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

March 4, 2024

Abstract

Claim reserving is primarily accomplished using macro-level models, with the Chain-Ladder method being the most widely adopted method. These methods are usually constructed heuristically and rely on oversimplified data assumptions, neglecting the heterogeneity of policyholders, and frequently leading to modest reserve predictions. In contrast, micro-level reserving leverages on stochastic modeling with granular information for improved predictions, but usually comes at the cost of more complex models that are unattractive to practitioners. In this paper, we introduce a simple macro-level type approach that can incorporate granular information from the individual level. To do so, we imply a novel framework in which we view the claim reserving problem as a population sampling problem and propose a reserve estimator based on inverse probability weighting techniques, with weights driven by policyholders' attributes. The framework provides a statistically sound method for aggregate claim reserving in a frequency and severity distribution-free fashion, while also incorporating the capability to utilize granular information via a regression-type framework. The resulting reserve estimator has the attractiveness of resembling the Chain-Ladder claim development principle, but applied at the individual claim level, so it is easy to interpret and more appealing to practitioners.

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

1 Introduction

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 viewd as an empirical IPW estimator, operating under a specified homogeneity assumption. 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. Indeed, actuaries have extensively researched statistical models that replicate the estimation of the Chain-Ladder method, often in a very artificial fashion. 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 micro-level 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 process (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. 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 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.