported and partially paid claims. While Equation (3) provides an estimator for the total reserve, it does not specify the allocation of the reserve to different types of payments. However, for accounting purposes, cash management, and risk assessment, actuaries need to specify the components of the overall reserve, commonly known as the IBNR (Incurred But Not Reported) reserve, the RBNS (Reported But Not Settled) reserve, and incremental payments over specific time periods.

In this section, we present how the population sampling framework can be adapted to decompose the estimation of the total reserve given by Equation (3) into these sub-components as per the actuary's requirements. To do so, we introduce the "change of population principle" as a general approach to accomplish these, and possibly other decompositions within the IPW framework, and then demonstrate its application in deriving the aforementioned reserves.

3.1 The change of population principle

In Section 2, we used the fact that the currently paid amount can be considered as a sample of the total amount of payments. As a result, we defined a sampling design within the total amount of payments along with its corresponding inclusion probabilities. However, it is important to recognize a simple yet crucial fact: the currently paid amount can also be regarded as a sample from various sub-populations within the total amount of payments.

Figure 1 demonstrates a method of partitioning the total liability at a given valuation time τ into sub-populations associated with specific reserves of interest. This figure provides a visual representation, akin to a run-off triangle, distinguishing reported and non-reported payments at τ . The x-axis represents the development time, which goes from t = T (i.e., the accident time) up to time $t = T + \omega$, being ω the maximum settlement time of a claim. This figure can be thought of as a screenshot of the classification of all the payments at a given valuation date τ .

To illustrate the concept, let's examine Figure 1. Combining all regions (A to G) yields the population discussed in Section 2, representing total liabilities. Focusing on the lower half of the figure, regions A to D, encompass payments associated with reported claims at τ . Additionally, by narrowing our focus to the lower half of the figure and considering payments up to a specific time, such as $t = t_2$ (the union of regions A, B, and C), we obtain a truncated version of payments. Specifically, this includes total payments made for currently reported claims, excluding those made after $t = t_2$. Note that the current paid amount is a sample from all of this subpopulation.

Adopting this approach, estimating the total liability for a specific subpopulation involves treating it as the main population from which current payments are sampled. Consequently, the IPW estimator, under a different sampling design, can be employed to estimate the liability. We refer to this approach as the "change of population principle". The different sample design under the change of population principle leads to different inclusion probabilities. Nevertheless, these probabilities can be easily determined using elementary conditional probability arguments. Specifically, when we limit the analysis to a subpopulation S, we denote the inclusion probability under this restriction as $\pi_i^S(\tau)$. This probability represents the likelihood of payment Y_j being reported at τ ,

Total liability for claims with accident time before au



Figure 1: Decomposition of the total liability as of valuation time τ into sub-populations

given its membership in subpopulation S. Note that these probabilities differ in meaning and in value from those defined in Section 2. Bayes' rule allows expressing this probability as

$$\pi_i^{\mathcal{S}}(\tau) = \frac{\pi_i(\tau)}{P_i(\mathcal{S})},$$

where $P_i(S)$ denotes the probability of payment j being sampled in subpopulation S according to the original sampling design, which is influenced by factors like reporting delay and claim evolution. It is important to note that in this context, we assume the subpopulation S is a subpopulation encompassing the current payments (region A in Figure 1).

We will observe that for the reserves of interest, these probabilities can be straightforwardly expressed in terms of the probabilities associated with the previously defined delay time random variables U_i and V_i . Therefore, no additional estimations are necessary.

3.2 Calculation of the RBNS reserve

The Reported But Not Settled (RBNS) reserve represents payments that are yet to be made for claims already reported at valuation time τ . This reserve corresponds to the combined regions B, C, and D in Figure 1.

To estimate the reserve, we apply the change of population principle and define the population S as the total payments associated with reported claims at τ . This population corresponds to the lower half of Figure 1, specifically $S = A \cup B \cup C \cup D$. In reserving terminology, this corresponds to the ultimate incurred losses for claims reported prior to τ .

Next, we determine the inclusion probabilities. The selection of the subpopulation depends on claim reporting, which occurs with probability $P_i(\mathcal{S}) = P(U_i \leq \tau - T_i | X_i)$. Utilizing the previously mentioned result derived from Bayes' rule, the new inclusion probabilities for this population are

$$\frac{\pi_i(\tau)}{P_i(\mathcal{S})} = \frac{P(U_i \le \tau - T_i | \mathcal{F}_i) \times P(V_i \le \tau - R_i | \mathcal{F}_i)}{P(U_i \le \tau - T_i | \mathcal{F}_i)} = P(V_i \le \tau - R_i | \mathcal{F}_i) = \pi_i^V(\tau).$$

This result is intuitive since the region only considers claims already reported at τ , and the remaining randomness pertains to the evolution of payment occurrences only. Consequently, we

can utilize the IPW estimator to obtain an estimator for the total payments of reported claims as

$$\sum_{i=1}^{N^P(\tau)} \frac{Y_i}{\pi_i^V(\tau)}.$$

Hence, the RBNS reserve of interest can be obtained by subtracting this quantity from the current paid amount:

$$\hat{L}^{RBNS}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \frac{Y_{i}}{\pi_{i}^{V}(\tau)} - L^{P}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \frac{1 - \pi_{i}^{V}(\tau)}{\pi_{i}^{V}(\tau)} Y_{i}.$$
(7)

3.3 Calculation of the pure IBNR reserve

To estimate the Incurred But Not Reported (IBNR) reserve, we cannot directly apply the change of population principle since the current paid amount is not a subpopulation of the not reported claims population (Figure 1). However, we can easily overcome this by considering the current paid amount as the difference between two populations: the total payments (all regions in Figure 1) and the total payments of currently reported claims (lower half of Figure 1). Estimations for the liabilities associated with these populations have been discussed in Sections 2 and 3.2, respectively. Therefore, the IBNR liability can be estimated as the difference between these two estimations.

$$\hat{L}^{IBNR}(\tau) = \hat{L}^{IBNS}(\tau) - \hat{L}^{RBNS}(\tau) = \sum_{i=1}^{N^{P}(\tau)} \left(\frac{1 - \pi_{i}(\tau)}{\pi_{i}(\tau)} - \frac{1 - \pi_{i}^{V}(\tau)}{\pi_{i}^{V}(\tau)}\right) Y_{i} = \sum_{i=1}^{N^{P}(\tau)} \left(\frac{1 - \pi_{i}^{U}(\tau)}{\pi_{i}^{U}(\tau)}\right) \frac{Y_{i}}{\pi_{i}^{V}(\tau)}.$$
(8)

We would like to highlight that this approach to estimating the IBNR is analogous to the conventional actuarial method using run-off triangles, where the total reserve is estimated using the incurred claims triangle and subtracting the reserve obtained from the paid claims triangle. Unlike aggregate approaches that may yield negative reserve estimates, our method ensures non-negative estimations.

3.4 Calculation of cumulative and incremental payments

The estimator we have presented so far provides the ultimate amount of liabilities, but insurance companies require projections of the reserve payments over specific periods. These payments, known as incremental claims, can be estimated within our framework as follows: we utilize the change of population principle to estimate cumulative claims for different periods and then calculate the incremental claims as the difference between these cumulative claims. Notably, our model is continuous rather than discrete, allowing for the accommodation of any desired periodicity for incremental claims. We will illustrate this process for the total reserve only, however, it can be similarly applied to the RBNS.

Let's consider an insurance company assessing claims whose accidents incurred before the valuation time τ and is interested in estimating the incremental claims associated with a future period between t_1 and t_2 ($\tau < t_1 < t_2$), denoted as $L(\tau, t_1, t_2)$. Visually, $L(\tau, t_1, t_2)$ corresponds to C and F in Figure 1.

We start by considering the population of cumulative claims up to time $t_1 \ge \tau$, where only payments made up to t_1 are included i.e., $L(\tau, 0, t_1)$ (regions A, B, and E in Figure 1). Using the change of population principle, a payment belongs to this population if its payment time is before t_1 , which occurs with probability $P_i(S) = \pi_i(t_1)$. Thus, the inclusion probability is:

$$\frac{\pi_i(\tau)}{P_i(\mathcal{S})} = \frac{\pi_i(\tau)}{\pi_i(t_1)},$$

and so the IPW estimator for cumulative claims is given by

$$\hat{L}(\tau, 0, t_1) = \sum_{i=1}^{N^P(\tau)} \frac{\pi_i(t_1)}{\pi_i(\tau)} Y_i .$$
(9)

The incremental claims between t_1 and t_2 are then given by $L(\tau, t_1, t_2) = L(\tau, 0, t_2) - L(\tau, 0, t_1)$, and so an unbiased estimator for incremental claims is

$$\hat{L}(\tau, t_1, t_2) = \hat{L}(\tau, 0, t_2) - \hat{L}(\tau, 0, t_1) = \sum_{i=1}^{N^P(\tau)} \frac{\pi_i(t_2) - \pi_i(t_1)}{\pi_i(\tau)} Y_i.$$
(10)

This is a very intuitive expression: the denominator, $\pi_i(\tau)$, scales the observed claims Y_i to the total amount, while the difference in probabilities in the numerator, $\pi_i(t_2) - \pi_i(t_1)$, captures the proportion of the total observed between t_1 and t_2 .

4 The Chain-Ladder as a Horvitz-Thompson estimator

This section constitutes a key contribution of the paper, demonstrating how the classical Chain-Ladder method and certain recent extensions inherently manifest as specific instances of the IPW estimator. Consequently, the theory of population sampling offers an alternative and statistically robust framework for understanding the Chain-Ladder method.

To proceed, it is essential to recognize that the Chain-Ladder provides discrete-time estimates of cumulative claims. For the purposes of the discussion that follows, let's assume we have predefined m periods for the development of the lifetime of a claim, from reporting to settlement. We denote these periods as $T = t_0 < t_1 < t_2 < \ldots < t_m = T + \omega$, where T represents the accident year and ω is the assumed maximum settlement time for a claim. To simplify the proof, we will only focus on the claims that occur in the same accident year T meaning that we will look only at a particular row in the Chain-Ladder triangle, say the k-th row, and we will assume that the accident year is T = 0, so that calendar time and development time match, notationly. This choice is made without any loss of generality, since each row of the triangle is treated similarly in terms of the development principle, and the total reserve is obtained by adding up the estimates row-wise. Furthermore, we also assume that the $\tau = t_k$ for simplicity.

To demonstrate the desired relationship, it is important to note that the classical Chain-Ladder method implicitly assumes that the development of claims is homogeneous, with all claims sharing the same development factors during the accident years. We formally state this within our context and notation in Assumption 1.

Assumption 1. [Homogeneous development assumption] The evolution of claims with identical accident years in the portfolio constitutes a homogeneous process, indicating that individual characteristics do not influence the patterns of reporting delays and settlement times. Consequently, the inclusion probabilities are the same for all claims with the same accident time, i.e., $\pi_i(\tau) = \pi(\tau)$.

We will show now that this simplified assumption (that we do not have to impose in general) suffices to derive the Chain-Ladder estimator and several of its extensions, as empirical versions of the IPW estimators.

4.1 Derivation for the Chain-Ladder

The Chain-Ladder method offers an iterative estimation of cumulative payments up to the ultimate. In line with our notation, we let $\hat{L}_{CL}(\tau, 0, t)$ denote the cumulative claim estimates from the Chain-Ladder method, with payments associated with accidents occurring before the valuation time τ and paid before time $t, t > \tau$. The Chain-Ladder estimates take the form

$$\hat{L}_{CL}(\tau, 0, \tau) = L^{P}(\tau)$$
$$\hat{L}_{CL}(\tau, 0, t_{n}) = \hat{L}_{CL}(\tau, 0, t_{n-1}) \times f(t_{n-1}, t_{n}), \quad n = k+1, \dots, m_{p}$$

where the factors $f(t_{n-1}, t_n)$ are the so-called development factors or link ratios and $L^P(\tau)$, represents the total payments made to time τ . These factors are estimated through Equation (13), as outlined in the literature on stochastic claim reserving e.g., Mack (1999). In a similar vein, employing the recursion, the ultimate liability is estimated as

$$\hat{L}_{CL}(\tau) = \hat{L}_{CL}(\tau, 0, t_m) = L^P(\tau) \times f(t_k, t_{k+1}) \times f(t_{k+1}, t_{k+2}) \times \ldots \times f(t_{m-1}, t_m),$$

and the reserve estimation is obtained by subtracting this estimation from the current paid amount.

Now let's proceed to connect this estimation with the IPW estimator. It is noteworthy that under Assumption 1, we can represent the IPW estimator for cumulative claims at any time t_n using Equation (9) as

$$\hat{L}(\tau, 0, t_n) = \sum_{j=1}^{N^P(\tau)} \frac{\pi_j(t_n)}{\pi_j(\tau)} Y_j = \left(\sum_{j=1}^{N^P(\tau)} Y_j\right) \frac{\pi(t_n)}{\pi(\tau)} = L^P(\tau) \frac{\pi(t_n)}{\pi(\tau)},$$

and so the IPW estimators for the cumulative payments over the periods above would satisfy the recursive relationship

$$\hat{L}(\tau, 0, t_n) = \hat{L}(\tau, 0, t_{n-1}) \frac{\pi(t_n)}{\pi(t_{n-1})}, \quad n = k+1, \dots, m$$

This equation provides an estimate of the reserve, following the same recursive principle as the Chain-Ladder method. In fact, one can easily observe that the implied development factors or link ratios of the IPW estimator are given by

$$f(t_{n-1}, t_n) = \frac{\pi(t_n)}{\pi(t_{n-1})}.$$
(11)

This expression provides an explicit interpretation of the development factors as the ratio of two probabilities. More specifically, it can be interpreted as the inverse of the conditional probability of a payment occurring by time t_{n-1} , given it occurs by time t_n . Although not explicitly stated in the notation (as we assume T = 0 for simplification purposes, it is worth noting that the development factors might potentially be influenced by the accident time, given that the inclusion probabilities could exhibit such dependence.

Observe that the expression above provides a "theoretical version" of the development factors because the inclusion probabilities are unknown a priori. We therefore continue by demonstrating how the classical Chain-Ladder method naturally provides an estimator for Equation (11). To accomplish this, it is important to note that in the classical Chain-Ladder method, the development factors remain constant across rows, meaning they do not vary with accident time. Hence, we need to make the assumption outlined below regarding the progression of claims over time, enabling us to utilize the information from all claims collectively in the estimation process, regardless of their accident time. It is worth mentioning that while such an assumption can be relaxed, it would necessitate a more explicit structure outlining the evolution of claims over time.

Assumption 2. [Stationary development assumption] The evolution of claims is a stationary process through time, indicating that the accident and reporting dates do not influence the patterns of reporting delays and settlement times. Consequently, the distribution associated with the inclusion probabilities does not change over time.

That said, consider a non-parametric approach where we use the membership indicators of the payments as the data, denoted as $\mathbf{1}_i(t)$ for $t \in t_0, \ldots, t_m$. In our simplified setup, we can interpret $\mathbf{1}_i(t_n)$ as the indicator of whether a payment is made or not before the development period t_n . Note that under Assumption 2, the membership indicators are all Bernoulli random variables with a common mean $\pi(t)$ that do not change through time. Assumption 2 enables us to use the membership indicators of all claims regardless of the accident time, as they have the same distribution. Along those lines, consider a Weighted Maximum Likelihood Estimation (WMLE) for the inclusion probabilities $\pi(t)$ using the membership indicators along with pre-specified weights γ_i . It is straightforward to show that for a fixed value t and treating the payments as exchangeable, the WMLE is given by the weighted empirical proportion of payments falling before period t out of the total $\hat{\pi}(t) = \frac{\sum_{i=1}^{N(\tau)} \gamma_i \mathbf{1}_i(t)}{\sum_{i=1}^{N(\tau)} \gamma_i}$. Therefore, by the invariance property, the WMLE estimators of the development factors are given by

$$\hat{f}(t_{n-1}, t_n) = \frac{\hat{\pi}(t_n)}{\hat{\pi}(t_{n-1})} = \frac{\sum_{i=1}^{N(\tau)} \gamma_i \mathbf{1}_i(t_n)}{\sum_{i=1}^{N(\tau)} \gamma_i \mathbf{1}_i(t_{n-1})}.$$
(12)

The choice of weights influences the principles governing the estimation of development factors on run-off triangles. For example, if we define the weights as $\gamma_i = Y_i$, the estimator simplifies to

$$\hat{f}(t_{n-1}, t_n) = \frac{\sum_{i=1}^{N(\tau)} Y_i \mathbf{1}_i(t_n)}{\sum_{i=1}^{N(\tau)} Y_i \mathbf{1}_i(t_{n-1})} = \frac{\text{Cumulative payments up to time } t_n}{\text{Cumulative payments up to time } t_{n-1}},$$
(13)

which is the same estimator traditionally employed in the Chain-Ladder method. In the case of $\gamma_i = 1$, we revert to the scenario of an unweighted likelihood, and the estimator takes the form:

$$\hat{f}(t_{n-1}, t_n) = \frac{\sum_{i=1}^{N(\tau)} \mathbf{1}_i(t_n)}{\sum_{i=1}^{N(\tau)} \mathbf{1}_i(t_{n-1})} = \frac{\text{Number of payments up to time } t_n}{\text{Number of payments up to time } t_{n-1}},$$

which coincides with the development factors obtained by applying the Chain-Ladder method to the triangle of the number of claims. Similarly, if we set up $\gamma_i = Y_i^2$, it is not difficult to show that we recover a version of the least squares estimator of the linear regression through the origin approach, akin to a model described in Mack (1999). Other choices of weights can be employed to derive different development principles over run-off triangles, as discussed, for instance, in (Asmussen and Steffensen, 2020, p. 437).

These examples demonstrate that we can derive various types of triangle development models as particular cases of weighted likelihood estimation of inclusion probabilities. As such, the development factors of the Chain-Ladder (and some alternatives) aim to estimate the associated ratio of inclusion probabilities in Equation (11). It is known from the theory of weighted likelihood estimation that, under some regularity conditions on the weights, the aforementioned WMLE of the inclusion probabilities is consistent Wang et al. (2004), and so will be the associated estimators of the development factors in Equation (12). In practice, one might prefer one weighting approach over another based on the finite sample size properties of the estimator, such as a better bias-variance trade-off or the actuary's experience. Finally, it is pertinent to note that while we have justified the estimation using maximum likelihood, alternative valid estimation frameworks are available. For example, one could opt for a method of moments approach for the inclusion probabilities. However, it is important to highlight that this would result in estimators that coincide with the ones presented above.

In conclusion, we observe that the Chain-Ladder method is essentially a particular case of the IPW estimator, where the inclusion probabilities are estimated using a specific empirical weighted proportion. We emphasize that this construction of the Chain-Ladder naturally follows from the probabilistic interpretation of the reserving problem as a population sampling problem, the IPW principle, and Assumptions 1 and 2, rather than artificially constructing an estimator that replicates the Chain-Ladder estimate.

Furthermore, it is important to highlight that this derivation of the Chain-Ladder method deviates significantly from previous modeling approaches, as illustrated in e.g., England and Verrall (2002). Martínez-Miranda et al. (2012), Taylor and McGuire (2016), Engler and Lindskog (2024) and subsequent literature. Specifically, we emphasize that the sampling framework offers a distribution-free construction for the Chain-Ladder method, where the assumptions pertain to the delay of payments rather than the claim amounts themselves. It is crucial to emphasize that our intention is not to discredit previous research on stochastic claim reserving, but rather to illustrate that the favorable properties of the Chain-Ladder method extend beyond the necessity for a stochastic model based solely on frequency and severity.

4.2 Derivations for some extensions of the Chain-Ladder

Over the past decade, there has been significant research aimed at comprehending the nature of the development factors in the Chain-Ladder model, with the goal of creating generalizations. In this section, we demonstrate how several of these extensions naturally arise in an interpretable fashion from the understanding of IPW estimators. This highlights the versatility and convenience of the population sampling framework for claim reserving.

Double Chain-Ladder

Let's now concentrate on the associated incremental payments for a given time window (t_1, t_2) , specifically for claims on a given accident date. Under the homogeneity Assumption 1, we can express the IPW estimator for incremental claims in Equation (10) as:

$$\hat{L}(\tau, t_1, t_2) = \sum_{i=1}^{N^P(\tau)} \frac{\pi(t_2) - \pi(t_1)}{\pi^U(\tau) \pi^V(\tau)} Y_i = \left(\frac{N^P(\tau)}{\pi^U(\tau)}\right) (\pi(t_2) - \pi(t_1)) \left(\frac{1}{N^P(\tau)} \sum_{i=1}^{N^P(\tau)} \frac{Y_i}{\pi^V(\tau)}\right)$$
$$= \hat{N}(\tau) p(t_1, t_2) \bar{Y}(\tau)$$

where $\hat{N}(\tau) = N^P(\tau)/\pi^U(\tau)$ is the IPW estimator of the total number of claims whose accident year is before valuation time τ , $p(t_1, t_2) = \pi(t_2) - \pi(t_1)$ is the expected proportion of claims that will be paid in the interval (t_1, t_2) , and $\bar{Y}(\tau) = \frac{1}{N^P(\tau)} \sum_{i=1}^{N^P(\tau)} \frac{Y_i}{\pi^V(\tau)}$ is the estimated average claim amount per claim for claims reported by valuation time τ .

This estimator is, in spirit, similar to the so-called double Chain-Ladder estimator described by Verrall et al. (2010) and Martínez-Miranda et al. (2012). The IPW estimator $N^P(\tau)/\pi^U(\tau)$ represents the (first) Chain-Ladder estimation of the total number of claims when applied to the run-off triangle of the number incremental claims, and $\sum_{i=1}^{N^P(\tau)} \frac{Y_i}{\pi^V(\tau)}$ is the (second) Chain-Ladder estimation of the total claim amount of reported claims when applied to such run-off triangle, see Section 3.2. Therefore, the IPW implicitly leverages information from both the triangles of paid claims and the number of claims, a modeling approach shown to be desirable as described in Verrall et al. (2010). Note that the decomposition of the reserve into IBNR and RBNS, which is naturally achieved in the IPW framework, is analogous to the one obtained from the double Chain-Ladder.

It is worth noting that this derivation of the double Chain-Ladder principle differs from other approaches, such as the collective reserving model in Wahl et al. (2019) and subsequent extensions. However, the IPW does so without imposing assumptions on frequency or severity, in contrast to the collective reserving model. Indeed, analogously to the traditional Chain-ladder method, the IPW framework also provides reasoning that leads to a double Chain-Ladder type estimation in a distribution-free fashion. Thus, the IPW and the collective reserving frameworks are not particular cases of one another but rather complementary from a theoretical standpoint. Lastly, it is important to note that the point estimation between these methods may differ depending on the actual estimator selected for the different components.

Continuous Chain-Ladder

Martínez-Miranda et al. (2013) and Bischofberger et al. (2020) introduced what they termed the continuous Chain-Ladder, utilizing data from observations in a continuous time instead of discrete time as given by cell partition of the triangle. They argue that the Chain-Ladder method can be viewed as a structured density approach on the run-off triangle of claims, denoted as $\tilde{f}_{Z,T}(z,t)$, and derived an estimate for outstanding claims of the form

$$\frac{\int_0^\omega \int_{\omega-t}^\omega \widetilde{f}_{Z,T}(z,t) \mathrm{d}z \mathrm{d}t}{\int_0^\omega \int_0^{\omega-t} \widetilde{f}_{Z,T}(z,t) \mathrm{d}z \mathrm{d}t} \sum_{i=1}^{N^P(\tau)} Y_i.$$

We can identify such an estimator as the IPW estimator of the reserve in Equation (3), under the homogeneity Assumption 1, and with a specific inclusion probability. Their approach implicitly derives an estimator of the average inclusion probability by estimating the joint density function of the (Z,T) variables, which is then integrated out on the upper half of the triangle of reported claims, i.e., $\pi = \int_0^{\omega} \pi(t) \tilde{f}_T(t) dt = \int_0^{\omega} \int_0^{\omega-t} \tilde{f}_{Z,T}(z,t) dz dt$. This estimation is then incorporated into an estimator of the form $\frac{1-\pi}{\pi} \sum_{i=1}^{N^P(\tau)} Y_i$.

The authors acknowledge that they do not provide any estimation theory justifying the structure of the estimator above. Nevertheless, note that the framework of IPW estimators might provide such a theoretical foundation. The details, unfortunately, go beyond the scope of this paper.

Chain-Ladder via reverse hazard rates

A recent contribution by Hiabu (2017) demonstrated that the estimated development factors of the Chain-Ladder method have a one-to-one correspondence with the estimator of the hazard rate of a counting process defined in reversed time, as shown in the equation below. This correspondence makes it useful to define more general estimators based on the Chain-Ladder principle.

Such correspondence can be easily derived from the IPW framework. In our notation, Z is the random variable associated with the entire development time of the claim from its accident, and we can define the reversed time random variable as $\tilde{Z} = \omega - Z$. Consequently, the development

factors can be expressed as:

$$f(t_{n-1}, t_n) = \frac{P(Z \le t_n)}{P(Z \le t_{n-1})} = \frac{P(\tilde{Z} \ge \omega - t_n)}{P(\tilde{Z} \ge \omega - t_{n-1})} = \left(1 - \frac{P(\tilde{Z} \ge \omega - t_n) - P(\tilde{Z} \ge \omega - t_{n-1})}{P(\tilde{Z} \ge \omega - t_n)}\right)^{-1}$$
$$= (1 - \delta_n \alpha_{t_n})^{-1}$$

where $\delta_n = t_n - t_{n-1}$ and $\alpha_{t_n} = \frac{P(\tilde{Z} \ge \omega - t_n) - P(\tilde{Z} \ge \omega - t_{n-1})}{\delta_n P(\tilde{Z} \ge \omega - t_n)}$ is the hazard rate of the process in reverse time \tilde{Z} over the interval $(\omega - t_n, \omega - t_{n-1})$.

This expression is employed for modelling development factors through a kernel estimator of the hazard rate. Further research also proposes more general estimation approaches for the reversed hazard rate using multiplicative models, as demonstrated in Hiabu et al. (2021). It is noteworthy that such approaches can be viewed as applications of survival models for the estimation of development factors, and, in a very implicit manner, models for the inclusion probabilities.

Chain-Ladder via demographic methods

Recent preprints, such as Pittarello et al. (2023) and Hiabu et al. (2024), have expanded on the idea of reverse hazard rates and established a correspondence between development factors and central death rates. Such a correspondence can also be derived from the IPW framework perspective. Using the analogy of the reversed time counting process of claims as a pure death process, we have:

$$f(t_{n-1}, t_n) = \frac{P(\tilde{Z} \ge \omega - t_n)}{P(\tilde{Z} \ge \omega - t_{n-1})} = \left(P(\tilde{Z} \ge \omega - t_{n-1} | \tilde{Z} \ge \omega - t_n)\right)^{-1} = \left(1 - \frac{m_{t_n}}{1 + m_{t_n}(1 - \eta_{t_n})}\right)^{-1}$$
$$= \frac{1 + m_{t_n}(1 - \eta_{t_n})}{1 - \eta_{t_n}m_{t_n}}$$

where m_{t_n} represents the reversed central death rate associated with the process in reverse time over the interval $(\omega - t_n, \omega - t_{n-1})$, and η_{t_n} is the so-called separation factor, which denotes the average fraction of the interval $(\omega - t_n, \omega - t_{n-1})$ lived through by those who died in the interval, akin to an exposure term. The second-to-last equality is derived from well-known results from life tables relating survival probabilities and central death rates. See for e.g., (London, 1997, p. 53-55).

In this manner, mortality models inspired by demography literature can be applied to describe such central death rates and implicitly the development factors of the Chain-Ladder method. See the two references mentioned earlier for illustration. It is crucial to note that such mortality models are specific instances of a survival model and are closely related to the modeling of inclusion probabilities in this paper.

We conclude this section by re-emphasizing that the previous cases represent particular simplifications obtained under the sampling theory framework. In fact, the IPW estimator in its full generality will allow us to further generalize the classical Chain-Ladder method by allowing the inclusion probabilities to be dependent on each policyholder's risk characteristics, as well as other sources of heterogeneity.

5 Estimation of the inclusion probabilities and model validation

The key element required to implement the IPW estimator is the unknown inclusion probabilities $\pi_i(\tau)$. These probabilities are associated with the evolution of a claim, i.e., reporting and settlement delays, and depend on various attributes of the nature of the claim and the policyholder, denoted

as X_i , as well as the payment amount Y_i itself. In this section, we outline briefly the methods and considerations to take into account for estimation.

As explained in Section 2, the inclusion probabilities consist of two separate components: the probability of reporting $\pi_i^U(\tau)$, and the probability of payment $\pi_i^V(\tau)$. Each one is estimated separately, so we discuss different strategies in Sections 5.1 and 5.2.

For the estimation, it is crucial to consider the possibility of having more than one payment per claim. For the sake of notation, we will index the different claims themselves using the index j and denote with \mathcal{I}_j the set of indexes of the payments Y_i associated with the j-th claim. In other words, the sets \mathcal{I}_j create a partition of the entire set of indexes $i = 1, \ldots, N(\tau)$. Additionally, following common practice in reserving applications, we will use ω to denote the maximum settlement time of a claim. This value is defined based on the nature of the contract, such as the final date to file a claim for an accident after its occurrence, or according to the policies of the insurance company.

5.1 Estimation of the reporting delay times probabilities $\pi_i^U(\tau)$

As mentioned in Section 2.1, a value U_i may be repeated several times depending on the number of payments per claim. Since reporting occurs only once, we consider only one realization of this variable per claim. Hence, we will use the notation U_j to refer to these unique values, forming the working dataset.

To estimate the inclusion probabilities, we require the cumulative distribution function of the reporting delay times. It is common in reserving literature (see, for example, Verrall and Wüthrich (2016)) to assume that the reporting delay times U_j , conditioned on claim attributes X_j , follow a common distribution function $F_{U|X_j}(u)$. This distribution function is the target of estimation in this section, using some regression framework to incorporate dependence on covariates.

To achieve this, we note that the reporting delay time is a time-to-event random variable commonly studied in survival modeling. Therefore, approaches already implemented in statistical software can be utilized to estimate the overall distribution function and the desired probabilities. For instance, hazard-based models, particularly the Cox regression model, are by far the most popular techniques used in survival analysis for this task (e.g., George et al. (2014)), and are widely implemented in packages, e.g. Bender et al. (2018). Nevertheless, note that there are alternative approaches that can offer different and flexible structures, as inspired by the machine learning literature. For instance, Fung et al. (2022) propose a flexible model based on mixture of experts, with implementations available for insurance data Tseung et al. (2021); Wiegrebe et al. (2024) consider non-linear regression on covariates via deep learning approaches; and Sonabend et al. (2021) propose approaches based on survival trees, both implementable via R packages as in Sonabend et al. (2021). The choice of the model must be achieved in a data-driven fashion aiming for the best fit to the data.

Consequently, the desired inclusion probabilities for all the payments $i \in I_j$ associated with the *j*-th claim are derived using the relationship:

$$\pi_j^U(\tau) = Pr(U_j \le \tau - T_i | X_j) = F_{U|X_j}(\tau - T_j) =$$
(14)

A crucial aspect in the estimation of the model above is accounting for the right truncation of the data. Indeed, due to the delay in the reporting times, our observations are limited to the conditional random variables: $U_j|U_j \leq \tau - T_j$, and ignoring this fact would result in a downward bias in the overall distribution. Fortunately, the literature on survival analysis has widely explored this issue and provided solutions that the user can adopt for the estimation of the model, such as back-censoring strategies or weighted likelihood methods. See for e.g., Verbelen et al. (2015), Fung et al. (2022)), and Vakulenko-Lagun et al. (2020) for the specific case of Cox regression models. Moreover, as motivated in Section 4, one may also include a weighting strategy that accounts for the claim sizes on the likelihood to better resemble the Chain-Ladder estimation.

5.2 Estimation of the payment times probabilities $\pi_i^V(\tau)$

Estimating this probability differs significantly from the previous case due to the potential for multiple payment events per claim, unlike the one-time event of claim reporting. The payment process is a *recurrent event process*, requiring an adapted modeling approach. See for e.g., Cook et al. (2007) and Andersen and Ravn (2023).

Recurrent events relate closely to counting processes, where the former focuses on event times, and the latter on the number of events. The modeling approach required for this paper involves selecting a suitable counting process (incorporating attributes) to data-driven model the number of payments per claim. The fitted model is then used to determine the desired inclusion probabilities.

That being said, assume the arrivals of payments for the *j*-th claim, conditioned on claim and policyholder attributes, are governed by a counting process $\{M_j(v); v \in (0, \omega)\}$, with times $V_i \in \mathcal{I}_j$ as the associated occurrence times. For convenience, we use the convention of v = 0 as the actual time of reporting the claim, which initiates the payment counting process.

To compute the inclusion probabilities, we first express the distribution of arrival times within the interval $(0, \omega)$ in terms of the process $M_j(v)$. To do this consider the reversed-time version of the process, denoted as $\tilde{M}_j(v) = M_j(\omega - v)$, in which the reversed time is defined as $\tilde{v} = \omega - v$ (e.g. Klein and Moeschberger (2003)). This reversed-time process can be seen as a pure death process, where $\tilde{M}_j(0) = M_j(\omega)$ plays the role of the initial number of lives and $\tilde{M}_j(v)$ the number of lives at time v. See for e.g., Hiabu (2017). The reversed times $\tilde{V}_i = \omega - V_i$ correspond to the event times of such process. These times can be seen as realizations from a lifetime random variable, denoted as \tilde{V} , following a specific distribution based on the original counting process's evolution. This random variable defines the distribution of the location of the occurrences within the interval $[0, \omega)$, and therefore, the distribution of the variable $V := \omega - \tilde{V}$, denoted as $F_{V|X_j,U_j}(v)$, determines the distribution of occurrence times of the original process $M_j(t)$. The distribution of this random variable determines the inclusion probabilities.

Using the analogy with the reversed time counting process, we can easily compute the distribution of the time occurrences as:

$$F_{V|\mathcal{F}_{j}}(v) = P(V \le v|\mathcal{F}_{j}) = P(\tilde{V} \ge \omega - v|\mathcal{F}_{j}) = E\left(P(\tilde{V} \ge \omega - v|\tilde{M}(0), \mathcal{F}_{j})\right)$$
$$= E\left(\frac{E(\tilde{M}_{j}(\omega - v)|\tilde{M}_{j}(0), \mathcal{F}_{j})}{\tilde{M}_{j}(0)}\right) = E\left(\frac{\tilde{M}_{j}(\omega - v)}{\tilde{M}_{j}(0)} \mid \mathcal{F}_{j}\right) = E\left(\frac{M_{j}(v)}{M_{j}(\omega)} \mid \mathcal{F}_{j}\right)$$
(15)

where the fourth equality holds by the life-table relationship $\omega_{-v}p_0 = l_{\omega-v}/l_0$ (see for e.g., relationship (3.51) in (London, 1997, p. 42)), where $l_{\omega-v} = E(\tilde{M}_j(\omega-v)|\tilde{M}_j(0), \mathcal{F}_j)$ and $l_0 = \tilde{M}_j(0)$; and the second last equality is the tower property of conditional expectations.

Equation 15 provides an intuitive expression for the desired probability linked to the claim's evolution up to settlement. Essentially, the right-hand side of Equation (15) represents the expected proportion of the number of payments made by time v out of the total number of payments. Another interpretation is as the inverse of a development factor for the number of payments from time t to the ultimate value at time ω . Note that this expression can be analytically computed only for certain processes, such as the Poisson process and some extensions, as illustrated in Example 1 below.

Example 1. Suppose M(t) is a non-homogeneous Poisson process with intensity $\lambda(v)$, then

$$F_{V|\mathcal{F}_j}(v) = E\left(\frac{M_j(v)}{M_j(\omega)} \mid \mathcal{F}_j\right) = E\left(\frac{E(M_j(v)|M_j(\omega), \mathcal{F}_j)}{M_j(\omega)}\right) = E\left(\frac{M_j(\omega)\frac{\int_0^v \lambda_{M|\mathcal{F}_j}(s)ds}{\int_0^\omega \lambda_{M|\mathcal{F}_j}(s)ds}}{M_j(\omega)}\right) = \frac{\int_0^v \lambda_{M|\mathcal{F}_j}(s)ds}{\int_0^\omega \lambda_{M|\mathcal{F}_j}(s)ds}$$

where we use that $M_j(v)|M_j(\omega) \sim Binom\left(n = M_j(\omega), p = \frac{\int_0^v \lambda_{M|X_i,U_i}(s)ds}{\int_0^\omega \lambda_{M|X_i,U_i}(s)ds}\right)$ in the third equality.

As a result, the desired inclusion probabilities can be derived using Equation (15). In this paper, we shall focus on the case of a non-homogeneous Poisson process, and so the inclusion probabilities will be computed using the expression in Example 1.

6 Numerical study with real data

In this section, we showcase the application of the IPW estimator using a real data set obtained from a European automobile insurance company. The data set comprises information on Body Injury (BI) claims from January 2009 to December 2012.

In line with our methodology's primary objective of serving as an alternative to traditional macro-level models, while being simpler than fitting a micro-level model, we maintain a simplified approach to emphasize the practicality of the method in real-world applications.

6.1 Description of the data

The dataset contains detailed information about claim settlements, policyholder attributes, and automobile characteristics within the aforementioned time period. This information encompasses factors such as car weight, engine displacement, engine power, fuel type (gasoline or diesel), car age, policyholder age, and region (a total of five regions). Furthermore, the dataset includes details related to the accidents themselves, such as the time of occurrence and the type of accident (type 1 and type 2). Additionally, information regarding the progression of claim payments is available, including reporting time, settlement amounts, and the corresponding occurrence times.

Our statistical modeling study focuses on the evolution of claims, specifically related to reporting delay time and payment times. We illustrate some characteristics of these quantities, such as the distributions in Figure 2 and summarize key statistics in Table 1.

Variable	Year	Mean	Std. Dev.	Min.	1st Qu.	Median	3rd Qu.	Max.
	2009	1.03	4.09	0.00	0.05	0.10	0.28	43.28
Reporting	2010	0.45	1.85	0.00	0.04	0.07	0.13	25.66
delay time	2011	0.50	2.16	0.00	0.04	0.07	0.17	24.09
	2012	0.57	2.34	0.00	0.04	0.08	0.18	32.69
	2009	3.32	7.04	0.10	0.93	1.43	2.52	48.05
Payment	2010	2.31	4.37	0.04	0.81	1.31	1.92	35.37
delay time	2011	2.82	5.41	0.03	0.91	1.41	2.14	35.38
	2012	2.49	4.95	0.04	0.98	1.44	2.14	40.23

Table 1: Summary statistics per year of the reporting delay time and the payment delay time (both in months)

Upon reviewing the information presented in Figure 2 and Table 1, it is evident that the reporting delay tends to be relatively short, with an average duration of less than a month. However,



Figure 2: Density functions by year of the reporting delay time (left panel) and the payment time (right panel) in months. Plots are in the log scale.

there is notable variability in the tail behavior of this variable. In contrast, the progression of claim payments typically spans a few months on average, but there are instances where settlement times can extend over several years. It is important to note that the distributions of these variables exhibit complexities that are challenging to capture using simple parametric models. Specifically, they display significant temporal fluctuations, indicating that historical data may not adequately represent future events. To account for this, we define the maximum development time for future analysis as $\omega = 24$ months, with approximately 99.95% of claims being settled within this timeframe.

6.2 Model fitting

To model the reporting delay time, we employ a Cox regression model. Likewise, we adopt a nonhomogeneous Poisson process for the recurrent payments process. As the maximum development time is 24 months, then the time window from 2009 to 2010 can be only used for training purposes, while the time window from 2011 to 2012 will be used for testing.

For the sake of illustration, we calculate reserves on a monthly basis i.e., several valuation dates each one month apart from the other. To ensure that the estimation captures the most recent evolution of claims as much as possible, we re-fit the models employing a rolling window approach, where only the last two years of data preceding a valuation date are used for the fitting. This approach aligns with the practices employed by industry professionals in their daily work. We do not consider a time-series model via correlated frailties due to the short time window of the training sets i.e., only 2 years. Although there may be some variations in model parameters across different dates, the overall fit behaves similarly across time. We proceed to present the fitting processes for only one of the valuation dates (September 2012) in the testing period.

Reporting delay time

For the reporting delay time, U, we fit a Cox regression model to describe the hazard function $\lambda_{U|\mathcal{F}}(u)$ using the attributes of the policyholder (contained in \mathcal{F}) as covariates as:

 $\log(\lambda_{U|\mathcal{F}}(u)) = \log(\lambda_0(u)) + \beta_1 \text{Car-Weight} + \beta_2 \text{Engine-Power} + \beta_3 \text{Fuel-Type} + \beta_4 \text{Age} + \beta_5 \text{Car-Age} + \beta_6 \text{Accident-type} + \beta_7 \text{Claim-Amount} + S_8(\text{Accident-day}) + \beta_9 \text{Region}$

The log of the baseline hazard function denoted as $\log(\lambda_0(u))$, is estimated using a B-Spline representation. Additionally, we incorporate non-linear effects for the covariate "Accident-day" with the term S_8 (Accident-day), which is also estimated using a B-Spline representation. Note that the non-linear effect associated with calendar time provides the model with a dynamic-like structure, allowing it to account for some temporal changes. Lastly, for the sake of interpretability, we work with the standardized version of the continuous covariates.

To estimate the parameters while considering the right truncation of the data, we utilize a generalized additive model implementation via the piece-wise exponential modeling approach, as described by Bender et al. (2018). This approach is available in R packages such as flexsurvreg, pammtools or GJRM. Moreover, we use the claim sizes (or case estimates if unavailable) as weights for each observation. The results of the estimation are presented in Table 2 and Figure 3.

Coefficient	β_1	β_2	β_3	β_4	β_5	β_6	β_7
Value	0.022	-0.001	0.082	0.074	-0.071	-0.057	-0.069
Std. Err.	0.0008	0.0007	0.0012	0.0005	0.0005	0.0017	0.0002
p-val	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001

Table 2: Estimated coefficients of the Cox regression model for the reporting delay time

Table 2 shows that all the policyholder attributes included in the model are statistically relevant to describe the behaviour of the reporting delay time. The same conclusion applies to the categorical variable "Region", although the detailed results are not presented in the table due to its numerous categories. The left panel of Figure 3 displays the baseline hazard function, indicating a large hazard rate during the initial months, suggesting a concentration of reporting delays within this period. The hazard rate then decreases rapidly but remains nonzero for large delay times, indicating a heavy tail. Furthermore, the right panel of Figure 3 illustrates the non-linear effect associated with the accident day within a year. Time 0 is the beginning of the year i.e., January 1 and time 1 indicates the end of the year i.e., December 31. We can appreciate a seasonal pattern. Specifically, the hazard rate for the reporting delay time is higher at the end and beginning of the year when compared to the hazard rate at the middle of the year.

To assess the adequacy of the model fit, we employ the so-called normal pseudo-residuals. These are constructed using the probability integral transform Rüschendorf (2009), along with the quantile function of the normal distribution. These quantities should follow a normal distribution if the fitted model resembles properly the distribution of the data. In the case of truncation, these pseudo-residuals are computed as:

$$r_i^U = \Phi^{-1} \left(\frac{\hat{F}_{U|\mathcal{F}_i}(U_i)}{\hat{F}_{U|\mathcal{F}_i}(\tau - T_i)} \right)$$

where $\hat{F}_{U|\mathcal{F}_i}$ is the estimated version of $F_{U|\mathcal{F}_i}$, and Φ^{-1} is the quantile function of the standard normal distribution.

Along those lines, the left-hand side of Figure 6 presents a QQ plot comparing these normal pseudo residuals against the theoretical normal distribution. From the plot, it is evident that the



Figure 3: Fitted baseline hazard function for the reporting delay time (left panel) and fitted effect of the accident day on the hazard of the reporting delay time (right panel)

normal pseudo residuals exhibit no significant deviations from their expected theoretical counterparts. Hence, there is no evidence suggesting a significant lack of fit in the fitted distribution function.

Lastly, we use the model to obtain the inclusion probabilities, which can be computed by plugging the estimated hazard function in the expression below:

$$\pi_i^U(\tau) = Pr(U_i \le \tau - T_i | \mathcal{F}_i) = F_{U|\mathcal{F}_i}(\tau - T_j) = 1 - \exp\left(-\int_0^{\tau - T_j} \lambda_{U|\mathcal{F}_i}(u) du\right).$$
(16)

Payment times

Next, we present the model for claim evolution utilizing a non-homogeneous Poisson process structure that depends on covariates. In this case, we consider a maximum settlement time of $\omega = 24$ months and proceed to model the intensity function of the process via a proportional hazard model (also known as extended Cox regression or as the *Andersen-Gill* model (e.g., Andersen et al. (1985))) with intensity described as follows:

$$log(\lambda_{V|\mathcal{F}_i}(v)) = log(\lambda_0(v)) + \alpha_1 \text{Car-Weight} + \alpha_2 \text{Engine-Power} + \alpha_3 \text{Fuel-Type} + \alpha_4 \text{Age} + S_5(v) \text{Car-Age} + \alpha_6 \text{Accident-type} + \alpha_7(\text{Payment-Amount}) + S_8(\text{Reporting-day}) + \alpha_9 \text{Region} + S_{10}(\text{Reporting-delay-time})$$

Once again, we work with the standardized version of the continuous covariates for interpretability. In this instance, our model incorporates non-linear effects for the Reporting-day and Reportingdelay-time through the smooth functions S_8 and S_{10} , respectively, along with a time-varying regression coefficient for Car-Age denoted by S_5 . These non-linear effects, as well as the logarithm of the baseline intensity function, $\log(\lambda_0(v))$, are estimated using B-Spline representations. The inclusion of these non-linear terms in the regression was based on standard model selection approaches for regression modeling i.e., the significance of regression coefficients, AIC, etc. The estimation is performed with a similar implementation to the previous model. The fitted parameters of the model are detailed in Table 3, and visual representations are provided in Figures 4 and 5.

This modeling approach mirrors the methodology presented in the previous section, and thus, we refrain from delving into further details. However, we highlight three key distinctions in this regression model. Firstly, the regression structure considers multiple events per claim instead of a single event per claim. Secondly, we incorporate the reporting day as a non-linear effect aiming to capture time-related effects. Thirdly, we include the observed reporting delay time as a covariate, as the counting process is conditional on its realization.

Coefficient	α_1	α_2	α_3	α_4	α_6	α_7
Value	-0.005	0.011	-0.007	0.031	0.025	-0.028
Std. Err.	0.0008	0.0007	0.0011	0.0005	0.0017	0.0002
p-val	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001	< 0.0001

Table 3: Estimated coefficients of the extended Cox regression model for the payment counting process

The results display similarities to the previous case, as depicted in Table 3, where all policyholder attributes included in the model are statistically significant in describing the behavior of the number of payments per claim. The left panel of Figure 4 illustrates the baseline intensity function of the process. It is apparent that the intensity is initially high during the first couple of months, indicating a significant number of payments occurring within this period. Subsequently, the hazard rate decreases, approaching zero, albeit not rapidly, suggesting the occurrence of some payments after several months from reporting. We also highlight the observation that certain payments tend to occur around 6 months and 1 year, as illustrated in the graph.



Figure 4: Fitted baseline intensity function for the payment counting process time (left panel) and fitted effect of the car age on the intensity of the payment counting process(right panel)

The right panel of Figure 4 displays the time-varying coefficient associated with the Age of the Car. This effect is practically negligible in the intensity function for the occurrence of payments within approximately the first year, but it negatively influences the frequency of payments after

that period. In this scenario, the older the car, the fewer the number of payments after the first year. Similarly, claims from older cars tend to settle more quickly than those from newer cars.

This observation can be explained by factors such as an increased likelihood of total loss events, where major incidents prompt a single lump-sum payment. Additionally, the reduced economic viability of repairing older vehicles and limitations in coverage contribute to insurers opting for settlements over multiple payments. The lower perceived replacement value of older cars further encourages insurers to streamline the process with a lump-sum settlement, minimizing the overall number of payments before claim resolution.



Figure 5: Fitted effect of the reporting day on the intensity function for the payment counting process (left panel) fitted effect of the reporting delay time on the intensity of the payment counting process (right panel)

The left panel of Figure 5 illustrates the non-linear effect associated with the reporting day, suggesting a quarterly seasonal pattern. In this context, claims reported around the end and beginning of the year tend to be settled with fewer payments than those reported throughout the year. On the other hand, the right panel of Figure 5 presents the effect of the reporting delay time on the intensity of the number of payments. In brief, we observe that claims with shorter reporting delay times are slightly associated with claim settlement involving more payments than those with longer delay times.

Several reasons can explain such a phenomenon. For instance, swiftly reported claims often involve more complex or severe incidents, necessitating prolonged settlement periods for meticulous investigation and processing. In the case of body injury-related claims, an early report might trigger immediate payments for initial expenses, with subsequent adjustments accounting for evolving situations or additional treatments. Additionally, policyholders exhibiting prompt reporting may engage in proactive behavior, contributing to a higher frequency of payments as they actively seek additional coverage or claim multiple aspects of their policy.

To assess the goodness of fit, we once again employ the normal pseudo residuals (using an analogous expression as the one shown above) and validate them accordingly. The right panel of Figure 6 presents a QQ plot comparing these normal pseudo residuals against the theoretical normal distribution. Similar to the previous analysis, no significant deviation is observed between the normal pseudo residuals and their expected theoretical counterparts. While the points do not

align as closely to the 45-degree line as it was in the previous case, there is no evidence indicating a considerable lack of fit in the fitted distribution function.



Figure 6: Q-Q plots of the normal pseudo-residuals. The left panel is for the Cox model for the reporting delay time and the right panel is for the extended Cox model for the payment counting process.

Lastly, we use the model to compute the inclusion probabilities using an expression from Example 1, as the model is a type of non-homogenous Poisson process:

$$\pi_i^V(\tau) = P(V_i \le \tau - R_i | \mathcal{F}_i) = F_{V|\mathcal{F}_i}(\tau - R_i) = \frac{\int_0^{\tau - R_i} \lambda_{M|\mathcal{F}_i}(s) ds}{\int_0^\omega \lambda_{M|\mathcal{F}_i}(s) ds}$$
(17)

6.3 Results of the estimation of the reserves for a single date

Here we show the estimation of the outstanding claims (IBNS), RBNS and IBNR reserves for the same date of the testing period (September 2012). To ease the visualization and comparison, we present the estimation of the reserves in the classical run-off triangle format in Tables 5 and 6, yet however, recall that our method does not rely on a given periodicity for its calculation or the construction of a triangle.

The inclusion probabilities $\pi_i(\tau)$ and $\pi_i^V(\tau)$ for the outstanding not settled claims are estimated directly from the models from the previous section using Equations (14), (17) and (1). Figure 7 displays the histogram of such probabilities and Table 4 displays some summary statistics. Briefly, we observe that the inclusion probabilities vary drastically from one claim to another due to the heterogeneity of the claims. Note that the probabilities tend to be closer to 1 than to 0 due to the low average reporting delay and payment time, and therefore only the most recently reported claims have a small probability.

Tables 5 and 6 present cumulative run-off triangles for total outstanding claims and reported but not settled claims, respectively, as of September 2012 The incurred but not reported claims reserve estimation is derived from the difference between these triangles, which is not shown to avoid redundancy. We completed the lower half of the triangles using the true reserve values, the IPW estimator (using Equation (9)), and the Chain-Ladder method, on a monthly basis, for comparison



Figure 7: Histogram of the inclusion probabilities used for the IPW estimator of the total reserve (left panel) and the RBNS reserve (right panel)

Probability	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
$\pi(\tau)$	0.110	0.734	0.878	0.802	0.926	0.982
$\pi^V(au)$	0.239	0.806	0.901	0.858	0.943	0.984

Table 4: Summary statistics of the inclusion probabilities

purposes. To maintain readability and practicality, the table is limited to 13 months, representing approximately 98% of settled claims within this period. To differentiate between the RBNS and IBNR claims components in the Chain-Ladder method, we utilize the double Chain-Ladder method by Martínez-Miranda et al. (2012).

Analyzing the lower half of the triangles in Tables 5 and 6, we observe that the IPW estimator provides cumulative payment estimations that exhibit similar trends and magnitudes as the actual cumulative claims. No evident patterns of under or over-estimation are observed. Additionally, the IPW estimator exhibits different behavior to the Chain-Ladder method, indicating the impact of using the individual level information in the estimation.

The overall findings indicate that, at the cell level, the IPW estimator performs similarly to the Chain-Ladder method for this date. We do note that the IPW tends to better capture the changes in the reserve on the most recent dates as a result of accounting for the composition of the portfolio in the estimation. However, it is important to note that the IPW estimator does not consistently outperform the Chain-Ladder method in all cells of the triangle. We emphasize that, even though the IPW estimator can provide such estimations at the cell level, it may not possess the same precision as the estimation of the reserve as a whole. The more granular the desired estimation (i.e., the smaller the subpopulation of interest), the lower the level of accuracy.

Along those lines, instead of focusing on cell-level comparisons, our emphasis lies now on the aggregation of cells to determine the actual reserve value, which is the ultimate objective of estimation. It is noteworthy that the IPW estimator directly provides an estimation of the total reserves using Equations (3), (7) and (8), eliminating the need for constructing the run-off triangle in comparison to the Chain-Ladder method. Table 7 presents the aggregated reserve values obtained by

1 5,544 5,544 5,544 9,356 9,356 9,356 39,787 39,787 65,812	$\begin{array}{c} 65,812\\ 2 \\ 72,508\\ 67,687\\ \hline 2 \\ 464,871\\ 1 \\ 432,108\\ 0 \\ 427,539\\ 9 \\ 521,505\\ 5 \\ 500,189\\ \end{array}$
1 5,544 5,544 9,356 9,356 9,356 9,356 39,787 39,787 39,787 65,812	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
	$\begin{array}{r cccc} & 67,687 \\ \hline 2 & 464,871 \\ 1 & 432,108 \\ 0 & 427,539 \\ \hline 9 & 521,505 \\ 5 & 500,189 \\ \end{array}$
	$\begin{array}{c ccccc} 2 & 464,871 \\ 1 & 432,108 \\ 0 & 427,539 \\ \hline 9 & 521,505 \\ 5 & 500,189 \end{array}$
	$\begin{array}{c cccc} 1 & 432,108 \\ 0 & 427,539 \\ \hline 9 & 521,505 \\ 5 & 500,189 \\ \end{array}$
$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	$\begin{array}{c ccc} 0 & 427,539 \\ 9 & 521,505 \\ 6 & 500,189 \end{array}$
	$\begin{array}{c cccc} 9 & 521,505 \\ 6 & 500,189 \end{array}$
471,409 471,	6 500,189
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	
472,267 476,	9 491,963
559,824 559,824 612,	7 645,614
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	6 629,708
<u>563,068</u> 577,972 584,	0 607,539
403,159 407,692 407,692 409,	5 412,925
$ \begin{tabular}{ c c c c c c c c c c c c c c c c c c c$	1 456,839
404,157 417,295 427,383 432,	9 447,641
492,646 492,646 550,816 551,	7 554,910
$\begin{bmatrix} 6 & 7,925 & 123,729 & 285,715 & 342,347 & 446,330 & 456,970 & 471,253 & 492,646 & 496,290 & 513,169 & 527,684 & 553,505 & 563, 563, 564, 564, 564, 564, 564, 564, 564, 564$	0 591,803
499,643 518,644 534,253 547,258 553,	4 573,918
454,720 455,515 457,373 457,737 459,	0 482,794
$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	1 541,259
450,614 464,207 477,651 490,427 501,720 505,	4 521,377
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	0 382,743
8 21,450 129,410 266,031 333,922 334,447 336,823 339,792 354,857 363,281 375,890 382,666 390,574 395,	3 430,756
337,586 347,283 355,723 365,547 375,033 383,383 386	9 397,049
0 10 056 014 040 000 746 000 0572 046 750 0572 382,372 383,440 388,655 390,737 390,737 391,097 480,000 401	1 506,410
9 19,856 214,340 292,746 323,878 346,572 355,058 396,709 425,225 452,507 462,365 473,363 482,900 491,	2 549,739
3/5/254 404/109 419/255 432/30 446/150 458/600 408/61/ 4/33	<u>9 489,991</u>
10 10 500 141 007 050 017 010 071 000 175 06 100 000 500 170 000 500 500 500 500 500 500 500 500 5	7 387,255
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	J 668,998
325,099 300,201 409,408 424,871 436,130 432,003 404,319 414,899 480,	<u>/ 497,105</u>
	2 438,539
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	2 041,400 2 417.156
205,9/1 300,304 333,012 305,025 304,294 313,000 804,294 313,099 385,949 395,323 402,351 400, 90,750 90,750 90,750 90,012 90,012 90,012 90,0000 90,0000 90,0000 90,0000 90,0000 90	$\frac{3}{9}$ $\frac{411,100}{214.647}$
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	5 514,045 9 959 446
$\begin{array}{ c c c c c c c c c c c c c c c c c c c$	2 202,440 0 128 621
13,121 33,002 103,103 112,312 110,300 121,040 124,301 127,103 131,101 133,305 134, 17,210 100 670 112,971 100 029 102 544 240 702 249 660 2472 248,974 255 000 255 000 255	$\frac{3}{0}$ 130,031
$\begin{array}{c c c c c c c c c c c c c c c c c c c $	5 095,148 7 101 716
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	60 765

Table 5: Monthly cumulative run-off triangle for all outstanding claims for September 2012. In black is the actual value, in blue is the estimation using the IPW estimator, and in red is the traditional Chain-Ladder

summing the ultimate values for each accident date, along with the corresponding estimation errors. Our findings reveal that the IPW estimator yields reserve values that closely align with their true counterparts for all reserve types, exhibiting significantly lower estimation errors compared to the Chain-Ladder method.

Furthermore, to evaluate the predictive quality of these estimates from a probabilistic standpoint, Figure 8 illustrates the predictive distributions of the reserves juxtaposed with the actual observed values. These are constructed based on the log-transformation, and the approximate sampling distribution of the IPW estimators as discussed in Section 2.4 . Notably, we observe that the true values consistently fall within the central region of the distribution, closely aligning with the corresponding modes, which represent the predicted reserve values. Consequently, the IPW-based predictions exhibit consistency with the observed reality.

	12	13	ULT
			65,812
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	65,812	65,812	67,577
			65,812
		408,962	439,571
2 46,283 187,790 287,157 327,979 331,239 341,457 375,575 386,305 386,305 394,726 408,965	408,962	410,211	428,203
		409,388	424,855
	471,409	471,409	471,409
$ \begin{vmatrix} 3 & 23,813 & 218,029 & 381,885 & 413,203 & 457,516 & 464,451 & 470,453 & 471,409 & 471,4$	473,056	479,076	491,644
	472,266	476,759	488,865
559,824	559,824	612,687	645,614
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	578,585	588,820	612,891
	571,447	577,620	596,900
403,159 407,693	407,692	409,845	412,352
$\begin{bmatrix} 5 & 61,068 & 167,822 & 297,202 & 348,347 & 370,234 & 379,379 & 396,235 & 396,235 & 403,159 & 405,374 & 416,018 &$	425,558	432,532	446,039
403,879 413,700	419,768	424,508	436,897
492,646 492,646 492,646	550,816	551,367	551,516
$\begin{bmatrix} 6 & 7,925 & 123,729 & 285,715 & 342,347 & 446,330 & 456,970 & 471,253 & 492,646 & 496,290 & 513,169 & 525,413 & 525,513 & $	538,032	546,020	564,057
499,467 517,292 528,723	536,975	542,727	559,589
454,720 454,720 456,578 456,943	456,942	458,335	462,642
$\begin{bmatrix} 7 & 14,738 & 160,911 & 314,757 & 359,414 & 443,251 & 447,872 & 454,720 & 452,884 & 466,792 & 475,395 & 485,127 &$	494,178	500,544	511,826
450,548 462,084 475,002 484,279	490,159	494,052	506,155
336,823 336,970 338,189 338,189 338,189	338,189	360,450	361,875
8 21,450 129,410 266,031 333,922 334,447 336,823 339,792 354,857 363,281 369,985 376,650	384,445	388,954	397,548
337,545 346,725 353,592 362,745 369,43	374,010	376,963	384,903
	390,737	390,737	406,591
9 19,850 214,340 292,740 323,878 340,572 353,058 390,709 419,029 433,802 443,238 453,71	462,879	468,211	481,965
37(4,215) 401,905 410,906 427(811) 440,107 448,90	454,714	459,227	472,401
	383,137	384,757	387,200
10 13,500 141,231 252,815 318,311 324,381 301,421 399,014 423,228 435,138 445,998 454,831 423,228 435,138 445,998 454,831 423,228 435,138 445,998 454,831 445,998 455,138 445,998 455,138 445,998 455,138 445,198 455,138	462,960	470,552	485,388
<u>323,010</u> 312,080 400,327 414,538 425,943 435,909 447,450	403,120	408,089	471,391
	379,554	379,554	421,809
11 23,920 138,409 208,993 213,024 312,400 329,884 300,104 301,205 311,423 383,008 389,92 399,914 3	398,312	403,495	413,103
201,210 229,001 310,122 331,021 341,230 334,139 330,330 310,139 \$22,750 327,500 320,000 320,412 300,200 300,900 300,900 772 300,277 300,200	014,010	202.076	205 580
12 8 104 67 608 72 66 106 101 122 23,413 230,500 239	146 201	147 062	295,580
12 0,104 01,005 10,005 100,200 110,131 122,261 123,353 130,001 140,312 140,312 140,014 140,312 140,014 140,312 140,014 140,312 140,014 140,312 140,014 140,312 140,014 140,312 140,014 140,312 140,312 140,314 140,	191 599	199 515	125 105
17 210 07 020 127 000 162 012 166 424 129 004 105 004 105 004 125 004	121,088	104 227	104 227
12 12 214 14 091 22 761 30 325 42 105 44 376 44 557 48 407 40 306 55 304 51 304 45 55 304 51 304 55 500 56 55 304 55 500 56 55 500 56 55 500 56 55 500 56 55 500 56 55 500 56 55 500 56 55 500 56 55 500 56 55 500 500	52 057	52 446	194,221 52 976
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	51 351	51 683	52 538

Table 6: Monthly cumulative run-off triangle for only reported claims for September 2012. In black is the actual value, in blue is the estimation using the IPW estimator, and in red is the double Chain-Ladder estimation



Figure 8: Predictive distributions of the IPW method for the IBNS, RBNS and IBNR reserves for September 2012. The vertical line in black is the true observed value.

Reserve Type	Method	Ultimate	Reserve	Error	% Error
	True value	$5,\!871,\!171$	$1,\!643,\!872$	-	-
IBNS	IPW	$5,\!869,\!469$	$1,\!642,\!169$	1,702	0.1%
	CL	$5,\!138,\!422$	$911,\!123$	732,749	44.6%
RBNS	True value	5,116,312	$908,\!598$		-
	IPW	$5,\!104,\!503$	896,789	$11,\!810$	1.3%
	CL	$4,\!971,\!626$	$763,\!912$	$144,\!686$	15.9%
	True value	$754,\!859$	735,273		-
IBNR	IPW	764,966	$745,\!381$	-10,107	-1.4%
	CL	166,796	147,210	588,063	80.0%

Table 7: Error metrics for the aggregate reserves for September 2012.

6.4 Estimation of the reserve for several dates

Here we present the estimation of reserves for all 24 months in the testing period. Figure 9 illustrates the estimations for the IBNS compared to the true value of the reserve at the corresponding month. Additionally, Figures 10 and 11 depict the estimation for RBNS and IBNR, respectively. To provide a comprehensive analysis, we include 95% confidence intervals for the estimations and include the estimates of the Chain-Ladder (CL) method or comparison. We use the double Chain-Ladder approach to distinguish between IBNR and RBNS, and we use the approach of Section 2.4 to construct the confidence intervals. Furthermore, Table 8 presents error metrics to assess the disparities between the estimations across all dates. We would like to note that this kind of temporal analysis is often overlooked in the claim-reserving literature due to its inherent challenges for having consistent estimations. Our aim is not to boast about the complexity of the analysis but rather to show the behavior of the IPW method in distinct valuation dates.

Reserve Type	Method	ME	RMSE	MAE	MAPE
IDNG	IPW	21,908	$268,\!386$	$292,\!419$	17%
IDING	CL	$170,\!363$	$749,\!437$	$647,\!976$	49%
DDNG	IPW	133,738	226,994	$225,\!052$	23%
NDNS	CL	$144,\!503$	305,765	$311,\!635$	33%
IDND	IPW	-111,830	$182,\!154$	$210,\!047$	43%
	CL	$25,\!860$	414,344	$386,\!959$	76%

Table 8: Error metrics for the total of the reserves over the testing period. ME: Mean error, RMSE: Root mean square error, MAE: Mean absolute error, MAPE: Mean absolute percentage error

With respect to the IBNS i.e., the total reserve, Figure 9 demonstrates that the IPW estimator produces predictions that closely align with the true value of the reserve for the majority of the observed periods, exhibiting no discernible pattern of under or overestimations. Additionally, the actual reserve value consistently falls within the associated confidence intervals, indicating a consistent fit with the predicted value. Notably, the IPW prediction proves to be more accurate than the traditional Chain-Ladder method on almost all dates of the considered period. This observation is further supported by the results in Table 8, where the error metrics for the IPW over the 24-month period outperform those of the Chain-Ladder method. Therefore, the IPW along with the use of individual claims information has more predictive power than the macro reserving method. It is worth noting that the confidence intervals for the IPW method typically tend to be wider than those for the Chain-Ladder method. Essentially, this arises from the bias-variance trade-off, as the IPW method factors in the heterogeneity of the data in its predictions, whereas the Chain-Ladder



Figure 9: Estimation for the IBNS reserve per month

method does not.

With respect to the RBNS, Figure 10 reveals that the IPW estimator produces estimations exhibiting a similar pattern to those generated by the Chain-Ladder method. Both methods initially underestimate the reserve for the first year, followed by estimations that fluctuate around the true value of the reserve for the second year. While the estimated patterns align for these methods, the IPW estimator tends to be closer to the true value of the reserves in almost all cases, and when not, the estimation is comparable. Indeed, the IPW consistently outperforms the Chain-Ladder method on average throughout the entire period, as indicated by the lower error metrics in Table 8.

From a global perspective, it is noteworthy that the estimations capture the general trend of the true RBNS reserve, demonstrating congruent patterns of fluctuations and variations in a comparable manner with the true value. This is evident, for instance, in the increasing trend during the first year of the reserve and its subsequent decrease in a portion of the second year.

Regarding the IBNR reserve, Figure 11 illustrates that the IPW estimator provides a reasonable estimation throughout the considered period, fluctuating around the true reserve. It is noteworthy that the IPW estimator displays a more variable behavior compared to previous scenarios. This variability is expected due to the relatively low reporting delay time, resulting in the IBNR representing a smaller proportion of the subpopulation. Generally, as the size of the subpopulation decreases, the variance of the IPW increases.

Despite this variability, the IPW estimator consistently outperforms the traditional Chain-Ladder method over the entire 24-month period, as indicated in Table 8. It is important to highlight that the Chain-Ladder estimations of the reserves follow a consistent trend across all reserves, i.e., underestimation at the beginning of the first year, overestimation at the end of the first year and the beginning of the second, and then underestimation again at the end. In contrast, the IPW



Figure 10: Estimation for the RBNS reserve per month

estimation does not exhibit the same pattern. This discrepancy is attributed to the fact that the Chain-Ladder assumes a homogeneous portfolio, while the IPW does not.

Observation 1. We encountered instability in the behavior of the IPW estimator (i.e., absurdly abnormal large reserves) when performing the estimation on some dates, along the same lines as the behavior described in Section 2.5. To address this issue, we implemented the adjusted version of the IPW estimator, as described in Algorithm 1, and compared it to the raw IPW estimator. If the percentual difference between the two estimations exceeded a certain threshold (e.g., more than 3%), we retained the adjusted estimation. For cases where the difference was not significant, we kept the original estimation. As a result, the estimations become more stable across dates.

7 Conclusions

Macro-level reserving models, particularly the Chain-Ladder method, overlook the underlying heterogeneity within the portfolio of policyholders, treating all claims equally, and providing most of the time modest estimations. Therefore, the estimation of the reserve does not benefit from the use of the individual attributes of the policyholders, which have been shown to provide a significant improvement in the accuracy of these methods in the literature of micro-level reserving.

In this paper, we address the limitation of macro-level reserving models by proposing a statistically justified macro-level reserve estimator based on Inverse Probability Weighting (IPW). Unlike traditional macro-level models, our method incorporates individual-level information in the weights to improve the accuracy of reserve estimation. Moreover, such incorporation is achieved within a less complex framework compared to micro-level models, in the sense that no explicit assumptions on claim frequency or severity are made.



Figure 11: Estimation for the IBNR reserve per month

The IPW estimator serves as a hybrid approach that bridges the gap between macro and microlevel methods. It assigns attribute-driven weights to each claim, allowing for a development factor specific to each claim's settlement, similar to the familiar principles of the Chain-Ladder method when applied at the granular level. This method represents an initial step towards obtaining more precise reserves from macro-level models and serves as an intermediate stage in the development of a customized micro-level reserving model. We hope practitioners find this method appealing as it is a natural extension of the traditional Chain-Ladder method, accounting for portfolio heterogeneity in a statistically justified fashion.

The IPW and population sampling framework offers a new theoretical insight into the Chain-Ladder method, revealing it as a specific case. Several properties of the IPW estimator directly apply to the Chain-Ladder, allowing the derivation of both known and potentially unknown characteristics. For instance, the well-established unbiasedness of the Chain-Ladder stems from the unbiasedness of the IPW estimator and Assumption 1. A similar result applies to the distributional properties of the estimator. It is worth noting that specific assumptions might differ from those in other frameworks like Mack (1999). Although we don't explicitly outline all connections here, recognizing their potential is crucial.

Future research exploring the interplay between claim reserving and population sampling techniques holds promise for advancing reserve estimation. This involves investigating inherited properties of the Chain-Ladder as an IPW estimator and exploring implications. Additionally, other population sampling-motivated methods, like double robust estimators, offer new avenues for addressing the reserving problem. Indeed, our ongoing work has revealed that population sampling has the potential to offer a broader framework for reserving, encompassing both general macrolevel and micro-level models. Additionally, adopting a sampling perspective may yield insights into improving model estimation techniques, thereby striving for enhanced outcomes. Other research directions should explore alternative approaches for estimation, potentially through tailored models for the development of claims that are specifically designed for the estimation of the inclusion probabilities in Equation (15). In fact, one may consider alternative inclusion probabilities to use as weights, aiming for better reserve estimations. Moreover, the IPW framework and the change of population principle could be applied beyond the discussed reserves. For instance, the incurred but not paid (IBNP) reserve, and the unearned premium reserve (UPR), among others.

Acknowledgments

This work was partly supported by Natural Sciences and Engineering Research Council of Canada [RGPIN 284246, RGPIN-2017-06684]. Sebastián Calcetero Vanegas acknowledges the Mountain Pygmy Possum, an endangered species in Australia, for inspiring research on population sampling and also this project.

Conflicts of interest or Competing interests

The authors declare no conflicts of interest or competing interests in this paper, with no financial or personal affiliations that could compromise the objectivity or integrity of the presented work.

References

- Andersen, P. K. and Ravn, H. (2023). Models for Multi-State Survival Data: Rates, Risks, and Pseudo-Values. CRC Press.
- Andersen, P. K., Ørnulf Borgan, Hjort, N. L., Arjas, E., Stene, J., and Aalen, O. (1985). Counting process models for life history data: A review [with discussion and reply]. *Scandinavian Journal* of Statistics, 12(2):97–158.
- Antonio, K. and Plat, R. (2014). Micro-level stochastic loss reserving for general insurance. Scandinavian Actuarial Journal, 2014(7):649–669.
- Arnab, R. (2017). Survey Sampling Theory and Applications. Academic Press.
- Asmussen, S. and Steffensen, M. (2020). Risk and Insurance. Springer.
- Bender, A., Groll, A., and Scheipl, F. (2018). A generalized additive model approach to time-toevent analysis. *Statistical Modelling*, 18(3-4):299–321.
- Berger, Y. G. (1998). Rate of convergence for asymptotic variance of the horvitz-thompson estimator. Journal of Statistical Planning and Inference, 74(1):149–168.
- Bischofberger, S. M., Hiabu, M., and Isakson, A. (2020). Continuous chain-ladder with paid data. Scandinavian Actuarial Journal, 2020(6):477–502.
- Boumezoued, A. and Devineau, L. (2017). Individual claims reserving: a survey. https://hal.science/hal-01643929. Preprint. Accessed: 2023-06-04.
- Chauvet, G. and Vallée, A.-A. (2020). Inference for two-stage sampling designs. Journal of the Royal Statistical Society Series B: Statistical Methodology, 82(3):797–815.

- Chen, Q., Elliott, M. R., Haziza, D., Yang, Y., Ghosh, M., Little, R. J. A., Sedransk, J., and Thompson, M. (2017). Approaches to improving survey-weighted estimates. *Statistical Science*, 32(2):227–248.
- Cook, R. J., Lawless, J. F., et al. (2007). The Statistical Analysis of Recurrent Events. Springer.
- 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.
- England, P. and Verrall, R. (2002). Stochastic claims reserving in general insurance. *British* Actuarial Journal, 8(3):443–518.
- Engler, N. and Lindskog, F. (2024). Mack's estimator motivated by large exposure asymptotics in a compound poisson setting. ASTIN Bulletin: The Journal of the IAA, 54(2):310–326.
- Fattorini, L. (2006). Applying the horvitz-thompson criterion in complex designs: a computerintensive perspective for estimating inclusion probabilities. *Biometrika*, 93(2):269–278.
- Fattorini, L. (2009). An adaptive algorithm for estimating inclusion probabilities and performing the horvitz-thompson criterion in complex designs. *Computational Statistics*, 24:623–639.
- Francis, L. (2016). Astin working party releases report on reserving practices for general insurance worldwide. https://ar.casact.org/astin-working-party-releases-report-on-reserving-practices-forgeneral-insurance-worldwide. Accessed: 2023-06-04.
- Fung, T. C., Badescu, A. L., and Lin, X. S. (2021). A new class of severity regression models with an application to ibnr prediction. North American Actuarial Journal, 25(2):206–231.
- Fung, T. C., Badescu, A. L., and Lin, X. S. (2022). Fitting censored and truncated regression data using the mixture of experts models. North American Actuarial Journal, 26(4):496–520.
- George, B., Seals, S., and Aban, I. (2014). Survival analysis and regression models. Journal of Nuclear Cardiology, 21(4):686–694.
- Hiabu, M. (2017). On the relationship between classical chain ladder and granular reserving. Scandinavian Actuarial Journal, 2017(8):708–729.
- Hiabu, M., Hofman, E., and Pittarello, G. (2024). A machine learning approach based on survival analysis for ibnr frequencies in non-life reserving. arXiv preprint arXiv:2312.14549.
- Hiabu, M., Mammen, E., Martínez-Miranda, M. D., and Nielsen, J. P. (2021). Smooth backfitting of proportional hazards with multiplicative components. *Journal of the American Statistical* Association, 116(536):1983–1993.
- Hirano, K., Imbens, G. W., and Ridder, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189.
- Hulliger, B. (1995). Outlier robust horvitz-thompson estimators. Survey Methodology, 21(1):79–87.
- Kim, J. K. and Shao, J. (2021). Statistical Methods for Handling Incomplete Data. Chapman and Hall/CRC.
- Klein, J. P. and Moeschberger, M. L. (2003). Survival Analysis: Techniques for Censored and Truncated Data. Springer.

London, D. (1997). Survival Models and Their Estimation. Actex Publications, 3th revised edition.

- Ma, X. and Wang, J. (2020). Robust inference using inverse probability weighting. *Journal of the American Statistical Association*, 115(532):1851–1860.
- Mack, T. (1994). Which stochastic model is underlying the chain ladder method? Insurance: Mathematics and Economics, 15(2-3):133–138.
- 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.
- Martínez-Miranda, M. D., Nielsen, J. P., Sperlich, S., and Verrall, R. (2013). Continuous chain ladder: Reformulating and generalizing a classical insurance problem. *Expert Systems with Applications*, 40(14):5588–5603.
- Martínez-Miranda, M. D., Nielsen, J. P., and Verrall, R. (2012). Double chain ladder. ASTIN Bulletin: The Journal of the IAA, 42(1):59–76.
- Pittarello, G., Hiabu, M., and Villegas, A. M. (2023). Chain ladder plus: a versatile approach for claims reserving. arXiv preprint arXiv:2301.03858.
- Quarg, G. and Mack, T. (2004). Munich chain ladder. Blätter der DGVFM, 26(4):597–630.
- Rüschendorf, L. (2009). On the distributional transform, sklar's theorem, and the empirical copula process. *Journal of Statistical Planning and Inference*, 139(11):3921–3927.
- Salditt, M. and Nestler, S. (2023). Parametric and nonparametric propensity score estimation in multilevel observational studies. *Statistics in Medicine*, 42(23):4147–4176.
- Särndal, C.-E., Swensson, B., and Wretman, J. (2003). Model Assisted Survey Sampling. Springer Science & Business Media.
- Schmidt, K. (2017). A bibliography on loss reserving. https://tudresden.de/mn/math/stochastik/ressourcen/dateien/schmidt/dsvm/reserve.pdf?lang=en. Preprint. Accessed: 2023-06-04.
- Sonabend, R., Király, F. J., Bender, A., Bischl, B., and Lang, M. (2021). mlr3proba: an r package for machine learning in survival analysis. *Bioinformatics*, 37(17):2789–2791.
- Taha, A., Cosgrave, B., Rashwan, W., and McKeever, S. (2021). Insurance reserve prediction: Opportunities and challenges. In 2021 International Conference on Computational Science and Computational Intelligence (CSCI), pages 290–295.
- Tan, Z. (2006). A distributional approach for causal inference using propensity scores. Journal of the American Statistical Association, 101(476):1619–1637.
- Taylor, G. and McGuire, G. (2016). Stochastic loss reserving using generalized linear models. CAS Monograph, 3:1–112.
- Taylor, G., McGuire, G., and Sullivan, J. (2008). Individual claim loss reserving conditioned by case estimates. Annals of Actuarial Science, 3(1-2):215–256.
- Thompson, S. K. (2012). Sampling. John Wiley & Sons.

- Tseung, S. C., Badescu, A. L., Fung, T. C., and Lin, X. S. (2021). Lrmoe. jl: a software package for insurance loss modelling using mixture of experts regression model. Annals of Actuarial Science, 15(2):419–440.
- Vakulenko-Lagun, B., Mandel, M., and Betensky, R. A. (2020). Inverse probability weighting methods for cox regression with right-truncated data. *Biometrics*, 76(2):484–495.
- Verbelen, R., Gong, L., Antonio, K., Badescu, A., and Lin, S. (2015). Fitting mixtures of erlangs to censored and truncated data using the em algorithm. ASTIN Bulletin: The Journal of the IAA, 45(3):729–758.
- Verrall, R., Nielsen, J., and Jessen, A. (2010). Including count data in claims reserving. ASTIN Bulletin: The Journal of the IAA, 40(2):871–887.
- Verrall, R. J. and Wüthrich, M. V. (2016). Understanding reporting delay in general insurance. *Risks*, 4(3):25.
- Wahl, F., Lindholm, M., and Verrall, R. (2019). The collective reserving model. Insurance: Mathematics and Economics, 87:34–50.
- Wang, X., van Eeden, C., and Zidek, J. V. (2004). Asymptotic properties of maximum weighted likelihood estimators. *Journal of Statistical Planning and Inference*, 119(1):37–54.
- Wiegrebe, S., Kopper, P., Sonabend, R., Bischl, B., and Bender, A. (2024). Deep learning for survival analysis: a review. *Artificial Intelligence Review*, 57(3):65.
- 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.
- Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., and Zhang, A. (2021). A survey on causal inference. ACM Transactions on Knowledge Discovery from Data (TKDD), 15(5):1–46.