



HAL
open science

Machine learning, bias correction and plug-in estimators for an accurate microlevel reserving

Olivier Lopez, Xavier Milhaud

► **To cite this version:**

Olivier Lopez, Xavier Milhaud. Machine learning, bias correction and plug-in estimators for an accurate microlevel reserving. 2019. hal-02292377v1

HAL Id: hal-02292377

<https://hal.science/hal-02292377v1>

Preprint submitted on 19 Sep 2019 (v1), last revised 9 Apr 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Machine learning, bias correction and plug-in estimators for an accurate microlevel reserving

Olivier Lopez¹, Xavier Milhaud²

¹Sorbonne Université, CNRS, Laboratoire de Probabilités, Statistique et Modélisation, LPSM, 4 place Jussieu, F-75005 Paris, France

²Université de Lyon, Université Claude Bernard Lyon1, ISFA, LSAF, F-69007, Lyon, France

September 19, 2019

Abstract

Thanks to nonparametric estimators coming from machine learning, microlevel claim reserving has become more and more popular in actuarial sciences. Recent research has focused on how to integrate the whole information one can have on claims to predict individual reserves, with varying success due to incomplete observations. In this paper, we introduce three extensions to comparable existing works: how to deal with censoring and truncation present in such type of data, how to cope with inflation when the inflation factor is unknown, and how to implement an adequate strategy leading to robust personalized reserve estimates. Using independent test sets, our results - on guarantees with typical long development times - indicate the importance of using the total claim development time to predict the reserves with acceptable accuracy. To remain close to reality, our applications are based on two open portfolios based on real-life datasets.

Keywords : reserving, censoring, truncation, CART, Kaplan-Meier weights.

1 Introduction

In non-life insurance, the final cost is rarely known when the claim occurs. In some cases, many years may pass before the amounts are settled. Standard procedures to predict the related reserve, typically Chain-Ladder (CL) approaches, rely on aggregated cumulative claim amounts. These techniques have the disadvantage not to exploit additional information one may have on the claims, and that may help to improve the prediction. Our

paper aims at using this information so as to improve the prediction of claim amounts for still open claims, the so-called Reported But Not Settled (RBNS) claims. We propose a regression-tree approach based on a particular weighting scheme in order to compensate the presence of censoring and truncation, naturally present in such data.

Right-censoring and left-truncation are known to have a significant impact when there is an underlying time phenomenon, see e.g. Fleming and Harrington [2011a]. In our framework, the claim lifetime (i.e. the difference between the occurrence date and the date at which the claim is fully settled) plays a crucial role to explain the final claim amount. Typically, claims that take a lot of time to be settled are more likely to be expensive, as pointed by Lopez [2019], Maegebier [2013], Spierdijk and Koning [2011], Pitt [2007], or Bluhm [1993]. Hence, calibrating a model only based on closed (or fully observed) claims usually conducts to an underestimation of the average amount. Indeed, such procedure is based on a set of observations in which there is an overrepresentation of claims with low lifetime, and therefore of claims with low value. The link between unsettled claims and censoring has been considered, for example, by Lopez et al. [2016], Albrecher et al. [2017] or Lopez et al. [2019]. We here extend such existing works to the estimation of claim amounts, and determine proper weights to be attributed to each observation to cancel the bias caused by censoring and truncation. A weighted regression tree is fitted, adapted from the CART algorithm initially developed by Breiman et al. [1984]. This algorithm has gained much popularity in actuarial sciences in the last decade due to its flexibility; see for instance Wüthrich [2018a] and Baudry and Robert [2017] for recent papers related to CART and reserving (note that other interesting works based on other machine learning techniques were recently published, see for instance Duval and Pigeon [2019] and Wüthrich [2018b]). CART easily allows to deal with data heterogeneity and nonlinearities between responses and explanatory factors, while preserving a relative simplicity of interpretation in the obtained model. The last step then consists of using the model to predict the claim reserves for RBNS claims.

The rest of the paper is organized as follows. Section 2 describes the structure of the observations, the general setting, and necessary treatments to be made on the data. Section 3 is devoted to the description of our approach to make the predictions, including the bias correction that removes the impact of censoring and truncation. Real data analyses are then performed in section 4 to assess the robustness of our technique. Results are compared to two widely used competing approaches in the industry, namely Chain Ladder and Cox modelling.

2 General framework

2.1 Observations

Consider n claims, with deflated amounts $(M_i)_{1 \leq i \leq n}$. For some of them, the final amount M_i is not observed, since the claim is still open. Indeed, the (random) time before the claim i is fully settled (the so-called claim lifetime), denoted by T_i , is censored. Introducing some censoring variables $(C_i)_{1 \leq i \leq n}$, define $(Y_i, \delta_i, N_i, \mathbf{X}_i)$ as i.i.d. replications of

$$\begin{cases} Y &= \inf(T, C), \\ \delta &= \mathbf{1}_{T \leq C}, \\ N &= \delta M, \\ \mathbf{X} &= (X^{(1)}, \dots, X^{(d)}) \in \mathbb{R}^d, \end{cases}$$

where \mathbf{X} are the claim features.

Moreover, let us denote by d_i the date at which the i -th claim occurs, s_i the date at which it is fully settled, and f_i the date at which the claim stops being observed. We have

$$\begin{aligned} T_i &= s_i - d_i, \\ C_i &= f_i - d_i. \end{aligned}$$

In practice f_i is often the same date whatever the claim, since most of time the cause of censoring is the database extraction. However, if some claims have been transferred to another company, information is lost at the transfer date, and f_i may vary between claims. Notice also that the database only contains claims that have been communicated, the Incurred But Not Reported (IBNR) claims being absent. This phenomenon is linked to left truncation. Introducing

$$\tau_i = r_i - d_i,$$

where r_i is the reporting date of claim i , only claims such that $C_i \geq \tau_i$ are observed. We assume that $(\tau_i)_{1 \leq i \leq n}$ are i.i.d., with the same distribution as a random variable τ . This truncation phenomenon is not completely standard, compared to the classical left-truncation model in survival analysis. Indeed, it is often assumed, in the literature, that observation only occurs when $T \geq \tau$ (see e.g. Tsai et al. [1987]) or when $Y \geq \tau$ (see e.g. Sellero et al. [2005]). Here, the situation is slightly different: $T < \tau$ means that the claim has been solved before reporting - which is the case, for example, if the indemnity has been fixed in advance - and the claim is present in the database.

To sum up, the observations are made of $(Y_i, \delta_i, N_i, \tau_i, \mathbf{X}_i)_{1 \leq i \leq n}$ i.i.d. with the same distribution as $(Y, \delta, N, \tau, \mathbf{X})$ conditionally to $C \geq \tau$. We assume hereafter that τ is independent from the other variables and, additionally, that C is independent from T, M, \mathbf{X} .

As already mentioned, censoring and truncation induce bias in the analysis if they are not taken into account. The former leads to under-represent claims with large times before settlement (large values of T_i) among the observations such that $\delta_i = 1$. Due to the positive correlation between T and M , this also has consequences on the distribution of M . The latter (left-truncation) clearly impacts claims that occur just before the extraction. In this paper, our aim is twofold: first, we want to understand the impact of T_i and \mathbf{X}_i on M_i . Typically, we aim to approximate the final cost of a claim with characteristics \mathbf{X}_i and lifetime T_i , that is $E[M_i|T_i, \mathbf{X}_i]$. Second, we would like to predict the final amount of unsettled claims. This consists, for one given claim with current observed duration Y_i , to estimate $E[M_i|T_i \geq Y_i, \delta_i = 0, \mathbf{X}_i]$.

Remark 2.1. *In many cases, one may have a preliminary idea on M_i even in the situation when the claim is still open ($\delta_i = 0$). A first way to take this information into account is to consider that the last evaluation of the claim is contained in the covariates \mathbf{X}_i . If partial payments have already been made, another possibility is to consider bivariate censoring. This leads to modify N_i : instead of taking $N_i = \delta_i M_i$, N_i is the cumulated amount already paid, and we know that $M_i \geq N_i$ (see also Remark 2.1 in Lopez [2019]).*

2.2 Preliminary step: removing inflation

Recall that classical statistical techniques rely on the assumption that observations are i.i.d. Hence, M_i must be corrected from inflation. Various methodologies can be used for such a task. Among them, using external data sources may be useful. Otherwise, another possibility is to follow the procedure by Lopez [2019].

Let M'_i denote the claim amount before removing inflation. In general, the data reports M'_i instead of M_i . We assume that

$$\log M'_i = \beta d_i + \log M_i, \tag{2.1}$$

where d_i is the date at which the i -th claim occurred, and β is a drift coefficient that is going to be estimated using our data. We assume that $(M_i)_{1 \leq i \leq n}$ are i.i.d. and independent from $(d_i)_{1 \leq i \leq n}$. The dates d_i take their value in $\{0, \dots, D\}$ (where $D+1$ years are observed). We then proceed in the following way:

- Compute $m'_{i,j}$, defined as the average of the fully observed claims that occurred on the fiscal year d_i , and which are settled after j years (for (i, j) such that $d_i + j \leq t_1$, where t_1 is the last observed date). Let $n_{i,j}$ denote the number of such claims.
- Under (2.1), we know that $\log m'_{i,j} \approx \beta_j d_i + \alpha_j$, where $\alpha_j = E[\log M_i | T_i = j]$. For each j , we compute $\hat{\beta}_j$ the weighted least-square estimator of the slope β based on the points $(m'_{i,j}, d_i)_{i:d_i+j \leq t_1}$. More precisely, one solves

$$(\hat{\alpha}_j, \hat{\beta}_j) = \arg \min_{\alpha_j, \beta_j} \sum_{i:d_i+j \leq t_1} n_{i,j} (\log m'_{i,j} - \alpha_j - \beta_j d_i)^2.$$

- It yields to our final estimator of β , given by $\hat{\beta} = \frac{\sum_j n_j^{1/2} \hat{\beta}_j}{\sum_j n_j^{1/2}}$, with $n_j = \sum_i n_{i,j}$.

Finally, for each claim i such that $\delta_i = 1$ (settled claims), we consider $\hat{M}_i = M'_i e^{-\hat{\beta} d_i}$ as an estimator of the final claim amount M_i , once removed the inflation effect. In practice, one will use this amount in applications.

3 Weighted regression-tree procedure

3.1 Statistical issues: technical details and useful references

From a statistical viewpoint, the main difficulty when dealing with censored and truncated data is the fact that the classical empirical means are biased. If one wishes to estimate $q_\phi = E[|\phi(T, M, \mathbf{X})|] < \infty$, one may write q_ϕ as an integral, that is

$$q_\phi = \int \phi(t, m, \mathbf{x}) dF(t, m, \mathbf{x}),$$

where $F(t, m, \mathbf{x}) = \mathbb{P}(T \leq t, M \leq m, \mathbf{X} \leq \mathbf{x})$, and plug a consistent estimator of F . In absence of censoring and truncation, a natural estimator of F is the empirical distribution function, leading to an estimation of q_ϕ through an empirical mean.

Consider first the simplest problem of estimating the marginal distribution function $F_T(t) = \mathbb{P}(T \leq t) = 1 - S(t)$, where $S(t)$ is the survival function. We propose to use the following estimator

$$\hat{F}_T(t) = 1 - \prod_{Y_i \leq t} \left(1 - \frac{\delta_i \mathbf{1}_{\tau_i < Y_i}}{\sum_{j=1}^n \mathbf{1}_{\tau_j < Y_i \leq Y_j}} \right),$$

where we assumed that there is no ties among the variables $(Y_i)_{1 \leq i \leq n}$ (an alternative formula in case of ties is proposed in Appendix A). This estimator slightly differs from

the classical one proposed by Tsai et al. [1987] under left-truncation, since, in our case, truncation occurs when $C \leq \tau$. The legitimacy of using \hat{F}_T is shown in section A.

Note that \hat{F}_T is piecewise constant, and can therefore be rewritten as

$$\hat{F}_T(t) = \sum_{i=1}^n w_{i,n} \mathbf{1}_{Y_i \leq t},$$

where

$$w_{i,n} = \frac{\delta_i \mathbf{1}_{\tau_i < Y_i}}{\sum_{j=1}^n \mathbf{1}_{\tau_j < Y_i \leq Y_j}} \prod_{Y_k < Y_i} \left(1 - \frac{\delta_k \mathbf{1}_{\tau_k < Y_k}}{\sum_{j=1}^n \mathbf{1}_{\tau_j < Y_k \leq Y_j}} \right). \quad (3.1)$$

Practically speaking, $w_{i,n}$ equals 0 when the observation is censored. The quantity $w_{i,n}$ can be thought as the weight to be put at observation i to correct the bias caused by censoring and truncation in our setting. Based on the idea of Sellero et al. [2005], one could use the same weight when it comes to estimate the joint distribution function F . Hence, we define

$$\hat{F}(t, m, \mathbf{x}) = \sum_{i=1}^n w_{i,n} \mathbf{1}_{Y_i \leq t, N_i \leq m, \mathbf{X}_i \leq \mathbf{x}}.$$

A natural estimator of q_ϕ thus yields

$$\hat{q}_\phi = \int \phi(t, m, \mathbf{x}) d\hat{F}(t, m, \mathbf{x}) = \sum_{i=1}^n w_{i,n} \phi(Y_i, N_i, \mathbf{X}_i).$$

3.2 Combination of our weighting scheme with CART algorithm

We here explain how to modify CART to suit our needs, i.e. the prediction of the individual claim amount M given the corresponding lifetime T and features \mathbf{X} . For the paper to be self-contained, we follow the description by Lopez et al. [2019]. However, the weighting scheme slightly differs, mostly because of the presence of truncation.

CART is a convenient way to estimate a regression function. Suppose that we want to estimate $\pi(t, \mathbf{x}) = E[\phi(M) | T = t, \mathbf{X} = \mathbf{x}]$: at each step of the algorithm, one determines "rules" $\mathbf{z} = (t, x^{(1)}, \dots, x^{(d)}) \rightarrow R_j(t, \mathbf{x})$ to split the data and create partitions of the covariate space. That is, for each possible value of (t, \mathbf{x}) , $R_j(t, \mathbf{x})$ equals 1 or 0 depending on whether some conditions are satisfied by $\mathbf{z} = (t, \mathbf{x})$. In addition, $R_j(t, \mathbf{x})R_{j'}(t, \mathbf{x})$ equals 0 when $j \neq j'$ and $\sum_j R_j(t, \mathbf{x})$ equals 1 to guarantee that created subsets of the partition are exhaustive and disjoint.

Our weighted CART algorithm can then be expressed as follows, where T and M are respectively replaced by their observed version Y and N .

Algorithm: weighted CART (denoted further by wCART)

Step 1: $R_1(\mathbf{z}) = 1$ for all $\mathbf{z} = (y, \mathbf{x})$, and $n_1 = 1$ (corresponds to the root node).

Step $k+1$: Let (R_1, \dots, R_{n_k}) denote the rules obtained at step k . For $j = 1, \dots, n_k$,

- if all observations such that $\delta_i R_j(Y_i, \mathbf{X}_i) = 1$ have the same characteristics, then keep rule j as it is no longer possible to segment the population;
- else, rule R_j is replaced by two rules R_{j1} and R_{j2} determined in the following way: for each component $Z^{(l)}$ of $\mathbf{Z} = (Y, \mathbf{X})$ ($l = 1, \dots, d+1$), define the best threshold $z_\star^{(l)}$ to split the data, such that $z_\star^{(l)} = \arg \min_{z^{(l)}} m(R_j, z^{(l)})$, with

$$m(R_j, z^{(l)}) = \sum_{i=1}^n w_{i,n} (\phi(N_i) - \bar{n}_{l-}(z^{(l)}, R_j))^2 \mathbf{1}_{Z_i^{(l)} \leq z^{(l)}} R_j(\mathbf{Z}_i) + \sum_{i=1}^n w_{i,n} (\phi(N_i) - \bar{n}_{l+}(z^{(l)}, R_j))^2 \mathbf{1}_{Z_i^{(l)} > z^{(l)}} R_j(\mathbf{Z}_i),$$

where

$$\bar{n}_{l-}(z, R_j) = \frac{\sum_{i=1}^n w_{i,n} \phi(N_i) \mathbf{1}_{Z_i^{(l)} \leq z} R_j(\mathbf{Z}_i)}{\sum_{k=1}^n w_{k,n} \mathbf{1}_{Z_k^{(l)} \leq z} R_j(\mathbf{Z}_k)}, \quad \bar{n}_{l+}(z, R_j) = \frac{\sum_{i=1}^n w_{i,n} \phi(N_i) \mathbf{1}_{Z_i^{(l)} > z} R_j(\mathbf{Z}_i)}{\sum_{k=1}^n w_{k,n} \mathbf{1}_{Z_k^{(l)} > z} R_j(\mathbf{Z}_k)}.$$

Then, select the best component to consider, that is $\hat{l} = \arg \min_l m(R_j, z_\star^{(l)})$.

Define the two new rules $R_{j1}(\mathbf{z}) = R_j(\mathbf{z}) \mathbf{1}_{z^{(\hat{l})} \leq z_\star^{(\hat{l})}}$, and $R_{j2}(\mathbf{z}) = R_j(\mathbf{z}) \mathbf{1}_{z^{(\hat{l})} > z_\star^{(\hat{l})}}$.

- Let n_{k+1} denote the new number of rules.

Stopping rule: stop if $n_{k+1} = n_k$.

A binary tree structure underlies this algorithm, since it is based on successive decompositions into two groups. The final obtained tree is called the *maximal tree*. Compared to the classical CART algorithm by Breiman et al. [1984], the splitting criterion (minimized at each step to split the population) is a weighted quadratic loss (instead of a quadratic loss) so as to compensate for censoring and truncation, as explained in section 3.1. Each set of rules $\mathcal{R} = (R_1, \dots, R_K)$ is associated with an estimator of the regression function, that is $\hat{\pi}^{\mathcal{R}}(t, \mathbf{x}) = \sum_{j=1}^K \hat{\pi}_j R_j(\mathbf{z})$, where

$$\hat{\pi}_j = \frac{\sum_{i=1}^n w_{i,n} \phi(N_i) R_j(\mathbf{Z}_i)}{\sum_{i=1}^n w_{i,n} R_j(\mathbf{Z}_i)}.$$

Of course, the maximal tree is not satisfactory to estimate the regression function $\pi(t, \mathbf{x})$. Indeed, the terminal nodes of this tree define clusters which are trivial (for one given leaf

R_j , there is either only one observation i such that $\delta_i R_j(\mathbf{Z}_i) = 1$, or all of the observations are associated to the same value of \mathbf{z}). A pruning step must then be performed to extract an appropriate subtree (the so-called *optimal tree*), in order to achieve some trade-off between adequacy and complexity. This can be understood as a model selection procedure. Let $K(\mathcal{R})$ denote the number of leaves of a subtree: the pruning approach proposed by Breiman et al. [1984], adapted to our framework, consists of minimizing

$$\sum_{i=1}^n w_{i,n} (\phi(N_i) - \hat{\pi}^{\mathcal{R}}(\mathbf{Z}_i))^2 + \alpha \frac{K(\mathcal{R})}{n},$$

where $\alpha > 0$ is a tuning parameter, usually chosen through cross-validation. Theoretical consistency properties of these building and pruning strategies are straightforward thanks to the properties of Kaplan-Meier estimators, and were shown in Lopez et al. [2016].

Remark 3.1. *Sometimes, the amount M may be a deterministic function of the lifetime T . The algorithm is thus easily adapted by replacing M by T , and \mathbf{Z} by \mathbf{X} . In this simplest situation where only one single censored and truncated variable has to be predicted, competing approaches include survival trees and survival forests, see e.g. Ishwaran et al. [2008] and Molinaro et al. [2004].*

3.3 Strategies to predict the final cost of RBNS claims

In absence of censoring or truncation, classical regression tree techniques can be used. Otherwise, extensions such as wCART should be considered to estimate quantities like $\pi(t, \mathbf{x}) = E[\phi(M) | T = t, \mathbf{X} = \mathbf{x}]$ (or $E[\phi(M) | \mathbf{X} = \mathbf{x}]$). From the selected tree estimator, it is then possible to deduce a predictor of $\phi(M)$ for claims where both T and \mathbf{X} are fully observed. Unfortunately, in our framework, we are interested in predicting the final amount ($\phi = id$) of RBNS claims, for which T is not fully observed. Say that the current observed duration equals y , we therefore need to adapt the algorithm:

- the best prediction (in terms of quadratic loss) of M given the available information is $E[M | T \geq y, \mathbf{X} = \mathbf{x}]$, which can also be written as

$$E[M | T \geq y, \mathbf{X} = \mathbf{x}] = \frac{E[M \mathbf{1}_{T \geq y} | \mathbf{X} = \mathbf{x}]}{E[\mathbf{1}_{T \geq y} | \mathbf{X} = \mathbf{x}]}. \quad (3.2)$$

Hence, two trees built using wCART can be used to estimate the numerator and the denominator of (3.2). This is the strategy described in Lopez et al. [2016], which requires the computation of a large number of trees (two trees per open claim). We refer to this procedure as strategy (A).

- build one single tree $\hat{\pi}$ by wCART to estimate $\pi(t, \mathbf{x}) = E[M | T = t, \mathbf{X} = \mathbf{x}]$. Then fit a model for $T | T \geq y, \mathbf{X} = \mathbf{x}$, from which a prediction $\hat{T}(y, \mathbf{x})$ can be computed, and predict M by $\hat{\pi}(\hat{T}(y, \mathbf{x}), \mathbf{x})$. Let us denote by (B) this strategy.

In the second approach (B), several prediction models for T can be used. A weighted regression tree is a possibility, with no obligation. A classical (semi)parametric model, or any machine learning prediction model adapted to censoring and truncation, can be used instead. In the sequel, we consider three different cases, among which weighted regression trees for the conditional distribution of T , and a simplified prediction of T assuming that T does not depend on \mathbf{X} . In the latter case, the prediction of T knowing that $T \geq y$ can be performed using

$$\hat{T} = \frac{\int t \mathbf{1}_{t > y} d\hat{F}_T(t)}{\int \mathbf{1}_{t > y} d\hat{F}_T(t)} = \frac{\sum_{i=1}^n w_{i,n} Y_i \mathbf{1}_{Y_i \geq y}}{\sum_{j=1}^n w_{j,n} \mathbf{1}_{Y_j \geq y}}. \quad (3.3)$$

We respectively refer to these methods as strategies (B)a), (B)b), and (B)c); summarized below. First estimate π by wCART to obtain $\hat{\pi}$, then

(B)a) : Build two trees $\hat{r}_{1,y}(\mathbf{x})$ and $\hat{r}_{2,y}(\mathbf{x})$ to estimate $r_{1,y}(\mathbf{x}) = E[T \mathbf{1}_{T \geq y} | \mathbf{X} = \mathbf{x}]$ and $r_{2,y}(\mathbf{x}) = E[\mathbf{1}_{T \geq y} | \mathbf{X} = \mathbf{x}]$ using wCART, and compute $\hat{T} = \hat{r}_{1,y}(\mathbf{x}) \hat{r}_{2,y}(\mathbf{x})^{-1}$; or

(B)b) : Predict T by \hat{T} from (3.3); or

(B)c) : Predict T by \hat{T} from the Cox model.

Finally, the prediction of M follows $\widehat{M} = \hat{\pi}(\hat{T}, \mathbf{x})$.

3.4 Expectations and comments on these strategies

Strategy (A) looks computationally intensive, but reveals quite simple to implement. The main pitfall appears for high thresholds of duration: indeed, when y gets high, the tree is built on very few events. As every classification technique, it becomes tricky to get robust estimators in this case, which does not ensure an accurate prediction of the denominator of (3.2). As this quantity tends to zero, predictions might overestimate the reality.

Concerning the procedure (B)b), it is obviously faster than (B)a). The computation of the weights $(w_{i,n})_{1 \leq i \leq n}$ is done once for all at the beginning, then the computation of (3.3) is very fast for each threshold y . However, strategy (B)b) is expected to be biased. Indeed, M is supposed to depend on \mathbf{X} , and T and M are expected to be dependent. Thus, assuming that T is independent from \mathbf{X} does not really make sense. Here, T is not predicted using all the information, but only through approaching $E[T | T \geq y]$, instead of

trying to estimate $E[T | T \geq y, \mathbf{X} = \mathbf{x}]$. Nevertheless, when summing the predictions over all open claims, one may hope that the bias vanishes, since $E[E[T | T \geq y, \mathbf{X}] | T \geq y] = E[T | T \geq y]$. We do not know by advance the bias amplitude for such a prediction, but due to the simplicity of the estimation technique compared to competing ones, one may expect a gain in terms of estimation variance. Hence, although asymptotically incorrect, this approach may lead, in some particular cases, to some accurate prediction.

Using Cox modelling (strategy (B)c)), one has to check the validity of the proportional hazards assumption. The latter is not fulfilled for most of covariates in our application. We thus expect to have mixed results. Finally, we challenge the previous strategies with two (biased) competitors, where the input information on T as a covariate is not realistic:

(C) : estimate π by $\hat{\pi}$ using wCART ; then predict M by $\widehat{M} = \hat{\pi}(y, \mathbf{x})$.

(D) : estimate π by $\hat{\pi}$ using wCART ; then estimate $r(\mathbf{x}) = E[T | \mathbf{X} = \mathbf{x}]$ by $\hat{r}(\mathbf{x})$ based on the wCART algorithm and predict M by $\widehat{M} = \hat{\pi}(\hat{r}(\mathbf{x}), \mathbf{x})$.

In the first situation, T is replaced by Y . Hence, compared to (A) and (B), the time before settlement is under estimated. This may lead to underestimations of M , if T is positively correlated with M . On the other hand, (D) uses a prediction of T given the available information, but does not take into account that $T \geq y$. Hence, some of the predictions will be smaller than y , and the overall result should also be underestimated.

4 Applications: reserve estimate for RBNS claims

We compare here the methodologies proposed in section 3.3 with benchmarks techniques, widely used for reserving purpose. Concretely, the global reserve is also estimated thanks to Chain Ladder on the one hand (on aggregated data), and Cox modelling on the other hand (as recommended for some life insurance risks such as short-term disability).

4.1 Context of our studies

To be in line with regulatory constraints (disclosure of regular reports) and current practices in insurance companies, the estimations of claim amounts and reserves are made on a quarterly basis. Building the database and running the estimation procedure is thus performed many times (every quarter), with updated policyholders' features and claim characteristics (in particular if it is open, closed, or new). Each time, we split our data

into two independent subsamples to assess the prediction power of the methods: the learning set which represents two third of the initial dataset, and the test set containing the remaining observations (from which the predictions are made). Moreover, we backtest our reserve predictions: only the claims closed at the last observed date are considered. This way, the whole information is known a posteriori, and one can compare predictions to the actual reserves finally needed.

Remark 4.1. *We also tried to change the time step between estimations (from quarterly to annually, bi-annually, monthly and weekly). We do not present the results for the sake of conciseness, but the conclusions were similar to those obtained hereafter.*

4.2 First application: income protection

Short-term disability insurance was designed to protect the policyholders against the loss of some revenue. In our context this coverage can last up to three years, meaning that the duration of payments is capped. In this application, predicting the final claim amount is thus similar to predicting the residual lifetime in the disability state (given the individual features \mathbf{X}). For the sake of simplicity, let us say that the insurer has to pay 1€ for each insured day, i.e. $M = T$ with our notations. We wish to predict the global reserve at various settlement dates (each quarter between 01/01/2008 and 10/01/2009) by aggregating individual predicted remaining claim lifetimes, using strategy (A) and Remark 3.1. Other strategies suggested in section 3.3 are useless since there is no need to plug in estimators for T , since T is not an explanatory variable in this case.

4.2.1 About the database

Our database reports 65 670 claims related to income protection guarantees over six years, from 01/01/2006 to 12/31/2011. For each claim, we know the gender of the policyholder (14 455 males, 51 215 females), her socio-professional category (2 406 managers, 62 799 employees and 465 others), her age when the claim occurred, the duration in the disability state (sometimes right-censored), the commercial network (three kinds of brokers: 28 662 “*Net-A*”, 4 890 “*Net-B*” and 32 118 “*Net-C*”), and the cause (57 131 sicknesses, and 8 539 accidents) that triggered the coverage. The overall censoring rate equals 7.2% on 12/31/2011. The mean observed duration in the disability state is 100 days (beyond a deductible of 30 days), with a median of 42 days and a standard deviation of 162 days. If necessary, more details about the data can be found in Lopez et al. [2019], section 3.

4.2.2 Claim lifetime predictions

All the results are stored on the top part of Table 2. Let us analyze and comment them.

First notice that the global reserve is strongly underestimated for all quarters of year 2008, with both Chain Ladder (CL) and strategy (A). On 01/01/2008, CL's error approximates 60%, whereas (A)'s error is roughly about 44%. CL does not take into account the censoring phenomenon, meaning that the method considers that all the claims are fully observed. Not surprisingly, this causes large underestimations of the final claim amounts. Although (A) is supposed to appropriately deal with censoring, largest observed lifetimes only equal two years at the beginning of 2008. The asymptotic properties of our tree estimator are guaranteed provided that observations (almost) entirely map the domain of possible values for T , which is not the case here. Indeed, some of the claims will last up to three years in practice, naturally resulting in an underestimation of the global reserve.

Second, the errors of CL and (A) decrease as time flies. This is coherent, as more and more information becomes available. One year later (on 01/01/2009), (A)'s error strongly lowers to reach 0.5%, whereas it still equals 39% with CL. The estimation of the global reserve thus improves much faster using (A) than in the CL case. Strategy (A) benefits from the fact that larger observed lifetimes (up to 3 years) are now included in the learning set, allowing to significantly improve the quality of our estimator. On the contrary, CL still suffers from the proportion of censored lifetimes (almost 11.9%).

Third, when looking at most recent settlement dates, (A) tends to overestimate the global reserve whereas underestimation by CL persists (but lowers). This makes sense since the portfolio is observed until 12/31/2011: getting closer to this date, the percentage of fully observed claims increases (for instance, the censoring rate equals 8.25% in the learning sample on 10/01/2009). On the one hand, this is good news for the CL approach, even though the underestimation remains high (-24%). However, such improvements for CL would not be experienced in reality: recall that the observation period terminates on 12/31/2011, and that backtesting implies that all considered claims are closed at this date. On the other hand, strategy (A) faces a selection bias when predicting the reserves, due to an overrepresentation of claims with short developments (whatever the features \mathbf{X}). There are RBNS claims to pay for, but the backtesting procedure leads to lower the actual global remaining lifetime, and thus the reserve. Indeed, our estimator is based on past information, and thus anticipates longer developments on average for the considered claims. Hopefully, this overestimation is therefore a non-issue since this is due to backtesting, and is thus artificial. To make sure about that, we ran the estimations

on the whole database (without only selecting closed claims on 12/31/2011) and ensured that the prediction error remained low and stable, whatever the settlement date.

Cox modelling globally shows better performance, especially for earliest settlement dates. Despite some constraints related to the semiparametric relation between the lifetime T and the covariates \mathbf{X} , Cox's error never exceeds 7.3% (except for the last settlement date where the selection bias plays the same role as previously). Still, these outstanding results should be moderated since they originate from very favourable circumstances: low censoring rate, no particular indication that the proportional hazard assumption may not be reasonable, and bounded lifetimes that cover a somewhat narrow interval.

To sum up, (A) seems to be as efficient as Cox to estimate the reserve, provided that the main underlying characteristics of the insured risk have been observed. Concretely, the insurer needs at least three years of historical data in order to apply it with high accuracy (because short-term disability can last up to three years). Moreover, it is important to mention that both Cox model and (A) accord with designating the policyholder's age as the most discriminant risk factor to explain T , and that weighted tree estimators remain consistent when the censoring rate gets much higher (see section 4 in Lopez et al. [2016]).

4.3 Second application: Third Party Liability (TPL) insurance

The first application focused on predicting remaining claim lifetimes, it is now time to move to the prediction of residual claim amounts for RBNS claims. To this aim, we use the strategies introduced in Section 3.3, and see whether the results significantly differ. Such a problem is more complex than the former, since there is an additional stochastic layer impacting the prediction (represented by the random claim lifetime). We consider an open source non-life insurance dataset providing claims in motor insurance, named `ausautoBI8999` and available in the R package `CASdatasets`¹. A lot of claims have long development times, causing specific claim management processes and atypical triangles (see Appendix B.2) when applying loss triangle techniques (like Chain Ladder).

4.3.1 Brief description of the database under study

The dataset is made of 22 036 settled automobile bodily injury claims in Australia. These claims arose from accidents occurring from July 1989 to January 1999. The database contains event dates (accident, reporting, closing), operational time (indicator of claim

¹See the following webpage: <http://dutangc.free.fr/pub/RRepos/web/CASdatasets-index.html>

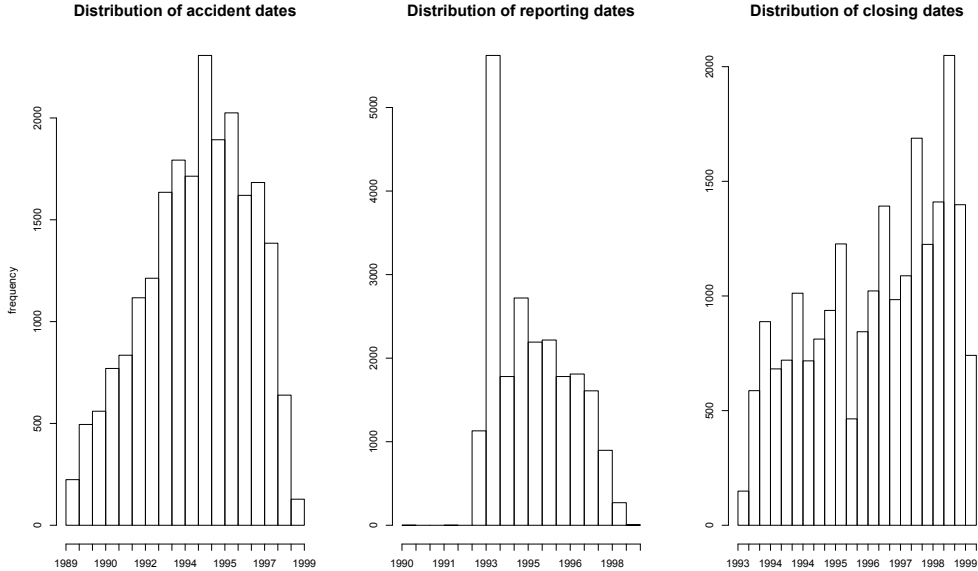


Figure 1: Distribution of dates of events in `ausautoBI8999` dataset.

management difficulties), type of injury, number of injured people, potential legal representation of the policyholder, and aggregated settled claim amount. Table 3 in Appendix B.1 summarizes descriptive statistics on these variables, as well as other created by-hand variables useful for our study (e.g. reporting delay, or claim duration). Figure 1 shows that the data is highly skewed in terms of reporting and closing dates, which is very likely due to data collection. In particular, note that there is no settled claim before 01/01/1993. Knowing that the mean claim lifetime equals 558 days and that accident dates start in 1989, it looks necessary to omit the data before 01/01/1993 in our study to avoid issues related to data quality. Finally, our database thus reports 16 822 claims.

4.3.2 Estimation of reserves for RBNS claims

First, claim amounts have to be inflated to the most recent date (01/01/1999). The annual inflation rate, estimated considering the whole observation period (see section 2.2), equals 0.39%. Recall that we wish to estimate the expected cost of RBNS claims, i.e. $E[M | T > y, \mathbf{X}]$, where the claim lifetime T is now a censored explanatory variable. We use the six strategies (A), (B)a), (B)b), (B)c), (C) and (D) to do so; which obviously lead to different estimators of the individual reserves. Strategies (B)a), (B)b), (B)c), (C) and (D) give the same tree estimator $\hat{\pi}$, but reserve predictions will differ for one simple reason: rebuilding the information on T is not considered in the same way. The estimator $\hat{\pi}$ highlights that the claim lifetime has rightly been detected as the risk factor with the

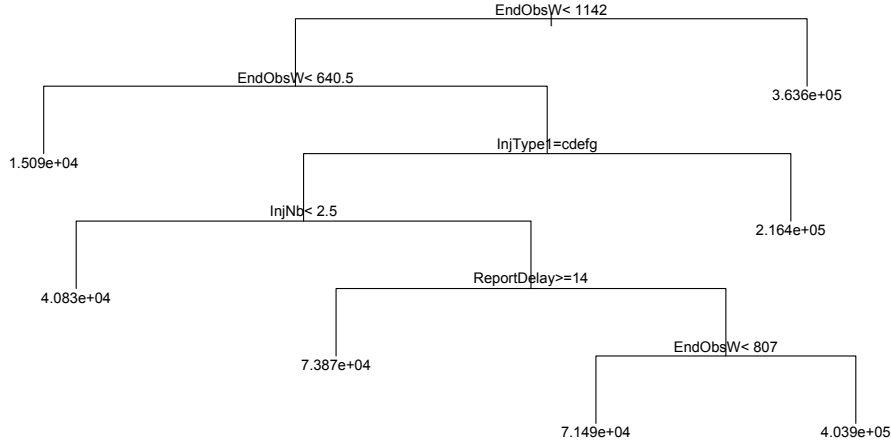


Figure 2: Optimal tree following strategy (D). The claim lifetime ('EndObsW') clearly appears as the most important explanatory variable to predict claim amounts.

strongest impact on the settled claim amounts, for all settlement dates considered. As an illustration, on 03/31/1997, Figure 2 shows that the most discriminant threshold for claim duration slightly exceeds three years, meaning that claims that last more than three years before being closed are expected to cause significantly higher final claim amounts. Let us furthermore point out that this threshold is stable, whatever the settlement date.

Concerning the reserve predictions, they are globally in line with the expectations given in section 3.4. The bottom part of Table 2 shows that the assessment error of the global reserve never falls down 0.5%, whatever the strategy and the settlement date considered. However, our reserve estimates are more stable and accurate than insurance practice (Chain Ladder and Cox), especially when looking at strategy (B)a). Chain Ladder always underestimates the reserve, as it does not take into account the censoring phenomenon (improvements related to the last settlement dates are fictive, as already explained at the end of Section 4.2.2).

To go further and detail our results, let us start with strategy (C). It considers the claim lifetime as a fully observed input in the modelling, yet it is not the case for censored claims in reality. Combined to the fact that the backtest sample is only made of censored cases, and due to the positive correlation between lifetimes and amounts (Kendall's tau equals 0.36), reserves are therefore always underestimated. The quality of predictions improves as time flies, which makes sense since the estimator is based on information containing lower bias for more recent settlement dates. For example, on 06/30/1998, the censoring rate decreases to 22% and the estimator error reaches about 15%.

With (D), we still have huge underestimations but the reason is different. We replace

T by $\hat{T} = E[\widehat{T} | \mathbf{X}]$ to make the predictions, where \hat{T} comes from the wCART algorithm. The predicted lifetime \hat{T} underestimates $E[T | T \geq y, \mathbf{X}]$, because it does not integrate the information on elapsed time. Knowing that this strategy is used for all individuals in the backtest sample, (D) obviously underestimates the individual reserves. However, no improvement is observed here for most recent settlement dates, which is not surprising because getting more and more experience does not guarantee that \hat{T} better estimates T . Strategy (D) is thus worse than (C) to make predictions when the censoring rate is low.

Strategy (B)a) seems to be reasonable: in terms of computation power, it is slightly more demanding than (A). However, it is easier to estimate $E[T | T > y, \mathbf{X} = \mathbf{x}]$ than $E[M | T > y, \mathbf{X} = \mathbf{x}]$, since T is a low-dispersed random variable as compared to M . As expected, strategy (A) tends to overestimate the global reserve because it suffers from the lack of data related to high values of y (added to the high dispersion of M). This effect is even more pronounced for the last settlement dates, where there are fewer claims whose lifetimes exceed high threshold of y (due to backtesting).

On 09/30/1996, (B)a) strongly underestimates the actual reserve needed. This is due to the distribution of the claim lifetimes in the created subsets: Table 1 shows that $r_{2,y}(\mathbf{x})$ (Section 3.3) is greater in the learning sample than in the validation one (for a given value of y), which should lead to underestimate $\hat{T} = E[T | T > y, \mathbf{X} = \mathbf{x}]$ when applying the model to the backtest sample, and thus to underestimate the individual reserves. The same story applies for strategy (B)b). Hopefully, the distributions of observed durations are similar between learning and validation samples for other settlement dates.

The overestimations of the global reserve in 1998 are caused by the selection bias due to our backtesting approach, see the discussion in Section 4.2.2. For other settlement dates, (B)a) seems to be more accurate than (B)b). As already mentioned, (B)b) is biased since it considers that T does not depend on \mathbf{X} . In this case, this assumption seems to be too strong to get robust estimates of the reserve. Concerning (B)a), the mean error rate is about $\pm 7\%$. Unless this may seem substantial, the explanation is threefold : i) learning samples are of limited size, ii) the censoring rate is huge (between 30% and 55% depending on the settlement date), and iii) there is an additional stochastic layer (represented by T as an explanatory variable) impacting predictions of M .

Quantile level:	0	10%	15%	20%	25%	30%	35%	40%
Learning set	0	91	121	152	181	212	242	273
Validation set	0	61	92	123	153	184	214	244

Table 1: Quantiles of claim durations between learning and validation sets on 09/30/1996.

	01/01/2008	04/01/2008	07/01/2008	10/01/2008	01/01/2009	04/01/2009	07/01/2009	10/01/2009
Income protection (section 4.2)								
(1) Size of the i)learning / ii)validation sets	20 542/10 271	23 370/11 686	26 214/13 107	28 740/14 371	31 962/15 982	34 796/17 399	37 700/18 850	40 344/20 172
(2) Censoring rate in validation set	16.24%	13.66%	12.8%	11.4%	11.89%	10.32%	9.27%	8.25%
⇒ Number of backtested claims : $(2) \times (1)ii)$	1 688	1 596	1 677	1 638	1 900	1 795	1 747	1 664
(3) Total paid amount at settlement date	818 079	955 809	1 115 449	1 259 591	1 448 942	1 608 799	1 771 356	1 955 760
(4) Paid at settlement date (censored claims)	278 230	286 354	323 982	336 883	378 083	388 346	387 616	399 445
(5) Final paid amount (censored claims)	657 047	650 253	708 685	719 172	778 448	780 152	768 116	741 743
(6) Exact (backtested) global reserve : $(5) - (4)$	378 817	363 899	384 703	382 289	400 365	391 806	380 500	342 298
(7) Global reserve by Chain Ladder	151 017	166 614	193 593	207 677	243 701	242 688	254 947	259 834
(8) <i>Error of Chain Ladder</i> : $((7) - (6))/(6)$	-60.1%	-54.2%	-50%	-45%	-39%	-38%	-33%	-24%
(9) Global reserve by Cox modelling	406 559	359 710	386 701	366 381	414 068	388 272	389 268	378 820
(10) <i>Error of Cox modelling</i> : $((9) - (6))/(6)$	7.3%	-1.2%	0.5%	-4.2%	3.4%	-0.9%	2.3%	9.5%
(11) Global reserve by strategy (A)	211 357	227 088	263 030	312 400	402 398	384 361	387 525	374 133
(12) <i>Error of strategy (A)</i> : $((11) - (6))/(6)$	-44.2%	-42%	-31.6%	-18.3%	0.5%	-1.9%	1.8%	9.3%
TPL insurance (section 4.3)								
(1) Size of the i)learning / ii)validation sets	7 960/3 981	8 513/4 257	9 022/4 511	9 515/4 758	9 995/4 998	10 436/5 219	10 782/5 392	11 033/5 517
(2) Censoring rate in validation set	54.43%	50.69%	47.21%	42.45%	37.99%	33.01%	28.52%	22.77%
⇒ Number of backtested claims : $(2) \times (1)ii)$	2 167	2 116	2 099	2 046	1 995	1 785	1 595	1 230
(3) Total paid amount at settlement date	35 852 729	38 609 774	51 151 746	56 014 781	68 835 620	78 851 031	93 701 263	102 519 887
(4) Paid at settlement date (censored claims)	unknown	unknown	unknown	unknown	unknown	unknown	unknown	unknown
(5) Exact (backtested) global reserve :	97 257 223	102 257 131	96 664 801	91 507 030	77 662 687	69 683 516	65 002 213	50 562 874
(6) Global reserve by Chain Ladder	22 928 280	22 965 831	32 037 675	34 391 312	56 800 817	64 518 373	79 882 662	94 678 896
(7) <i>Error of Chain Ladder</i> : $((6) - (5))/(5)$	-76.4%	-77.5%	-66.9%	-62.4%	-26.9%	-7.4%	+22.9%	+87.2%
(8) Global reserve by strategy (A)	115 270 629	105 774 935	104 192 632	109 865 208	102 158 129	113 866 731	95 394 567	75 973 350
(9) <i>Error of strategy (A)</i> : $((8) - (5))/(5)$	+18.5%	+3.4%	+7.8%	+20.1%	+31.5%	+63.4%	+46.7%	+50.2%
(10) Global reserve by strategy (B)a)	70 118 268	107 898 314	97 258 540	86 771 554	88 493 582	64 516 565	54 414 821	64 722 217
(11) <i>Error of strategy (B)a)</i> : $((10) - (5))/(5)$	-27.9%	+5.5%	+0.6%	-5.2%	+13.9%	-7.5%	-16.3%	+28%
(12) Global reserve by strategy (B)b)	69 488 566	133 367 406	121 172 197	98 547 137	131 446 851	69 330 511	58 843 816	86 412 081
(13) <i>Error of strategy (B)b)</i> : $((12) - (5))/(5)$	-28.5%	+30.4%	+25.3%	+7.7%	+69.2%	-0.5%	-9.5%	+70.9%
(14) Global reserve by strategy (B)c)	101 130 170	112 344 821	112 207 661	136 491 137	118 179 879	76 968 555	58 125 322	77 492 663
(15) <i>Error of strategy (B)c)</i> : $((14) - (5))/(5)$	+4%	+9.9%	+16.1%	+49.1%	+52.2%	+10.5%	-10.6%	+53.3%
(16) Global reserve by strategy (C)	62 906 056	60 324 551	57 297 234	61 678 665	60 692 333	60 133 761	52 562 268	43 341 206
(17) <i>Error of strategy (C)</i> : $((16) - (5))/(5)$	-35.3%	-41%	-40.1%	-32.6%	-21.8%	-13.7%	-19.1%	-14.3%
(18) Global reserve by strategy (D)	60 422 123	34 157 053	39 025 819	52 726 684	39 484 387	56 791 202	49 962 130	22 723 368
(19) <i>Error of strategy (D)</i> : $((18) - (5))/(5)$	-37.9%	-66.6%	-59.6%	-42.4%	-49.2%	-18.5%	-23.1%	-55.1%

Table 2: Quarterly assessment of the global reserve at various settlement dates, with corresponding error rates.

Looking at the results from Cox modelling (strategy (B)c)), they are very unstable. The quality of the predictions for the overall reserve can reveal either very good or very bad. As compared to the application of Section 4.2, the Cox model strongly deteriorates when applied to predict \hat{T} (before plugging it into the model on M , built by wCART). Looking more deeply to the results of the Cox step, we understand that the Cox estimator is affected by both a higher censoring rate and the violation of the proportional hazards assumption. Besides, T is a random variable with higher dispersion than in the previous application, which is suspected to make its modelling more complex.

In a nutshell, we recommend to use the strategy (B)a) to make the reserve predictions, as it outperforms all other methods and shows stable results in terms of prediction error. Once the claim lifetime T is consistently estimated for RBNS claims, running the wCART algorithm to predict M (based on \hat{T} and \mathbf{X}) seems to be acceptable to predict the claim amounts with reasonable accuracy.

5 Conclusion

We proposed different methodologies to perform individual claim reserving, based on regression trees. The main features of these approaches are the possibility to use all available information on a claim to predict its final state; the nonlinear and flexible structure of regression trees; and the correction of truncation and censoring phenomenons. In other words, the information on the time since occurrence of the claim is appropriately and fully integrated in the model in our framework. Our applications are mainly a picture of the reserve at some point of time. In particular, no dynamic readjustment of the reserve - due to new information or events that affect the claim - is considered. Nevertheless, our technique may be easily modified to incorporate this, as long as the required information is available. Among other possible improvements, let us mention the possibility to use random forests (i.e. aggregations of regression trees) to stabilize the results, since the CART algorithm is known to be sometimes sensitive to the introduction of new data. The drawback would be a loss of intelligibility of the obtained model.

References

- H. Albrecher, J.L. Teugels, and J. Beirlant. *Reinsurance: Actuarial and Statistical Aspects*. Statistics in practice. Wiley, 2017. ISBN 9780470772683. URL <https://books.google.fr/books?id=HZjObwAACAAJ>.
- M. Baudry and C.Y. Robert. Non parametric individual claim reserving in insurance. Working Paper, 2017.
- W.F Bluhm. Duration-based policy reserves. *Transactions of Society of Actuaries*, 45: 11–53, 1993.
- L Breiman, J Friedman, R A Olshen, and C J Stone. *Classification and Regression Trees*. Chapman and Hall, 1984.
- F. Duval and M. Pigeon. Individual loss reserving using a gradient boosting-based approach. *Risks*, 7(79):1–19, 2019. doi: 10.3390/risks7030079.
- Thomas R Fleming and David P Harrington. *Counting processes and survival analysis*, volume 169. John Wiley & Sons, 2011a.
- Thomas R Fleming and David P Harrington. *Counting processes and survival analysis*, volume 169. John Wiley & Sons, 2011b.
- Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, Michael S Lauer, et al. Random survival forests. *The annals of applied statistics*, 2(3):841–860, 2008.
- Olivier Lopez. A censored copula model for micro-level claim reserving. *Insurance: Mathematics and Economics*, 87:1 – 14, 2019. ISSN 0167-6687. doi: <https://doi.org/10.1016/j.insmatheco.2019.04.001>. URL <http://www.sciencedirect.com/science/article/pii/S0167668718300441>.
- Olivier Lopez, Xavier Milhaud, and Pierre-Emmanuel Therond. Tree-based censored regression with applications in insurance. *Electronic Journal of Statistics*, 10:2685–2716, 2016. URL [dx.doi.org/10.1214/16-EJS1189](https://doi.org/10.1214/16-EJS1189).
- Olivier Lopez, Xavier Milhaud, and Pierre-E. Théron. A tree-based algorithm adapted to microlevel reserving and long development claims. *ASTIN Bulletin*, page 1–22, 2019. doi: 10.1017/asb.2019.12.
- A. Maegebier. Valuation and risk assessment of disability insurance using a discrete time trivariate markov renewal reward process. *Insurance: Mathematics and Economics*, 53 (3):802–811, 2013. doi: 10.1016/j.insmatheco.2013.09.013.
- Annette M. Molinaro, Sandrine Dudoit, and Mark J. van der Laan. Tree-based multivariate regression and density estimation with right-censored data. *Journal of Multivariate Analysis*, 90(1):154 – 177, 2004. ISSN 0047-259X. doi: <https://doi.org/10.1016/j.jmva.2004.01.001>.

1016/j.jmva.2004.02.003. URL <http://www.sciencedirect.com/science/article/pii/S0047259X04000296>. Special Issue on Multivariate Methods in Genomic Data Analysis.

D.G.W. Pitt. Modelling the claim duration of income protection insurance policyholders using parametric mixture models. *Annals of Actuarial Science*, 2(1):1–24, 2007.

César Sánchez Sellero, Wenceslao González Manteiga, and Ingrid Van Keilegom. Uniform representation of product-limit integrals with applications. *Scandinavian Journal of Statistics*, 32(4):563–581, 2005. ISSN 03036898, 14679469. URL <http://www.jstor.org/stable/4616898>.

L. Spierdijk and R.H. Koning. Calculating loss reserves in a multistate model for income insurance. Working Paper, 2011.

Wei-Yann Tsai, Nicholas P. Jewell, and Mei-Cheng Wang. A note on the product-limit estimator under right censoring and left truncation. *Biometrika*, 74(4):883–886, 1987. ISSN 00063444. URL <http://www.jstor.org/stable/2336484>.

M.V. Wüthrich. Machine learning in individual claims reserving. *Scandinavian Actuarial Journal*, 2018(6):465–480, 2018a. doi: 10.1080/03461238.2018.1428681.

M.V. Wüthrich. Neural networks applied to chain-ladder reserving. *European Actuarial Journal*, 8(2):407–436, 2018b. doi: 10.1007/s13385-018-0184-4.

A Estimator of $S(t)$ in our censoring-truncation model

Let us recall that, for a discrete variable A taking value at point $\{a_1, \dots, a_k\}$, its survival function $S_A(t) = \mathbb{P}(A \geq a)$ can be written as

$$S_A(t) = \prod_{j=1}^k (1 - \lambda_A(a_j)), \quad (\text{A.1})$$

with

$$\lambda_A(t) = -\frac{dS_A(t)}{S_A(t)}.$$

A way to determine an estimator of S_A hence reduces to replace λ_A in (A.1) by a consistent estimator obtained from the data. If the variable A is not discrete, it can still be approximated by a discrete distribution where the $(a_j)_{1 \leq j \leq k}$ are replaced by the value of the complete observations (in our case, the uncensored observations). This is the basis of the construction of Kaplan-Meier and other product-limit based estimator, see for example Fleming and Harrington [2011b].

Hence, our aim is to determine a consistent estimator of $\lambda_T(t) = -\frac{dS_T(t)}{S_T(t)}$.
Let $L(t) = \mathbb{P}(\tau \leq t)$, $S_C(t) = \mathbb{P}(C \geq t)$, $\alpha = P(\tau < C)$, and

$$S_1(t) = E[\delta \mathbf{1}_{\tau \leq Y} \mathbf{1}_{Y \leq t} | \tau < C].$$

By basic computations, we obtain $dS_1(t) = -\alpha^{-1} S_C(t) L(t) dS_T(t)$.
Indeed,

$$\begin{aligned} S_1(t) &= \alpha^{-1} E[\mathbf{1}_{\tau \leq t} \mathbf{1}_{T \leq T} E[\mathbf{1}_{\max(\tau, T) < C} | \tau, T]] \\ &= \alpha^{-1} E[\mathbf{1}_{\tau \leq T} \mathbf{1}_{T \leq t} S_C(\max(\tau, T))] \\ &= \alpha^{-1} E[\mathbf{1}_{\tau \leq T} \mathbf{1}_{T \leq t} S_C(T)] \\ &= -\alpha^{-1} \int_0^t S_C(y) L(y) dS_T(y), \end{aligned}$$

where we used that (T, τ) is independent from C for the second line, and the independence between T and C for the last line. On the other hand, let

$$S_2(t) = E[\mathbf{1}_{\tau < t < Y} | \tau < C].$$

We have $S_2(t) = \alpha^{-1} S_T(t) S_C(t) L(t)$. Since

$$\begin{aligned} S_2(t) &= \alpha^{-1} E[\mathbf{1}_{\tau < t} \mathbf{1}_{t < T} \mathbf{1}_{\max(t, \tau) < C}] \\ &= \alpha^{-1} E[\mathbf{1}_{\tau < t} E[\mathbf{1}_{t < T} \mathbf{1}_{\max(t, \tau) < C} | \tau]] \\ &= \alpha^{-1} E[\mathbf{1}_{\tau < t} S_T(t) S_C(\max(t, \tau))] \\ &= \alpha^{-1} E[\mathbf{1}_{\tau < t} S_T(t) S_C(t)], \end{aligned}$$

hence

$$-\frac{dS_T(t)}{S_T(t)} = \frac{dS_1(t)}{S_2(t)}.$$

The quantities S_1 and S_2 can be estimated consistently by

$$\hat{S}_1(t) = \frac{1}{n} \sum_{i=1}^n \delta_i \mathbf{1}_{\tau_i < Y_i} \mathbf{1}_{Y_i \leq t} \quad \text{and} \quad \hat{S}_2(t) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau_i < t \leq Y_i}.$$

This leads to the following estimator of the survival function,

$$\hat{S}_T(t) = 1 - \hat{F}_T(t) = \prod_{Y_i \leq t} \left(1 - \frac{d\hat{S}_1(Y_i)}{S_2(Y_i)} \right),$$

when there is no ties. In case of ties, let (t_1, \dots, t_k) denote the distinct values taken by $(Y_i)_{1 \leq i \leq n}$, leading to

$$\hat{S}_T(t) = \prod_{t_i \leq t} \left(1 - \frac{d\hat{S}_1(t_i)}{S_2(t_i)} \right).$$

B Application on TPL insurance

B.1 Descriptive statistics for the motor insurance dataset

Depending on the type of the variable, we give different classical indicators: for categorical variables, exposure for each category is provided. Concerning numerical variables, the minimum, the maximum, the median, the mean, and the standard deviation are given.

Variable:	Type	Min.	Median	Mean	Std.	Max.
AccDate	date	07/01/1989	10/01/1994	08/02/1994		01/01/1999
ReportDate	date	09/01/1990	03/01/1995	05/12/1995		02/01/1999
FinDate	date	07/01/1993	01/01/1997	10/11/1996		03/01/1999
Reporting delay	numerical	0	59	113	173	1 430
Claim duration	numerical	0	486	558	381	2 069
Operational time	numerical	0.1	45.9	46.33	27.1	99.1
InjNb	numerical	1	2	2.13	1.37	5
AggClaim	numerical	10	13 854	38 367	90 981	4 485 797
Legal	boolean	No: 8 008	Yes: 14 028			
InjType1	categorical		Fatal: High: Medium: Minor: Severe: Small: Not recorded:			
	frequency		256 189 1 133 15 638 188 3 376 1 256			

Table 3: Descriptive statistics on available information for `ausautoBI8999` dataset.

B.2 Loss triangle: non cumulated payments for RBNS claims on 12/31/1996 (settlement date)

	dev1	dev2	dev3	dev4	dev5	dev6	dev7	dev8	dev9	dev10	dev11	dev12	dev13	dev14	dev15	dev16
01/01/93	109	26992	57383	236791	342703	199194	184154	524182	562563	381179	427698	250028	449081	651059	605897	18632
04/01/93	0	36044	170751	161602	90869	332125	323779	454470	207602	389262	191875	862233	344286	275318	118890	
07/01/93	4856	60631	227743	394798	210000	370244	352686	500941	804994	1876735	330368	234574	848961	357564		
10/01/93	2122	49047	139021	168804	98346	571368	326922	647178	655088	686878	829336	342203	87614			
01/01/94	4932	111734	100917	209810	208715	299321	223234	728861	263702	855347	370049	120106				
04/01/94	2025	15266	235556	372367	240441	371492	792235	136263	651922	603713	175600					
07/01/94	8358	152299	688247	350385	436593	363827	579805	326972	506667	227678						
10/01/94	19386	124653	303847	346686	157300	439116	451604	300136	142562							
01/01/95	61320	102081	253811	165662	154319	413768	491512	686257								
04/01/95	10535	81881	128657	224975	310139	476608	148182									
07/01/95	15125	79378	67132	245390	360690	719365										
10/01/95	902	43090	234033	228407	139189											
01/01/96	0	112493	262520	41871												
04/01/96	5382	91241	52514													
07/01/96	6592	16329														
10/01/96	0															