Stockholms universitet

Reserve Risk Modelling under Solvency II: Established Models and a Gradient Boosting Approach

Jimmy Eriksson Widfors

Masteruppsats i försäkringsmatematik
Master Thesis in Actuarial Mathematics

Masteruppsats 2026:4
Försäkringsmatematik
Februari 2026

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Stockholm University

Mathematical Statistics
Stockholm University
Master Thesis **2026:4**
http://www.math.su.se

Reserve Risk Modelling under Solvency II: Established Models and a Gradient Boosting Approach

Jimmy Eriksson Widfors*

February 2026

Abstract

Under the Solvency II framework, insurers are required to hold capital equal to the 99.5% Value-at-Risk (VaR) of their one-year loss distribution, with reserve risk representing a key component. This study compares several models for estimating reserve risk. A bootstrap one-year risk framework is applied to established stochastic reserving models, including the Mack model under different distributional assumptions and the Over-Dispersed Poisson (ODP) model, as well as to a machine-learning extension of the ODP model based on a Gradient Boosting Machine (GBM). All models are evaluated using synthetic claims data generated for several Lines of Business (LoBs).

The results reveal clear differences across models. For some LoBs, the GBM–ODP model produces narrower loss distributions and lower VaR estimates, whereas the Mack models yield more conservative results, with the ODP model positioned between them. Although the results suggest potential advantages of the GBM–ODP model in terms of both accuracy and capital efficiency, certain results indicate a need for further analysis to improve robustness in reserve risk applications.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: jimmyew1@gmail.com. Supervisor: Mathias Millberg Lindholm.

Acknowledgements

I would like to thank my supervisor, Mathias Millberg Lindholm, for his guidance and support throughout this thesis, my former manager, Jonas Hermansson, for encouraging me to complete it, and my current manager, Carolina Björn, for her support during the final stages of this work.

AI Usage Statement

Artificial intelligence tools were used to assist with language refinement in this thesis. The development of the methodology, analysis, and conclusions was carried out independently by the author.

Contents

1 Introduction

Since the implementation of Solvency II in European legislation, strict capital requirements have been imposed on insurance companies to ensure their ability to meet future obligations. According to Directive 2009/138/EC (EU, 2009), the Solvency Capital Requirement (SCR) corresponds to a 99.5% Value-at-Risk (VaR) over a one-year time horizon. This means that an insurance company must be able to cover its annual losses with 99.5% probability. Conversely, there is a 0.5% probability that the company's capital will be insufficient, a concept often referred to as the "200-year risk".

A key component of the overall capital requirement is the Basic Solvency Capital Requirement (BSCR), which consists of several risk categories, including market risk, life insurance risk, and non-life insurance risk. These categories are further divided into subcategories. In this study, the focus is on reserve risk, a key subcategory of non-life insurance risk. Reserve risk is defined as the risk that the claims reserve is insufficient over a one-year time horizon. The claims reserve itself is typically measured by the Best Estimate (BE), representing the expected value of future outstanding claim payments.

This study evaluates different models for estimating reserve risk and compares their performance. A bootstrap framework, as discussed by England and Verrall (2006), is combined with the one-year risk approach of Ohlsson and Lauzenings (2009) and applied jointly to the models under consideration.

The models examined in this study are the Mack model, under various distributional assumptions, and the Over-Dispersed Poisson (ODP) model, both of which are described in England and Verrall (2006). In addition, the ODP model is extended using a tree-based Gradient Boosting Machine (GBM), following the approach of Lindholm et al. (2020).

The data used in this study are generated by the stochastic simulation machine introduced by Gabrielli and Wüthrich (2018).

2 Data

The data in this study are provided by the stochastic simulation machine, introduced by Gabrielli and Wüthrich (2018). The simulation machine is based on neural networks, whose parameters are calibrated on historical non-life claims data consisting of $9,977,298$ individual claims during the period 1994-2005. Each generated individual claim by the simulation machine provides the paid amount during each development year and the following features:

- the claims number $ClNr$, which serves as a distinct claims identifier
- the line of business LoB, which is categorical with labels in $\{1,\ldots,4\}$
- the claims code cc, which is categorical with labels in $\{1,\ldots,53\}$ and denotes the labor sector of the injured
- the accident year AY, which is in $\{1994,\ldots,2005\}$
- the accident quarter AQ, which is in $\{1,\ldots,4\}$
- the age of the injured age (in 5-year age buckets), which is in $\{15,20,\ldots,70\}$
- the injured part inj_part, which is categorical with labels in $\{10,\ldots,99\}$ and denotes the part of the body injured
- the reporting year RY, which is in $\{1994,\ldots,2016\}$.

The analysis is based on the features LoB, AY, and RY, where RY will hereafter be referred to as the Development Year (DY). For reproducibility purposes, the number of expected simulated claims is 1,200,000, using seed 12345.

	LoB 1	LoB 2	LoB 3	LoB 4
Number of claims	298,770	359,796	239,700	300,129
Total payments	336,382	824,308	688,339	523,831
Total outs. payments	49,422	77,917	73,443	85,023

Table 1: Summary statistics of generated individual claims per LoB

3 Theory and Methodology

3.1 Claims reserving method

Reserving in actuarial science is mainly divided into two different types of reserves. The first is based on claims that have already occurred, and is referred to as the claims reserve. The other is based on future claims, and is referred to as the premium reserve. Since this study focuses on the risk linked to the claims reserve, this section will describe one of the most common methods for calculating the claims reserve.

In order to calculate the claims reserve, historical data over claims development is needed. Typically, the data is summarized in a claims development triangle, as shown in Table 2. The cells in the triangle usually represent cumulative (or incremental) amount paid, incurred amount or number of reported claims. In this

case, the cells c_{ij} are defined as the cumulative payment at development year j that occurred in accident year i. The column corresponding to development year $j = 0$ describes the total payment made during the same year as the accident year, $i = 1, 2, \ldots, k$. When development year $j = m$ is reached, the accident year is assumed to be fully developed, and no further payments will be made. The corresponding cell $c_{i,m}$ is therefore the total amount paid for accident year i, and is referred to as the ultimo claim cost.

Accident	Development year					
year	0	1	2	$\ldots$	$m-1$	m
1	$c_{1,0}$	$c_{1,1}$	$c_{1,2}$	$\ldots$	$c_{1,m-1}$	$c_{1,m}$
2	$c_{2,0}$	$c_{2,1}$	$c_{2,2}$	$\ldots$	$c_{2,m-1}$	
3	$c_{3,0}$	$c_{3,1}$	$c_{3,2}$	$\ldots$		
$\vdots$	$\vdots$	$\vdots$	$\vdots$			
$k-1$	$c_{k-1,0}$	$c_{k-1,1}$				
k	$c_{k,0}$					

Table 2: Claims development triangle

Assuming that the valuation takes place at the end of accident year $i = k$, all cumulative payments c_{ij} with $i + j \leq k$ are known and form the *upper triangle*. The future payments $C_{i,j}$, i.e., those with $k < i + j \leq k + m$, must be predicted in order to calculate the claims reserve. These cells are referred to as the *lower triangle* and are shown in Table 3.

Accident	Development year					
year	0	1	2	$\ldots$	$m-1$	m
1						
2						$C_{2,m}$
3					$C_{3,m-1}$	$C_{3,m}$
$\vdots$					$\vdots$	$\vdots$
$k-1$			$C_{k-1,2}$	$\ldots$	$C_{k-1,m-1}$	$C_{k-1,m}$
k		$C_{k,1}$	$C_{k,2}$	$\ldots$	$C_{k,m-1}$	$C_{k,m}$

Table 3: Future claims development triangle

By combining the cells in Tables 2 and 3, the outstanding payments R_k at the end of accident year $i = k$ can be defined as

$$R_k = \sum_{i=2}^{k} (C_{i,m} - c_{i,k+1-i}) \tag{1}$$

The quantity R_k is to be estimated in order to calculate the claims reserve. One

of the most common methods for this purpose is the chain ladder method. In Mack (1993), three assumptions regarding the claims data are made for the method to be valid in theory. The first assumption is that there exist development factors $f_1, \ldots, f_m > 0$ such that

$$\mathbb{E}[C_{ij} \mid C_{i,1}, ..., C_{i,j-1}] = f_j C_{i,j-1} \tag{2}$$

The second one states that the set of random variables

$$\{C_{ij}; j = 1, 2, ...\} \tag{3}$$

are independent, i.e., the development of the accident years is independent. The third one states that

$$\text{Var}(C_{i,j} \mid C_{i,1}, ..., C_{i,j-1}) = C_{i,j-1}\sigma_j^2, \tag{4}$$

where σ_j is an unknown parameter.

Given these assumptions, the development factors can be estimated as

$$\hat{f}_j = \frac{\sum_{i=1}^{k-j} c_{i,j}}{\sum_{i=1}^{k-j} c_{i,j-1}} \tag{5}$$

By introducing $F_j = f_j f_{j+1} ... f_m$, the claims reserve can be estimated as

$$\hat{R}_k = \sum_{i=2}^{k} c_{i,k+1-i}(\hat{F}_{m+1-i} - 1) \tag{6}$$

3.2 One-year risk

Under the Solvency II regulation, the time horizon of the capital requirement is one year. In Ohlsson and Lauzeningks (2009), it is stated that the one-year reserve risk is the risk linked to the one-year run-off result.

Consider first the reserve at the beginning of the risk year, referred to as the opening reserve and denoted by R_0. During the risk year, payments related to historical claims are made, with a total amount of X_1. At the end of the risk year, given

the observed payments X_1, the reserve is re-estimated, yielding the closing (or outgoing) reserve R_1. The run-off result, denoted by T, is then defined as

$$T = R_0 - X_1 - R_1. \tag{7}$$

The one-year reserve risk is described by the distribution of T conditioned on the information available at time 0. By analysing the loss distribution $-T$, one can assess the extent to which the actual outcome may deviate from the initial claims reserve. Since the focus is on reserve risk, interest lies in the right tail of this distribution. To quantify this tail risk, an appropriate risk measure is required. In the following subsection, a risk measure is introduced and applied to determine an estimate of the one-year reserve risk.

3.3 Value-at-Risk

Value-at-Risk (VaR) is a metric for measuring the risk in financial assets. According to the Solvency II regulation, the capital requirement must correspond to 99.5% Value-at-Risk over a one-year horizon. This is referred to as the Solvency Capital Requirement (SCR). Mathematically, it can be described as follows. Let L denote the loss variable associated with claims reserve risk over a one-year horizon.

For a confidence level $\alpha \in (0,1)$, VaR_α is defined as

$$\mathrm{VaR}_\alpha(L) = \min\{l \in \mathbb{R} : P(L > l \mid \mathscr{F}_0) \leq 1-\alpha\}. \tag{8}$$

In words, the VaR of L at confidence level α is the smallest loss l such that the probability of L exceeding l is less than or equal to $1-\alpha$.

Let $L = -T$. By combining (7) and (8), the SCR for reserve risk can be expressed as

$$\mathrm{VaR}_{0.995}(L) = \mathrm{VaR}_{0.995}(X_1 + R_1 \mid \mathscr{F}_0) - R_0, \tag{9}$$

where $\mathscr{F}_0$ denotes the information available at time 0, i.e., at the beginning of the risk year.

The distribution of the one-year reserve risk loss is not available in closed form and must therefore be obtained through simulation. In the following subsection, a bootstrap-based framework for constructing predictive distributions in claims reserving is introduced.

3.4 Framework for bootstrapped predictive distributions in claims reserving

Bootstrapping is a simulation-based method that is very useful for obtaining predictive distributions. In this thesis, the framework described by England and Verrall (2006) for bootstrapped predictive distributions is used to estimate VaR. The framework provides a practical way to assess reserve uncertainty arising from both parameter uncertainty and future claim variability.

After fitting a reserving model to the observed triangle data, residuals are resampled to generate multiple pseudo triangles. The model is then re-estimated for each pseudo triangle, and future claims (one year ahead in this case) are simulated based on the fitted parameters. Repeating this procedure yields a distribution of predicted outstanding claims, from which measures like VaR can be obtained. A schematic illustration of the procedure is shown in Figure 1.

In the following sections, the different model types estimated using this framework are presented. The Mack bootstrap and ODP models follow England and Verrall (2006), whereas the ODP–GBM extension is based on Lindholm et al. (2020).

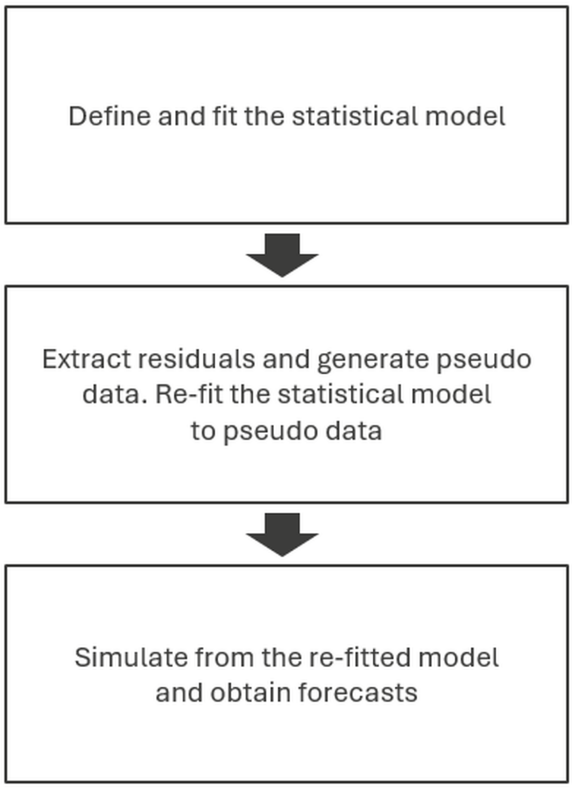

Figure 1: Framework for obtaining bootstrapped predictive distributions

3.5 Mack bootstrap

In Section 3.1, an established method for calculating the claims reserve based on triangular data was described. Building on Mack's assumptions together with the framework in 3.4, a bootstrap approach can be used to generate a predictive distribution for the one-year reserve risk. As described in Section 3.2, at the beginning of the risk year the opening reserve R_0 is known, whereas the payments during the year X_1 and the outgoing reserve R_1 are unknown. Since the estimation of R_1 depends on the realised payments X_1, the first step is to generate a stochastic estimate of X_1, corresponding to a new diagonal in the claims triangle. This procedure can be carried out by following the six steps described below.

1. Estimate f_j for $1 \leq j \leq m$ by the same method used for estimating the claims reserve.

2. Estimate σ_j^2 as

$$\hat{\sigma}_j^2 = \frac{1}{n_j - 1} \sum_i C_{i,j-1}(f_{i,j} - \hat{f}_j)^2 \tag{10}$$

   where n_j equals the number of terms in the summation, the summation runs over all i used for estimating f_j, and $f_{i,j} = C_{i,j}/C_{i,j-1}$ denotes the individual development factor.

3. Compute independent and identically distributed residuals. Based on the assumptions in (2) and (4), the standardized residuals can be defined as

$$r_{i,j} = \frac{\sqrt{C_{i,j-1}}}{\hat{\sigma}_j}(f_{i,j} - \hat{f}_j). \tag{11}$$

4. Draw a random sample of residuals $r_{i,j}^*$, with replacement, for all $i + j \leq k$. Then compute bootstrap development factors as

$$f_{i,j}^* = \hat{f}_j + r_{i,j}^* \frac{\hat{\sigma}_j}{\sqrt{C_{i,j-1}}} \tag{12}$$

   in order to preserve the conditional mean and variance structure implied by the Mack model. Repeat this step B times. Choose B sufficiently large for the results to converge.

5. For each random sample, compute a new estimate $\hat{f}_j^*$ by the same method used for estimating the claims reserve.

6. For $i+j=k+1$, compute $C_{i,j+1}$, i.e., the new diagonal in the triangle, as $C_{i,j}f_{i,j}^{**}$, by assuming that

$$f_{i,j}^{**} \sim N(\hat{f}_j^*, \hat{\sigma}_j^2/C_{i,j}), \tag{13}$$

which follows directly from the conditional mean and variance assumptions of the Mack model applied to the next development period. The normality assumption may be replaced by a gamma or lognormal distribution with the same conditional mean and variance.

The Mack bootstrap results in B simulated claims triangles, including the diagonal of the risk year corresponding to the simulated payments X_1. For each simulated triangle, the outgoing reserve R_1 is calculated using the same reserving method as for the original data. Finally, the VaR is derived as described in Section 3.3.

3.6 The over-dispersed Poisson model

Using regression-based reserving models is another possible method for estimating the claims reserve. England and Verrall (2002) show that a Poisson regression model with a log-link and accident year and development year effects reproduces the classical chain-ladder reserve estimates. Assuming that the incremental payments $I_{i,j} = C_{i,j} - C_{i,j-1}$ are Poisson distributed and allowing for over-dispersion, it follows that

$$I_{i,j} \sim \text{ODP}(\mu_{i,j}, \phi), \tag{14}$$

with

$$\mathbb{E}(I_{i,j}) = \mu_{i,j}, \qquad \text{Var}(I_{i,j}) = \mu_{i,j}\phi, \tag{15}$$

where $\phi > 0$ denotes the dispersion parameter.

The mean structure is specified using a log-link function. Let $\psi_{i,j}$ denote the linear predictor, such that

$$\mu_{i,j} = \exp\{\psi_{i,j}\}. \tag{16}$$

The linear predictor is defined parametrically as

$$\psi_{i,j} = \beta_i^{(1)} + \beta_j^{(2)}, \tag{17}$$

where $\beta_i^{(1)}$ and $\beta_j^{(2)}$ represent accident year and development year effects, respectively.

Estimation is performed using a quasi-likelihood approach. The parameters in the linear predictor are estimated by maximizing the Poisson log-likelihood, treating ϕ as fixed. After fitting the mean model, Pearson residuals

$$r_{i,j}^{(P)} = \frac{I_{i,j} - \hat{\mu}_{i,j}}{\sqrt{\hat{\mu}_{i,j}}} \tag{18}$$

are used to estimate the dispersion parameter according to

$$\hat{\phi} = \frac{1}{n-p} \sum_{i,j} \left(r_{i,j}^{(P)} \right)^2, \tag{19}$$

where n is the number of observed incremental payments and p is the number of parameters in the linear predictor.

By simulating from this model in accordance with Section 3.2 and Section 3.4, VaR is obtained as described in Section 3.3. The ODP model provides a parametric baseline for the machine-learning–based reserving models presented in the following section.

3.7 Tree-based gradient boosting extension of the over-dispersed Poisson model

In Section 3.6, the ODP model for estimating the claims reserve was introduced. Building on this framework, the model parameters may be estimated using machine learning techniques. Lindholm et al. (2020) present various machine learning–based reserving models. Using the ODP model as a base building block, they estimated the mean structure by means of tree-based gradient boosting machines (GBM) and neural networks (NN). Here, the theory of the tree-based GBM approach used in their paper is briefly described.

Let the linear predictor $\psi_{i,j}$ be explained by a regression function. Using the log-link function, the model for incremental payments $I_{i,j}$ is specified as

$$I_{i,j} \sim \text{ODP}(\exp\{\psi_{i,j}\}, \phi), \tag{20}$$

with

$$\mathbb{E}(I_{i,j}) = \exp\{\psi_{i,j}\}, \;\; \text{Var}(I_{i,j}) = \exp\{\psi_{i,j}\}\phi. \tag{21}$$

Furthermore, the linear predictor $\psi_{i,j}$ is approximated by using some function $f(\boldsymbol{x};\boldsymbol{\gamma})$, where $\boldsymbol{x}$ is a covariate vector and $\boldsymbol{\gamma}$ a parameter vector. The optimization of $f(\boldsymbol{x};\boldsymbol{\gamma})$ will be done by minimizing a loss function L equivalent to the negative Poisson log-likelihood.

In this case, the function $f(\boldsymbol{x};\boldsymbol{\gamma})$ will be defined in accordance with a binary regression tree of depth k. This can be defined as follows. Let $\boldsymbol{X}$ be the feature space and $\boldsymbol{x} \in \boldsymbol{X}$. In the present application, $\boldsymbol{x}$ consists of accident year and development year indices, so that $\boldsymbol{X}$ corresponds to the observed cells of the upper triangle. A binary regression tree of depth k will then divide $\boldsymbol{X}$ into at most 2^k different regions $A_j, j = 1,2,\ldots,2^k$, where each region will be assigned a value δ_j. This leads us to

$$f(\boldsymbol{x};\boldsymbol{\gamma}) = \sum_{j=1}^{2^k} \delta_j \mathbb{1}_{\{\boldsymbol{x}\in A_j\}}. \tag{22}$$

Here, $\boldsymbol{\gamma}$ is the vector of all parameters needed to define A_j, including δ_j, for all j. The optimization problem can then be represented as

$$\delta_j := \arg\min_{\delta} \sum_{i:\boldsymbol{x}_i\in A_j} L(y_i, f(\boldsymbol{x}_i;\boldsymbol{\gamma})), \tag{23}$$

where

$$f(\boldsymbol{x}_i;\boldsymbol{\gamma}) = \delta_j, \;\; i : \boldsymbol{x}_i \in A_j.$$

The estimation of all δ_j will be done using a greedy algorithm. That means, the splitting point of the covariate generating the smallest value of the loss function L for each binary decision will be chosen.

The next step is to introduce the theory of gradient boosting, which is an ensemble machine learning technique widely used in the field of supervised learning, see Friedman (2001). By combining an ensemble of weak learners, e.g. binary regression trees, the GBM algorithm results in a single strong learner. In gradient boosting, the predictor is constructed iteratively as a sum of such base learners. Let $\hat{\mathrm{H}}_0(\boldsymbol{x}) = 0$. At iteration b, the predictor is updated according to

$$\hat{H}_b(\boldsymbol{x}) = \hat{H}_{b-1}(\boldsymbol{x}) + \alpha f(\boldsymbol{x};\hat{\boldsymbol{\gamma}}_b), \tag{24}$$

where α is a hyper-parameter, usually referred to as the learning rate.

The regression tree $f(\boldsymbol{x};\hat{\boldsymbol{\gamma}}_b)$ is obtained by fitting to the negative gradient of the loss function. When it comes to tree-based gradient boosting, the algorithm proceeds as follows. For each iteration $b = 1,2,\ldots,B$, compute

$$g_i = -\frac{\partial}{\partial z} L(y_i, z)\bigg|_{z=\hat{H}_{b-1}(\boldsymbol{x}_i)}, \quad i = 1,2,\ldots,m, \tag{25}$$

and fit a regression tree by solving

$$\hat{\boldsymbol{\gamma}}_b = \arg\min_{\boldsymbol{\gamma}} \sum_{i=1}^{m} (g_i - f(\boldsymbol{x}_i;\boldsymbol{\gamma}))^2. \tag{26}$$

In the Poisson case, the loss function corresponding to the negative log-likelihood is

$$L(y_i, z) = \exp(z) - y_i z, \tag{27}$$

where y_i denotes the observed incremental payment. Substituting this loss function into the expression above yields the negative gradient

$$g_i = y_i - \exp\big(\hat{H}_{b-1}(\boldsymbol{x}_i)\big), \quad i = 1, 2, \ldots, m. \tag{28}$$

In this analysis, the learning rate is fixed at 0.1 throughout. Moreover, after B iterations, the fitted predictor can be expressed in terms of the base learners as

$$\hat{H}_B(\boldsymbol{x}) = \alpha \sum_{b=1}^{B} f(\boldsymbol{x}; \hat{\gamma}_b). \tag{29}$$

In the GBM implementation used in this study, the fitted mean structure from the ODP model is incorporated as an offset. Specifically, the linear predictor is written as

$$\psi_{i,j} = \psi_{i,j}^{\text{ODP}} + \hat{H}_B(\boldsymbol{x}_{i,j}), \tag{30}$$

so that the corresponding mean becomes

$$\mathbb{E}(I_{i,j}) = \exp\{\psi_{i,j}^{\text{ODP}}\} \exp\big(\hat{H}_B(\boldsymbol{x}_{i,j})\big). \tag{31}$$

Using the ODP mean as an offset, the gradient boosting model is restricted to capturing systematic deviations from the parametric baseline, rather than re-estimating the entire mean structure.

When it comes to the estimation of the dispersion parameter, the procedure differs from the classical ODP approach. Since GBM is a non-parametric model and its degrees of freedom are not well defined, the dispersion parameter ϕ is estimated directly as the mean of the squared Pearson residuals, rather than using a degrees-of-freedom adjustment as in (19).

In the next subsection, the procedure used to estimate the optimal number of trees and prevent overfitting is described.

3.7.1 Model estimation and selection of the number of trees

It is well known that due to the complex nature of machine learning techniques, the algorithms may lead to overfitting. This means that during the training procedure,

the model captures not only the underlying structure but also the noise in the data. A common approach to mitigate this risk is to divide the available data into two sets: a training set used for parameter estimation, and a validation set used to evaluate the predictive performance of the model.

In this study, the optimal number of trees, denoted by B^*, is determined through a two-fold cross-validation, as described by Hastie et al. (2009). Each fold contains distinct claims, ensuring that observations belonging to the same claim never appear in both the training and validation samples. This allows the model to be evaluated on entirely unseen claims and provides a realistic assessment of its predictive performance. As with all other model estimation procedures in this study, the selection of B^* is based only on the observed part of the triangle data, that is, the upper triangle shown in Table 2, with both training and validation samples drawn exclusively from these observed incremental payments.

For each fold, a GBM model is trained on one subset and validated on the other. At each boosting iteration $i = 1, 2, \ldots, B$, the Poisson deviance is computed on the out-of-fold predictions and averaged across folds.

Formally, the optimal number of trees is selected as

$$B^* = \arg\min_i D(i),$$

where $D(i)$ denotes the mean out-of-fold deviance at iteration i.

In practice, this procedure is implemented by first training two GBM models, one per fold, using a predefined maximum number of trees (for example 1,000). The Poisson deviance is then evaluated at each iteration, and the iteration that minimizes the mean out-of-fold deviance is selected.

4 Results

4.1 Selecting the number of trees

Choosing an appropriate number of trees is crucial, as it directly affects the stability and predictive performance of the GBM–ODP model. As noted in Section 3.7.1, the number of trees is determined using two-fold cross-validation, selecting the point at which the cross-validated deviance reaches its minimum.

Figure 2 shows the deviance as a function of the number of trees for each Line of Business (LoB). For LoB 1 and LoB 2, the deviance initially decreases as the number of trees increases, before flattening and eventually increasing due to overfitting.

For LoB 3, the deviance rises almost immediately, indicating that the model be-

gins to overfit at a very early stage. In contrast, the deviance curve for LoB 4 flattens after an initial decrease, indicating that additional trees offer only marginal improvements and that overfitting occurs at a relatively large number of trees.

For consistency across the LoBs, the minimum-deviance selection method is applied, and the resulting number of trees for each LoB is reported in Table 4. Alternative tree-selection approaches are explored further in Section 5.1.

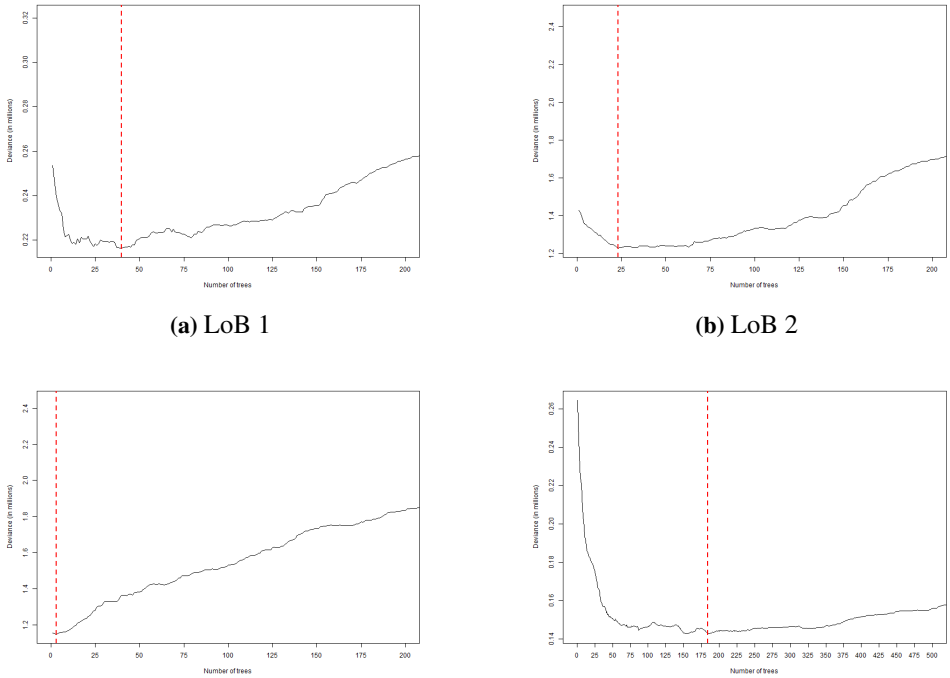

(a) LoB 1 **(b)** LoB 2

(c) LoB 3 **(d)** LoB 4

Figure 2: Deviance as a function of the number of trees for the GBM–ODP model across all LoBs.

	LoB 1	**LoB 2**	**LoB 3**	**LoB 4**
Optimal number of trees	40	23	3	184

Table 4: Optimal number of trees for the GBM–ODP model across the LoBs, based on the minimum-deviance selection method.

4.2 Bootstrap simulations

Table 5 shows the Best Estimate (BE) and VaR based on 100,000 simulations for each model and LoB. The BE corresponds to an estimate of the expected value of future outstanding payments and therefore represents the opening reserve R_0 in the one-year risk framework described in Section 3.2.

The True Reserve (TR), which is also reported in the table, is defined as the sum of the outstanding payments, corresponding to the lower triangle of the synthetic data, see Section 2.

As shown in the table, the BE values are identical for the Mack and ODP models. This is expected, since the Mack model is based on the chain-ladder approach and the ODP model's BE is analytically equivalent to the chain-ladder estimate, as mentioned in Section 3.6. The GBM–ODP model, however, produces different results. For LoB 1 and LoB 2, the GBM–ODP model outperforms the other models by producing estimates closer to the TR.

Model		**LoB 1**	**LoB 2**	**LoB 3**	**LoB 4**
	TR	49,422	77,917	73,443	85,023
Mack–N	BE	44,340 (-10.3%)	72,168 (-7.4%)	78,913 (+7.4%)	88,304 (+3.9%)
	VaR	2,674 (6.0%)	7,574 (10.5%)	5,519 (7.0%)	4,152 (4.7%)
Mack–LN	BE	44,340 (-10.3%)	72,168 (-7.4%)	78,913 (+7.4%)	88,323 (+3.9%)
	VaR	2,664 (6.0%)	7,654 (10.6%)	5,526 (7.0%)	4,146 (4.7%)
Mack–G	BE	44,340 (-10.3%)	72,168 (-7.4%)	78,913 (+7.4%)	88,323 (+3.9%)
	VaR	2,707 (6.1%)	7,674 (10.6%)	5,569 (7.1%)	4,189 (4.7%)
ODP	BE	44,340 (-10.3%)	72,168 (-7.4%)	78,913 (+7.4%)	88,323 (+3.9%)
	VaR	2,288 (5.2%)	5,504 (7.6%)	5,244 (6.6%)	3,899 (4.4%)
GBM–ODP	BE	45,575 (-7.8%)	75,066 (-3.7%)	78,943 (+7.5%)	94,946 (+11.7%)
	VaR	936 (2.1%)	2,349 (3.1%)	4,513 (5.7%)	-1,853 (-2.0%)

Table 5: BE and VaR for each model across LoBs. The first row shows the TR. BE percentage deviations from TR and VaR percentages relative to BE are shown in parentheses in smaller font.

The observed variation in VaR for the Mack models appears to be influenced, to a small extent, by the choice of distribution. Specifically, the gamma distribution consistently produces slightly higher VaR estimates across all LoBs compared to the lognormal and normal distributions. Between the lognormal and normal

distributions, neither consistently yields higher nor lower VaR estimates, as the differences vary across LoBs.

Another conclusion from the table is that the ODP models produce lower VaR estimates compared to the Mack models. Extending the ODP models with gradient boosting decreases the VaR estimates even further for LoB 1 and LoB 2, which may partly be explained by a reduction in residual variability achieved by the GBM. Some of the differences may also be attributed to the models' different approaches to estimating the dispersion parameter ϕ. To enable a fairer comparison between the ODP and GBM–ODP models, the simulations will therefore be repeated using aligned estimators of the dispersion parameter ϕ. This will be carried out in Section 5.2.

In addition to the numerical results, histograms of the bootstrapped loss distributions for LoB 1 are shown in Figure 3. The Mack model under the normality assumption is excluded from the figure, as its distribution is very similar to those obtained under the gamma and lognormal assumptions. All four distributions appear approximately symmetric around zero and resemble a bell-shaped normal distribution. The GBM–ODP model produces the narrowest distribution, which is consistent with its comparatively low VaR estimate in Table 5. The Mack models, by contrast, produce substantially wider distributions. The ODP model lies between these extremes, displaying more variability than the GBM–ODP but less than the Mack models. The same conclusions apply to LoB 2 across all models. For LoB 3 and LoB 4, the conclusions remain valid for all models except the GBM–ODP model.

For LoB 3, the effect of extending the ODP model with gradient boosting is limited. The cross-validation procedure selects a very small optimal number of trees, see Section 4.1. With such a small number of boosting iterations, the fitted model remains close to the ODP offset, resulting in BE and VaR estimates that are similar to those obtained under the ODP model.

Examining the GBM–ODP results for LoB 4 more closely, the table reveals a negative VaR. This is clearly illustrated by the histogram of the GBM–ODP model in Figure 4. The loss distribution is expected to be centered around zero, with both positive and negative realizations reflecting outcomes above and below the BE. In this case, however, all simulated outcomes are negative, suggesting a bias in the opening reserve, as it is consistently higher than the sum of payments during the risk year and the closing reserve. This implies a negative loss in every scenario and, consequently, a negative VaR. This is not consistent with Solvency II, where the risk measure is required to reflect a potential loss and therefore cannot be negative.

The following section further investigates the effects of the alternative tree-selection method and the choice of dispersion estimator. In addition, the influence of the

number of simulations on the results is examined. The analysis presented so far has been based on 100,000 simulations, and by repeating the procedure with varying simulation counts, the sensitivity of the outcomes can be assessed. Finally, the impact of the simulation seed is explored to determine the extent to which the choice of seed affects the results.

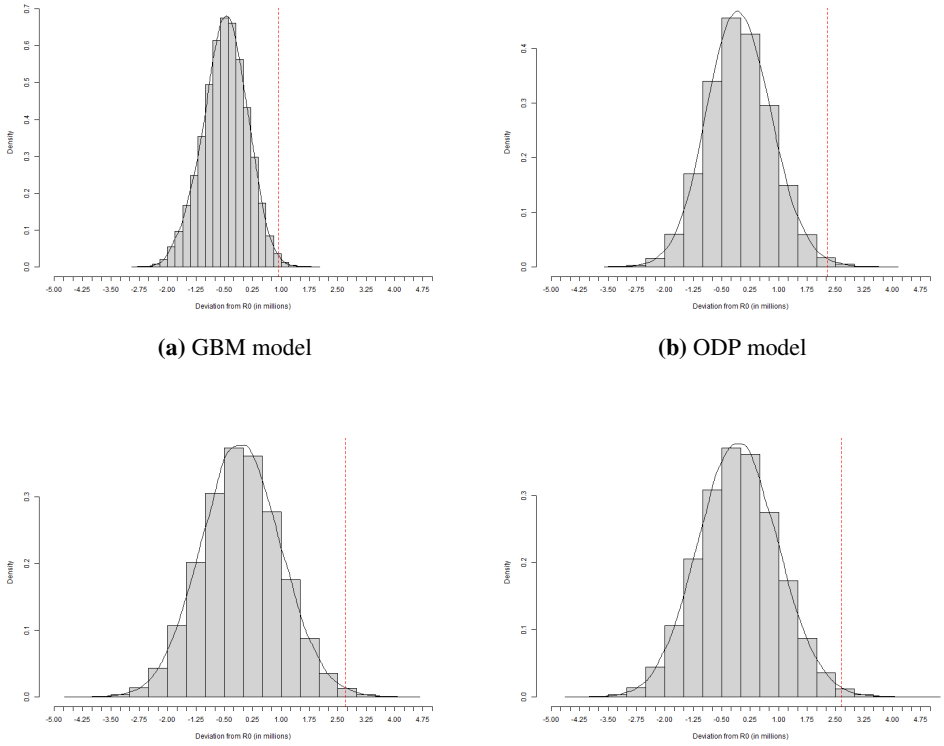

(a) GBM model

(b) ODP model

(c) Mack-gamma model

(d) Mack-lognormal

Figure 3: Histograms illustrating the loss distribution for LoB 1 under the respective models. The vertical red dashed line indicates the 99.5th percentile in each distribution.

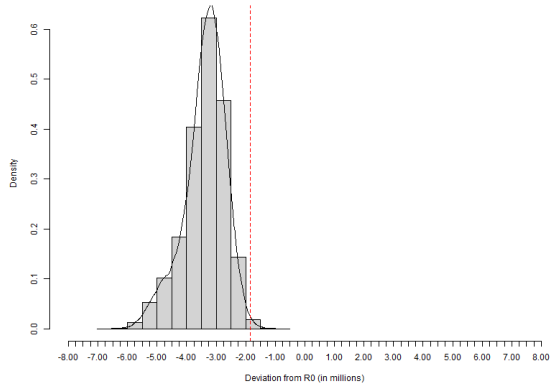

Figure 4: Histogram illustrating the loss distribution for LoB 4 under the GBM–ODP model. The vertical red dashed line indicates the 99.5th percentile.

5 Analysis

5.1 Alternative tree selection methods

In Table 5, the number of trees for the GBM–ODP models has been chosen based on where the cross-validated deviance reaches its minimum. In this subsection, two alternative selection criteria are explored. The first is the *1-SE rule*, which selects the least complex model whose cross-validated deviance lies within one standard error of the minimum. The second approach considered is the *elbow method*, which chooses the point where the deviance curve begins to flatten, indicating diminishing returns from adding additional trees. As this criterion is based on a visual assessment of the curve's shape, the resulting choice is approximate and may vary slightly depending on interpretation.

These two approaches are applied to all LoBs where they are suitable, allowing comparison with the minimum-deviance method. For LoB 3, however, the elbow method is not applicable because the deviance increases almost immediately.

Method	**LoB 1**	**LoB 2**	**LoB 3**	**LoB 4**
Minimum-Deviance	40	23	3	184
1-SE-rule	24	22	1	152
Elbow method	10	20	-	50

Table 6: Selected number of trees based on different selection methods across the LoBs.

Table 7 shows that, for LoB 1 and LoB 2, the minimum-deviance criterion yields BE values that are closer to the TR than those obtained using the 1-SE rule or

the elbow method, indicating no improvement from applying alternative selection criteria in these cases. For LoB 3, the 1-SE rule selects a single tree, yielding BE and VaR that are slightly closer to the ODP model, which is expected, given that the ODP estimate enters the GBM–ODP model as an offset. For LoB 4, the elbow method yields a positive VaR estimate, unlike the negative VaR obtained under the minimum-deviance and 1-SE criteria, and both alternative selection methods produce BE values closer to the TR.

Method		**LoB 1**	**LoB 2**	**LoB 3**	**LoB 4**
	TR	49,422	77,917	73,443	85,023
Minimum deviance	BE	45,575 (-7.8%)	75,066 (-3.7%)	78,943 (+7.5%)	94,946 (+11.7%)
	VaR	936 (2.1%)	2,349 (3.1%)	4,513 (5.7%)	-1,853 (-2.0%)
1-SE rule	BE	45,138 (-8.7%)	75,033 (-3.7%)	78,946 (+7.5%)	94,031 (+10.6%)
	VaR	1,120 (2.5%)	2,375 (3.2%)	4,813 (6.1%)	-871 (-0.9%)
Elbow method	BE	44,579 (-9.8%)	74,965 (-3.8%)	-	91,184 (+7.3%)
	VaR	1,604 (3.6%)	2,381 (3.2%)	-	1,323 (1.5%)

Table 7: TR together with BE and VaR for each selection method across all LoBs. BE percentage deviations from TR and VaR percentages relative to BE are shown in parentheses in smaller font.

5.2 Alignment of dispersion parameter estimation

In Sections 3.6 and 3.7, the ODP and GBM–ODP models rely on different approaches to estimate the dispersion parameter ϕ. To enable a more meaningful comparison between the two models, both are fitted using the mean of the squared Pearson residuals, which was previously applied only to the GBM–ODP model. This alignment ensures that any remaining differences in performance can be attributed to the model structures themselves rather than to differences in the estimation of the dispersion parameter ϕ.

Table 8 presents the VaR estimates for the ODP model using the different approaches. When aligning the methods, the VaR estimates decrease by roughly 10% across all LoBs compared to the original approach. This indicates that part of the difference previously observed between the ODP and GBM–ODP models can be explained by the choice of estimator.

Model	Estimator	LoB 1	LoB 2	LoB 3	LoB 4
ODP	original ϕ	2,288	5,504	5,244	3,899
ODP	sq. Pearson mean ϕ	2,058	4,900	4,752	3,480
GBM–ODP	sq. Pearson mean ϕ	936	2,349	4,513	-1,853

Table 8: VaR estimates by model and dispersion-parameter estimation method across LoBs.

5.3 Convergence test of simulated VaR estimates

Observing convergence as the number of simulations increases is important for assessing both the reliability and efficiency of the results. When the estimated quantities stabilize and show minimal variation with additional simulations, it indicates that random fluctuations are diminishing and that the estimates are approaching their true theoretical values. From a computational perspective, examining convergence is also useful for determining an appropriate balance between accuracy and calculation time.

To evaluate the convergence of the VaR simulations for the different models, the incremental percentage change in VaR between successive simulation sizes is calculated for each LoB. Since perfect convergence is difficult to achieve in practice, an appropriate threshold must be chosen to determine when convergence is considered sufficient. In this analysis, thresholds of $\pm 1\%$ and $\pm 2\%$ were considered.

Figure 5 presents the results, showing the incremental percentage change in VaR as the number of simulations increases. It is difficult to draw far-reaching conclusions when comparing the models, as the variation across LoBs and simulation levels remains substantial and the convergence patterns are not fully consistent. Nevertheless, a general tendency can be observed in the reduction in variability with increasing simulation size and the overall stabilisation within the $\pm 2\%$ threshold beyond 40,000 – 50,000 simulations. At the same time, only a limited number of model–LoB combinations remain within the $\pm 1\%$ threshold at 80,000 – 100,000 simulations. A corresponding convergence plot for the Mack model under the normality assumption is provided in the appendix.

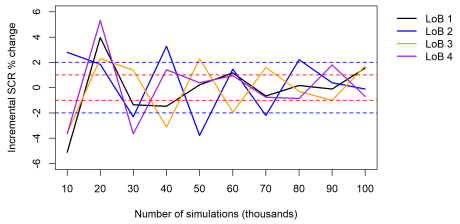

(a) Mack (Gamma)

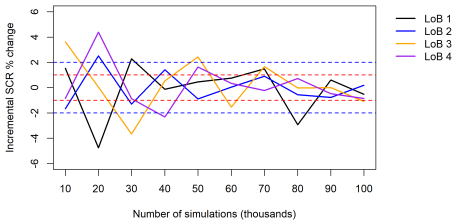

(b) Mack (Lognormal)

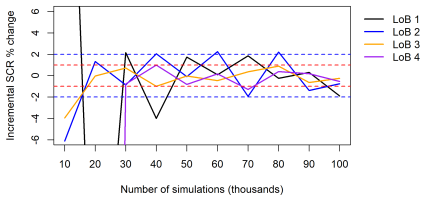

(c) GBM–ODP

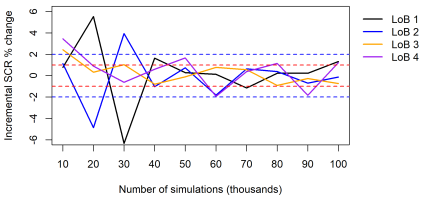

(d) ODP

Figure 5: Comparison of VaR percentage increases across simulation sizes for the Mack (gamma and lognormal), GBM–ODP, and ODP models. The red dashed lines denote the ±1% deviation threshold, while the blue dashed lines indicate the ±2% threshold.

5.4 Effect of the data simulation seed

This subsection examines how the model outcomes change when the underlying claims data are regenerated using different seeds in the simulation machine described in Section 2. Only the Mack-gamma model is included among the Mack models, both to keep the comparison manageable and because earlier results showed that the choice of distributional assumption within the Mack framework had only a limited impact on the outcomes.

As shown in Table 9, which includes seed 12345 used in Section 4.2, the GBM–ODP model yields BE values for LoB 1 and LoB 2 that are consistently closer to the TR across all seeds than the corresponding Mack–gamma and ODP estimates.

The GBM–ODP model also generates negative VaR estimates for seeds 100 and 200, as previously observed and discussed for seed 12345 in Section 4.2. Apart from this, the GBM–ODP model produces the lowest VaR estimates in most cases, while the Mack–gamma model consistently yields the highest VaR estimates and the ODP model generally falls in between.

			BE / VaR		
Seed	**LoB**	**TR**	**Mack–gamma**	**ODP**	**GBM–ODP**
12345	1	49,422	-10.3% / 6.11%	-10.3% / 5.16%	-7.8% / 2.05%
12345	2	77,917	-7.4% / 10.63%	-7.4% / 7.63%	-3.7% / 3.13%
12345	3	73,443	+7.4% / 7.06%	+7.4% / 6.65%	+7.5% / 5.72%
12345	4	85,023	+3.9% / 4.74%	+3.9% / 4.41%	+11.7% / -1.95%
100	1	48,462	-2.73% / 6.16%	-2.73% / 4.34%	-0.65% / 3.02%
100	2	80,492	-2.90% / 9.36%	-2.90% / 7.12%	+1.31% / 5.93%
100	3	89,532	-13.26% / 11.76%	-13.26% / 7.94%	-8.79% / 4.15%
100	4	91,779	-3.77% / 5.62%	-3.77% / 4.24%	+4.86% / -3.03%
200	1	47,707	-6.63% / 5.61%	-6.63% / 4.78%	-0.56% / -1.21%
200	2	82,761	-2.73% / 11.46%	-2.73% / 8.50%	-2.68% / 7.54%
200	3	68,912	+9.68% / 6.17%	+9.68% / 5.98%	+11.07% / 3.81%
200	4	84,926	-6.46% / 4.03%	-6.46% / 3.25%	-6.10% / 2.34%

Table 9: TR, BE deviations from TR, and VaR as a percentage of BE across LoBs for three different random seeds.

5.5 Discussion

The purpose of this study has been to compare the performance and behaviour of different reserve risk models using synthetic claims data generated from a stochastic simulation machine. The analysis focused on two key outputs, the BE and the VaR, and examined how these quantities vary across LoBs.

The results indicate systematic differences in the dispersion of the loss distributions produced by the models. The GBM–ODP approach tends to generate narrower distributions, which in turn leads to lower VaR estimates. Its BE values are also, in many cases, closer to the TR. By contrast, the ODP and Mack models generally exhibit wider distributions, with the Mack model typically producing the largest dispersion and the highest VaR estimates.

Although these findings point to potential advantages of the GBM–ODP approach in terms of both accuracy and capital efficiency, the results also highlight the need for further investigation. In particular, the occurrence of negative VaR estimates in some cases shows that the predictive distribution is not centred around zero, as expected. This behaviour may be related to aspects of the fitting procedure. Further work could therefore investigate the robustness of the GBM–ODP approach under alternative tuning choices, regularisation strategies and validation procedures.

In addition, it could be informative to further decompose the results by accident year and development period. Such a breakdown would provide more granular insight into how the different models behave across the triangle and could help identify whether differences in performance are driven by specific calendar effects, early or late development periods, or particular accident years.

Overall, the study suggests that the GBM–ODP approach has promising properties compared to more established reserve risk models, but that additional analysis is required before drawing firm conclusions about its suitability for reserve risk modelling.

References

European Parliament and Council of the European Union (2009). *Directive 2009/138/EC (Solvency II)*. Official Journal of the European Union, L 335, 17 December 2009, 1–155.

A. Gabrielli and M. V. Wüthrich (2018). *An Individual Claims History Simulation Machine*. Risks, 6(2), 29.

T. Mack (1993). *Distribution-free Calculation of the Standard Error of Chain Ladder Reserve Estimates*. ASTIN Bulletin, 23(2), 213–225.

E. Ohlsson and J. Lauzeningks (2009). *The one-year non-life insurance risk*. Insurance: Mathematics and Economics, 45(2), 203–208.

P. D. England and R. J. Verrall (2006). *Predictive Distributions of Outstanding Liabilities in General Insurance*. Annals of Actuarial Science, 1(II), 221–270.

P. D. England and R. J. Verrall (2002). *Stochastic claims reserving in general insurance*. British Actuarial Journal, 8(3), 443–544.

M. Lindholm, R. J. Verrall, F. Wahl and H. Zakrisson (2020). *Machine Learning, Regression Models, and Prediction of Claims Reserves*. Casualty Actuarial Society E-Forum, Summer 2020.

J. H. Friedman (2001). *Greedy Function Approximation: A Gradient Boosting Machine*. Annals of Statistics, 29(5), 1189–1232.

T. Hastie, R. Tibshirani and J. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York.

A Appendix

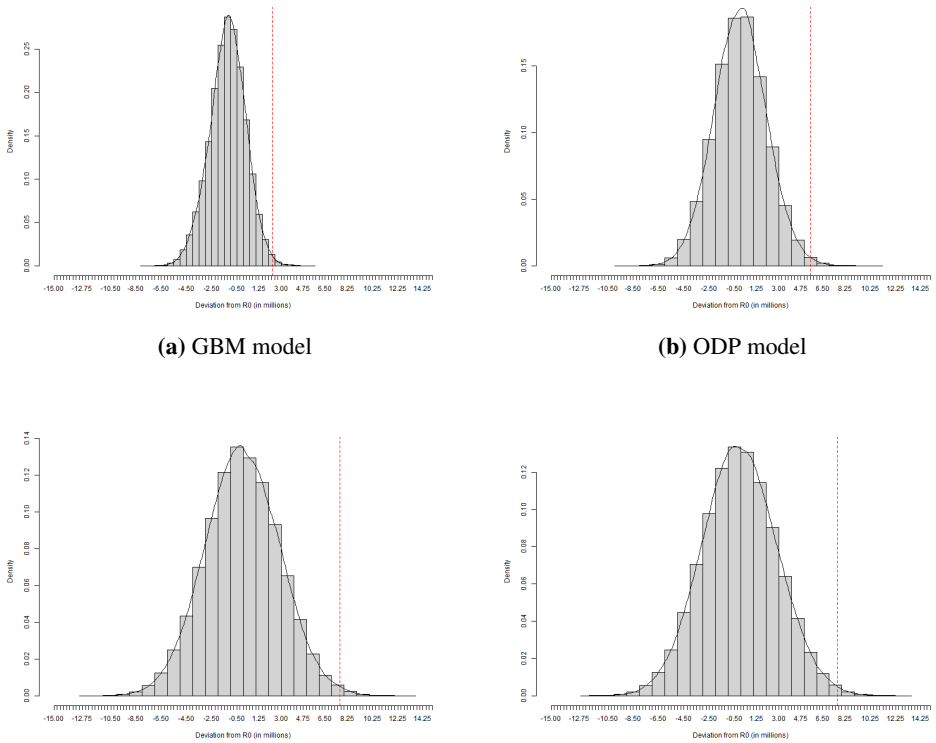

(a) GBM model

(b) ODP model

(c) Mack-gamma model

(d) Mack-lognormal

Figure 6: Histograms illustrating the distribution of the deviation from R_0 for LoB 2 under the respective models. The vertical red dashed line indicates the 99.5th percentile in each distribution.

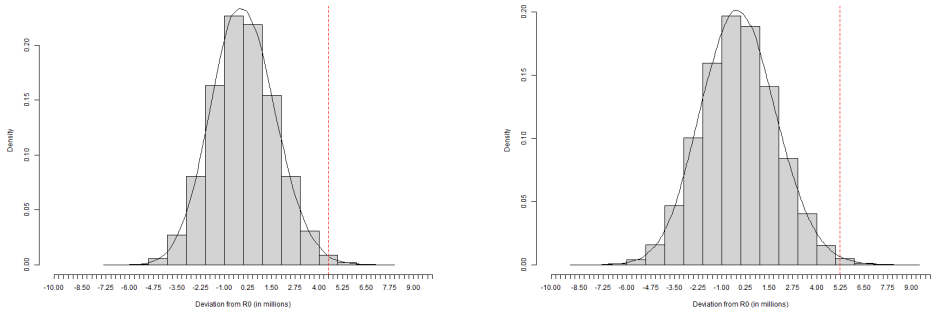

(a) GBM model

(b) ODP model

(c) Mack-gamma model

(d) Mack-lognormal

Figure 7: Histograms illustrating the distribution of the deviation from R_0 for LoB 3 under the respective models. The vertical red dashed line indicates the 99.5th percentile in each distribution.

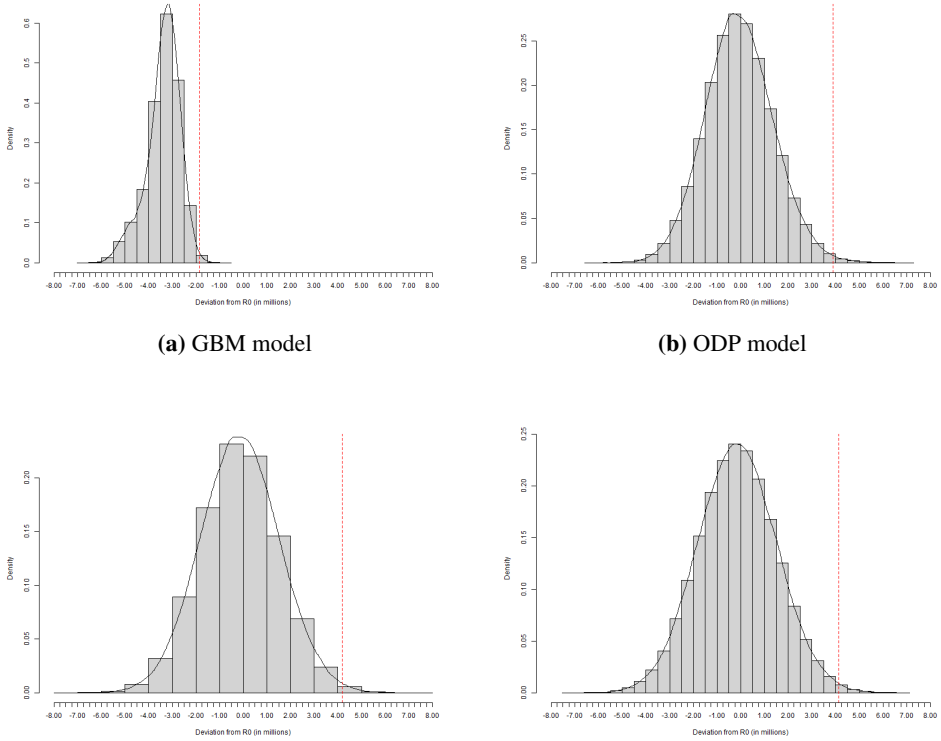

(a) GBM model

(b) ODP model

(c) Mack-gamma model

(d) Mack-lognormal

Figure 8: Histograms illustrating the distribution of the deviation from R_0 for LoB 4 under the respective models. The vertical red dashed line indicates the 99.5th percentile in each distribution.

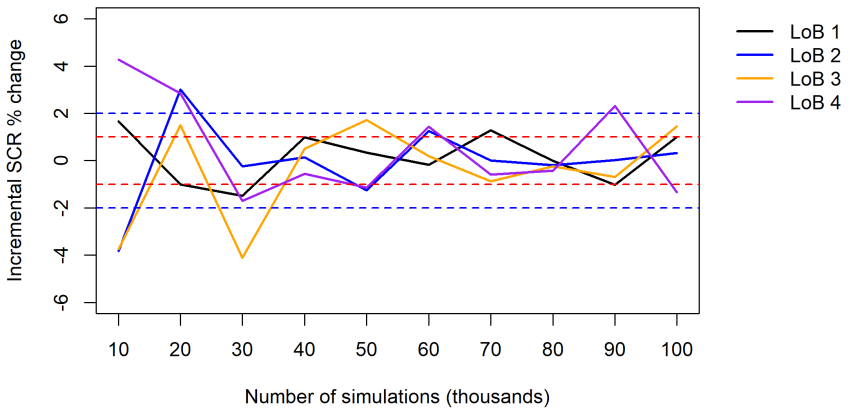

Figure 9: Comparison of SCR percentage increases across simulation sizes for the Mack model under normality assumption. The red dashed lines denote the ±1% deviation threshold, while the blue dashed lines indicate the ±2% threshold.