



HAL
open science

Boosting insights in insurance tariff plans with tree-based machine learning methods

Roel Henckaerts, Marie-Pier Côté, Katrien Antonio, Roel Verbelen

► To cite this version:

Roel Henckaerts, Marie-Pier Côté, Katrien Antonio, Roel Verbelen. Boosting insights in insurance tariff plans with tree-based machine learning methods. *North American Actuarial Journal*, 2021, 25 (2), 10.1080/10920277.2020.1745656 . hal-04015739

HAL Id: hal-04015739

<https://hal.science/hal-04015739v1>

Submitted on 6 Mar 2023

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.

Boosting insights in insurance tariff plans with tree-based machine learning methods

Roel Henckaerts^{2,4}, Marie-Pier Côté¹, Katrien Antonio^{2,3,4}, and Roel Verbelen^{2,4}

¹*École d'actuariat, Université Laval, Canada.*

²*Faculty of Economics and Business, KU Leuven, Belgium.*

³*Faculty of Economics and Business, University of Amsterdam, The Netherlands.*

⁴*LRisk, Leuven Research Center on Insurance and Financial Risk Analysis, KU Leuven, Belgium.*

Abstract

Pricing actuaries typically operate within the framework of generalized linear models (GLMs). With the upswing of data analytics, our study puts focus on machine learning methods to develop full tariff plans built from both the frequency and severity of claims. We adapt the loss functions used in the algorithms such that the specific characteristics of insurance data are carefully incorporated: highly unbalanced count data with excess zeros and varying exposure on the frequency side combined with scarce, but potentially long-tailed data on the severity side. A key requirement is the need for transparent and interpretable pricing models which are easily explainable to all stakeholders. We therefore focus on machine learning with decision trees: starting from simple regression trees, we work towards more advanced ensembles such as random forests and boosted trees. We show how to choose the optimal tuning parameters for these models in an elaborate cross-validation scheme, we present visualization tools to obtain insights from the resulting models and the economic value of these new modeling approaches is evaluated. Boosted trees outperform the classical GLMs, allowing the insurer to form profitable portfolios and to guard against potential adverse risk selection.

Key words: cross-validation, deviance, frequency–severity modeling, gradient boosting machine, interpretable machine learning, model lift

1 Introduction

Insurance companies bring security to society by offering protection against financial losses. They allow individuals to trade uncertainty for certainty, by transferring the risk to the insurer in exchange for a fixed premium. An insurer sets the price for an insurance policy before its actual cost is revealed. Due to this phenomenon, known as the reverse production cycle of the insurance business, it is of vital importance that an insurer properly assesses the risks in its portfolio. To this end, tools from predictive modeling come in handy.

The insurance business is highly data driven and partly relies on algorithms for decision making. In order to price a contract, property and casualty (P&C, or: general, non-life) insurers predict the loss cost y for each policyholder based on his/her observable characteristics \mathbf{x} . The insurer therefore develops a predictive model f , mapping the risk factors \mathbf{x} to the predicted loss cost \hat{y} by setting $\hat{y} = f(\mathbf{x})$. For simplicity, this predictive model is usually built in two stages by considering separately the frequency and severity of the claims. Generalized linear models (GLMs), introduced by [Nelder and Wedderburn \(1972\)](#), are the industry standard to develop state-of-the-art analytic insurance pricing models ([Haberman and Renshaw, 1997](#); [De Jong and Heller, 2008](#); [Frees, 2015](#)). Pricing actuaries often code all risk factors in a categorical format, either based on expert opinions ([Frees and Valdez, 2008](#); [Antonio et al., 2010](#)) or in a data-driven way ([Henckaerts et al., 2018](#)). GLMs involving only categorical risk factors result in predictions available in tabular format, that can easily be translated into interpretable tariff plans.

Technological advancements have boosted the popularity of machine learning and big data analytics, thereby changing the landscape of predictive modeling in many business applications. However, few papers in the insurance literature go beyond the actuarial comfort zone of GLMs. [Dal Pozzolo \(2010\)](#) contrasts the performance of various machine learning techniques to predict claim frequency in the Allstate Kaggle competition. [Guelman \(2012\)](#) compares GLMs and gradient boosted trees for predicting the accident loss cost of auto at-fault claims. [Liu et al. \(2014\)](#) approach the claim frequency prediction problem using multi-class AdaBoost trees. [Wüthrich and Buser \(2019\)](#) and [Zöchbauer et al. \(2017\)](#) show how tree-based machine learning techniques can be adapted to model claim frequencies. [Yang et al. \(2018\)](#) predict insurance premiums by applying a gradient boosted tree algorithm to Tweedie models. [Pesantez-Narvaez et al. \(2019\)](#) employ XGBoost to predict the occurrence of claims using telematics data. [Ferrario et al. \(2018\)](#) and [Schelldorfer and Wüthrich \(2019\)](#) propose neural networks to model claim frequency, either directly or via a nested GLM. Machine learning techniques have also been used in other insurance applications, such as policy conversion or customer retention ([Spedicato et al., 2018](#)), renewal pricing ([Krasheninnikova et al., 2019](#)) and claim fraud detection ([Wang and Xu, 2018](#)).

Insurance pricing models are heavily regulated and they must meet specific requirements before being deployed in practice, posing some challenges for machine learning algorithms. Firstly, the European Union’s General Data Protection Regulation ([GDPR, 2016](#)), effective May 25, 2018, establishes a regime of “algorithmic accountability” of decision-making machine algorithms. By law, individuals have the right to an explanation of the logic behind the decision ([Kaminski, 2018](#)), which means that pricing models must be transparent and easy to communicate to all stakeholders. Qualified transparency ([Pasquale, 2015](#)) implies that customers, managers and the regulator should receive information in different degrees of scope and depth. Secondly, every policyholder should be charged a fair premium, related to his/her risk profile, to minimize the potential for adverse selection ([Dionne et al., 1999](#)). If the heterogeneity in the portfolio is not carefully reflected in the pricing, the good risks will be prompt to lapse and accept a lower premium elsewhere, leaving the insurer with an inadequately priced portfolio. Thirdly, the insurer has the social role of creating solidarity among the policyholders. The use of machine learning for pricing should in no way lead to an extreme “personalization of risk” or discrimination, e.g., in the form of extremely high premiums ([O’Neil, 2017](#)) for some risk profiles that actually entail no risk transfer. By finding a trade-off between customer segmentation and risk pooling, the insurer avoids adverse selection while offering an effective insurance product involving a risk transfer for all policyholders. In a regime of algorithmic accountability, insurers should be held responsible for their pricing models in terms of transparency, fairness and solidarity. It is therefore very important to be able to “look under the hood” of machine learning algorithms and the resulting pricing models. That is exactly one of the goals of this paper.

In this paper, we study how tree-based machine learning methods can be applied to insurance pricing. The building blocks of these techniques are decision trees, covered in [Friedman et al. \(2001\)](#). These are simple predictive models that mimic human decision-making in the form of yes-no questions. In insurance pricing, a decision tree partitions a portfolio of policyholders into groups of homogeneous risk profiles based on some characteristics. The partition of the portfolio is directly observable, resulting in high transparency. For each subgroup, a constant prediction is put forward, automatically inducing solidarity among the policyholders in a subgroup (as long as the size of this subgroup is large enough). These aspects of decision trees make them good candidates for insurance pricing. However, the predictive performance of such simple trees tends to be rather low. We therefore consider more complex algorithms that combine multiple decision trees in an ensemble, i.e., tree-based machine learning. These ensemble techniques usually provide better predictive performance, but at the cost of less transparency. We employ model interpretation tools to understand these “black boxes”, allowing us to interpret the

resulting models and underlying decision process. Tree-based machine learning techniques are often praised for their ability to discover interaction effects in the data, a very useful insight for insurers that will be explored in this paper.

Insurance claim data typically entails highly imbalanced count data with excess zeros and varying exposure-to-risk on the frequency side, combined with long- or even heavy-tailed continuous data on the severity side. Standard machine learning algorithms typically deal with data that is more normal-like or balanced. [Guelman \(2012\)](#) models the accident loss cost by simplifying the frequency count regression problem into a binary classification task. This however cannot factor in varying exposure-to-risk and leads to a loss of information regarding policyholders who file more than one claim in a period. [Wüthrich and Buser \(2019\)](#) and [Zöchbauer et al. \(2017\)](#) show how the specific data features on the frequency side can be taken into account for regression.

We extend the existing literature by also putting focus on the severity side of claims and obtaining full tariff plans on real-world claims data from an insurance portfolio. We develop an elaborate cross-validation scheme instead of relying on built-in routines from software packages and we take into account multiple types of risk factors: categorical, continuous and spatial information. The goal of this paper is to investigate how tree-based pricing models perform compared to the classical actuarial approach with GLMs. This comparison puts focus on statistical performance, interpretation and business implications. We go beyond a purely statistical comparison, but acknowledge the fact that the resulting pricing model has to be deployed, after marketing adjustments, in a business environment with specific requirements.

The rest of this paper is structured as follows. Section 2 introduces the basic principles and guidelines for building a benchmark pricing GLM. Section 3 consolidates the important technical details on tree-based machine learning. Section 4 presents interpretations from the optimal frequency and severity models fitted on a Belgian insurance data set, together with an out-of-sample performance comparison. Section 5 reviews the added value from a business angle and Section 6 concludes this paper. In an accompanying online supplement, available at <https://github.com/henckr/sevtree>, we provide more details on the construction and interpretation of tree-based machine learning methods for the severity.

2 State-of-the-art insurance pricing models

To assess the possible merits of tree-based machine learning for insurance pricing, we first have to establish a fair benchmark pricing model that meets industry standards. GLMs are by far the most popular pricing models in today's industry. This section outlines the basic principles and steps for creating a benchmark pricing GLM with the strategy from [Henckaerts et al. \(2018\)](#).

A P&C insurance company is interested in the total loss amount L per unit of exposure-to-risk e , where L is the total loss for the N claims reported by a policyholder during the exposure period e . P&C insurers usually opt for a so-called frequency-severity strategy to price a contract ([Denuit et al., 2007](#); [Frees et al., 2014](#); [Parodi, 2014](#)). Claim frequency F is the number of claims N filed per unit of exposure-to-risk e . Claim severity S refers to the cost per claim and is defined by the average amount per claim filed, that is the total loss amount L divided by the number of claims N . The technical price π (or: pure/risk premium) then follows as:

$$\pi = \mathbb{E}\left(\frac{L}{e}\right) \stackrel{\text{indep.}}{=} \mathbb{E}\left(\frac{N}{e}\right) \times \mathbb{E}\left(\frac{L}{N} \mid N > 0\right) = \mathbb{E}(F) \times \mathbb{E}(S),$$

assuming independence between the frequency and the severity component of the premium ([Klugman et al., 2012](#)). Alternatives, allowing dependence between F and S , are investigated in the literature ([Gschlößl and Czado, 2007](#); [Czado et al., 2012](#); [Garrido et al., 2016](#)).

Predictive models for both F and S are typically developed within the framework of GLMs. Let Y , the response variable of interest, follow a distribution from the exponential family. The structure of a GLM with all explanatory variables in a categorical format is:

$$\eta = g(\mu) = \mathbf{z}^\top \boldsymbol{\beta} = \beta_0 + \sum_{j=1}^q \beta_j z_j, \quad (1)$$

with η the linear predictor, $g(\cdot)$ the link function and μ the expectation of Y . The $q + 1$ dimensional 0/1-valued vector \mathbf{z} contains a 1 for the intercept together with the q dummy variables expressing the policyholder's risk profile. In a claim frequency model, the response variable N typically follows a count distribution such as the Poisson. Assuming $g(\cdot) = \ln(\cdot)$, the model may account for exposure-to-risk through an offset $\ln(e)$ such that the risk premium is proportional to the exposure. In a claim severity model, the response variable L/N typically follows a right skewed distribution with a long right tail such as the gamma or log-normal. Only policyholders filing at least one claim, i.e., $N > 0$, contribute to the severity model calibration and the number of claims N is used as a case weight in the regression (Denuit and Lang, 2004).

Henckaerts et al. (2018) details a data-driven strategy to build a GLM with all risk factors in a categorical format. This strategy aligns the practical requirements of a business environment with the statistical flexibility of generalized additive models (GAMs, documented by Wood, 2006). GAMs extend the linear predictor in Eq. (1) with (multidimensional) smooth functions. After an exhaustive search, the starting point for both frequency and severity is a flexible GAM with smooth effects for continuous risk factors, including two-way interactions, and a smooth spatial effect. These smooth effects are used to bin the continuous and spatial risk factors, thereby transforming them to categorical variables. Figure 1 schematizes how decision trees and unsupervised clustering are applied to achieve this binning. The output of this framework is an interpretable pricing GLM, which serves as benchmark pricing model in this study.

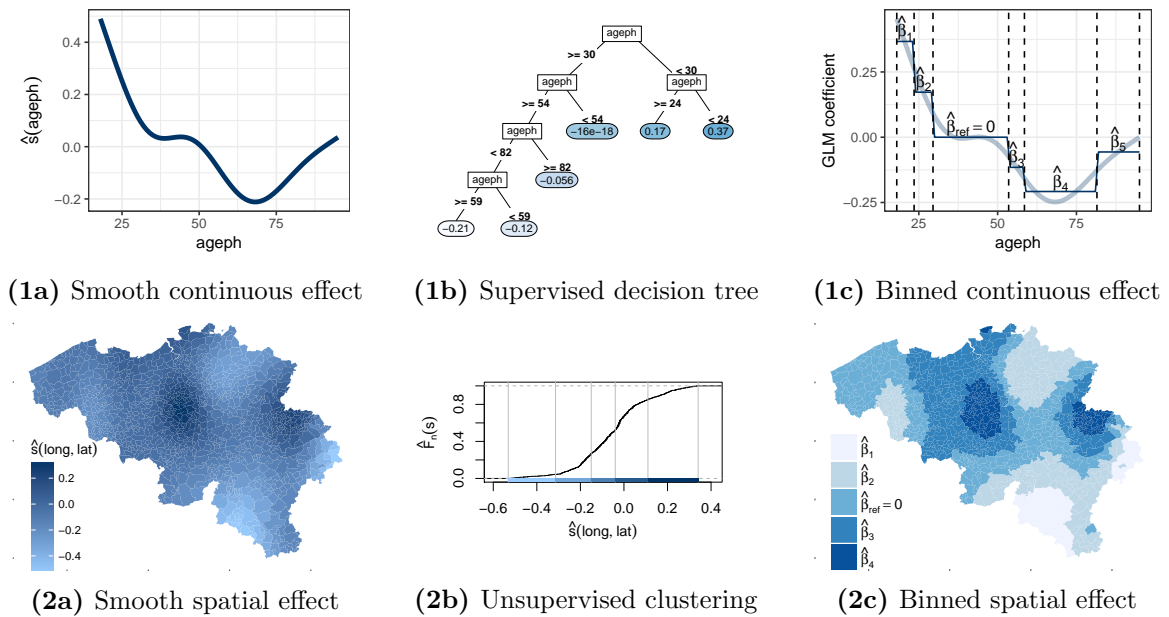


Figure 1: Schematic overview of the binning strategy of Henckaerts et al. (2018) for a continuous risk factor (1a-1c) and a spatial risk factor (2a-2c). The smooth GAM effect of a continuous risk factor (1a) is fed as a response to a decision tree (1b), which splits the continuous risk factor into bins (1c). The smooth spatial effect (2a) is clustered in an unsupervised way (2b), resulting in groups of postcode areas (2c). These categorical risk factors are used in a GLM.

3 Tree-based machine learning methods for insurance pricing

Section 3.1 introduces the essential algorithmic details needed for understanding the tree-based modeling techniques used in this paper. We consider regression trees (Breiman et al., 1984), random forests (Breiman, 2001) and gradient boosting machines (Friedman, 2001) as alternative modeling techniques for insurance pricing. These models rely on the choice of a loss function, which has to be tailored to the characteristics of insurance data as we motivate in Section 3.2. In Section 3.3 and 3.4, we explain our tuning strategy and present interpretation tools.

3.1 Algorithmic essentials

Regression tree Decision trees partition data based on yes-no questions, predicting the same value for each member of the constructed subsets. A popular approach to construct decision trees is the Classification And Regression Tree (CART) algorithm, introduced by Breiman et al. (1984). The predictor space R is the set of possible values for the p variables x_1, \dots, x_p , e.g., $R = \mathbb{R}^p$ for p unbounded, continuous variables and $R = [\min(x_1), \max(x_1)] \times [\min(x_2), \max(x_2)]$ for two bounded, continuous variables. A tree divides the predictor space R into J distinct, non-overlapping regions R_1, \dots, R_J . In the j th region, the fitted response \hat{y}_{R_j} is computed as a (weighted) average of the training observations falling in that region. The regression tree predicts a (new) observation with characteristics \mathbf{x} as follows:

$$f_{\text{tree}}(\mathbf{x}) = \sum_{j=1}^J \hat{y}_{R_j} \mathbb{1}(\mathbf{x} \in R_j). \quad (2)$$

The indicator $\mathbb{1}(A)$ equals one if event A occurs and zero otherwise. As the J regions are non-overlapping, the indicator function differs from zero for exactly one region for each \mathbf{x} . A tree therefore makes the same constant prediction \hat{y}_{R_j} for the entire region R_j .

It is computationally impractical to consider every possible partition of the predictor space R in J regions, therefore CART uses a top-down greedy approach known as recursive binary splitting. From the full predictor space R , the algorithm selects a splitting variable x_v with $v \in \{1, \dots, p\}$ and a cut-off c such that $R = R_1(v, c) \cup R_2(v, c)$ with $R_1(v, c) = \{R \mid x_v \leq c\}$ and $R_2(v, c) = \{R \mid x_v > c\}$. This forms two nodes in the tree, one containing the observations satisfying $x_v \leq c$ and the other containing the observations satisfying $x_v > c$. For a categorical splitting variable, the corresponding factor levels are replaced by their empirical response averages, see Section 8.8 in Breiman et al. (1984). These averages are sorted from low to high and a cut-off c is chosen such that the factor levels are split into two groups. The splitting variable x_v and cut-off c are chosen such that their combination results in the largest improvement in a carefully picked loss function $\mathcal{L}(\cdot, \cdot)$. For $i = \{1, \dots, n\}$, where n is the number of observations in the training set, the CART algorithm searches for x_v and c minimizing the following summations:

$$\sum_{i: \mathbf{x}_i \in R_1(v, c)} \mathcal{L}(y_i, \hat{y}_{R_1}) + \sum_{i: \mathbf{x}_i \in R_2(v, c)} \mathcal{L}(y_i, \hat{y}_{R_2}).$$

A standard loss function is the squared error loss, but we present more suitable loss functions for claim frequency or severity data in Section 3.2. In a next iteration, the algorithm splits R_1 and/or R_2 in two regions and this process is repeated recursively until a stopping criterion is satisfied. This stopping criterion typically puts a predefined limit on the size of a tree, e.g., a minimum improvement in the loss function (Breiman et al., 1984), a maximum depth of the tree or a minimum number of observations in a node of the tree (Friedman et al., 2001).

A large tree is likely to overfit the data and does not generalize well to new data, while a small tree is likely to underfit the data and fails to capture the general trends. This is related to the bias-variance tradeoff (Friedman et al., 2001) meaning that a large tree has low bias and high variance while a small tree has high bias but low variance. To prevent overfitting, the performance of a tree is penalized by the number of regions J as follows:

$$\sum_{j=1}^J \sum_{i: \mathbf{x}_i \in R_j} \mathcal{L}(y_i, \hat{y}_{R_j}) + J \cdot cp \cdot \sum_{i: \mathbf{x}_i \in R} \mathcal{L}(y_i, \hat{y}_R), \quad (3)$$

where the first part assesses the goodness of fit and the second part is a penalty measuring the tree complexity. The strength of this penalty is driven by the complexity parameter cp , a tuning parameter (see Section 3.3 for details on the tuning strategy). A large (small) value for cp puts a high (low) penalty on extra splits and will result in a small (large) tree. The complexity parameter cp is usually scaled with the loss function evaluated for the root tree, which is exactly the last summation in Eq. (3); see Remark 3.8 in Zöchbauer et al. (2017). This ensures that $cp = 1$ delivers a root tree without splits capturing an overall y estimate (denoted \hat{y}_R in Eq. (3)) and $cp = 0$ results in the largest possible tree allowed by the stopping criterion.

Figure 2 depicts an example of a regression tree for claim frequency data. The rectangles are internal nodes which partition observations going from top to bottom along the tree. The top node, splitting on the bonus-malus level bm , is called the root node. The ellipses are leaf nodes, containing prediction values for the observations ending in that specific leaf. Going from left to right, the leaf nodes are ordered from low (light blue) to high (dark blue) prediction values. Decision trees have many advantages because their predictions are highly explainable and interpretable, both very important criteria for regulators. The downside of trees however is that the level of predictive accuracy tends to be lower compared to other modeling techniques. This is mainly driven by the high variance of a tree, e.g., slight changes in the data can result in very different trees and therefore rather different predictions for certain observations. The predictive performance can be substantially improved by aggregating multiple decision trees in ensembles of trees, thereby reducing the variance. That is the idea behind popular ensemble methods, such as random forests and gradient boosting machines, which are discussed next.

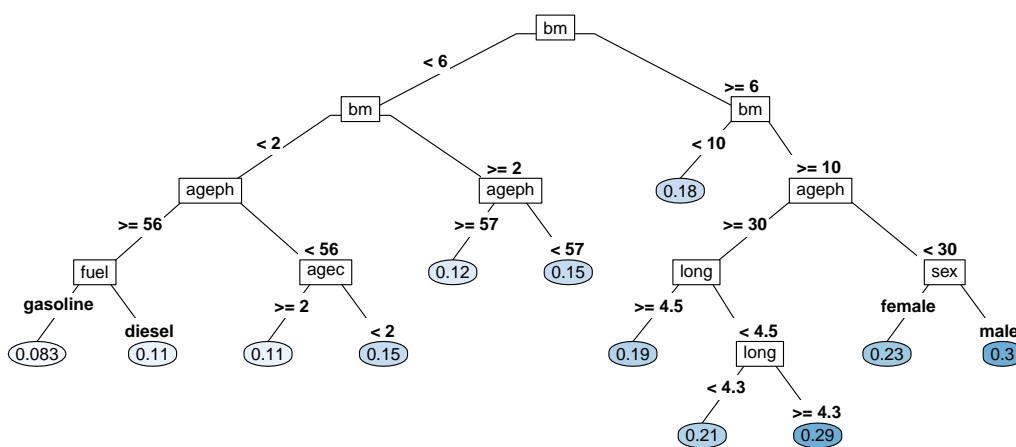


Figure 2: Visual representation of a regression tree for claim frequency with nodes (rectangles) containing the splitting variable x_v , edges (lines) representing the splits with cut-off c and leaf nodes (ellipses) containing the prediction values \hat{y}_{R_j} . The variable names are defined in Table 7.

Random forest Bagging, which is short for **bootstrap aggregating** (Breiman, 1996), and random forests (Breiman, 2001) are similar ensemble techniques combining multiple decision trees. Bagging reduces the variance of a single tree by averaging the forecasts of multiple trees on bootstrapped samples of the original data. This stabilizes the prediction and improves the predictive performance compared to a single decision tree. Starting from the data set \mathcal{D} , the idea of bagging is to take bootstrap samples $\{\mathcal{D}_t\}_{t=1,\dots,T}$ and to build T decision trees, one for each \mathcal{D}_t independently. The results are then aggregated in the following way:

$$f_{\text{bagg}}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T f_{\text{tree}}(\mathbf{x} | \mathcal{D}_t), \quad (4)$$

where the condition $(| \mathcal{D}_t)$ indicates that the tree was developed on the sample \mathcal{D}_t .

The performance improvement through variance reduction gets bigger when there is less correlation between the individual trees, see Lemma 3.25 in Zöchbauer et al. (2017). For that reason, the trees are typically grown deep (i.e., $cp = 0$ in Eq. (3)), until a stopping criterion is satisfied. Taking bootstrap samples of smaller sizes $\delta \cdot n$, with n the number of observations in \mathcal{D} and $0 < \delta < 1$, decorrelates the trees further and reduces the model training time. However, a lot of variability remains within a bagged ensemble because the trees built on the bootstrapped data samples are still quite similar. This is especially the case when some explanatory variables in the data are much more predictive than the others. The important variables will dominate the first splits, causing all trees to be similar to one another. To prevent this, a random forest further decorrelates the individual trees by sampling variables during the growing process. At each split, m out of p variables are randomly chosen as candidates for the optimal splitting variable. Besides this adaptation, a random forest follows the same strategy as bagging and predicts a new observation according to Eq. (4). The random forest procedure is detailed in Algorithm 1 where T and m are treated as tuning parameters (see Section 3.3 for details on the tuning strategy).

for $t = 1, \dots, T$ **do**

generate bootstrapped data \mathcal{D}_t of size $\delta \cdot n$ by sampling with replacement from data \mathcal{D} ;

while *stopping criterion not satisfied* **do**

randomly select m of the p variables;

find the optimal splitting variable x_v from the m options together with cut-off c ;

$f_{\text{rf}}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T f_{\text{tree}}(\mathbf{x} | \mathcal{D}_t)$;

Algorithm 1: Procedure to build a random forest model.

A random forest improves the predictive accuracy obtained with a single decision tree by using more, and hopefully slightly different, trees to solve the problem at hand. However, the trees in a random forest are built independently from each other (i.e., the **for** loop in Algorithm 1 can be run in parallel) and do not share information during the training process.

Gradient boosting machine In contrast to random forests, boosting is an iterative statistical method that combines many weak learners into one powerful predictor. Friedman (2001) introduced decision trees as weak learners; each tree improves the current model fit, thereby using information from previously grown trees. At each iteration, the pseudo-residuals are used to assess the regions of the predictor space for which the model performs poorly in order to improve the fit in a direction of better overall performance. The pseudo-residual $\rho_{i,t}$ for observation i in iteration t is calculated as the negative gradient of the loss function $-\partial \mathcal{L}\{y_i, f(\mathbf{x}_i)\} / \partial f(\mathbf{x}_i)$, evaluated at the current model fit. This typical approach called stepwise gradient descent ensures that a lower loss is obtained at the next iteration, until convergence. The boosting

method learns slowly by fitting a small tree of depth d (with a squared error loss function) to these pseudo-residuals, improving the model fit in areas where it does not perform well. For each region R_j of that tree, the update \hat{b}_j is calculated as the constant that has to be added to the previous model fit to minimize the loss function, namely b that minimizes $\mathcal{L}\{y_i, f(\mathbf{x}_i) + b\}$ over this region. A shrinkage parameter λ controls the learning speed by shrinking updates for $\mathbf{x} \in R_j$ as follows: $f_{new}(\mathbf{x}) = f_{old}(\mathbf{x}) + \lambda \cdot \hat{b}_j$. A lower λ usually results in better performance but also increases computation time because more trees are needed to converge to a good solution. Typically, λ is fixed at the lowest possible value within the computational constraints (Friedman, 2001). The collection of T trees at the final iteration is used to make predictions.

Stochastic gradient boosting, introduced by Friedman (2002), injects randomness in the training process. In each iteration, the model update is computed from a randomly selected subsample of size $\delta \cdot n$. This improves both the predictive accuracy and model training time when $\delta < 1$. The (stochastic) gradient boosting machine algorithm is given in Algorithm 2 where T and d are treated as tuning parameters (see Section 3.3 for details on the tuning strategy).

```

initialize fit to the optimal constant model:  $f_0(\mathbf{x}) = \arg \min_b \sum_{i=1}^n \mathcal{L}(y_i, b)$ ;
for  $t = 1, \dots, T$  do
    randomly subsample data of size  $\delta \cdot n$  without replacement from data  $\mathcal{D}$ ;
    for  $i = 1, \dots, \delta \cdot n$  do
         $\rho_{i,t} = - \left[ \frac{\partial \mathcal{L}\{y_i, f(\mathbf{x}_i)\}}{\partial f(\mathbf{x}_i)} \right]_{f=f_{t-1}}$ 
    fit a tree of depth  $d$  to the pseudo-residuals  $\rho_{i,t}$  resulting in regions  $R_{j,t}$  for  $j = 1, \dots, J_t$ ;
    for  $j = 1, \dots, J_t$  do
         $\hat{b}_{j,t} = \arg \min_b \sum_{i: \mathbf{x}_i \in R_{j,t}} \mathcal{L}\{y_i, f_{t-1}(\mathbf{x}_i) + b\}$ 
    update  $f_t(\mathbf{x}) = f_{t-1}(\mathbf{x}) + \lambda \sum_{j=1}^{J_t} \hat{b}_{j,t} \mathbb{1}(\mathbf{x} \in R_{j,t})$ ;
 $f_{\text{gbm}}(\mathbf{x}) = f_T(\mathbf{x})$ ;

```

Algorithm 2: Procedure to build a (stochastic) gradient boosting machine.

3.2 Loss functions for insurance data

The machine learning algorithms discussed in Section 3.1 require the specification of a loss (or: cost) function that is to be minimized during the training phase of the model. We first present a general discussion on the loss function choice, followed by details on the R implementation.

Loss functions The standard loss function for regression problems is the squared error loss:

$$\mathcal{L}\{y_i, f(\mathbf{x}_i)\} \propto \{y_i - f(\mathbf{x}_i)\}^2,$$

where y_i is the observed response and $f(\mathbf{x}_i)$ is the prediction of the model for variables \mathbf{x}_i . However, the squared error is not necessarily a good choice when modeling integer-valued frequency data or right-skewed severity data. We use the concept of deviance to make this idea clear. The deviance is defined as $D\{y, f(\mathbf{x})\} = -2 \cdot \ln[\mathcal{L}\{f(\mathbf{x})\} / \mathcal{L}(y)]$, a likelihood ratio where $\mathcal{L}\{f(\mathbf{x})\}$ is the model likelihood and $\mathcal{L}(y)$ the likelihood of the saturated model (i.e., the model in which the number of parameters equals the number of observations). The condition $\mathcal{L}\{f(\mathbf{x})\} \leq \mathcal{L}(y)$ always holds, so the ratio of likelihoods is bounded from above by one. For competing model fits, the best one obtains the lowest deviance value on holdout data. We therefore use a loss function $\mathcal{L}(\cdot, \cdot)$ such that $D\{y, f(\mathbf{x})\} = \sum_{i=1}^n \mathcal{L}\{y_i, f(\mathbf{x}_i)\}$. This idea was put forward by Venables and Ripley (2002) for general classification and regression problems.

Assuming constant variance, the normal (or: Gaussian) deviance can be expressed as follows:

$$\begin{aligned} D\{y, f(\mathbf{x})\} &= 2 \ln \prod_{i=1}^n \exp \left\{ -\frac{1}{2\sigma^2} (y_i - y_i)^2 \right\} - 2 \ln \prod_{i=1}^n \exp \left[-\frac{1}{2\sigma^2} \{y_i - f(\mathbf{x}_i)\}^2 \right] \\ &= \frac{1}{\sigma^2} \sum_{i=1}^n \{y_i - f(\mathbf{x}_i)\}^2, \end{aligned}$$

which boils down to a scaled version of the sum of squared errors. This implies that a loss function based on the squared error is appropriate when the normal assumption is reasonable. More generally, the squared error is suitable for any continuous distribution symmetrical around its mean with constant variance, i.e., any elliptical distribution. However, claim frequency and severity data do not follow any elliptical distribution, as we show in Section 4.1. Therefore, in an actuarial context, [Wüthrich and Buser \(2019\)](#) and [Zöchbauer et al. \(2017\)](#) propose more suitable loss functions inspired by the GLM pricing framework from Section 2.

Claim frequency modeling involves count data, typically assumed to be Poisson distributed in GLMs. Therefore, an appropriate loss function is the Poisson deviance, defined as follows:

$$\begin{aligned} D(y, f(\mathbf{x})) &= 2 \ln \prod_{i=1}^n \exp(-y_i) \frac{y_i^{y_i}}{y_i!} - 2 \ln \prod_{i=1}^n \exp\{-f(\mathbf{x}_i)\} \frac{f(\mathbf{x}_i)^{y_i}}{y_i!} \\ &= 2 \sum_{i=1}^n \left[y_i \ln \frac{y_i}{f(\mathbf{x}_i)} - \{y_i - f(\mathbf{x}_i)\} \right]. \end{aligned} \quad (5)$$

When using an exposure-to-risk measure e_i , $f(\mathbf{x}_i)$ is replaced by $e_i \cdot f(\mathbf{x}_i)$ such that the exposure is taken into account in the expected number of claims. Thus, the Poisson deviance loss function can account for different policy durations. Predictions from a Poisson regression tree in Eq. (2) are equal to the sum of the number of claims divided by the sum of exposure for all training observations in each leaf node: $\hat{y}_{R_j} = \sum_{i \in I_j} N_i / \sum_{i \in I_j} e_i$ for $I_j = \{i : \mathbf{x}_i \in R_j\}$. This optimal estimate is obtained by setting the derivative of Eq. (5) with respect to f equal to zero. As a tree node without claims leads to a division by zero in the deviance calculation, an adjustment can be made to the implementation with a hyper-parameter that will be introduced in Section 3.3.

Right-skewed and long-tailed severity data is typically assumed to be gamma or log-normally distributed in GLMs. In Section 4, we present the results obtained with the gamma deviance as our preferred model choice, but a discussion on the use of the log-normal deviance is available in the supplementary material. The gamma deviance is defined as follows:

$$\begin{aligned} D\{y, f(\mathbf{x})\} &= 2 \ln \prod_{i=1}^n \frac{1}{y_i \Gamma(\alpha)} \left(\frac{\alpha y_i}{y_i} \right)^\alpha \exp \left(-\frac{\alpha y_i}{y_i} \right) - 2 \ln \prod_{i=1}^n \frac{1}{y_i \Gamma(\alpha)} \left\{ \frac{\alpha y_i}{f(\mathbf{x}_i)} \right\}^\alpha \exp \left\{ -\frac{\alpha y_i}{f(\mathbf{x}_i)} \right\} \\ &= 2 \sum_{i=1}^n \alpha \left\{ \frac{y_i - f(\mathbf{x}_i)}{f(\mathbf{x}_i)} - \ln \frac{y_i}{f(\mathbf{x}_i)} \right\}. \end{aligned} \quad (6)$$

The shape parameter α acts as a scaling factor and can therefore be ignored. When dealing with case weights, α can be replaced by the weights w_i . In severity modeling, the response is the average claim amount of the N_i observed claims and the number of claims N_i should be used as case weight. Predictions from a gamma regression tree in Eq. (2) are equal to the sum of the total loss amount divided by the sum of the number of claims for all training observations in each leaf node: $\hat{y}_{R_j} = \sum_{i \in I_j} L_i / \sum_{i \in I_j} N_i$ for $I_j = \{i : \mathbf{x}_i \in R_j\}$. This optimal estimate is obtained by setting the derivative of Eq. (6) with respect to f equal to zero.

Implementation in R Our results are obtained with two special purpose packages for tree-based machine learning in the statistical software R. For the regression trees and random forests, we developed our own package called `distRforest` (Henckaerts, 2019). For stochastic gradient boosting, we chose the implementation from Southworth (2015) of the `gbm` package, originally developed by Ridgeway (2014). Our `distRforest` package extends the `rpart` package by Therneau and Atkinson (2018) such that it is capable of developing regression trees and random forests with our specific desired loss functions for both claim frequency and severity. We had to go beyond the standard implementations especially because of the loss functions appropriate for actuarial applications. Although the `rpart` package supports the Poisson deviance for regression trees, it did not facilitate the use of a suitable loss function for severity data.

3.3 Tuning strategy

Tuning and hyper-parameters Table 1 gives an overview of the parameters used by the algorithms described in Section 3.1. Some of these are chosen with care (tuning parameters), while others are less influential and are set to a sensible predetermined value (hyper-parameters). Instead of relying on the built-in tuning strategies of the R packages mentioned in Section 3.2, we perform an extensive grid search to find the optimal values among a predefined tuning grid displayed in Table 8 in Appendix B. We prefer a grid search above other tuning strategies, such as Bayesian optimization (Xia et al., 2017), for its ease of implementation while being a sound approach. The hyper-parameter κ enforces a stopping criterion for trees used across the three algorithms, ensuring that a split is not allowed if a resulting node would contain less than 1% of the observations. The hyper-parameter δ in Algorithms 1 and 2 specifies to develop the trees in the ensemble techniques on 75% of the available training data. The shrinkage parameter λ in Algorithm 2 is set at a low value for which computation time is still reasonable, namely 0.01. The parameter γ helps to avoid division by zero when optimizing the Poisson deviance in Eq. (5). This parameter is therefore only used when growing a regression tree and random forest for claim frequency. We refer the reader to Section 8.2 in Therneau and Atkinson (2019) for details on the `rpart` implementation. In short, a gamma prior is assumed on the Poisson rate parameter to keep it from becoming zero when there is no claim in a node. With $I_j = \{i : \mathbf{x}_i \in R_j\}$, the prediction in a node is adapted as follows:

$$\hat{y}_{R_j}^\gamma = \frac{\gamma^{-2} + \sum_{i \in I_j} N_i}{\gamma^{-2}/\hat{y}_R + \sum_{i \in I_j} e_i}.$$

Note that $\hat{y}_{R_j}^\gamma = \hat{y}_R$ for $\gamma = 0$ and $\hat{y}_{R_j}^\gamma = \hat{y}_{R_j} = \sum_{i \in I_j} N_i / \sum_{i \in I_j} e_i$ for $\gamma = \infty$.

	Tuning parameters	Hyper-parameters
Regression tree	complexity parameter cp coefficient of variation gamma prior γ	$\kappa = 0.01$
Random forest	number of trees T number of split candidates m	$cp = 0$ $\gamma = 0.25$ $\kappa = 0.01$ $\delta = 0.75$
Gradient boosting machine	number of trees T tree depth d	$\lambda = 0.01$ $\kappa = 0.01$ $\delta = 0.75$

Table 1: Overview of the tuning and hyper-parameters for the different machine learning techniques.

Cross-validation Machine learning typically relies on training data to build a model, validation data to tune the parameters and test data to evaluate the out-of-sample performance of the model. In this paper, we develop an extensive cross-validation scheme, inspired by K -fold cross-validation (Friedman et al., 2001), that serves two purposes. First, we tune the parameters in the algorithms under study with a 5-fold cross-validation approach. Second, we evaluate the predictive performance of the algorithms investigated on multiple data sets, instead of on a single test set. Algorithm 3 outlines the basic principles of our approach and Figure 3 gives a schematic representation. The full data \mathcal{D} is split in six disjoint and stratified (Neyman, 1934) subsets $\mathcal{D}_1, \dots, \mathcal{D}_6$ by ordering first on claim frequency, then on severity. The ordered observations are assigned to each of the subsets in turn. Stratification ensures that the distribution of response variables is similar in the six subsets, as we illustrate in Table 2 for the data introduced in Section 4.1. The **foreach** and inner **for** loop in Algorithm 3 represent the typical approach to perform 5-fold cross-validation on data from which we already separated a hold-out test set \mathcal{D}_k . The **foreach** loop iterates over the tuning grid and the **for** loop allows the validation set \mathcal{D}_ℓ to vary. The optimal tuning parameters are those that minimize the cross-validation error, which is obtained by averaging the error on the validation sets. The outer **for** loop in Algorithm 3 allows the hold-out test set to vary and model performance is evaluated on this test set \mathcal{D}_k . Advantages of evaluating a trained model on multiple test sets are threefold. First, we obtain multiple performance measures per model class which results in a more accurate performance assessment. Second, it allows to perform sensitivity checks to assess the stability of different algorithms. Third, it exempts us from the choice of a specific test set which could bias results.

Input: model class (**mclass**) and corresponding tuning grid (**tgrid**)
split data \mathcal{D} into 6 disjoint stratified subsets $\mathcal{D}_1, \dots, \mathcal{D}_6$;
for $k = 1, \dots, 6$ **do**
 leave out \mathcal{D}_k as test set;
 foreach parameter combination in **tgrid** **do**
 for $\ell \in \{1, \dots, 6\} \setminus k$ **do**
 train a model $f_{k\ell}$ of **mclass** on $\mathcal{D} \setminus \{\mathcal{D}_k, \mathcal{D}_\ell\}$;
 evaluate the model performance on \mathcal{D}_ℓ using loss function $\mathcal{L}(\cdot, \cdot)$;
 valid_error $_{k\ell} \leftarrow \frac{1}{|\mathcal{D}_\ell|} \sum_{i \in \mathcal{D}_\ell} \mathcal{L}\{y_i, f_{k\ell}(\mathbf{x}_i)\}$;
 valid_error $_k \leftarrow \frac{1}{5} \sum_{\ell \in \{1, \dots, 6\} \setminus k} \text{valid_error}_{k\ell}$;
 optimal parameters from **tgrid** are those that minimize valid_error $_k$;
 train a model f_k of **mclass** on $\mathcal{D} \setminus \mathcal{D}_k$ using the optimal parameters;
 evaluate the model performance on \mathcal{D}_k using loss function $\mathcal{L}(\cdot, \cdot)$;
 test_error $_k \leftarrow \frac{1}{|\mathcal{D}_k|} \sum_{i \in \mathcal{D}_k} \mathcal{L}\{y_i, f_k(\mathbf{x}_i)\}$;

Output: optimal tuning parameters + performance measure for each of the six folds.

Algorithm 3: Cross-validation scheme for model tuning and performance evaluation.

	\mathcal{D}_1	\mathcal{D}_2	\mathcal{D}_3	\mathcal{D}_4	\mathcal{D}_5	\mathcal{D}_6
$\sum N_i / \sum e_i$	0.1391687	0.1391433	0.1392443	0.1392213	0.1391517	0.1393045
$\sum L_i / \sum N_i$	1296.165	1302.894	1324.667	1312.619	1330.884	1287.832

Table 2: Summary statistics of response variables in the different data subsets \mathcal{D}_1 to \mathcal{D}_6 .




Figure 3: Graphical representation of the cross-validation scheme. The holdout test set for data fold k is \mathcal{D}_k , indicated in red. Within data fold k , we tune the parameters by 5-fold cross-validation on $\mathcal{D} \setminus \mathcal{D}_k$ with the validation sets \mathcal{D}_ℓ in green and the training data $\mathcal{D} \setminus \{\mathcal{D}_k, \mathcal{D}_\ell\}$ in blue. After tuning, we train the model on $\mathcal{D} \setminus \mathcal{D}_k$ using the optimal parameters for data fold k .

3.4 Interpretability matters: opening the black box

The GDPR’s regime of “algorithmic accountability” and the resulting “right to explanation” highlight the vital importance of interpretable and transparent pricing models. However, machine learning techniques are often considered black boxes compared to statistical models such as GLMs. In a GLM, parameter estimates and their standard errors give information about the effect, uncertainty and statistical relevance of all variables. Such quick and direct interpretations are not possible with machine learning techniques, but this section introduces tools to gain insights from a model. A good source on interpretable machine learning is [Molnar \(2019\)](#). These tools are evaluated on the data used to train the optimal models, i.e., $\mathcal{D} \setminus \mathcal{D}_k$ for data fold k . A subset of the training data can be used to save computation time if needed.

Variable importance Variable selection and model building is often a time consuming and tedious process with GLMs ([Henckaerts et al., 2018](#)). An advantage of tree-based techniques is their built-in variable selection strategy, making a priori design decisions less critical. Unraveling the variables that actually matter in the prediction is thus crucial. For $\ell \in \{1, \dots, p\}$, [Breiman et al. \(1984\)](#) measure the importance of a specific feature x_ℓ in a decision tree t by summing the improvements in the loss function over all the splits on x_ℓ :

$$\mathcal{I}_\ell(t) = \sum_{j=1}^{J-1} \mathbb{1}\{v(j) = \ell\} (\Delta\mathcal{L})_j.$$

The sum is taken over all $J - 1$ internal nodes of the tree, but only the nodes where the splitting variable x_v is x_ℓ contribute to this sum. These contributions $(\Delta\mathcal{L})_j$ represent the

difference between the evaluated loss function before and after split j in the tree. The idea is that important variables appear often and high in the decision tree such that the sum grows largest for those variables. We normalize these variable importance values such that they sum to 100%, giving a clear idea about the relative contribution of each variable in the prediction.

We can easily generalize this approach to the ensemble techniques by averaging the importance of variable x_ℓ over the different trees that compose the ensemble:

$$\mathcal{I}_\ell = \frac{1}{T} \sum_{t=1}^T \mathcal{I}_\ell(t),$$

where the sum is taken over all trees in the random forest or gradient boosting machine.

Partial dependence plots Besides knowing which variables are important, it is meaningful to understand their effect on the prediction target. Partial dependence plots, introduced in [Friedman \(2001\)](#), show the marginal effect of a variable on the predictions obtained from a model. Hereto, we evaluate the prediction function in specific values of the variable of interest x_ℓ for $\ell \in \{1, \dots, p\}$, while averaging over a range of values of the other variables \mathbf{x}^* :

$$\bar{f}_\ell(x_\ell) = \frac{1}{n} \sum_{i=1}^n f_{\text{model}}(x_\ell, \mathbf{x}_i^*). \quad (7)$$

The vector \mathbf{x}_i^* holds the realized values of the other variables for observation i and n is the number of observations in the training data. Interaction effects between x_ℓ and another variable in \mathbf{x}^* can distort the effect ([Goldstein et al., 2015](#)). Suppose that half of the observations show a positive association between x_ℓ and the prediction outcome (higher x_ℓ leads to higher predictions), while the other half of the observations show a negative association between x_ℓ and the prediction outcome. Taking the average over all observations will cause the partial dependence plot to look like a horizontal line, wrongly indicating that x_ℓ has no effect on the prediction outcome. Individual conditional expectations can rectify such wrong conclusions.

Individual conditional expectation Individual conditional expectations, introduced by [Goldstein et al. \(2015\)](#), also show the effect of a variable on the predictions obtained from a model, but on an individual level. We evaluate the prediction function in specific values of the variable of interest x_ℓ for $\ell \in \{1, \dots, p\}$, keeping the values of the other variables \mathbf{x}^* fixed:

$$\tilde{f}_{\ell,i}(x_\ell) = f_{\text{model}}(x_\ell, \mathbf{x}_i^*), \quad (8)$$

where \mathbf{x}_i^* are the realized values of the other variables for observation i . We obtain an effect for each observation i , allowing us to detect interaction effects when some (groups of) observations show different behavior compared to others. For example, two distinct groupings will emerge when half of the observations have a positive association and the other half a negative association between x_ℓ and the prediction outcome. Individual conditional expectations can also be used to investigate the uncertainty of the effect of variable x_ℓ on the prediction outcome. The partial dependence plot can be interpreted as the average of this collection of individual conditional expectations, i.e., $\bar{f}_\ell(x_\ell) = \frac{1}{n} \sum_{i=1}^n \tilde{f}_{\ell,i}(x_\ell)$.

4 Case study: claim frequency and severity modeling

An insurer's pricing team uses proprietary data to deliver a fine-grained tariff plan for a portfolio. As a typical example of such data, we study a motor third party liability (MTPL) portfolio from

a Belgian insurer in 1997. This section puts focus on the claim frequency and severity models that are developed with the different modeling techniques. We briefly introduce the data and we report the optimal tuning parameters for the frequency and severity models. Afterwards, we use the tools from Section 3.4 to gain some insights in these optimal models. We conclude this section with an out-of-sample deviance comparison to assess the statistical performance of the different modeling techniques.

4.1 Quick scan of the MTPL data

The data used here is also analyzed in [Denuit and Lang \(2004\)](#), [Klein et al. \(2014\)](#) and [Henckaerts et al. \(2018\)](#). We follow the same data pre-processing steps as the aforementioned papers, e.g., regarding the exclusion of very large claims. Table 7 in Appendix A lists a description of the available variables. The portfolio contains 163,212 unique policyholders, each observed during a period of exposure-to-risk expressed as the fraction of the year during which the policyholder was exposed to the risk of filing a claim. Claim information is known in the form of the number of claims filed and the total amount claimed (in euro) by a policyholder during her period of exposure. The data set lists five categorical, four continuous and two spatial risk factors, each of them informing about specific characteristics of the policy or the policyholder. A detailed discussion on the distribution of all variables is available in [Henckaerts et al. \(2018\)](#). Regarding spatial information, we have access to the 4-digit postal code of the municipality of residence and the accompanying latitude/longitude coordinates of the center of this area. The GAM/GLM benchmarks employ spatial smoothing over the latitude/longitude coordinates. In line with this approach, we use the coordinates as continuous variables in the tree-based models.

Figure 4 displays the distribution of the claims information (`nclaims` and `amount`) and the exposure-to-risk measure (`expo`). Most policyholders in the portfolio are claim-free during their insured period, some file one claim and few policyholders file two, three, four or five claims. The majority of all these claims involve small amounts, but very large claims occur as well. Most policyholders are exposed to the risk during the full year, but there are policyholders who started the policy in the course of the year or surrendered the policy before the end of the year. Figure 4 motivates the use of loss functions which are not based on the squared error loss. We work with the Poisson and gamma distribution/deviance for frequency and severity respectively as our preferred distributional assumption (for GAM/GLM) and loss function choice (for tree-based techniques). Note that earlier work on this data, such as [Denuit et al. \(2007\)](#) and [Henckaerts et al. \(2018\)](#), assumed a log-normal distribution for severity. We illustrate the difference and motivate our choice in the supplementary material.




Figure 4: Distribution of the claim counts, amounts and the exposure-to-risk measure in the MTPL data.

4.2 Optimal tuning parameters

Table 3 lists the optimal tuning parameters for the different machine learning techniques in each of the six data folds. Comparing the number of splits in the trees and the number of trees in the ensembles, we conclude that the frequency models are more extensive compared to the severity variants. This is driven by the lower sample size for severity modeling and the fact that the severity of a claim is typically much harder to predict than the frequency (Charpentier, 2014).

The complexity parameter cp does not give much information about the size of a regression tree and therefore Table 3 also lists the number of splits in the tree. All frequency trees contain between 20 and 38 splits while the severity trees comprise of only one or two splits. The coefficient of variation for the gamma prior γ remains stable over the different data folds.

The number of trees T in the random forest is very unstable over the different folds for both frequency and severity. This shows that the size of the eventual model highly depends on the training data when simply averaging independently grown trees. In four out of six cases, the number of split candidates m is equal to 5 for frequency models and 2 for the severity models. The low value of m for the severity random forests indicates that the variance reduction is the main driver for reducing the loss, as opposed to finding the best split out of multiple candidates.

The number of trees T in the gradient boosting machine is more stable over the folds compared to the random forest. This shows that the sequential approach of growing a boosted model is less affected by the specific data fold. The tree depth d ranges from 3 to 5 in the frequency models. Table 3 reveals that the largest values of T correspond to the smallest values of d and vice versa, highlighting the interplay between these tuning parameters. In five out of six cases, the severity models use stumps (i.e., trees with only one split) as weak learners. A tree depth of $d = 1$ makes these models completely additive, without interaction by construct.

	data fold	Regression tree			Random forest		Boosting machine	
		cp	γ	splits	T	m	T	d
Frequency	1	7.3×10^{-5}	0.125	38	4900	5	2600	3
	2	1.4×10^{-4}	0.125	24	900	5	2000	4
	3	1.1×10^{-4}	0.125	31	400	8	1400	5
	4	1.2×10^{-4}	0.250	27	5000	5	1500	5
	5	1.8×10^{-4}	0.250	20	600	10	1900	4
	6	1.7×10^{-4}	0.250	23	100	5	2700	3
Severity	1	3.3×10^{-3}	-	2	4300	2	600	1
	2	5.8×10^{-3}	-	1	200	2	300	1
	3	3.7×10^{-3}	-	2	600	1	500	1
	4	7.3×10^{-3}	-	1	100	2	400	2
	5	5.4×10^{-3}	-	1	3600	2	600	1
	6	5.4×10^{-3}	-	1	100	1	600	1

Table 3: Overview of the optimal tuning parameters for the tree-based machine learning techniques.

We also tune the benchmark GLMs for each of the six data folds separately, i.e., we perform the binning strategy from Henckaerts et al. (2018) in each fold k such that the optimal bins are chosen for the training data at hand $\mathcal{D} \setminus \mathcal{D}_k$. A grid is used for the two tuning parameters involved, one for the continuous variables and one for the spatial effect, thereby avoiding the two-step binning procedure initially proposed in Henckaerts et al. (2018). Examples of the resulting benchmark GLMs for frequency and severity are presented in Appendix D.1.

4.3 Model interpretation

We will use the variable importance measure to find the most relevant variables in the frequency and severity models. Afterwards, we will make use of partial dependence plots and individual conditional expectations to gain understanding on a selection of interesting effects for the claim frequency. Similar results on claim severity can be found in the supplementary material.

Variable importance To learn which variables matter for predicting claim frequency and severity, we compare in Figure 5 the variable importance plots for the different machine learning techniques over the six data folds. The variables are ranked from top to bottom, starting with the most important one as measured by the average variable importance over the folds (multiple variables with zero importance are ordered alphabetically from Z to A). By contrasting the graphs in the left column of Figure 5, we see that the important variables (mostly bonus-malus scale and age) are similar across all methods for the frequency model. Other relevant risk factors are the power of the vehicle and the spatial risk factor (combining the longitude and latitude information). The frequency GLM, presented in Table 9 in Appendix D.1, contains the top seven variables together with `coverage`, which is ranked at the ninth place for all methods.

The right column of Figure 5 shows the variable importance for severity models. The ranking is very dissimilar across the different modeling techniques. The regression tree models for severity contain only one split using the type of coverage in four out of the six folds, while the other two trees have an additional split on the age of the car. The random forest and gradient boosting machine include more variables, but they both lead to rather different rankings of the importance of those variables. The severity GLM, presented in Table 10 in Appendix D.1, contains three variables: `coverage`, `ageph` and `agec`. An interesting observation is that the most important risk factors in the gradient boosting machines are those selected in the GLMs.




Figure 5: Variable importance in the optimal regression tree (top), random forest (middle) and gradient boosting machine (bottom) per data fold (color) for frequency (left) and severity (right).

Partial dependence plot of the age effect Figure 6 compares the partial dependence effect of the age of the policyholder in frequency models. The two top panels of Figure 6 show the GLM and GAM effects on the left and right respectively. As explained in Section 2, due to our proposed data-driven approach, the GLM effects are a step-wise approximation of the GAM effects. The risk of filing a claim is high for young policyholders and gradually decreases with increasing ages to stabilize around the age of 35. The risk starts decreasing again around the age of 50 and increases for senior policyholders around the age of 70. The bottom left panel of Figure 6 shows the age effect captured by the regression trees. The effect is less stable across the folds compared to the other methods, this is a confirmation and illustration of the variability of a single regression tree. There is also no increase in risk for senior policyholders in the regression trees. The bottom right panel of Figure 6 shows the age effect according to the gradient boosting machines. This looks very similar to the smooth GAM effect with one important distinction, namely the flat regions at the boundaries. This makes the tree-based techniques more robust with respect to extrapolation and results in less danger of creating very high premiums for risk profiles at edges. Note that the gradient boosting machine predicts a wider range of frequencies than the regression tree, namely 0.12 to 0.20 versus 0.12 to 0.165 respectively. The shape of the age effect in the random forest, available in Appendix C, is rather similar to the gradient boosting machine effect but on a slightly more compact range.

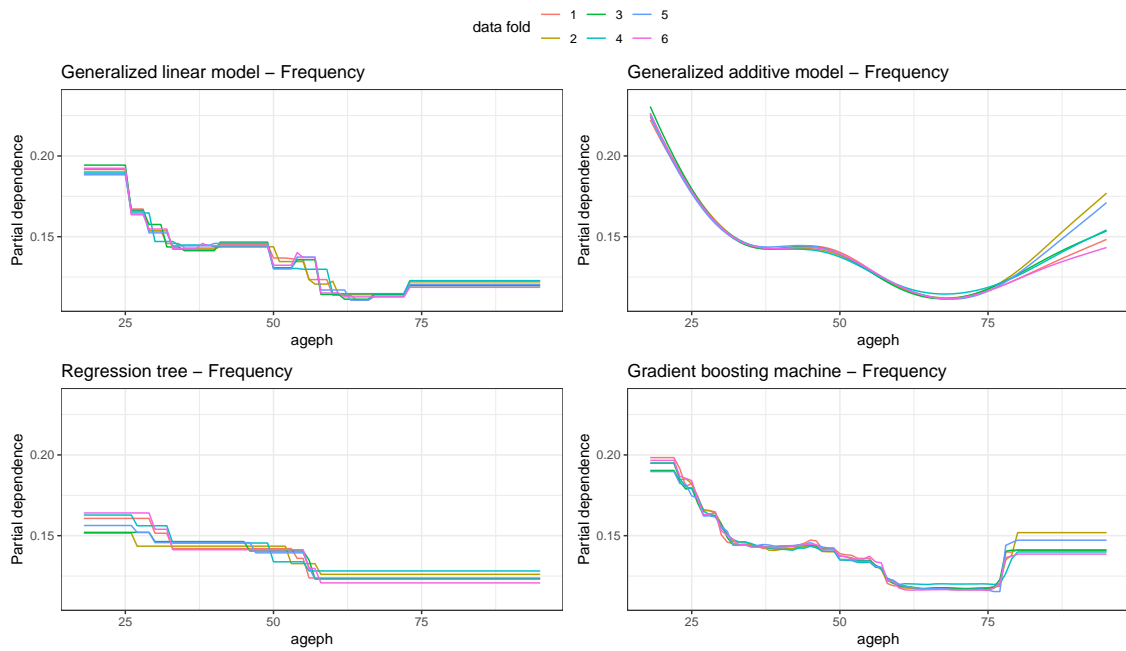


Figure 6: Partial dependence plot to visualize the effect of the age of the policyholder on frequency for the optimal model obtained per data fold (color) in a GLM (top left), GAM (top right), regression tree (bottom left) and gradient boosting machine (bottom right).

Partial dependence plot of the spatial effect Figure 7 compares the spatial effect in frequency models, more specifically the models trained on the data where fold \mathcal{D}_3 was kept as the hold-out test set. We choose a specific data fold because we otherwise need to show six maps of Belgium per method as opposed to overlaying six effects as in Figure 6. The two top panels show the GLM and GAM effects on the left and right respectively. Brussels, the capital of Belgium located in the center of the country, is clearly the most accident-prone area to live and drive a car because of heavy traffic. The southern and northeastern parts of Belgium are less risky because of sparser population and more rural landscapes. The bottom left panel of

Figure 7 shows the spatial effect as it is captured with a regression tree. Splitting on longitude and latitude coordinates results in a rectangular split pattern on the map of Belgium. The bottom right panel of Figure 7 shows the spatial effect for the gradient boosting machine. The underlying rectangular splits are still visible but in a smoother way compared to the regression tree. Brussels still pops out as the most risky area and the pattern looks similar to the GLM and GAM effects. The shape of the spatial effect in the random forest, available in Appendix C, is again similar to that of the gradient boosting machine on a more compact range.




Figure 7: Partial dependence plot to visualize the effect of the municipality of residence on frequency in a GLM (top left), GAM (top right), regression tree (bottom left) and gradient boosting machine (bottom right).

Figures 6 and 7 teach us that a single tree is not able to capture certain aspects in the data, resulting in a coarse approximation of the underlying risk effect. The ensemble techniques are able to capture refined risk differences in a much smoother way. The differences in the range of the predicted effects imply that the gradient boosting machine performs more segmentation while the random forest puts more focus on risk pooling.

Individual conditional expectation of the bonus-malus effect Figure 8 compares the bonus-malus effect captured with a regression tree (left) and a gradient boosting machine (right) for frequency data. As in Figure 7, we show the effect for the models trained on the data where fold \mathcal{D}_3 was kept as the hold-out test set. The gray lines are individual conditional expectations for 1000 random policyholders and the blue line shows the partial dependence curve. The values for \mathbf{x}^* in Eq. (8) are those registered for the selected policyholders. On average, we observe an increase in frequency risk as the blue line surges over the bonus-malus levels, which is to be expected because higher bonus-malus levels indicate a worse claim history. We can get an idea about the sensitivity of the bonus-malus effect across the different policyholders in the portfolio by comparing the steepness of the gray lines. Keeping all other risk factors fixed, a steeper effect indicates that a policyholder's risk is more sensitive to changes in the bonus-malus scale. This effect is driven by the combination of all risk factors registered for this policyholder.

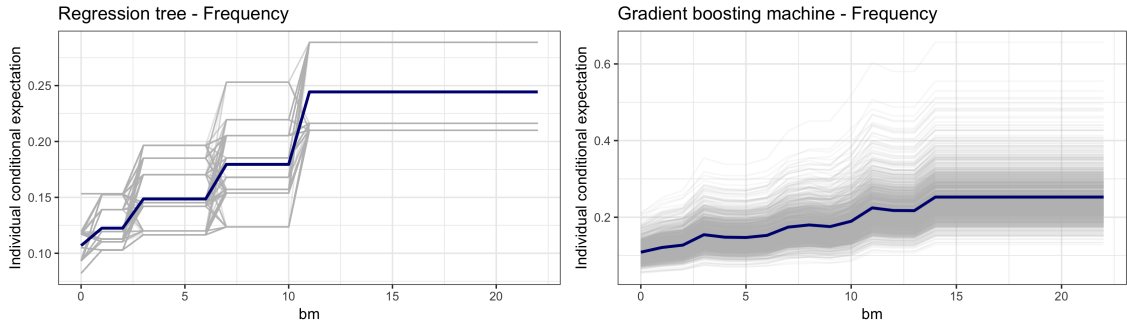


Figure 8: Effect of the bonus-malus scale on frequency in a regression tree (left) and gradient boosting machine (right) as partial dependence (blue) and individual conditional expectations (gray).

The previous figures show some counterintuitive results regarding the monotonicity of a fitted effect. For example, the bonus-malus individual conditional expectations for the regression tree in Figure 8 reveal decreases in risk over increasing bonus-malus levels for certain policyholders. This poses problems for the practical implementation of such a tariff because it assigns a lower premium to policyholders with a worse claim history. In practice, an actuary would specify monotonicity constraints on such a risk factor, either by an a posteriori smoothing of the resulting effect or by using an implementation that allows to specify such constraints a priori, e.g., the `gbm` package has this functionality. Our analysis does not enforce such constraints.

4.4 Hunting for interaction effects

Tree-based models are often praised for their ability to model interaction effects between variables (Buchner et al., 2017; Schiltz et al., 2018). The predictions of a model can not be expressed as the sum of the main effects when interactions are present, because the effect of one variable depends on the value of another variable. Friedman’s H -statistic, introduced by Friedman and Popescu (2008), estimates the interaction strength by measuring how much of the prediction variance originates from the interaction effect. We will put focus on two-way interactions between variables x_k and x_ℓ , but in theory this measure can be applied to arbitrary interactions between any number of variables. Let $\bar{f}_k(x_k)$ and $\bar{f}_\ell(x_\ell)$ represent the one-dimensional partial dependence of the variables as defined in Section 3.4 and $\bar{f}_{kl}(x_k, x_\ell)$ the two-way partial dependence, defined analogously to Eq. (7). The H -statistic is expressed as:

$$H_{kl}^2 = \frac{\sum_{i=1}^n \{\bar{f}_{kl}(x_k^{(i)}, x_\ell^{(i)}) - \bar{f}_k(x_k^{(i)}) - \bar{f}_\ell(x_\ell^{(i)})\}^2}{\sum_{i=1}^n \bar{f}_{kl}^2(x_k^{(i)}, x_\ell^{(i)})},$$

where $x_k^{(i)}$ indicates that the partial dependence function is evaluated at the observed value of x_k for policyholder i . Assuming the partial dependence is centered at zero, the numerator measures the variance of the interaction while the denominator measures the total variance. The ratio of both therefore measures the interaction strength as the amount of variance explained by the interaction. The H -statistic ranges from zero to one, where zero indicates no interaction and one implies that the effect of x_k and x_ℓ on the prediction is purely driven by the interaction.

Table 4 shows the fifteen highest two-way H -statistics among the variables available in the data set (as listed in Table 7 in Appendix A) for the frequency gradient boosting machine trained on the data where fold \mathcal{D}_3 was kept as the hold-out test set. The strongest interaction is found between the longitude and latitude coordinates, which is not a surprise seeing how these two variables together encode the region where the policyholder resides.

Variables	H -statistic	Variables	H -statistic	Variables	H -statistic
(lat, long)	0.2687	(agec, coverage)	0.1185	(bm, power)	0.0800
(fuel, power)	0.1666	(ageph, power)	0.1062	(ageph, lat)	0.0799
(agec, power)	0.1319	(ageph, bm)	0.0961	(agec, ageph)	0.0785
(ageph, sex)	0.1293	(power, sex)	0.0829	(long, sex)	0.0732
(coverage, long)	0.1203	(fuel, long)	0.0828	(agec, bm)	0.0678

Table 4: H -statistic of the 15 strongest two-way interactions between all the variables in the gradient boosting machine for frequencies, trained on data with \mathcal{D}_3 as hold-out test set.

The H -statistic informs us on the strength of the interaction between two variables, but gives us no idea on how the effect behaves. Figure 9 shows grouped partial dependence plots to investigate the interactions highlighted in gray in Table 4. The partial dependence plots of a specific variable are grouped into five equally sized groups based on the value of another variable. Interaction effects between both variables can be discovered by comparing the evolution of the curves over the five different groups. An interaction is at play when this evolution is different for policyholders in different groups. In order to focus purely on the evolution of the effect, we let all the curves in Figure 9 start at zero by applying a vertical shift.




Figure 9: Grouped partial dependence plots for the gradient boosting machine on frequency, trained on data with \mathcal{D}_3 as hold-out test set. The effect is binned in five equally sized groups. The left column shows the effects for the power of the car grouped by the age of the policyholder (top), the type of fuel (middle) and the bonus-malus scale (bottom). The right column shows the effects for the sex of the policyholder (top), age of the car (middle) and type of coverage (bottom), grouped by the age of the policyholder or car.

An important and well-known effect in insurance pricing is the interaction between the age of the policyholder and the power of the vehicle. Our benchmark GLM and GAM use this interaction effect in the predictor and Figure 11 in [Henckaerts et al. \(2018\)](#) shows that young policyholders with high power vehicles form an increased risk for the insurer regarding claim frequency. The top left panel of Figure 9 shows the partial dependence of the power of the vehicle, grouped by the age of the policyholder. The power effect is steepest for young policyholders, indicated by the red line. The steepness of the effect decreases for increasing ages. The difference in the steepness of the effect between young and old policyholders is a visual confirmation of the interaction at play between the variables `ageph` and `power`. The top right panel of Figure 9 shows the partial dependence for the sex of the policyholders, grouped by their age. For young policyholders, aged 18 to 33, we observe that males are on average more risky drivers compared to females, while for the other age groups the female drivers are perceived more risky than males. European insurers are not allowed to use gender in their tariff structure nowadays, implying that young female drivers might be partly subsidizing their male peers.

The middle left panel of Figure 9 shows the partial dependence of the power of the vehicle, grouped by the type of fuel. We observe that the steepness of the power effect is slightly higher for gasoline cars. Drivers typically choose a diesel car when their annual mileage is above average, which would justify their choice of buying a bigger and more comfortable car with higher horsepower. However, drivers who own a high powered gasoline car might choose such a car to accommodate for a more sportive driving style, making them more prone to the risk of a claim. The middle right panel of Figure 9 shows the partial dependence of the age of the vehicle, grouped by the policyholder's age. We observe a big bump for young policyholders in the vehicle age range from 5 to 15. This could indicate an increased claim frequency risk for starting drivers who buy their first car on the second-hand market. The sharp drop around 19 could relate to vintage cars that are used less often and are thus less exposed to the claim risk.

The bottom left panel of Figure 9 shows the partial dependence of the power of the vehicle, grouped by the bonus-malus scale. We observe that the power effect grows steeper for increasing levels occupied in the bonus-malus scale. This indicates that driving a high powered car becomes more risky for policyholders in higher bonus-malus scales. The bottom right panel of Figure 9 shows the partial dependence of the type of coverage, grouped by the age of the vehicle. For vehicles in the age range zero to three, we observe that adding material damage covers decreases the claim frequency risk less compared to other age ranges. This might indicate that policyholders who buy a new car add a material damage cover because they worry about damaging their newly purchased vehicle, while policyholders with older cars who still add damage covers are more risk-averse and also less risky drivers.

4.5 Statistical out-of-sample performance

Figure 10 compares the out-of-sample performance of the different models investigated over the six data folds. We evaluate the Poisson deviance for frequency models and the gamma deviance for severity models, see Eq. (5) and (6) respectively, on the holdout test data. In the left panel of Figure 10, we observe a clear ranking of the out-of-sample Poisson deviance among the different methods. The gradient boosting machine is most predictive, consistently leading to the lowest deviance values. The performance of GLMs and GAMs is very similar, which is expected because the GLM is a data driven approximation of the GAM, as explained in Section 2. Next in line is the random forest, and the regression tree is the least predictive for frequency. These results are very stable over the six data folds. The right panel of Figure 10 shows the out-of-sample gamma deviance for severity. The methods perform rather similarly

and there is no clear winner or loser over the different folds. The peak at the fourth fold reveals a weakness of the GAM: the extrapolation of smooth effects, see Figure 6, combined with out-of-sample testing can lead to huge deviance values. In the severity GAM trained on $\mathcal{D} \setminus \mathcal{D}_4$ and evaluated on \mathcal{D}_4 , the problem occurred with the age of the car. Specifically, the maximal value for `agec` in $\mathcal{D} \setminus \mathcal{D}_4$ for severity is 32 while the maximal value in \mathcal{D}_4 is 37, thus requiring an extrapolation of the calibrated smooth effect. This motivates to cap continuous variables at a certain cut-off in a pre-processing stage for a GAM. A tree-based method automatically deals with this problem thanks to the flat region at the outer ends of the effect, see Figure 6.




Figure 10: Out-of-sample Poisson deviance for frequency (left) and gamma deviance for severity (right), each color representing a different modeling technique.

This comparison only puts focus on the statistical performance of the frequency and severity models. In the next section, we combine both in a pricing model and compare the different tariff structures using practical business metrics relevant for an insurance company.

5 Model lift: from analytic to managerial insights

Choosing between two tariff structures is an important business decision. This creates the need to translate our model findings to evaluation criteria that capture a manager’s interest. We evaluate the economic value of a tariff using tools proposed in the literature to measure model lift (Frees et al., 2013; Goldburd et al., 2016, Section 7.2). In this context, model lift refers to the ability of a model to prevent adverse selection. An insurer might become victim hereof when a competitor refines its tariff structure via innovation such that good risks switch to the competitor and the insurer is left with the bad risks which are more prone to high losses.

We combine the claim frequency and severity models from Section 4 to obtain the technical premium for each policy under consideration, allowing us to compare different tariff structures. As Figure 3 illustrates, each observation is out-of-sample in exactly one of the data folds, more specifically the observations in \mathcal{D}_k are out-of-sample for data fold k . We then use the optimal model trained on $\mathcal{D} \setminus \mathcal{D}_k$ to predict the policyholders in holdout test set \mathcal{D}_k . Following this strategy, we obtain one premium per modeling technique for each policyholder in the full data \mathcal{D} . Table 5 shows a comparison between the predicted premium totals and the observed losses, both on the portfolio level and split by data fold. On average, every method is perfectly capable of replicating the total losses. Section 5.2 compares the model lift measures from Section 5.1 on the portfolio level. We also analyzed each of the data folds separately (not shown), leading to the same ranking of models as in Section 5.2, thereby validating the consistency of our results.

Data fold	GLM	CART	RF	GBM	Losses	GLM	CART	RF	GBM
1	4,396,698	4,341,397	4,407,389	4,376,619	4,365,483	1.01	0.99	1.01	1.00
2	4,420,933	4,339,615	4,419,903	4,384,513	4,388,147	1.01	0.99	1.01	1.00
3	4,369,876	4,313,768	4,380,972	4,337,848	4,461,478	0.98	0.97	0.98	0.97
4	4,370,502	4,374,666	4,383,748	4,324,014	4,422,213	0.99	0.99	0.99	0.98
5	4,405,369	4,374,368	4,399,226	4,357,937	4,485,079	0.98	0.98	0.98	0.97
6	4,397,372	4,375,151	4,412,852	4,363,588	4,342,569	1.01	1.01	1.02	1.01
portfolio	26,360,750	26,118,966	26,404,091	26,144,519	26,464,970	1.00	0.99	1.00	0.99

Table 5: Comparison of the predicted premiums and observed losses on portfolio level and by data fold. We show the premium and loss totals (left) and the ratio of premiums to losses (right).

5.1 Tools to measure model lift

Suppose that an insurance company has a tariff structure P^{bench} in place and a competitor introduces a tariff structure P^{comp} based on a new modeling technique or a different set of rating variables. We define the relativity r_i as the ratio of the competing premium to the benchmark premium for policyholder i :

$$r_i = \frac{P_i^{\text{comp}}}{P_i^{\text{bench}}}. \quad (9)$$

A small relativity indicates a profitable policy which can potentially be lost to a competitor offering a lower premium. A high relativity reveals an underpriced policy which could benefit from better loss control measures such as renewal restrictions. These statements make the assumption that P^{comp} is a more accurate reflection of the true risk compared to P^{bench} .

Loss ratio lift The loss ratio (LR) is the ratio of total incurred claim losses and total earned premiums. Following [Goldburd et al. \(2016\)](#), we assess the loss ratio lift in the following way:

1. sort the policies from smallest to largest relativity r_i ;
2. bin the policies into groups containing the same amount of total exposure e ;
3. within each bin, calculate the overall loss ratio using the benchmark premium P^{bench} .

The bins should have loss ratios around 100% if the benchmark tariff is a technically accurate reflection of the risk. However, an upward trend in the loss ratios would indicate that policies with a lower (higher) premium under the competing tariff are those with a lower (higher) loss ratio in the benchmark tariff, pointing out that the competing tariff better aligns the risk.

Double lift A double lift chart facilitates a direct comparison between two potential tariff structures. Following [Goldburd et al. \(2016\)](#), this chart is created in the following way:

1. sort the policies from smallest to largest relativity r_i ;
2. bin the policies into groups containing the same amount of total exposure e ;
3. within each bin, calculate the average actual loss amount (L) and the average predicted pure premium for both the models (P^{bench} and P^{comp});
4. within each bin, calculate the percentage error for both models as $P/L - 1$.

The best tariff structure is the one with the percentage errors closest to zero, indicating that those premiums match the actual losses more closely.

Gini index [Frees et al. \(2013\)](#) introduced the ordered Lorenz curve to compare the tariffs P^{bench} and P^{comp} by analyzing the distribution of losses versus premiums, where both are ordered by the relativities r from Eq. (9). The ordered Lorenz curve is defined as follows:

$$\left(\frac{\sum_{i=1}^n L_i \mathbb{1}\{F_n(r_i) \leq s\}}{\sum_{i=1}^n L_i}, \frac{\sum_{i=1}^n P_i^{\text{bench}} \mathbb{1}\{F_n(r_i) \leq s\}}{\sum_{i=1}^n P_i^{\text{bench}}} \right),$$

for $s \in [0, 1]$ where $F_n(r_i)$ is the empirical cumulative distribution function of the relativities r . This curve will coincide with the 45 degree line of equality when the technical pricing is done right by the benchmark premium. However, the curve will be concave up when P^{comp} is able to spot tariff deficiencies in P^{bench} . The cumulative distributions are namely taken from the most overpriced policies towards the most underpriced policies in P^{bench} . The Gini index, introduced by [Gini \(1912\)](#) and computed as twice the area between the ordered Lorenz curve and the line of equality, has a direct economic interpretation. A tariff structure P^{comp} that yields a larger Gini index is likely to result in a more profitable portfolio because of better differentiation between good and bad risks. The insurer can decide to only retain the policies with a relativity value below a certain threshold. Averaging this decision over all possible thresholds, [Frees et al. \(2013\)](#) show that the average percentage profit for an insurer equals one half of the Gini index.

5.2 Adverse selection and profits

The panels in the left column of Figure 11 show the loss ratio lift charts for the regression tree, random forest and gradient boosting machine respectively with the GLM as benchmark tariff (i.e., the GLM premium is the denominator in Eq. (9)). All tree-based methods show an increasing trend in the loss ratios. This implies that policies which would receive a lower premium under the competing tariff, those in the first bins, are policies with favorable loss ratios. At the same time, policies having a higher premium under the competing tariff, those in the last bins, exhibit detrimental loss ratios. The tree-based techniques are therefore able to spot deficiencies in the GLM benchmark tariff. One should not draw conclusions from these graphs too fast however. The middle panels of Figure 11 show the loss ratio lifts for the GLM with the three respective tree-based techniques as a benchmark tariff. Comparing these lift charts side by side, we can observe that the upwards trend is now steeper in the cases of the regression tree and random forest. Thus, the GLM is better in spotting deficiencies in those tree-based tariffs compared to the other way around. The gradient boosting machine and GLM result in rather complementary tariffs. The GLM is very good in the three middle relativity bins, but the gradient boosting machine is clearly outperforming in the first and last bin.

These findings are confirmed by the double lift charts in the right panels of Figure 11. These show the double lift charts obtained with the GLM as the benchmark tariff in the relativities. The red and turquoise line respectively show the percentage error for the tree-based model and the GLM. For both the regression tree and the random forest, the percentage error for the GLM benchmark tariff is more closely centered around zero compared to the competitor percentage error. We again notice the complementarity of the gradient boosting machine and GLM tariffs. The percentage error for the gradient boosting machine is closer to zero for the first and last relativity bin, but the GLM is closer to zero for the other three relativity bins.

The gradient boosting machine tariff clearly holds economic value over the GLM benchmark. However, in the bottom left panel of Figure 11, we observe that the gradient boosting machine is slightly over-correcting the GLM premium in the extreme ends of the tariff. The loss ratio in the first bin is 0.84 while the average relativity in that bin is equal to 0.78. Likewise, the average loss ratio in the last bin is 1.21 while the average relativity in that bin is equal to 1.30.




Figure 11: Assessment of model lift for the regression tree (top), random forest (middle) and gradient boosting machine (bottom). The left column shows the loss ratio lift for the tree-based techniques with the GLM as benchmark. The middle column shows the loss ratio lift for the GLM with the tree-based techniques as benchmark. The right column shows the double lift chart for the tree-based techniques with the GLM as benchmark.

Table 6 shows a two-way comparison of Gini indices for the machine learning methods and the GLM. The row names indicate the model generating the benchmark tariff structure P^{bench} while the column names indicate the model generating the competing tariff structure P^{comp} . The row-wise maximum values are indicated in bold. We observe that the gradient boosting machine achieves the highest Gini index when the benchmark is either the GLM, the regression tree or the random forest. When the gradient boosting machine serves as benchmark, the GLM attains the highest Gini index. We use the mini-max strategy of Frees et al. (2013) where we search for the benchmark model with the minimal maximum Gini index. In other words, we look for the benchmark model with the lowest value in bold in Table 6. The gradient boosting machine achieves this minimal maximum Gini index, indicating that this approach leads to a tariff structure that is the least probable to suffer from adverse selection. Note that the GLM tariff achieves the second place, before the random forest and regression tree.

Frees et al. (2013) explain that the average profit for an insurer is equal to half the Gini index. Let us assume that the insurance company uses state-of-the-art GLMs to develop their current tariff structure on this specific data. This implies that developing a competing tariff structure with gradient boosting machines would result in an average profit of around 3.3% for the insurer.

The average is taken over all possible decision-making strategies that the insurance company can take to retain policies based on the relativities. Therefore, by following an optimal strategy, the profit can even be higher for a specific choice of portfolio. We suspect that the improvement in profits could be even greater if there were more explanatory variables in the data.

		competitors:	GLM	CART	RF	GBM
benchmark	GLM			4.07	5.99	6.57
	CART		10.86		10.10	12.02
	RF		7.07	0.53		7.59
	GBM		3.93	0.67	2.30	

Table 6: Two-way comparison of Gini indices for the different tree-based techniques and GLM.

5.3 Solidarity and risk differentiation

From a social point of view, it is crucial for everybody to be able to buy insurance cover at a reasonable price. A tariff structure that follows from a complex machine learning algorithm should not lead to the “personalization of risk” with excessively high premiums for some policyholders. Figure 12 shows violin plots of the annual (i.e., exposure equals one) premium distribution in both the gradient boosting machine tariff P^{gbm} and the GLM tariff P^{glm} . We only consider the gradient boosting machine because Sections 4.5 and 5.2 teach us that only this method holds added value over the GLM. The left panel shows the annual premium amounts and we observe that both distributions look very similar. The minimum, median and maximum premium is 43, 155 and 1138 Euro in P^{gbm} and 41, 156 and 1230 Euro in P^{glm} respectively. The right panel shows the relative difference between both premiums, namely $(P^{\text{gbm}} - P^{\text{glm}})/P^{\text{gbm}}$. The difference is centered around zero and for half the policyholders the difference lies in the range $[-12\%, +12\%]$. This implies that, overall, P^{gbm} and P^{glm} trade off segmentation and risk pooling in a similar way, thereby finding a balance between differentiation and solidarity. For a small selection of policyholders, the gradient boosting machine leads to considerable discounts compared to the GLM.

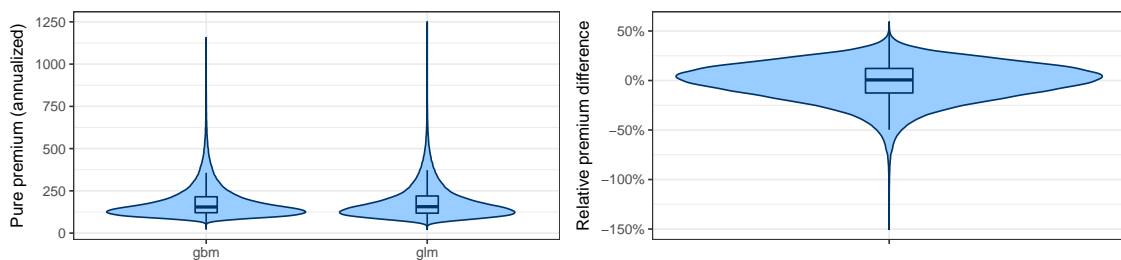


Figure 12: Comparison of the annual premium distribution in the gradient boosting machine tariff P^{gbm} and the GLM tariff P^{glm} : absolute premiums (left) and relative differences (right).

The gradient boosting machine and GLM result in similar premiums on a portfolio level (see Table 5), but on a coarser scale, they could lead to different approaches for targeting specific customer segments. Figure 13 are the relative premium differences between P^{gbm} and P^{glm} over the age of the policyholder in the left panel and the power of the car in the right panel. The blue dots show the average premium difference for policyholders with that specific characteristic, e.g., all policyholders aged 25. We observe that younger policyholders obtain a slightly lower

premium in the gradient boosting machine tariff, while senior policyholders obtain a slightly higher premium compared to the GLM tariff. For middle aged policyholders there are some fluctuations which can be explained by analyzing the age effects in Figure 6. For the group of dots around the age of 75, P^{gbm} gives an average 30% premium discount over P^{glm} . Figure 6 shows that the age effect starts increasing before the age of 75 in the GLM (top left), but only after the age of 75 for the gradient boosting machine (bottom right). Therefore, policyholders around the age of 75 obtain a better deal in the gradient boosting machine tariff. In the right panel of Figure 13, the premium differences increase roughly monotonically with the power of the car. Low powered cars obtain a lower premium in P^{gbm} while high powered cars get a lower premium in P^{glm} . In between the differences are close to zero, indicating that both tariffs treat those cars in a similar fashion.




Figure 13: Comparison of premium differences between P^{gbm} and P^{glm} over the age of the policyholder (left) and the power of the car (right).

We conclude that gradient boosting machines can be a valuable tool for the insurer, while the other tree-based techniques under investigation show little added value on our specific portfolio. The gradient boosting machine is able to deliver a tariff that assesses the underlying risk in a more accurate way, thereby guarding the insurer against adverse selection risks which eventually can result in a profit. The gradient boosting machine also honored the principle of solidarity in the portfolio, offering affordable insurance cover for all policyholders with premiums in the same range as the benchmark GLM.

6 Conclusions and outlook

In this study, we have adapted tree-based machine learning to the problem of insurance pricing, thereby leaving the comfort zone of both traditional ratemaking and machine learning. State-of-the-art GLMs are compared to regression trees, random forests and gradient boosting machines. These tree-based techniques can be used on insurance data, but care has to be taken with the underlying statistical assumptions in the form of the loss function choice. This paper brings multiple contributions to the existing literature. First, we develop complete tariff plans with tree-based machine learning techniques for a real-life insurance portfolio. In this process, we use the Poisson and gamma deviance because the classical squared error loss is not appropriate for a frequency-severity problem. Second, our elaborate cross-validation scheme gives a well thought and careful tuning procedure, allowing us to assess not only the performance of different methods, but also the stability of our results across multiple data folds. Third, we go beyond a purely statistical comparison and also focus on business metrics used in insurance companies to evaluate different tariff strategies. Fourth, we spend a lot of attention on the interpretability of the resulting models. This is a very important consideration for insurance companies within the GDPR's regime of algorithmic accountability. Fifth, our complete analysis is available in

well-documented R functions, readily applicable to other data sets. This includes functions for training, predicting and evaluating models, running the elaborate cross-validation scheme, interpreting the resulting models and assessing the economic lift of these models. Sixth, we extended the `rpart` package such that it is now possible to build regression trees with a gamma deviance as loss function and random forest with both the Poisson and gamma deviance as loss functions. This package is available at <https://github.com/henckr/distRforest>.

The gradient boosting machine is consistently selected as best modeling approach, both by out-of-sample performance measures and model lift assessment criteria. This implies that an insurer can prevent adverse selection and generate profits by considering this new modeling framework. However, this might be impossible because of regulatory issues, e.g., filing requirements (see Appendix D). In that case, an insurance company can still learn valuable information on how to form profitable portfolios from an internal, technical model and translate this to a commercial product which is in line with all those requirements. A possible approach would be to approximate a gradient boosting machine with a GLM, much in line with the strategy to develop the benchmark pricing GLM in this study. The gradient boosting machine can be used to discover the important variables and interactions between those variables, which can then be included in a GLM for deployment. Although we present the tools to detect potentially interesting variables and interactions, we leave for future work the building of a competitive GLM inspired by the gradient boosting machine.

Acknowledgments

This research was supported in part by the Research Foundation - Flanders (FWO) [SB grant 1S06018N] and by the Society of Actuaries James C. Hickman Scholar Program. Furthermore, Katrien Antonio acknowledges financial support from the Ageas Research Chair at KU Leuven and from KU Leuven's research council (project COMPACT C24/15/001). We are grateful for the access to the computing resources of the Department of Mathematics and Statistics at McGill University and the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government - department EWI. Thanks are also due to Christian Genest for supporting this project in various ways and to Michel Denuit for sharing the data analyzed in this paper. Finally, we would like to thank the editor and reviewers for their comments which substantially improved the paper.

References

- K. Antonio, E. W. Frees, and E. A. Valdez. A multilevel analysis of intercompany claim counts. *ASTIN Bulletin*, 40(1):151–177, 2010.
- L. Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen. *Classification and Regression Trees*. The Wadsworth and Brooks-Cole statistics-probability series. Taylor & Francis, New York, 1984.
- F. Buchner, J. Wasem, and S. Schillo. Regression trees identify relevant interactions: can this improve the predictive performance of risk adjustment? *Health economics*, 26(1):74–85, 2017.
- A. Charpentier. General insurance pricing. In *Computational Actuarial Science with R*, pages 507–542. Chapman and Hall/CRC, New York, 2014.
- C. Czado, R. Kastenmeier, E. C. Brechmann, and A. Min. A mixed copula model for insurance claims and claim sizes. *Scandinavian Actuarial Journal*, 2012(4):278–305, 2012.
- A. Dal Pozzolo. Comparison of data mining techniques for insurance claim prediction (MSc thesis). *Università degli Studi di Bologna*, 2010.

- P. De Jong and G. Z. Heller. *Generalized linear models for insurance data*. Cambridge University Press, Cambridge, 2008.
- M. Denuit and S. Lang. Non-life rate-making with Bayesian GAMs. *Insurance: Mathematics and Economics*, 35(3):627–647, 2004.
- M. Denuit, X. Maréchal, S. Pitrebois, and JF. Walhin. *Actuarial modelling of claim counts: Risk classification, credibility and bonus-malus systems*. John Wiley & Sons Ltd, West Sussex, 2007.
- G. Dionne, C. Gouriéroux, and C. Vanasse. Evidence of adverse selection in automobile insurance markets. In *Automobile Insurance: Road Safety, New Drivers, Risks, Insurance Fraud and Regulation*, pages 13–46. Springer, New York, 1999.
- A. Ferrario, A. Noll, and M. V. Wuthrich. Insights from inside neural networks. 2018. URL <https://ssrn.com/abstract=3226852>.
- E. W. Frees. Analytics of insurance markets. *Annual Review of Financial Economics*, 7:253–277, 2015.
- E. W. Frees and E. A. Valdez. Hierarchical insurance claims modeling. *Journal of the American Statistical Association*, 103(484):1457–1469, 2008.
- E. W. Frees, G. Meyers, and A. D. Cummings. Insurance ratemaking and a Gini index. *Journal of Risk and Insurance*, 81(2):335–366, 2013.
- E. W. Frees, R. A. Derrig, and G. Meyers. Predictive modeling in actuarial science. *Predictive Modeling Applications in Actuarial Science: Volume 1, Predictive Modeling Techniques*, 2014.
- J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 29(5):1189–1232, 2001.
- J. H. Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- J. H. Friedman and B. E. Popescu. Predictive learning via rule ensembles. *The Annals of Applied Statistics*, 2(3):916–954, 2008.
- J. H. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer, New York, 2001.
- J. Garrido, C. Genest, and J. Schulz. Generalized linear models for dependent frequency and severity of insurance claims. *Insurance: Mathematics and Economics*, 70:205–215, 2016.
- GDPR. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). *O.J. (L 119)*, 1:1–88, 2016.
- C. Gini. Variabilità e mutabilità (variability and mutability). *Cuppini, Bologna*, 1912.
- M. Goldburd, A. Khare, and D. Tevet. *Generalized linear models for insurance rating*. Casualty Actuarial Society, 2016.
- A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin. Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics*, 24(1):44–65, 2015.
- S. Gschlöbl and C. Czado. Spatial modelling of claim frequency and claim size in non-life insurance. *Scandinavian Actuarial Journal*, 2007(3):202–225, 2007.
- L. Guelman. Gradient boosting trees for auto insurance loss cost modeling and prediction. *Expert Systems with Applications*, 39(3):3659–3667, 2012.
- S. Haberman and A. E. Renshaw. Generalized linear models and actuarial science. *Insurance Mathematics and Economics*, 2(20):142, 1997.
- R. Henckaerts. *distRforest: Distribution-based Random Forest*, 2019. URL <https://www.github.com/henckr/distRforest>. R package version 1.0.
- R. Henckaerts, K. Antonio, M. Clijsters, and R. Verbelen. A data driven binning strategy for the construction of insurance tariff classes. *Scandinavian Actuarial Journal*, 2018(8):681–705, 2018.
- M. E. Kaminski. The right to explanation, explained. *Berkeley Technology Law Journal*, 34(1), 2018.
- N. Klein, M. Denuit, S. Lang, and T. Kneib. Nonlife ratemaking and risk management with Bayesian generalized additive models for location, scale, and shape. *Insurance: Mathematics and Economics*, 55:225–249, 2014.
- S. A. Klugman, H. H. Panjer, and G. E. Willmot. *Loss models: from data to decisions*. John Wiley & Sons, Hoboken, 2012.
- E. Krasheninnikova, J. García, R. Maestre, and F. Fernández. Reinforcement learning for pricing strategy optimization in the insurance industry. *Engineering Applications of Artificial Intelligence*, 80:8–19,

- 2019.
- J. Lemaire. *Bonus-malus systems in automobile insurance*. Springer, New York, 1995.
- Y. Liu, B. Wang, and S. Lv. Using multi-class AdaBoost tree for prediction frequency of auto insurance. *Journal of Applied Finance and Banking*, 4(5):45, 2014.
- C. Molnar. Interpretable machine learning: A guide for making black box models explainable. 2019. URL <https://christophm.github.io/interpretable-ml-book/>.
- J. A. Nelder and R. W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- J. Neyman. On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society*, 97(4):558–625, 1934.
- C. O’Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, 2017.
- P. Parodi. *Pricing in General Insurance*. Chapman and Hall/CRC, New York, 2014.
- F. Pasquale. *The black box society: The secret algorithms that control money and information*. Harvard University Press, Cambridge, 2015.
- J. Pesantez-Narvaez, M. Guillen, and M. Alcañiz. Predicting motor insurance claims using telematics data—XGBoost versus logistic regression. *Risks*, 7(2):70, 2019.
- G. Ridgeway. *gbm: Generalized Boosted Regression Models*, 2014. URL <https://CRAN.R-project.org/package=gbm>. R package version 2.1-06.
- J. Schelldorfer and M. V. Wüthrich. Nesting classical actuarial models into neural networks. 2019. URL <https://ssrn.com/abstract=3320525>.
- F. Schiltz, C. Masci, T. Agasisti, and D. Horn. Using regression tree ensembles to model interaction effects: a graphical approach. *Applied Economics*, 50(58):6341–6354, 2018.
- H. Southworth. *gbm: Generalized Boosted Regression Models*, 2015. URL <https://github.com/harrysouthworth/gbm>. R package version 2.1-06.
- G. A. Spedicato, C. Dutang, and L. Petrini. Machine learning methods to perform pricing optimization. A comparison with standard GLMs. *Variance*, 12(1):69–89, 2018.
- T. M. Therneau and B. Atkinson. *rpart: Recursive Partitioning and Regression Trees*, 2018. URL <https://CRAN.R-project.org/package=rpart>. R package version 4.1-13.
- T. M. Therneau and E. J. Atkinson. An introduction to recursive partitioning using the rpart routines (vignette), 2019. URL <https://cran.r-project.org/package=rpart>.
- W. N. Venables and B. D. Ripley. Tree-based methods. In *Modern Applied Statistics with S*, pages 251–269. Springer, New York, 2002.
- Y. Wang and W. Xu. Leveraging deep learning with LDA-based text analytics to detect automobile insurance fraud. *Decision Support Systems*, 105:87–95, 2018.
- Simon N Wood. *Generalized additive models: An introduction with R*. Chapman and Hall/CRC, 2006.
- M. V. Wüthrich and C. Buser. Data analytics for non-life insurance pricing (lecture notes). 2019.
- Y. Xia, C. Liu, Y. Li, and N. Liu. A boosted decision tree approach using Bayesian hyper-parameter optimization for credit scoring. *Expert Systems with Applications*, 78:225–241, 2017.
- Y. Yang, W. Qian, and H. Zou. Insurance premium prediction via gradient tree-boosted Tweedie compound Poisson models. *Journal of Business & Economic Statistics*, 36(3):456–470, 2018.
- P. Zöchbauer, M. V. Wüthrich, and C. Buser. Data science in non-life insurance pricing (MSc thesis). *ETH Zürich*, 2017.

A List of variables in the MTPL data

Claim information and exposure-to-risk measure	
<code>nclaims</code>	The number of claims filed by the policyholder.
<code>amount</code>	The total amount claimed by the policyholder in euro.
<code>expo</code>	The fraction of the year during which the insurer was exposed to the risk.

Categorical risk factors	
<code>coverage</code>	Type of coverage provided by the insurance policy: TPL, TPL+ or TPL++. TPL = only third party liability, TPL+ = TPL + limited material damage, TPL++ = TPL + comprehensive material damage.
<code>fuel</code>	Type of fuel of the vehicle: gasoline or diesel.
<code>sex</code>	Gender of the policyholder: male or female. (As from 21 December 2012, the European Court of Justice prohibited the use of gender in insurance tariffs to avoid discrimination between males and females, known as the Test-Achats Ruling. Gender is therefore only investigated for use within an internal technical tariff, but can not be used in a commercial product.)
<code>use</code>	Main use of the vehicle: private or work.
<code>fleet</code>	The vehicle is part of a fleet: yes or no.

Continuous risk factors	
<code>ageph</code>	Age of the policyholder in years.
<code>power</code>	Horsepower of the vehicle in kilowatt.
<code>agec</code>	Age of the vehicle in years.
<code>bm</code>	Level occupied in the former compulsory Belgian bonus-malus scale. From 0 to 22, a higher level indicates a worse claim history, see Lemaire (1995) . (This variable is typically not used as an a priori rating factor, but rather as an a posteriori correction in a credibility framework or bonus-malus scheme. We however keep <code>bm</code> in the data to assess the amount of information contained in this variable and to investigate the resulting effect.)

Spatial risk factor	
<code>long</code>	Longitude coordinate of the center of the municipality where the policyholder resides.
<code>lat</code>	Latitude coordinate of the center of the municipality where the policyholder resides.

Table 7: Description of the available variables in the MTPL data.

B Search grid for the tuning parameters

Regression tree	$cp \in \{1.0 \times 10^{-5}, 1.1 \times 10^{-5}, \dots, 1.0 \times 10^{-2}\}$ $\gamma \in \{2^{-6}, 2^{-5}, 2^{-4}, 2^{-3}, 2^{-2}, 2^{-1}, 2^0\}^*$
Random forest	$T \in \{100, 200, \dots, 5000\}$ $m \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$
Gradient boosting machine	$T \in \{100, 200, \dots, 5000\}$ $d \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$

Table 8: Search grid for the tuning parameters in the different tree-based machine learning techniques.
 * Note that the γ tuning parameter is only used in frequency models for the Poisson deviance.

C Random forests for claim frequency data

The left and right panel of Figure 14 show the partial dependence plot of the age and spatial effect in the claim frequency random forests, respectively. These effects are discussed in Section 4.3. The six random forests for frequency contain rather different number of trees, ranging from 100 to 5000 (see Table 3 earlier). However, these random forests exhibit very similar effects for the age of the policyholder in Figure 14. This indicates that the underlying model structure does not change drastically after including a sufficient number of trees.




Figure 14: Effect of the age of the policyholder (left) and the municipality of residence (right) on frequency in a random forest.

D Filing requirements for pricing models

Insurance companies can be required by regulation to file their rating model on paper. This section presents such frequency and severity models, trained on the data where \mathcal{D}_3 was kept as hold-out test set. Section D.1 and D.2 show the GLMs and regression trees respectively. The other five GLMs and regression trees are not shown due to lack of space. This filing requirement is more difficult to satisfy for the ensemble methods, but it is possible by printing the individual trees. However, this would result in a large amount of pages which is not very practical or insightful for the regulator.

D.1 GLM

	Coefficient	SE (σ)	t -stat	p -value		Coefficient	SE (σ)	t -stat	p -value
Intercept	-2.19 ***	0.03	-64.09	0.00	power[10,35)	-0.16 ***	0.04	-3.95	0.00
coverageTPL+	-0.10 ***	0.02	-5.46	0.00	power[35,42)	-0.05	0.03	-1.54	0.12
coverageTPL++	-0.10 ***	0.03	-3.97	0.00	power[42,49)	-0.02	0.03	-0.55	0.58
fueldiesel	0.18 ***	0.02	10.00	0.00	power[59,73)	0.02	0.02	0.89	0.37
fleetyes	-0.13 ***	0.05	-2.72	0.01	power[73,92)	0.07 *	0.04	1.76	0.08
ageph[18,26)	0.32 ***	0.04	8.94	0.00	power[92,243]	0.13 ***	0.05	2.88	0.00
ageph[26,29)	0.16 ***	0.04	4.61	0.00	latlong[-0.46,-0.36)	-0.31	0.19	-1.62	0.11
ageph[29,32)	0.11 ***	0.04	3.11	0.00	latlong[-0.36,-0.26)	-0.33 ***	0.06	-5.97	0.00
ageph[32,35)	0.02	0.04	0.46	0.65	latlong[-0.26,-0.18)	-0.20 ***	0.03	-5.86	0.00
ageph[50,54)	-0.11 ***	0.03	-3.53	0.00	latlong[-0.18,-0.12)	-0.23 ***	0.04	-6.38	0.00
ageph[54,58)	-0.05	0.04	-1.28	0.20	latlong[-0.12,-0.061)	-0.13 ***	0.03	-4.41	0.00
ageph[58,62)	-0.23 ***	0.05	-4.93	0.00	latlong[-0.061,-0.017)	-0.05	0.03	-1.62	0.11
ageph[62,73)	-0.25 ***	0.04	-7.06	0.00	latlong[0.025,0.077)	0.06 **	0.03	2.42	0.02
ageph[73,95]	-0.21 ***	0.05	-4.42	0.00	latlong[0.077,0.14)	0.03	0.03	0.86	0.39
bm[1,2)	0.14 ***	0.03	5.46	0.00	latlong[0.14,0.23)	0.04	0.03	1.29	0.20
bm[2,3)	0.16 ***	0.04	4.48	0.00	latlong[0.23,0.33]	0.36 ***	0.03	12.37	0.00
bm[3,5)	0.37 ***	0.03	11.86	0.00	agephpower-0.05	-0.04	0.06	-0.61	0.54
bm[5,7)	0.32 ***	0.03	11.39	0.00	agephpower-0.02	-0.11 ***	0.04	-3.15	0.00
bm[7,9)	0.48 ***	0.03	14.77	0.00	agephpower-0.01	-0.03	0.03	-0.79	0.43
bm[9,11)	0.52 ***	0.03	17.30	0.00	agephpower0.01	0.00	0.03	-0.02	0.98
bm[11,22]	0.75 ***	0.03	26.07	0.00	agephpower0.02	-0.04	0.05	-0.67	0.50
agec[9,14)	0.04 *	0.02	1.91	0.06	agephpower0.04	0.04	0.04	1.07	0.29
agec[14,48]	-0.02	0.03	-0.75	0.46					

Significance codes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 9: Frequency GLM specification, trained on the data where \mathcal{D}_3 was kept as hold-out test set.

	Coefficient	SE (σ)	t -stat	p -value		Coefficient	SE (σ)	t -stat	p -value
Intercept	7.31 ***	0.08	88.50	0.00	ageph[61,66)	-0.20 *	0.11	-1.73	0.08
coverageTPL+	-0.23 ***	0.05	-4.62	0.00	ageph[66,72)	-0.04	0.11	-0.33	0.75
coverageTPL++	0.18 **	0.07	2.54	0.01	ageph[72,95]	0.11	0.11	0.93	0.35
ageph[18,25)	0.13	0.11	1.23	0.22	agec[0,2)	0.22 **	0.11	2.09	0.04
ageph[25,28)	-0.02	0.10	-0.15	0.88	agec[2,3)	-0.11	0.09	-1.11	0.27
ageph[28,30)	0.08	0.11	0.70	0.49	agec[3,4)	-0.08	0.10	-0.82	0.41
ageph[30,33)	0.07	0.10	0.65	0.52	agec[4,6)	-0.08	0.07	-1.09	0.28
ageph[33,36)	-0.19 *	0.10	-1.87	0.06	agec[6,7)	-0.10	0.08	-1.22	0.22
ageph[36,39)	-0.26 **	0.10	-2.57	0.01	agec[7,8)	-0.10	0.09	-1.21	0.23
ageph[39,42)	-0.19 *	0.10	-1.84	0.07	agec[10,11)	-0.12	0.09	-1.33	0.18
ageph[42,45)	-0.01	0.10	-0.12	0.91	agec[11,12)	-0.05	0.09	-0.52	0.60
ageph[49,52)	-0.02	0.10	-0.21	0.83	agec[12,14)	-0.14 *	0.08	-1.69	0.09
ageph[52,55)	-0.13	0.11	-1.18	0.24	agec[14,48]	-0.01	0.09	-0.06	0.95
ageph[55,61)	-0.17 *	0.10	-1.69	0.09					

Significance codes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 10: Severity GLM specification, trained on the data where \mathcal{D}_3 was kept as hold-out test set.

D.2 Regression tree

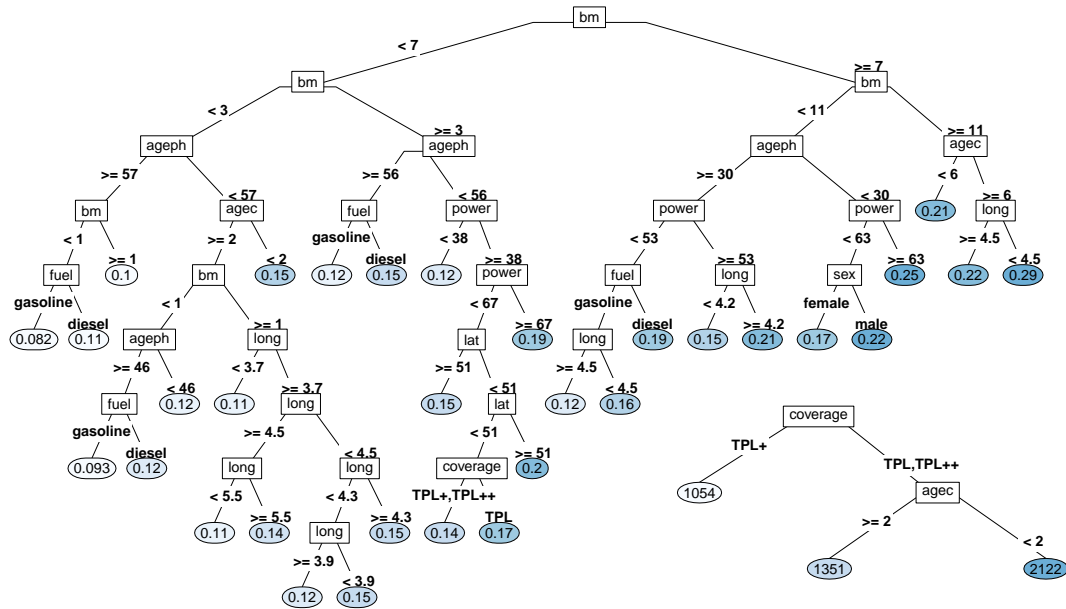


Figure 15: Regression trees for claim frequency (big, top left) and severity (small, bottom right), both trained on the data where \mathcal{D}_3 was kept as hold-out test set.

E Supplementary material

Supplementary material related to this article can be found online at <https://github.com/henckr/semtree>.