

MATRIX CALCULATION FOR ULTIMATE AND ONE-YEAR RISK IN THE SEMI-MARKOV INDIVIDUAL LOSS RESERVING MODEL

Carole Bettonville Louise d'Oultremont
Detralytics
Brussels, Belgium

Michel Denuit
Institute of Statistics, Biostatistics and Actuarial Science
Louvain Institute of Data Analysis and Modeling
UCLouvain
Louvain-la-Neuve, Belgium

Julien Trufin
Department of Mathematics
ULB
Brussels, Belgium

Robin Van Oirbeek
Belfius Insurance
Brussels, Belgium

November 4, 2020

Abstract

This paper proposes a multistate model with a Semi-Markov dependence structure describing the different stages in the settlement process of individual claims in general insurance. Every trajectory, from reporting to closure is combined with a modeling of individual link ratios to obtain the ultimate cost of each claim. Analytical expressions are derived for the moments of ultimate amounts whereas quantile risk measures can be obtained by simulation. In the one-year view, the proposed matrix calculations avoid the simulation-within-simulation issue and offer a tractable evaluation method. A case study illustrates the relevance of the proposed approach.

Keywords: IBNR, RBNP, RBNS, loss development, technical provisions, solvency calculation, financial reporting.

1 Introduction

In general insurance, it may take some time before the policyholder reports a claim to the insurer. Also, the claim sometimes cannot be finalized by the end of the accounting period but requires more time to be settled. Meanwhile, the insurer must constitute a reserve, or provision representing the future costs to be paid in relation to this claim.

Within the Solvency 2 prudential framework as well as in the new IFRS environment, the accurate evaluation of future cash flows and technical provisions has become increasingly important to assess the financial strength of insurance portfolios. It is also important to reconcile risk and reserving models so that all actuarial evaluations remain consistent. The individual claim development model proposed in this paper effectively addresses these challenges.

Individual loss reserving aims to model the settlement dynamics for each claim reported to the insurer. Compared to aggregate methods using triangular data only, which summarize all the experience with global figures cross-classified by accident year and development period, individual or micro-reserving requires granular data recording for each claim the sequence of cash-flows from reporting to closure. The literature devoted to individual loss reserving, or micro-reserving can be traced back to the 1990s with the development of a mathematical framework in continuous time by Arjas (1989) and Norberg (1993, 1999). Since then, many models have been proposed and we refer the interested reader to Charpentier and Pigeon (2016) for a comparison and Taylor (2018, 2019) for the evolution from aggregate to individual models. For a textbook treatment of loss reserving, see Wuthrich and Merz (2008, 2015).

In this paper, we adopt the multistate approach to loss reserving proposed by Hachemeister (1980) and further considered e.g. by Hesselager (1994), Hurlimann (2015), and Antonio et al. (2016). The goal is to design a stochastic model for trajectories until final settlement so that the associated cash flows can be used to determine the required reserve for outstanding claims.

Between occurrence and closure, a claim may go through several states reflecting successive stages in its development. Precisely, a claim is first said to be Incurred But Not Reported (IBNR, in short) between the occurrence of the insured event and its notification to the insurance company. During the IBNR period, the insurer is liable for the claim amount but is unaware of the claim's existence. Once notified, the company is aware of the claim but it may take some time before the first payment (if any) is made. The claim is then said to be Reported But Not Paid (RBNP, in short), meaning a reported claim for which no payments have been made yet. Unless the claim closes without payment, it enters the Reported But Not Settled (RBNS, in short) stage where it stays from the first payment until closure, as long as its final cost remains unknown.

In this paper, we adopt a discrete-time approach and model the reporting delay, the RBNP stage, and the times between consecutive payments by discrete random variables representing occupation times in a multistate process. The state space consists in an IBNR state, a RBNP state, a cascade of RBNS states, and two final states corresponding to closure with or without terminal payment. The occupation times in each state are studied with the help of discrete hazard rate that has initially been developed for two states, only, but can easily be extended to the multistate setting. The state-specific discrete hazard rate function

corresponds to the probability that a transition takes place given that the claim has stayed in the current state for some time.

As in Hesselager (1994), we assume that payments occur at the times of transition between states. Individual development factors, or link ratios together with an initial payment structure the cash-flows along the development pattern of each claim. Several candidate distributions are available to model the resulting development factors at an individual claim level. In this paper, we allow for general severity distributions, using the GAMLSS models that were successfully applied in an insurance context by Klein et al. (2014).

The reserving model proposed in this paper extends the approach proposed by Antonio et al. (2016) in several directions:

- firstly, we allow for duration dependence by switching from a Markov modeling to a Semi-Markov setting. The duration effect appears to be significant in the case study worked out in the present paper so that this extension improves the modeling of claim dynamics.
- secondly, we explicitly account for the IBNR part in loss reserving calculations by estimating the occurrence intensity and the distribution of reporting delay.
- thirdly, we explain how to obtain moments of outstanding losses with the help of matrix calculations in an augmented state space, avoiding simulation. This is especially useful when the one-year risk view is adopted because we do not have to deal with the simulation-within-simulation issue. Specifically, given the state occupied at the beginning of the evaluation year, the state reached by every claim at the end of that year is simulated together with the corresponding payments that emerge over the year, if any. At the end of the evaluation year, the expected ultimate loss is computed using matrix formulas starting from the states occupied by the claims and the cumulative payments at the end of the year obtained by simulation. In other words, an outer Monte Carlo simulation gathers payments that emerge over the evaluation year with the corresponding inner estimates (computed analytically) of conditional expected ultimate loss at the end of the year, making inner simulations unnecessary.

As an illustration, we analyze simulated data generated by the “Individual Claims History Simulation Machine” contributed by Gabrielli and Wuthrich (2018). This tool simulates individual cash-flows from occurrence to final settlement. Both the best estimate and the one-year risk of the resulting portfolio are considered. We also perform a case study based on real-life data from a motor insurance portfolio considered in Denuit and Trufin (2017, 2018). The multistate loss reserving model is embedded in the hybrid approach proposed in Denuit and Trufin (2018). Specifically, the individual multistate modeling is applied to those claims with longer path to settlement. The majority of claims that are rapidly reported and closed are still modeled collectively (fitting the actuarial model to individual observations) and the detailed multistate modeling is used for claims with a longer settlement process, only. This is particularly effective from a computational point of view since it avoids to simulate numerous short trajectories corresponding to claims that are rapidly reported and settled (during the accident year and the year after, in this example), concentrating the effort on the more complicated cases that generally also turn out to be more expensive. This hybrid approach appears to be a good compromise between accuracy and computing effort.

The remainder of this paper is organized as follows. We introduce the multistate loss reserving model in Section 2. Section 3 is devoted to numerical illustrations performed on simulated data to demonstrate the advantage of the proposed approach compared to the traditional Chain Ladder loss reserving procedure. Section 4 offers a case study based on individual claim developments in a motor third-party liability (MTPL) insurance portfolio. The final Section 5 briefly concludes the paper.

2 Model

2.1 Multistate approach

The different stages in the settlement process are modeled using a discrete-time multistate process. Time is measured in discrete steps, corresponding to periods $(t - 1, t)$, $t = 1, 2, \dots$, that can represent weeks, months, quarters, semesters or years. At the end of each time step $t = 1, 2, \dots$, the state occupied by the claim is recorded so that we model each specific settlement dynamics. Transitions are accompanied by cash-flows, payments or refunds. Because we work in discrete time, all the cash-flows of a period are aggregated into a single value. Henceforth, as in the majority of applications considered in the literature, we assume without loss of generality that time steps correspond to accounting, or calendar years.

Consider a claim related to an insured event occurring during a given accident year, taken as time origin 0. The stochastic process $\mathcal{S} = \{S_t, t = 0, 1, 2, \dots\}$ describes the claim settlement procedure, with $S_0 = \text{IBNR}$. The meaning is that the claim moves to state S_t during period t , that is, the transition takes place between time $t - 1$ and time t , $t = 1, 2, \dots$. Hence, we actually have two time scales: calendar time and time elapsed since accident year. Depending on the calculation, one time scale or the other may be more convenient. The state space and possible transitions are represented in Figure 2.1. Sometimes, several transitions take place within the same year (from IBNR to RBNP and then to RBNS, for instance) but we only record the state occupied at the end of the period. We assume that all claims are settled after a maximum number of payments n , such that we consider RBNS states indexed from 1 to n . Given that a claim can stay for several periods in each RBNS state, the total time to settlement may be longer than n . The value of n does not need to be specified in the application as letting n tend to infinity can be considered as a tail factor capturing unusually long developments.

2.2 Trajectories

The IBNR state is the initial one, where the claim stays until it is reported to the insurer. Once reported to the insurer, the claim leaves the IBNR state and there are several possibilities:

- either the claim is reported, is not terminated during the reporting period and the insurer does not make any payment during the reporting period. Then the claim moves from the IBNR state to the RBNP one. It remains there as long as it stays open but the insurer does not pay anything to beneficiaries. When leaving the RBNP state,

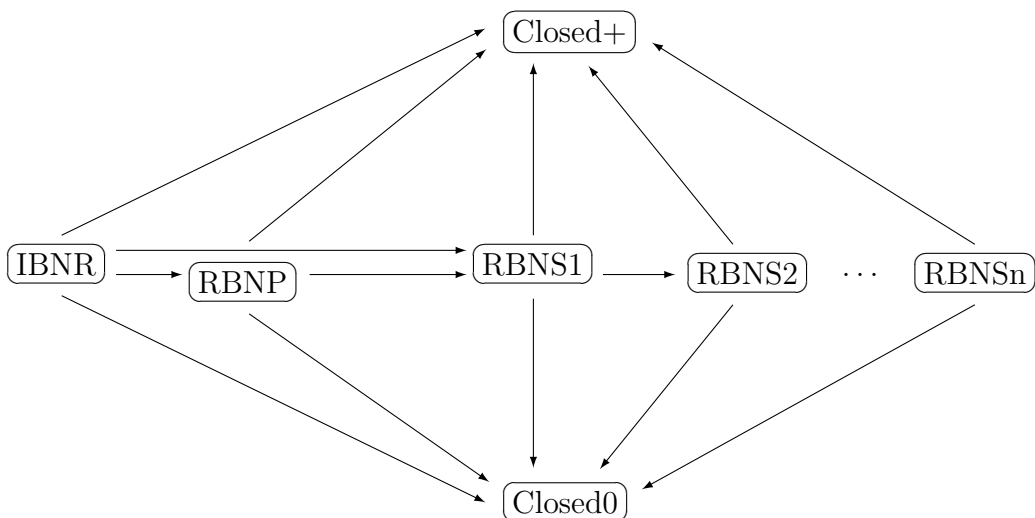


Figure 2.1: Multistate model describing the claim’s settlement process, from occurrence to closure.

the claim may terminate with or without payment, or payments can start so that it enters the first RBNS state.

- or it is reported, is not terminated during the reporting period and the insurer makes an initial payment during the reporting period. Then, the claim moves directly to the first payment state RBNS1. If the insurer makes at least one payment in addition to the terminal one then the claim enters the cascade of RBNS states shown in Figure 2.1. More specifically, the claim moves along the series RBNS1, RBNS2, ... of RBNS states until closure, with a transition occurring each time an additional payment is made.
- or the claim is terminated during the year it is reported and directly moves from the IBNR state either to the terminal state with payment Closed+ or to the terminal state without payment Closed0, both corresponding to closure. The difference between the two types of closure thus depends on the presence of a final payment upon termination.

Here, we do not allow for possible re-opening of a closed claim so that Closed0 and Closed+ are final states.

2.3 Claim occurrence rate and reporting lag

The numbers of insured events taking place in each accident year are assumed to be mutually independent, Poisson distributed. The means are allowed to vary between accident years. This specification is in line with claim occurrence times obeying a non-homogeneous Poisson process (after Hesselager, 1994). The time-varying Poisson means account for time trends or changes in the size of the portfolio. If the time step is less than one year, time-varying

Poisson means can also capture seasonality effects (more traffic accidents generally occur during winter periods in motor insurance, for instance).

The IBNR state is hidden: the insurer only observes claims when they are reported, that is, when they leave the IBNR state. At that time, the insurer gets information about the insured event, including its occurrence date. Thus, the reporting lag can be determined and the claim can be allocated to the right accident period. The next model can be used to estimate the distribution of the reporting delay and to infer the number of IBNR claims accounting for the censoring mechanism due to late reporting. It is based on a standard aggregated run-off triangle filled with numbers of claims cross-classified according to accident, or occurrence period and reporting lag.

Precisely, let us denote as N_{tj} the number of claims relating to insured events that occurred in accident year t , that is, from calendar time $t - 1$ to t and were reported to the insurer at development j , that is, during calendar year $t + j - 1$. Development 1 thus corresponds here to the accident year.

In line with the classical Chain Ladder model, we use the multiplicative specification

$$\mathbb{E}[N_{tj}] = \alpha_t \beta_j \text{ subject to } \sum_{j \geq 1} \beta_j = 1 \quad (2.1)$$

and we assume that the random variables N_{tj} are independent and Poisson distributed. Under these assumptions, the estimations agree with the output of the Chain Ladder algorithm. This also ensures that the total number $N_t = \sum_{j \geq 1} N_{tj}$ of claims for accident year t also obeys the Poisson distribution with mean $\mathbb{E}[N_t] = \alpha_t$, and that β_j is the probability that a claim is reported at lag j , that is $j - 1$ periods after its occurrence. The parameters α_t and β_j can be estimated from the triangle with observed N_{tj} using Poisson regression. But other models could also be used in case Poisson specification is not appropriate for the data under consideration, such as Negative Binomial for instance, and alternative mean functions can be considered in (2.1) in case simple products $\alpha_t \beta_j$ do not capture the reporting dynamics.

The expected claim number α_t for each accident year t allows the actuary to track possible trends (portfolio volume or market trends). Notice that a risk exposure (number of years of risk coverage, for instance) is useful to perform inference as α_t is the expected number of claims meaning the product between the claim rate per unit of exposure and the corresponding risk exposure. For the evaluation of the technical provisions, only the expected numbers of claims α_t for past occurrence periods matter. If available, a volume measure can be included as an offset in the Poisson regression analysis allowing to estimate the coefficients α_t and β_j .

Notice that here, we have assumed some stationarity in the reporting pattern, in that the same proportions β_j apply to all accident years. In case the analyst suspects that some change may have occurred, which speeds up or slows down the reporting of claims to the insurer, then different β_j can be specified depending on the accident year.

The distribution of the reporting lag, that is, of the sojourn time in state IBNR, can then be deduced from the resulting β_j . Precisely, the probability that the claim is reported at, or before development j is then given by

$$\mathbb{P}[S_j \neq \text{IBNR} | S_0 = \text{IBNR}] = \beta_1 + \dots + \beta_j, \quad j = 1, 2, \dots$$

This is the probability that a claim corresponding to an insured event that occurred in calendar year t is reported to the insurer in calendar year $t + j - 1$ at the latest.

2.4 Cash-flows

Cash-flows accompany transitions between states. At transition from IBNR or RBNP to the first RBNS state RBNS1, the insurer makes a first, non-terminal payment. Remember that we work in discrete time, such that several payments can actually be made during the development year and the total amount paid during this period is recorded as the first payment. Then, each time the insurer makes at least one payment in a period, the claim moves to the next RBNS state, until final settlement.

Cash-flows are modeled using an initial payment P_1 possibly multiplied with a sequence of link ratios describing the evolution of cumulative payments to the final cost. More specifically, when the claim enters the RBNS1 state, the insurer makes a first payment P_1 . This first payment may subsequently be followed by payments P_2, P_3, \dots along the RBNS cascade, where payment P_k is made when the claim enters the RBNS k state. The cumulative amount already paid by the insurer for a claim in state RBNS j is then given by

$$C_j = \sum_{k=1}^j P_k.$$

Notice that some of the payments P_k may be negative. For instance, this happens when the involved third party is finally recognized liable for the accident so that it must refund the insurer for the payments made so far. In such a case, because some of the costs for the insurer may never be reimbursed by the third party (such as the internal settlement expenses), the ultimate cumulative amount can still remain strictly positive. To account for negative payments, actuaries often rely on link ratios Λ_j defined as

$$\Lambda_j = \frac{C_{j+1}}{C_j}, \quad j = 1, 2, \dots$$

The advantage of working with such link ratios is that they can easily deal with negative payments by taking values smaller than one. Of course, this is not the only possible approach to model cash-flows related to each claim. Paid amounts P_2, P_3, \dots can also be modeled, instead of link factors Λ_j . Negative payments can be modeled with the help of finite mixture distributions (see Section 3 for more details).

Each claim is accompanied with its own sequence of cash-flows described by the random variables $P_1, \Lambda_1, \Lambda_2, \dots$ such that

$$C_1 = P_1 \text{ and } C_j = P_1 \prod_{k=1}^{j-1} \Lambda_k \text{ for } j \geq 2.$$

The link ratios Λ_j have a distribution specific to each RBNS state RBNS j . They are independent from one claim to another (so that amounts must have been corrected for inflation beforehand, for instance). For a given claim, we assume that these random variables are independent for different values of j , given the features included in the analysis. We will see

in the application that some dependence between P_1 and Λ_j is nevertheless needed to prevent final costs that are unrealistically large, without economic relevance. This is because a large initial payment P_1 is generally accompanied with smaller link ratios compared to smaller initial payments.

Notice that if the claim does not enter the RBNS states, being terminated with a single payment, then link ratios are not needed. In case the claim moves directly from IBNR or RBNP to Closed+ there is only a single payment P_1 made by the insurer. For convenience, link ratios are set to 1 for states with no associated cash-flow.

The multistate model described in this paper is a payment-to-payment model, in the sense that each transition corresponds to an actual payment made by the insurer. For this reason, there is no need to include a positive probability mass at the origin for the initial payment, or at 1 for the link ratios, which considerably eases the modeling.

2.5 Analytical calculation of moments

In this section, we show that all quantities that can be expressed as mathematical expectations can be computed analytically once the Semi-Markov model has been turned into a Markov one by adding artificial states. This is particularly useful when the reserve is calculated in terms of moments of the portfolio. However, this approach does not work for quantities like quantile for instance (simulation can be used, instead).

In the Markov setting, it is well documented that calculations can easily be performed using products of matrices containing the transition probabilities between each pair of states. In the Semi-Markov modeling proposed in this paper, calculations can still be performed in matrix form once artificial states are created to recover the Markov, memoryless property. This simple idea exploits the well-known fact that a process $\{S_t, t = 0, 1, 2, \dots\}$ is Semi-Markov if, and only if, the bivariate process $\{(S_t, D_t), t = 0, 1, 2, \dots\}$ is Markov, where D_t denotes the time spent in the current state S_t . If no duration effect is present in the data, matrix calculations can be performed directly, without augmenting the state space.

The idea is thus to split every Semi-Markov state, that is, every state where sojourn time D_t influences transition probabilities, into a series of artificial Markov states: the first one corresponds to occupying the physical state (RBNP or RBNS j) for 1 year, the second one corresponds to occupying the physical state (RBNP or RBNS j) between 1 and 2 years, and so on until no duration effect remains. Ultimately, exit probabilities do no more depend on the time spent in the current state and a Geometric tail distribution is used for the remaining sojourn time. A transition to the physical, Semi-Markov state RBNP or RBNS j means that the claim enters the first artificial state in the series. The first artificial states can only be occupied for one year but the last one corresponding to the Geometric tail distribution can be occupied for a longer time. The ultimate Geometric decay can thus be seen as a kind of tail factor. Alternatively, an upper bound can be specified on the length of stay in a given state (such as IBNR). The Semi-Markov model then becomes Markovian in this augmented state space.

Let \mathbf{M} denote the matrix containing the transition probabilities in the augmented state space. Also, denote as $\boldsymbol{\omega}$ the vector with 0 entries, except the one corresponding to the state currently occupied in the augmented state space which is equal to 1. In other words, the position with a 1 indicates the current state of the claim in the augmented state space, that

depends on both S_t and D_t . The probability to reach any state at horizon k can then be calculated as $\boldsymbol{\omega}^\top \mathbf{M}^k$. The corresponding probabilities in the Semi-Markov model describing the real settlement process are easily obtained by summing all the entries in the vector $\boldsymbol{\omega}^\top \mathbf{M}^k$ corresponding to the Semi-Markov state under consideration.

To obtain the corresponding costs, let us define the matrix $\widetilde{\mathbf{M}}$ where each entry of \mathbf{M} corresponding to a transition probability between states IBNR, RBNP, RBNS or Closed+ is multiplied with the corresponding $E[P_1]$ or $E[\Lambda_j]$. The element of \mathbf{M} is left unchanged if there is no corresponding transition between physical states or the transition is not accompanied with a cash flow. For a claim occupying a RBNS state, the observed initial payment and link factors are used instead of expected values. Summing the elements of $\boldsymbol{\omega}^\top \widetilde{\mathbf{M}}^k$ where k corresponds to the longest remaining time to settlement gives the expected ultimate cost, given the current development reflected in $\boldsymbol{\omega}$. Specifically,

- for a claim in state RBNS j at time t , with initial payment P_1 and link factors $\Lambda_1, \dots, \Lambda_{j-1}$, its expected ultimate cost is equal to

$$P_1 \left(\prod_{l=1}^{j-1} \Lambda_l \right) \left(\boldsymbol{\omega}_{\text{RBNS}j}^\top \widetilde{\mathbf{M}}^k \mathbf{1} \right)$$

where k is sufficiently large to reach Closed0 or Closed+ with certainty or with a high enough probability and $\boldsymbol{\omega}_{\text{RBNS}j}$ is the vector indicating the current state RBNS j of the claim in the augmented state space. In case there is no upper bound on the time to settlement, the choice of k can be based on the distribution of the absorption time in the multistate model, selecting it so that $P[S_{t+k} \in \{\text{Closed0}, \text{Closed+}\}] \geq 1 - \epsilon$ for some small enough ϵ , where this probability is obtained from $\boldsymbol{\omega}_{\text{RBNS}j}^\top \widetilde{\mathbf{M}}^k$.

- for a claim in state RBNP, its expected ultimate cost is equal to $\boldsymbol{\omega}_{\text{RBNP}}^\top \widetilde{\mathbf{M}}^k \mathbf{1}$.
- for all IBNR claims relating to accident year i , we add

$$\alpha_i \sum_{j \geq j_0+1} \beta_j \left(\boldsymbol{\omega}_{\text{IBNR}}^\top \widetilde{\mathbf{M}}^k \mathbf{1} \right)$$

if the evaluation date is $i + j_0$, $j_0 \geq 1$.

Replacing $E[P_1]$ and $E[\Lambda_j]$ with higher-order moments leads to higher moments of the ultimate loss. For instance, if the actuary is interested in the variance then $E[P_1^2]$ and $E[\Lambda_j^2]$ enter the matrix $\widetilde{\mathbf{M}}$ and if the skewness is needed then $E[P_1^3]$ and $E[\Lambda_j^3]$ enter $\widetilde{\mathbf{M}}$.

2.6 One-year risk view

Reserve risk has traditionally been discussed in terms of the risk that the estimated reserves will not be able to cover the claims payment during the full run-off of today's liabilities. This is the ultimate view of reserve risk. In the Solvency 2 framework however, the time horizon is one year. The claims development result studied by Solvency 2 is the sum of next year's

payments plus the difference between the reserves at the end and at the beginning of the calendar year.

The one-year risk can be effectively assessed with the help of the Semi-Markov multistate model proposed in this paper since it avoids the simulation-within-simulation issue. A predictive distribution can be obtained as follows from simulations of the claims development result in the multistate model under consideration:

- Payments that emerge over the next calendar period are simulated using the multistate model. Precisely, given the state occupied at the beginning of the evaluation year t^* , the state reached by every claim at the end of year t^* is simulated (as explained in Section 3) and the corresponding cumulative cost is updated by simulating the random variable P_1 or Λ_j associated to the transition, if any.
- At the end of the year, the expected ultimate loss is computed using the matrix formulas in the augmented state space starting from the states occupied by the claims and the cumulative payments at the end of the year.
- The claim development result is then obtained by subtracting the values generated from the two preceding steps from the total reserves available at the beginning of year t^* .

Repeating the whole procedure a large number of times generates realizations of the claim development result from which the amount of required capital can be determined.

3 Numerical illustration on simulated data

3.1 Database

The database used to illustrate this paper has been simulated from the “Individual Claims History Simulation Machine” contributed by Gabrielli and Wuthrich (2018). It records 50,313 claims with occurrence dates between 1994 and 2005. Claims are developed over a maximum of 11 years. A claim ID allows us to follow the whole trajectory, from reporting to closure. The database records the yearly amount of payments (equal to 0 if no payment has been made) together with the accident year, the current development year (the first development corresponding to accident year), the reporting year and the settlement year.

In order to allow interested readers to reproduce the analysis performed in this section, here are the parameters used in the simulation tool. The seed used for the simulation is 100. The total expected number of claims is 50,000. As far as the lines of business are concerned, all the weight is put on the first one. No inflation is applied on the expected annual number of claims, which remains constant over time. The chosen standard deviation parameter in the LogNormal distribution for the total claim amounts simulation is 1.85. Finally the chosen standard deviation parameter in the LogNormal distribution for the total recovery payments simulation is 1.5.

Following Allison (1982) and due to the definition of the multistate model, we build an augmented database with a single line for each development period, comprising the following new variables:

- a variable indicating the RBNP or RBNS state at the end of the considered time period, here calendar year.
- a variable representing the time spent in the state, indicated by the payment state variable.
- a binary variable equal to 1 if there is a transition and 0 otherwise.

The features included in this augmented database are then used to explain the transition probabilities, the mean initial payment and the expected link ratios. Since we work in a regression setting, any other explanatory feature could be included as well. For instance, it could be helpful to include circumstances surrounding the insured event, such the presence of bodily injuries for instance, as explanatory variables. Including circumstances surrounding the insured event of course requires predicting these circumstances for IBNR claims. Such a prediction can be based on past observed claims, including the reporting lag. We refer the reader to Bettonville et al. (2020) for more explanations about the construction of the augmented database.

3.2 Estimation of transition probabilities

The modeling proposed in Section 2.3 provides us with the distribution of the number of claims leaving the hidden IBNR state at each development (results are displayed in Table 3.11 and will be discussed later on). It remains to determine the destination state when a claim leaves the IBNR state. There are 4 possibilities: either the claims terminates and moves to the final state Closed0 or Closed+ or it further develops and then moves to the RBNP state or to the RBNS1 state. In the database under consideration, no transition from IBNR to Closed0 has been observed and we thus set the corresponding probability equal to 0, that is,

$$P[S_t = \text{Closed0} | S_{t-1} = \text{IBNR}] = 0 \text{ for all } t.$$

For claims which do not terminate during the reporting year, we then have to model the probability to move either to the RBNP state or to the RBNS1 state.

The output of Binomial GLM with canonical logit link function for transition probabilities from the IBNR state is displayed in Table 3.1. The analysis has been performed with the function `glm` comprised into the statistical software R. We only show the results of the final model, after insignificant effects have been removed and grouping has been performed. The explanatory variable is the time spent in the IBNR state which is treated as a categorical feature to allow for full flexibility. The reference level corresponds to claims reported during the accident year (that is, $S_1 \neq \text{IBNR}$). To obtain the fitted probabilities, it suffices to sum the estimated intercept and the other coefficients reported in the table and then to compute $\exp(s)/(1+\exp(s))$ where s denotes the aforementioned sum. A positive estimated regression coefficient thus increases the fitted probability compared to the reference one obtained from the intercept, whereas a negative one decreases it.

Considering the estimated regression coefficients displayed in Table 3.1, we see that the probability of being closed during the reporting year culminates at development 2 before declining at later developments. No effect was found to be significant for explaining the

	Coefficient	Standard deviation	p-value
	$P[S_t = \text{Closed+} S_{t-1} = \text{IBNR and } S_t \neq \text{IBNR}]$		
Intercept ($t = 1$)	0.231783	0.009717	$< 10^{-6}$
$t = 2$	1.290108	0.048304	$< 10^{-6}$
$t \geq 3$	-0.976222	0.197155	0.006712
	$P[S_t = \text{RBNP} S_{t-1} = \text{IBNR and } S_t \neq \text{IBNR and } S_t \neq \text{Closed+}]$		
Intercept	-1.20879	0.01697	$< 10^{-6}$

Table 3.1: Estimated regression coefficients for transition probabilities from the IBNR state.

	Coefficient	Standard deviation	p-value
	$P[S_t \neq \text{RBNP} S_{t-1} = \text{RBNP and } D_{t-1} = j]$		
Intercept	4.8981	0.1774	$< 10^{-6}$
$j \geq 2$	-5.3035	0.2891	$< 10^{-6}$
	$P[S_t \in \{\text{Closed+}, \text{Closed0}\} S_{t-1} = \text{RBNP and } D_{t-1} = j \text{ and } S_t \neq \text{RBNP}]$		
Intercept	0.48541	0.03144	$< 10^{-6}$
$j \geq 2$	-1.58403	0.40946	0.000109
	$P[S_t = \text{Closed+} S_{t-1} = \text{RBNP and } D_{t-1} = j \text{ and } S_t \in \{\text{Closed+}, \text{Closed0}\}]$		
Intercept	3.13196	0.09696	$< 10^{-6}$
$j \geq 2$	-5.07787	1.07343	0.005552

Table 3.2: Estimated regression coefficients for transition probabilities from the RBNP state.

probability of moving to the RBNP state if the claim does not terminate during the reporting year so that we end up with an intercept-only model.

The output of Binomial GLM for transition probabilities from the RBNP state is displayed in Table 3.2. The response is the binary variable whose mean is the exit probability depending on the time spent in the current state. Because data become scarce, time spent in RBNP state has been categorized to have a sufficient volume of data in each class. This is a common problem encountered when dealing with loss reserving data: at longer development, the volume of data becomes so small that it does not lend itself to modeling. This is accounted here by aggregating all data together after two years spent in the RBNP state. We acknowledge that this choice is partly subjective but it produces reliable results on the database under study. Thus, we only consider two categories for the sojourn time: “1 year” and “ ≥ 2 years” for the time spent in the current RBNP state. We can see in Table 3.2 that the probability of leaving the RBNP state decreases with sojourn time.

To determine the destination state when a claim leaves the current RBNP state, we consider a binary variable equal to 1 if the claim moves to the terminal state (Closed0 or Closed+) and 0 otherwise. If the response is 0 then the claim moves from RBNP to RBNS1. This part of the analysis thus gives the termination probabilities given that the claim leaves its current RBNP state. We see from Table 3.2 that these probabilities decrease with sojourn time.

The output of the Binomial GLM for the transition probabilities from the RBNS states is displayed in Table 3.3. To fit the exit probabilities from the current RBNS state, we use the binary variable whose mean is the exit probability from the current RBNS state, depending on the sojourn time. Because data become scarce at later development periods, data relating to RBNS states RBNS5, RBNS6, etc. are pooled together for estimation purposes (the corresponding results are displayed in the row RBNS5+). This means that sojourn times are identically distributed in each RBNS j state with $j \geq 5$. Time spent in state has been categorized as well to have a sufficient volume of data in each class. We consider the categories “1 year”, “2 years”, ..., and “ ≥ 6 years” for the time spent in the current RBNS state. As expected, exit probabilities are globally increasing with the time spent in the RBNS state.

	Coefficient	Standard deviation	p-value
$P[S_t \neq \text{RBNS} S_{t-1} = \text{RBNP and } D_{t-1} = j]$			
Intercept	2.08671	0.02409	$< 10^{-6}$
RBNS2	-1.22616	0.03800	$< 10^{-6}$
RBNS3	-0.54988	0.08232	0.000040
RBNS4	-0.45943	0.11993	0.000128
RBNS5+	0.31904	0.15020	0.033666
$j = 2$	-2.57492	0.05111	$< 10^{-6}$
$j = 3$	-2.53644	0.05853	$< 10^{-6}$
$j = 4$	-2.60606	0.06836	$< 10^{-6}$
$j = 5$	-2.30319	0.07445	$< 10^{-6}$
$j \geq 6$	-1.76885	0.05913	$< 10^{-6}$
$P[S_t \in \{\text{Closed+}, \text{Closed0}\} S_{t-1} = \text{RBNS and } D_{t-1} = j \text{ and } S_t \neq \text{RBNS}]$			
Intercept	1.02588	0.01890	$< 10^{-6}$
RBNS2	-1.26471	0.05506	$< 10^{-6}$
RBNS3	-2.13416	0.09996	$< 10^{-6}$
RBNS4	-2.45871	0.14996	$< 10^{-6}$
RBNS5+	-2.92400	0.15086	$< 10^{-6}$
$j = 3$	1.50839	0.12786	$< 10^{-6}$
$j = 4$	2.31393	0.19187	$< 10^{-6}$
$j = 5$	2.86358	0.25626	$< 10^{-6}$
$j \geq 6$	4.20136	0.33861	$< 10^{-6}$
$P[S_t = \text{Closed+} S_{t-1} = \text{RBNS and } D_{t-1} = j \text{ and } S_t \in \{\text{Closed+}, \text{Closed0}\}]$			
Intercept	-1.00399	0.02200	$< 10^{-6}$
RBNS2	0.68037	0.09184	0.000003
RBNS3	1.71307	0.17810	$< 10^{-6}$
RBNS4	2.45297	0.28064	$< 10^{-6}$
RBNS5+	1.91128	0.29829	0.000067
$j = 2$	-1.25159	0.14916	$< 10^{-6}$
$j = 3$	-1.91325	0.18390	$< 10^{-6}$
$j = 4$	-2.11965	0.22272	$< 10^{-6}$
$j = 5$	-2.77314	0.30428	$< 10^{-6}$
$j \geq 6$	-3.30163	0.27022	$< 10^{-6}$

Table 3.3: Estimated regression coefficients for transition probabilities from RBNS states.

Considering the destination state when a claim leaves the current RBNS state, we use a binary variable equal to 1 if the claim moves to the terminal state (Closed0 or Closed+) and 0 otherwise. If the response is 0 then the claim moves from RBNS j to RBNS $j+1$, depending on its current state. This part of the analysis thus gives the termination probabilities given that the claim leaves its current RBNS state. As expected, we see from Table 3.3 that termination probabilities increase with sojourn time but decrease when the claim moves along the series of RBNS states.

The probability that the claim terminates with a final payment when leaving the current RBNS state is finally modeled. The corresponding estimates are displayed in the bottom panel of Table 3.3.

3.2.1 Model adequacy

To check model adequacy, we compare the observed average time spent in each state and the corresponding expected time, calculated according to the fitted model. Table 3.4 displays empirical averages and the corresponding expected values given by the model. We can see there that model predictions are close to the observed averages.

As an out-of-sample check, we have simulated a new set of data from the simulation machine and evaluated how well the model fitted on the original data performed on it. The seed used for the simulation of the new database is 42 (all the other parameters are left unchanged). The results have been added to Table 3.4. We can see there that the accuracy of the model remains stable for the first three states (RBNP-RBNS1-RBNS2) when we switch to another set of simulated values. The accuracy then becomes worse for the new data, on which the model is not adjusted. This can be attributed to the much smaller volume of

observations in the states with more payments, resulting in higher volatility.

This basic check suggests that the model appropriately captures the claim settlement dynamics. For further methods assessing multistate model adequacy, we refer the reader to Willekens (2014).

Model	RBNP	RBNS1	RBNS2	RBNS3	RBNS4	RBNS5
Initial database	1.0185	1.2718	2.4332	1.5839	1.5155	1.1742
New simulated database	1.0187	1.2579	2.6302	1.7474	1.4338	1.2559
Initial database - Model	0.000015	-0.014215	0.199076	0.038132	0.015771	0.003565
New database - Model	0.000194	-0.003934	0.197031	0.163562	-0.081655	0.081721
Initial database - Model (%)	0.0014	1.1177	8.1818	2.4076	1.0407	0.3037
New database - Model (%)	0.0190	0.3093	8.0977	10.3268	5.3881	6.9600

Table 3.4: Average time spent in each state computed according to the model or from simulated data. Initial database is the one used for estimation and new database is another set of simulated data.

3.3 Estimation of the payment process

3.3.1 Discrete mixture setting

To model the payment process, we need appropriate distributions for both the first payment P_1 and the subsequent link ratios Λ_j . The LogNormal distribution is certainly the natural choice for the distribution of the initial payment and subsequent link ratios since ultimate amounts corresponding to products of link ratios are therefore also LogNormal. However, QQ-plots and Kolmogorov-Smirnov tests indicate that Lognormal nor any classical distribution is able to deliver an acceptable fit to data under consideration. This is why we consider discrete mixture models combining two (or more) distributions. For this purpose, the `gamlss` and `gamlss.mx` packages provide actuaries with a variety of distributions including the classical LogNormal, Pareto, and Gamma distributions but also their combinations with the help of finite mixtures. Precisely, the function `gamlssMXfits` comprised into the R package `gamlss.mx` is used to fit this model to the available data.

Of course, this is not the only way to model the payment process. Composite Lognormal-Pareto models proposed by Cooray and Ananda (2005) suitably modified by Scollnik (2007) combine a Lognormal density up to an unknown threshold value and a Pareto density thereafter and could also be relevant here, as well as a Generalized Pareto modeling for the peaks over a sufficiently high threshold as in Cebrian et al. (2003).

3.3.2 First payment

Consider the distribution of the first payment, which is associated to the transition from IBNR to RBNS1 or Closed+ or from RBNP to RBNS1 or Closed+. Several models have been tried out and the best-fitting one corresponds to a mixture of two LogNormal distributions with respective parameters (μ_1, σ_1^2) and (μ_2, σ_2^2) . Parameter estimates are displayed in Table 3.5. We can see there that when a claim goes through the RBNP state, the payment tends to be larger than if the claim comes directly from the IBNR state.

In order to assess the goodness of the resulting fit for first payments, several regression diagnostic checks are available. Here, we use PP-plots for visual inspection completed with

Probability	First Log-normal component			Second Log-normal component		
	0.51			0.49		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	5.02666	0.01258	$< 10^{-6}$	6.25634	0.01432	$< 10^{-6}$
Time in State 1	0.13280	0.04683	0.00458	-0.41359	0.05329	0.000003
Previous State RBNP	0.53054	0.05839	$< 10^{-6}$	1.13816	0.06644	$< 10^{-6}$
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	0.569153	0.004618	$< 10^{-6}$	0.676564	0.004721	$< 10^{-6}$

Table 3.5: Estimated coefficients for first payments

formal uniformity tests. These tools appear to be effective with finite mixture models involving explanatory features, as those used here. Precisely, we have transformed the observed payments P_1 using the distribution function F_{P_1} of the fitted mixture and check whether the resulting values were uniformly distributed over the unit interval (as they should be if the model was correct). Figure 3.1 displays the empirical distribution function of these values compared to the 45-degree line corresponding to the unit uniform distribution function. The graph clearly suggests a close agreement since the black curve corresponding to the data adheres to the main diagonal. The fact that the PP-plots are close to being perfect is perfectly reasonable since the simulation machine of Gabrielli and Wuthrich (2018) simulates LogNormal payments. We have nevertheless adopted here an agnostic procedure and it turned out that a mixture of two LogNormal distributions outperformed a single LogNormal one. Because this could also be an artifact of overfitting (since the payment distribution is obviously more flexible compared to the one implemented in the simulation machine), we also perform an out-of-sample analysis, as described next.

In order to evaluate how well the model fitted on the original data performs on a new set of data, we have produced a new database from the simulation machine (obtained with seed 42) and evaluated the adequacy for first payments. The result has been added to the PP-plot visible on Figure 3.1 where the red curve corresponds to newly simulated data. The close agreement of the red curve with the 45-degree line shows that the model also performs well on data that the model has not seen.

In addition to this visual inspection, a formal Kolmogorov-Smirnov test for uniformity has been performed using the `kolmogorov.unif.test` function included in the R package `uniftest`. This results in a p -value equal to 0.72, well above the usual 5% level. Considering newly simulated data, the p -value for the uniformity of payments slightly decreases to 0.694 but still remains well above the usual significance level of 5%.

3.3.3 Impact of P_1 on Λ_1 and of C_j on Λ_j

To prevent unrealistically large costs for some claims, there is a need to introduce some dependence between the first payment P_1 and the first link ratio Λ_1 . Figure 3.2 displays the observed pairs (P_1, Λ_1) recorded in the database. We clearly see there that a large P_1 cannot be followed by a large Λ_1 . Otherwise, the ultimate cost for that claim becomes so large that it has no economic relevance. This is because some claims have large payments at early developments which are followed by moderate ones (such that the corresponding link ratios remain small) whereas other claims exhibit an opposite behavior, with small payments at early developments followed by larger ones at later stages, such that the corresponding link ratios are pretty large.

Figure 3.2(a) displays all the pairs (P_1, Λ_1) recorded in the database. Some extreme


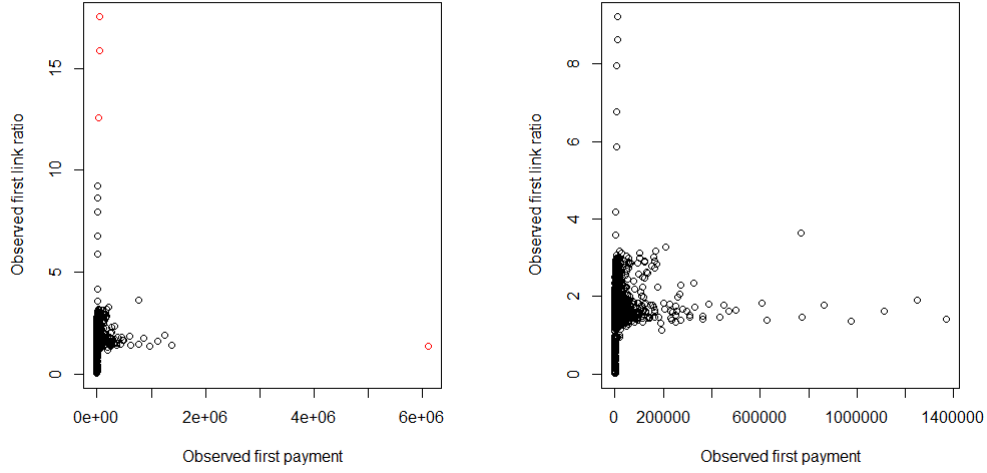


Figure 3.1: Empirical distribution function of the ranks in the two-component LogNormal mixture model for the first payment P_1 with parameter estimates displayed in Table 3.5. The curve printed in black corresponds to the initial database used for estimation, the curve printed in red corresponds to a newly generated database.

points appear printed in red. Removing these observations improves the readability of the graph, as it can be seen on the right panel in Figure 3.2(b). These extreme observations are nevertheless kept in the subsequent analysis. Based on Figure 3.2, we introduce a threshold of 100,000 for P_1 , with a distribution for Λ_1 that differs according to whether P_1 exceeds this threshold or not. Figure 3.2 shows that the majority of the very large link ratios are related to a first payment lower than 100,000. The latter observation also holds for the next link ratios where large cumulative amounts C_j are typically accompanied with smaller link ratios Λ_j .

Link ratios are modeled separately according to their association with the first, second, third, fourth and fifth or more payments. Figure 3.3 displays the histograms of the observed link ratios. Based on QQ-plots, no standard severity distribution provides a satisfactory fit for these data, with significant departures in the tails. As for the first payment P_1 , a mixture of two LogNormal distributions appears to provides a satisfactory fit to observed link ratios Λ_j .

The resulting fit is described in Tables 3.6-3.10 for link ratios Λ_1 , Λ_2 , Λ_3 , Λ_4 , and Λ_j , $j \geq 5$. Point estimates are given there, as well as standard errors. The intercept of the second LogNormal component for the the third link ratio was not significant and has thus been removed. Table 3.8 gives the results of the simpler model, putting this intercept to zero. These tables confirm a different dynamics when payments exceed 100,000 (the following link ratio being smaller compared to the case where payments remain below this threshold). Notice that we do not explicitly associate a link ratio to the transition to Closed+. Instead,



(a) All the data (b) With the three points removed

Figure 3.2: First link ratios function of the first payments.

if l payments are needed to settle the claim and there is a final payment, the last link ratio corresponds to the transition to Closed+.

Probability	First Log-normal component			Second Log-normal component		
	0.94			0.06		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	0.461831	0.015390	$< 10^{-6}$	1.0000	0.2228	0.01812
Time in State 1	0.025740	0.006046	0.01415	-	-	-
First Payment < 100000	0.032340	0.014229	0.0231	-0.5857	0.2278	0.0102
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	-2.151417	0.008696	$< 10^{-6}$	-0.04614	0.03448	0.181

Table 3.6: Estimated coefficients for the first link ratio.

Probability	First Log-normal component			Second Log-normal component		
	0.98			0.02		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	0.21349	0.01150	$< 10^{-6}$	-0.5535	0.1560	0.000403
First Payment < 100000	0.03728	0.01174	0.00154	-	-	-
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	-2.51581	0.02058	$< 10^{-6}$	-0.2754	0.1453	0.0583

Table 3.7: Estimated coefficients for the second link ratio.

Probability	First Log-normal component			Second Log-normal component		
	0.98			0.02		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	0.12317	0.01232	$< 10^{-6}$	0	-	-
First Payment < 100000	0.05713	0.01254	0.015963	-	-	-
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	-2.9527	0.03082	$< 10^{-6}$	-0.982	-	-

Table 3.8: Estimated coefficients for the third link ratio.

In order to assess the goodness of the resulting fit for link ratios, we proceed as for the first payment and transform the observation using the distribution function of the fitted

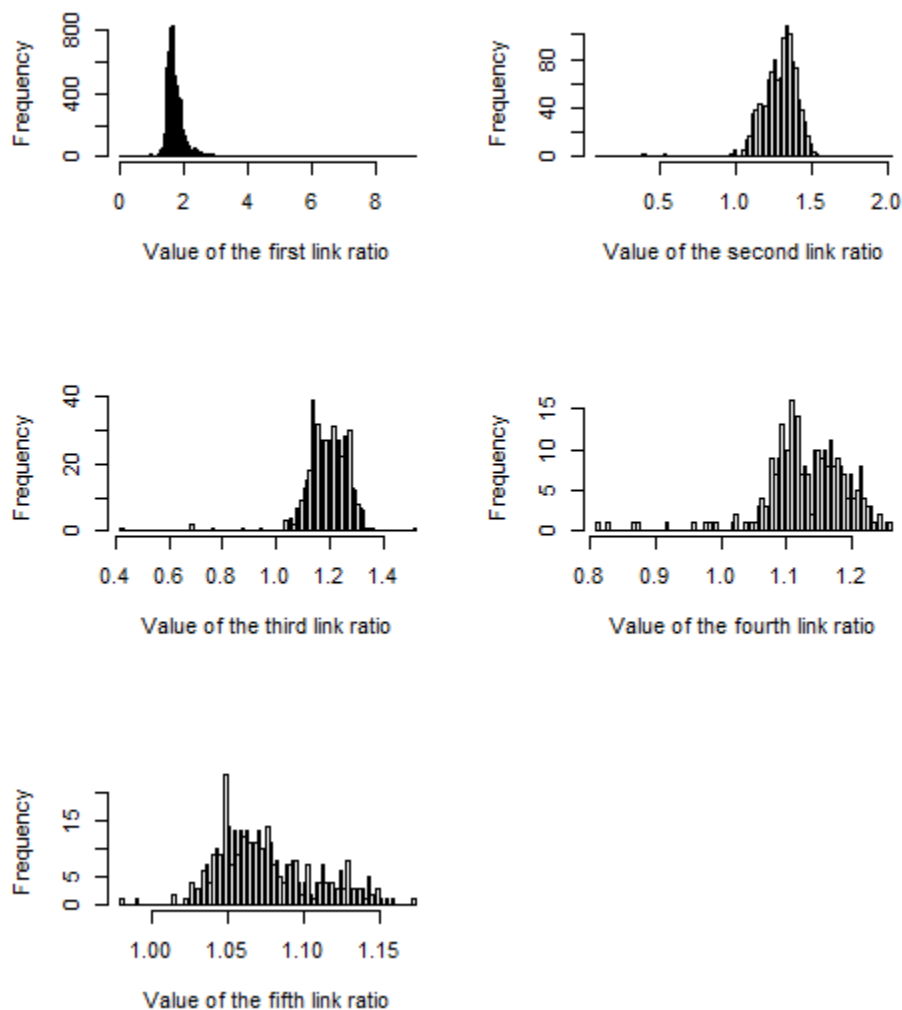


Figure 3.3: Histograms of observed link ratios.

Probability	First Log-normal component			Second Log-normal component		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	0.116128	0.018363	0.000452	-0.05063	0.02779	0.0695
Time in State 1	-0.027072	0.007868	0.000668	-	-	-
First Payment < 100000	0.037474	0.016799	0.026489	-	-	-
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	-3.20677	0.04283	$< 10^{-6}$	-2.3214	0.2002	$< 10^{-6}$

Table 3.9: Estimated coefficients for the fourth link ratio.

LogNormal mixture to check whether the resulting values were uniformly distributed over the unit interval. Figure 3.4 displays the empirical distribution function of these values compared to the 45-degree line corresponding to the unit uniform distribution function. The graph suggests close agreement, except for some moderate departure for both the initial data base used for estimation (curve printed in black) and the newly simulated data used for out-of-sample check (curve printed in red). In addition to visual inspection, A Kolmogorov-Smirnov

Probability	First Log-normal component			Second Log-normal component		
	0.51			0.49		
	μ_1	Std. Error	p-value	μ_2	Std. Error	p-value
Intercept	0.057128	0.008649	0.000056	0.111616	0.004208	$< 10^{-6}$
Time in State 1	0.034421	0.008907	0.000129	-0.055491	0.004330	$< 10^{-6}$
First Payment ≥ 100000	-0.035467	0.012241	0.003965	-	-	-
	$\ln \sigma_1$	Std. Error	p-value	$\ln \sigma_2$	Std. Error	p-value
Intercept	-3.55654	0.05043	$< 10^{-6}$	-4.22371	0.04806	$< 10^{-6}$

Table 3.10: Estimated coefficients for the fifth and following link ratio.

test for uniformity has been performed with the help of the function `kolmogorov.unif.test` available from the R package `uniftest`. Uniformity cannot be rejected at usual probability levels, as shown by the p -values equal to 96.7% for Λ_1 , 99.1% for Λ_2 , 97.2% for Λ_3 , 79.7% for Λ_4 , and 15.4% for Λ_j , $j \geq 5$. Considering the newly simulated data, the p -values for uniformity are respectively equal to 97.75%, 98.8%, 40.1%, 78.15%, and 78.3% so that model adequacy is not questioned.

3.4 Reserve calculations

Now that transition probabilities and payments have been modeled, we can use the multistate model to perform actuarial calculation.

3.4.1 Simulation

Assume that we are at time t^* and that we want to simulate the total cost of the portfolio. More specifically, we put the portfolio in run-off and we are interested in the remaining costs for all insured events that occurred in accident years up to t^* . We proceed separately for each type of claim, RBNS, RBNP or IBNR claims.

RBNS claims For all claims occupying a RBNS state, we have to simulate the last part of the trajectory until final settlement. First, we simulate the remaining sojourn time in the current RBNS state. More specifically, if the claim is in state RBNS j at time t^* and entered that state at time $t^* - w$, we determine the remaining duration d in state RBNS j as

$$d = \min \left\{ k \in \mathbb{N} \mid \mathbb{P}[S_{t^*+k} = \text{RBNS}j \mid S_t = \text{RBNS}j \text{ for } t = t^* - w, \dots, t^*] \leq u \right\}$$

where u is a realization of a random variable U obeying the unit uniform distribution and

$$\begin{aligned} & \mathbb{P}[S_{t^*+k} = \text{RBNS}j \mid S_t = \text{RBNS}j \text{ for } t = t^* - w, \dots, t^*] \\ &= \prod_{l=1}^k \mathbb{P}[S_{t^*+l} = \text{RBNS}j \mid S_t = \text{RBNS}j \text{ for } t = t^* - w, \dots, t^* + l - 1]. \end{aligned}$$

Each factor entering the product has been estimated from the data.

At time $t^* + d$, the claim leaves the state RBNS j and we simulate the transition to RBNS $j + 1$ or Closed0/Closed+. If the claim moves to state RBNS $j + 1$ then the procedure described above is repeated (with $w = 0$ since the claim just entered the next RBNS state) until the claim finally moves to Closed0/Closed+.

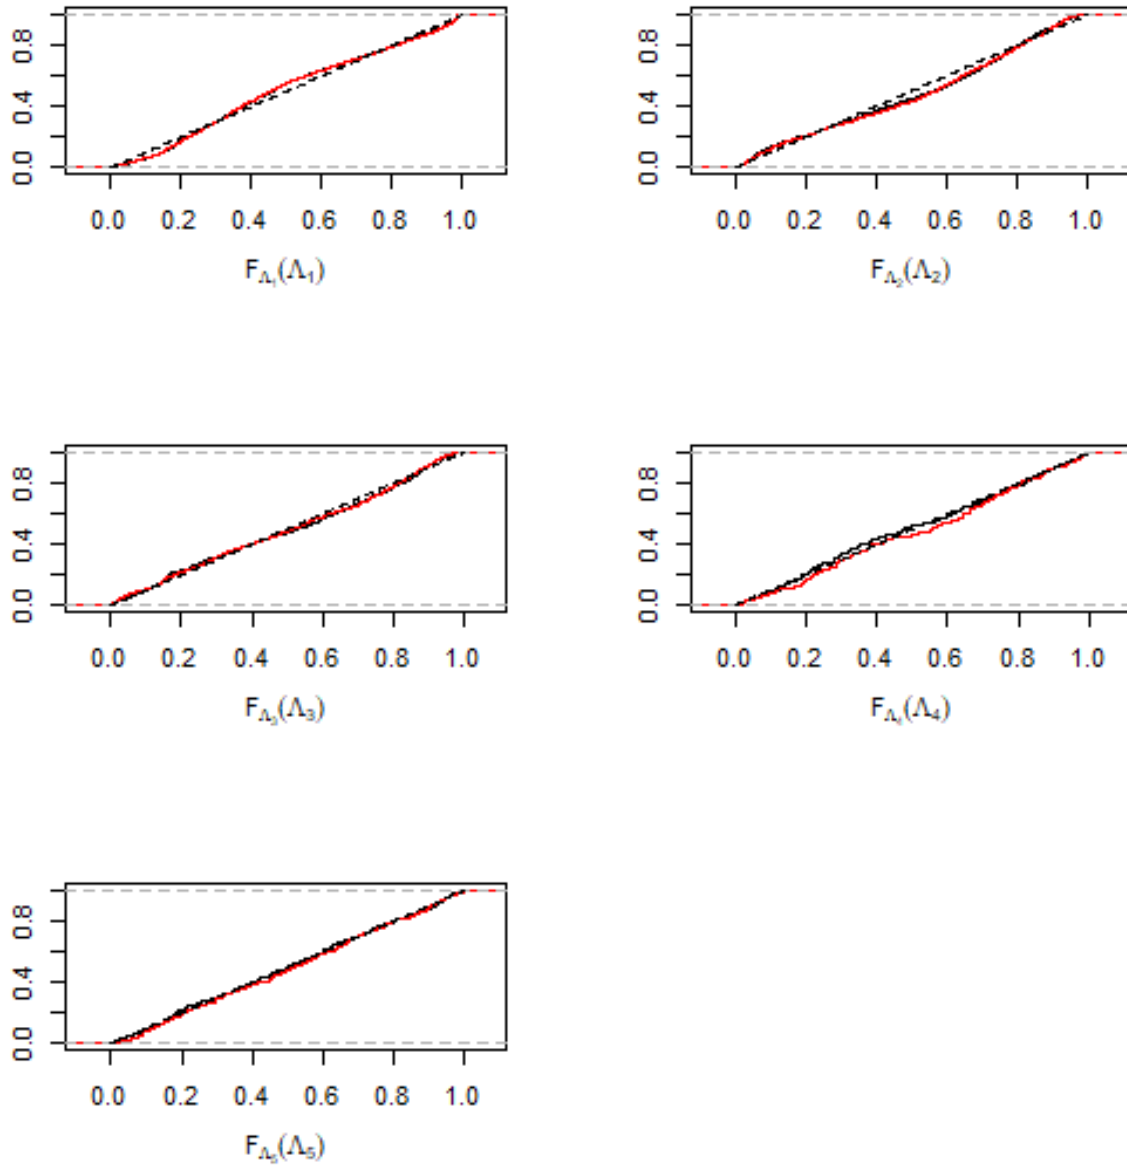


Figure 3.4: Empirical distribution function $F_{\Lambda_j}(\Lambda_j)$ of the ranks in the two-components LogNormal mixture model of each link ratio, $j \in \{1, 2, 3, 4, 5\}$. The curve printed in black corresponds to the initial database used for estimation, the curve printed in red corresponds to a newly generated database.

The corresponding link ratios $\Lambda_{j+1}, \Lambda_{j+2}, \dots$ are also simulated and applied to the current observed total payment C_j for the claim occupying state RBNS j , to obtain its ultimate cost. As described in Section 3.4.2, let us mention that the individual reserves can be analytically calculated to avoid the simulation-within-simulation approach.

RBNP claims For all claims occupying the RBNP state, we first simulate the remaining sojourn time in this state. More specifically, if the claim entered the state RBNP at time $t^* - w$, we determine the remaining duration d of the sojourn in state RBNP as

$$d = \min \left\{ k \in \mathbb{N} \mid \mathbb{P}[S_{t^*+k} = \text{RBNP} \mid S_t = \text{RBNP for } t = t^* - w, \dots, t^*] \leq v \right\}$$

where v is a realization of a random variable V obeying the unit uniform distribution and

$$\begin{aligned} & \mathbb{P}[S_{t^*+k} = \text{RBNP} \mid S_t = \text{RBNP for } t = t^* - w, \dots, t^*] \\ &= \prod_{l=1}^k \mathbb{P}[S_{t^*+l} = \text{RBNP} \mid S_t = \text{RBNP for } t = t^* - w, \dots, t^* + l - 1]. \end{aligned}$$

Each factor entering the product has been estimated from the data.

At time $t^* + d$, the claim leaves the state RBNP and we simulate the transition to RBNS1 or Closed0/Closed+. If the claim moves to state RBNS1 then we proceed as described above for RBNS claims (with $j = 1$ and $w = 0$ since the claim just entered the state RBNS1). The amount of the initial payment P_1 is simulated, as well as the subsequent link ratios $\Lambda_1, \Lambda_2, \dots$ and the product gives the ultimate cost of the claim.

IBNR claims Running a Poisson GLM on the claims aggregated by accident year and reporting year, we get the estimated parameters displayed in Table 3.11. These coefficients allow us to simulate the number of claims that still need to be reported in the following years.

	Coefficient	Standard deviation	p-value
Intercept	8.268839	0.004622	$< 10^{-6}$
RepDel 1Year	-2.626932	0.018540	$< 10^{-6}$
RepDel 2Years	-6.094087	0.106700	$< 10^{-6}$
RepDel 3Years	-7.270310	0.229460	$< 10^{-6}$
RepDel 4Years	-7.575691	0.353579	$< 10^{-6}$
RepDel 5Years	-7.981157	0.353583	$< 10^{-6}$

Table 3.11: Estimated regression coefficients for IBNR claims depending on accident year (AY) and reporting delay (RepDel).

For all accident years $t \leq t^*$ for which some claims may still be reported, i.e. such that

$$\sum_{j=1}^{t^*-t} \beta_j < 1$$

we first simulate the numbers $N_{t,j}$, $j \geq t^* - t + 1$, of claims originating from insured events which happened in year t but were reported at development j . This is done by generating independent realizations from Poisson random variables with means $\alpha_t \beta_j$. The trajectory of each of these $N_{t,j}$ claims is then simulated as explained before, based on the reporting lag j .

3.4.2 Analytical computations

Simulating individual trajectories to final settlement requires a lot of computation time (that can nevertheless be reduced by paralleling the process). The analytical approach considerably reduces the computation time and appears to be particularly effective to deal with the one-year risk in the Solvency 2 system because it avoids the so-called simulation-within-simulation issue.

As explained previously, each physical state is split into several artificial states combining the state occupied and the time already spent in it. Using these new artificial states we are back to the Markov setting for calculations. The matrix \mathbf{M} containing all transition probabilities from one artificial state to another has a block structure, each block corresponding to a physical, Semi-Markov state. The process is hierarchical (no return to a previous state is possible) so that \mathbf{M} is upper diagonal. As long as a duration effect is present, the element $m_{i,i}$ of \mathbf{M} is zero. When duration effects have disappeared, in the Geometric tail, then $m_{i,i}$ is the (constant) probability of remaining one additional year in the artificial state.

For RBNS claims, the first payment P_1 has been observed and the average link ratio to be included in the matrix $\widetilde{\mathbf{M}}$ is $E[\Lambda_j | P_1 < 100,000]$ or $E[\Lambda_j | P_1 \geq 100,000]$ depending on the value of P_1 . However, for IBNR and RBNP claims, the first payment is still unknown and the mean ultimate cost is then given by

$$\begin{aligned} E \left[P_1 \prod_{j=1}^n \Lambda_j \right] &= E[P_1 | P_1 < 100,000] \left(\prod_{j=1}^n E[\Lambda_j | P_1 < 100,000] \right) P[P_1 < 100,000] \\ &\quad + E[P_1 | P_1 \geq 100,000] \left(\prod_{j=1}^n E[\Lambda_j | P_1 \geq 100,000] \right) P[P_1 \geq 100,000]. \end{aligned}$$

3.4.3 Results

After completion of the whole process, we have cash-flows for each claim and thus also for the entire portfolio, in calendar years t^* , $t^* + 1$, $t^* + 2$, ... until all claims relating to accident year before t^* are finally settled. Repeating the procedure a large number of times (here 5,000 times), allows the actuary to obtain best-estimates by averaging the outcomes and to derive quantile risk measures, such as Value-at-Risk.

Table 4.7 displays the obtained results for the entire portfolio and compare them with the output of the classical Chain Ladder method. The results obtained with Chain Ladder have been produced with the `ChainLadder` package of `R`, based on the corresponding aggregated runoff triangle.

The best estimates displayed in the first row of Table 4.7 are obtained in the ultimate view. We see that the proposed Semi-Markov reserving model produces a total reserve slightly below the true one (since the simulation machine produces development to ultimate for all claims in the database) whereas Chain Ladder overestimates it to a large extent.

Under Solvency 2, the risk is assessed over a one-year horizon (resulting in the so-called one-year risk). The claim development result (CDR) is the difference between the reserve at the beginning of the year and the reserve at the end of the year supplemented with the payments made during the year. Precisely, the claim development result is obtained in four

steps. First, the reserve is computed using the analytical matrix method. We simulate all the cash flows generated during this year. The computation of the reserve at the end of the year is then possible using again the analytical matrix method. The three numbers can then be aggregated in the following way to give the claim development result :

$$\text{CDR} = \text{Reserve at the beginning of year } t - \text{Payments occurring during year } t - \text{Reserve at the end of year } t.$$

We repeat this calculation based on a large number of simulations (here 5,000 times) to obtain a sample of CDR values. This is compared with the CDR obtained from Chain Ladder by bootstrapping. Precisely, the results obtained with Chain-Ladder have been produced with the `CDR.BootChainLadder` function contained in the `ChainLadder` package of `R`, based on the corresponding aggregated runoff triangle. The mean CDRs are displayed in Table 4.7. Clearly, a brute-force simulation approach would not have been feasible here since for each 5,000 simulations performed to compute the reserve at the beginning of the year, another set of 5,000 simulations should be performed to predict the coming year.

	Database	Chain Ladder	Simulations	Analytical method
Reserve for open claims	26,550,428	33,862,957	24,059,974	23,398,774
Reserve for IBNR claims	391,428	-	647,094	647,096
Total Reserve	26,941,856	33,862,957	24,707,068	24,045,870
		Chain Ladder	Simulations + Analytical method	
Mean CDR		-486,340	-418,583	

Table 3.12: Results for the entire portfolio.

4 Case study with MTPL insurance data

4.1 Data set

The data set under consideration has been extracted from the MTPL insurance portfolio of an insurance company operating in the European Union. The observation period consists in calendar years 2004 till 2014. The database records all claims that originate from insured event occurred in accident years 2004 to 2013 so that we have observed developments up to 11 calendar years. We refer to Denuit and Trufin (2017, 2018) for a detailed description of the database and we only provide the reader here with relevant information for the analysis performed in the present paper.

The length of the settlement period appears to be a natural criterion to distinguish two types of claims:

- claims with short development to settlement, i.e. those claims that are reported and fully settled during the year of occurrence, or the year after;
- and claims with longer development to settlement, i.e. that are not yet reported or that are still open two years after occurrence.

Such a separation based on the time to settlement provides the actuary with a simple relevant criterion to isolate simple cases that are easily settled from more complicated ones

that develop over a longer period, making the multistate approach attractive. Indeed, the multistate model is not needed for the claims that are rapidly settled because trajectories with

$$\begin{aligned} S_0 &= \text{IBNR} \\ S_1 &\in \{\text{IBNR}, \text{RBNP}, \text{RBNS1}, \text{Closed0}, \text{Closed+}\} \\ S_2 &\in \{\text{Closed0}, \text{Closed+}\} \end{aligned}$$

are so short (with at most two states visited) that developing these claims on an individual basis does not bring any additional insight. This is why we isolate claims that are not yet reported or settled two years after occurrence and model their individual development using the proposed multistate model. Restricting accident years to 2004-2013 allows us to classify each claim into one of the two categories. Because we restrict the data to claims with longer developments, the settlement year varies from 2006 to 2014. All payments are expressed in euro currency and have been corrected for inflation.

4.2 Estimation of transition probabilities

As with simulated data in the preceding section, a Binomial GLM is used to model exit probabilities from the different states as well as the destination state when several transitions were possible. For the IBNR state, the estimated exit probabilities are

$$P[S_t \neq \text{IBNR} | S_{t-1} = \text{IBNR}] = \begin{cases} 0.903 & \text{for } t = 1 \\ 0.864 & \text{for } t \in \{2, 3\} \\ 1.000 & \text{for } t = 4. \end{cases}$$

In the database under consideration, no transition from IBNR to Closed0 has been observed such that we set the corresponding probability equal to 0, that is,

$$P[S_t = \text{Closed0} | S_{t-1} = \text{IBNR}] = 0 \text{ for all } t.$$

Furthermore,

$$P[S_t = \text{Closed+} | S_{t-1} = \text{IBNR}] = \begin{cases} 0.000 & \text{for } t \in \{1, 2\} \\ 0.507 & \text{for } t \geq 3. \end{cases}$$

and

$$P[S_t = \text{RBNS1} | S_{t-1} = \text{IBNR}, S_t \in \{\text{RBNP}, \text{RBNS1}\}] = \begin{cases} 0.301 & \text{for } t = 1 \\ 0.390 & \text{for } t = 1 \\ 0.108 & \text{for } t \geq 2. \end{cases}$$

Denoting as D_t the time spent in state occupied at time t , the estimated exit probabilities from RBNP and RBNS j states are as follows:

$$P[S_t \neq \text{RBNP} | D_t = j] = \begin{cases} 0.566 & \text{for } j = 1 \\ 0.646 & \text{for } j \geq 2 \end{cases}$$

$$P[S_t \neq \text{RBNS1} | D_t = j] = 0.684 \text{ for all } j \geq 1$$

$$P[S_t \neq \text{RBNS2} | D_t = j] = \begin{cases} 0.798 & \text{for } j = 1 \\ 0.684 & \text{for } j = 2 \\ 0.579 & \text{for } j \geq 3 \end{cases}$$

and, for any $l \geq 3$,

$$P[S_t \neq \text{RBNS}l | D_t = j] = \begin{cases} 0.844 & \text{for } j = 1 \\ 0.684 & \text{for } j = 2 \\ 0.579 & \text{for } j \geq 3. \end{cases}$$

The estimated termination probabilities from each state RBNP and RBNS are as follows:

$$P[S_t \in \{\text{Closed0}, \text{Closed+}\} | S_{t-1} = \text{RBNP}, D_{t-1} = j] = \begin{cases} 0.087 & \text{for } j = 1 \\ 0.794 & \text{for } j \geq 2, \end{cases}$$

$$P[S_t \in \{\text{Closed0}, \text{Closed+}\} | S_{t-1} = \text{RBNS1}, D_{t-1} = j] = \begin{cases} 0.556 & \text{for } j = 1 \\ 0.844 & \text{for } j = 2 \\ 0.893 & \text{for } j \geq 3, \end{cases}$$

$$P[S_t \in \{\text{Closed0}, \text{Closed+}\} | S_{t-1} = \text{RBNS2}, D_{t-1} = j] = \begin{cases} 0.556 & \text{for } j = 1 \\ 0.660 & \text{for } j = 2 \\ 0.792 & \text{for } j \geq 3, \end{cases}$$

$$P[S_t \in \{\text{Closed0}, \text{Closed+}\} | S_{t-1} = \text{RBNS3}, D_{t-1} = j] = \begin{cases} 0.462 & \text{for } j \in \{1, 2\} \\ 0.362 & \text{for } j \geq 3, \end{cases}$$

and for every $l \geq 4$

$$P[S_t \in \{\text{Closed0}, \text{Closed+}\} | S_{t-1} = \text{RBNS}l, D_{t-1} = j] = \begin{cases} 0.362 & \text{for } j = 1 \\ 0.462 & \text{for } j \geq 2. \end{cases}$$

The probability that the claim terminates with a final payment when leaving the current RBNP or RBNS state is estimated to

$$P[S_t = \text{Closed+} | S_{t-1} = \text{RBNP}, D_{t-1} = j, S_t \in \{\text{Closed0}, \text{Closed+}\}] = 1 \text{ for all } j$$

$$P[S_t = \text{Closed+} | S_{t-1} = \text{RBNS1}, D_{t-1} = j, S_t \in \{\text{Closed0}, \text{Closed+}\}] = \begin{cases} 0.437 & \text{for } j = 1 \\ 0.403 & \text{for } j = 2 \\ 0.478 & \text{for } j \geq 3 \end{cases}$$

and for all $l \geq 2$,

$$P[S_t = \text{Closed+} | S_{t-1} = \text{RBNS}l, D_{t-1} = j, S_t \in \{\text{Closed0}, \text{Closed+}\}] = 0.711 \text{ for all } j.$$

More details about the fitting procedure can be found in [Bettonville et al. \(2020\)](#).

The comparison of the observed average time spent in each state and the corresponding expected time calculated according to the fitted model reveals that model predictions are very close to the observed averages. This shows that the model adequately captures the claim settlement dynamics.

Probability	0.1154		0.8846	
	$\hat{\mu}_1$	Std. Error	$\hat{\mu}_2$	Std. Error
Time in state = 0	5.7932	0.1690	7.4481	0.0285
Time in state ≥ 1	4.1257	0.1828	6.9604	0.0328
	$\ln \hat{\sigma}_1$	Std. Error	$\ln \hat{\sigma}_2$	Std. Error
	0.9163	0.0351	0.1831	0.0127

Table 4.1: Estimated regression coefficients for the first payment from RBNP to RBNS1 to Closed+ in the two-component LogNormal mixture.

Probability	0.0416		0.9584	
	$\hat{\mu}_1$	Std. Error	$\hat{\mu}_2$	Std. Error
Time in state ≥ 0	5.8822	0.3168	7.4254	0.0296
	$\ln \hat{\sigma}_1$	Std. Error	$\ln \hat{\sigma}_2$	Std. Error
	1.0003	0.0824	0.1976	0.0172

Table 4.2: Estimated regression coefficients for the first payment from IBNR to RBNS1 or Closed+ in the two-component LogNormal mixture.

4.3 Estimation of the payment process

4.3.1 First payment

Here, we fit the distribution of the first payment, which is associated to the transition from IBNR to RBNS1 or Closed+ or from RBNP to RBNS1 or Closed+. Several models have been tried out and the best-fitting one corresponds to a mixture of two LogNormal distributions with respective parameters (μ_1, σ_1^2) and (μ_2, σ_2^2) . Parameter estimates are displayed in Table 4.1 for transition from RBNP to RBNS1 or Closed+ and in Table 4.2 for transition from IBNR to RBNS1 or Closed+ together with the corresponding standard errors.

In order to assess the goodness of the resulting fit for first payments, we have transformed the observation using the distribution function of the fitted mixture and check whether the resulting values were uniformly distributed over the unit interval (as they should be if the model was correct). Figures 4.1 and 4.2 display the empirical distribution function of these values compared to the 45-degree line corresponding to the unit uniform distribution function. The graph clearly suggests a close agreement. In addition to this visual inspection, a formal Kolmogorov-Smirnov test for uniformity has been performed, resulting in a p -value equal to 0.0747 above the usual 5% level.

4.3.2 Impact of P_1 on Λ_1 and of C_j on Λ_j

To prevent unrealistically large costs for some claims, there is a need to introduce some dependence between the first payment P_1 and the first link ratio Λ_1 . Figure 4.3 displays the observed pairs (P_1, Λ_1) recorded in the database. We clearly see there that a large P_1 cannot be followed by a large Λ_1 .

This is why we introduce a threshold for P_1 , with a distribution for Λ_1 that differs according to whether P_1 exceeds this threshold or not. To account for this phenomenon, the state RBNS1 is split into a state RBNS1a where P_1 is smaller than 5,000 and a state RBNS1b where P_1 exceeds this threshold. The threshold is determined graphically on the basis of Figure 4.3 which shows that the majority of the very large link ratios are related to a first payment lower than 5,000. The latter observation also holds for the next link ratios




Figure 4.1: Empirical distribution function of the ranks in the two-component LogNormal mixture model for the first payment P_1 from RBNP to RBNS1.




Figure 4.2: Empirical distribution function of the ranks in the two-component LogNormal mixture model for the first payment P_1 from IBNR to RBNS1.

where large cumulative amounts C_j are typically accompanied with smaller link ratios Λ_j .

4.3.3 Link ratios

The first link ratio Λ_1 is best modeled with a mixture of a LogNormal distribution with parameters (μ_1, σ_1^2) and a one-parameter Pareto distribution with parameter λ . The resulting fit is described in Table 4.3 where point estimates are given as well as standard errors.

The next link ratios are also modeled using a finite mixture with LogNormal and Pareto components. As mentioned before, the cumulative payment C_j is split into two categories based on a threshold of 5,000. Tables 4.4, 4.5 and 4.6 present the fitted models for link ratios Λ_2 , Λ_3 and Λ_j , $j \geq 4$.

In order to assess the goodness of the resulting fit for link ratios, we proceed as for the first payment and transform the observation using the distribution function of the fitted LogNormal-Pareto mixture to check whether the resulting values were uniformly distributed

First observed link ratio as a function of the first observed payment




Figure 4.3: Observed pairs (P_1, Λ_1) recorded in the database.

over the unit interval. Figure 4.4 displays the empirical distribution function of these values compared to the 45-degree line corresponding to the unit uniform distribution function. The graph suggests close agreement, except for some moderate departure in the left tail. In addition to visual inspection, A Kolmogorov-Smirnov test for uniformity has been performed resulting in p -values equal to 22.57% for Λ_1 , 45.02% for Λ_2 , 33.95% for Λ_3 , and 10.75% for $\Lambda_j, j \geq 4$. The fit thus appears to be acceptable for the data under consideration.




Figure 4.4: Empirical distribution function of the ranks of link factors Λ_j in the two-component mixture with a LogNormal distribution and a Pareto one.

Probability	LogNormal component		Pareto component	
	0.0889		0.9111	
	$\hat{\mu}$	Std. Error	$\hat{\lambda}$	Std. Error
Time in state 0 First payment $\leq 5,000$	2.6907	0.0758	0.1515	0.0316
Time in state 1 First payment $\leq 5,000$	2.6907	0.0758	0.4448	0.0615
Time in state ≥ 2 First payment $\leq 5,000$	2.6907	0.0758	0.7519	0.088
Time in state 0 First payment $> 5,000$	1.2612	0.1294	1.1328	0.0469
Time in state 0 First payment $> 5,000$	1.2612	0.1294	1.8506	0.1226
Time in state ≥ 2 First payment $> 5,000$	1.2612	0.1294	2.2911	0.1996
	$\ln \hat{\sigma}$	Std. Error		
	-0.1047	0.0514		

Table 4.3: Estimated regression coefficients for the first link ratio Λ_1 .

Probability	LogNormal component		Pareto component	
	0.093		0.907	
	$\hat{\mu}$	Std. Error	$\hat{\lambda}$	Std. Error
Time in state 1 payment $\leq 5,000$	1.5823	0.0619	0.833	0.0501
Time in state 2 payment $\leq 5,000$	2.5291	0.2258	1.0495	0.0895
Time in state 3 payment $\leq 5,000$	1.5823	0.0619	1.0495	0.0895
Time in state 1 payment $> 5,000$	0.942	0.0787	1.5647	0.0693
Time in state 2 payment $> 5,000$	0.5207	0.1246	3.1686	0.1855
Time in state 3 payment $> 5,000$	0.5207	0.1246	3.1686	0.1855
	$\ln \hat{\sigma}$	Std. Error		
	-0.9347	0.0801		

Table 4.4: Estimated regression coefficients for the second link ratio Λ_2 .

Probability	LogNormal component		Pareto component	
	0.104		0.896	
	$\hat{\mu}$	Std. Error	$\hat{\lambda}$	Std. Error
Time in state 1 payment $\leq 5,000$	0.0328	0.0037	0.7093	0.0784
Time in state 2 payment $\leq 5,000$	0.3785	0.0138	0.7093	0.0786
Time in state 3 payment $\leq 5,000$	0.4539	0.0162	1.7085	0.2577
Time in state 1 payment $> 5,000$	0.0231	0.0049	1.3572	0.1029
Time in state 2 payment $> 5,000$	0.5531	0.0105	1.3572	0.1029
Time in state 3 payment $> 5,000$	0.3008	0.0146	1.7085	0.2577
	$\ln \hat{\sigma}$	Std. Error		
	-4.1941	0.1259		

Table 4.5: Estimated regression coefficients for the third link ratio Λ_3 .

	LogNormal component		Pareto component	
Probability	0.1163		0.8837	
	$\hat{\mu}$	Std. Error	$\hat{\lambda}$	Std. Error
Time in state 1 payment $\leq 5,000$	0.2204	0.0018	0.754	0.0959
Time in state 2 payment $\leq 5,000$	0.0161	0.0024	0.754	0.096
Time in state 3 payment $\leq 5,000$	0.0962	0.0041	0.754	0.096
Time in state 1 payment $> 5,000$	0.0119	0.0015	1.143	0.0983
Time in state 2 payment $> 5,000$	0.0161	0.0024	1.143	0.0983
Time in state 3 payment $> 5,000$	0.9804	0.0052	1.143	0.0983
	$\ln \hat{\sigma}$	Std. Error		
	-5.2699	0.1338		

Table 4.6: Estimated regression coefficients for link ratios Λ_j , $j \geq 4$.

4.3.4 Reserving calculations

Now that transition probabilities and cash-flows have been modeled, we can use the multistate model to perform actuarial calculation. Simulations are conducted as explained in Section 3. Figures 4.5, 4.6 and 4.7 show empirical distributions for the total cost of IBNR claims, RBNP claims and RBNS claims, respectively, obtained from 10,000 simulations.




Figure 4.5: Empirical distribution for the total cost of IBNR claims.

On each trajectory, the actuary can apply the risk mitigation technique in force within the portfolio under consideration. For instance, if the portfolio is protected by an excess-of-loss reinsurance treaty, the part of the cost exceeding the priority will be paid by the reinsurer. Also, financial discounting can be introduced, if necessary.

After completion of the whole process, we have cash-flows for each claim and thus also for the entire portfolio. Repeating the procedure a large number of times (here 10,000 times) allows the actuary to derive risk measures, such as Value-at-Risk (VaR) and Tail-VaR.




Figure 4.6: Empirical distribution for the total cost of RBNP claims.




Figure 4.7: Empirical distribution for the total cost of RBNS claims.

Table 4.7 displays means and VaRs at 99.5% estimated from the 10,000 simulations for the three types of claims as well as for the entire portfolio. We notice that the difference between the sum of the VaR at 99.5% for the three types of claims and the VaR at 99.5% for the entire portfolio is given by 5,444,181, which can be interpreted as the diversification benefit resulting from the aggregation of the three types of claims.

The one-year risk can also be studied following the approach described in Section 3 using simulation combined with the proposed matrix-analytic method.

5 Conclusion

In this paper, a Semi-Markov model has been proposed to model the claim settlement process in general insurance. Cash-flows are associated to transitions between states and the total outstanding losses can easily be simulated while analytical calculation is possible for the

	IBNR	RBNP	RBNS	All
Mean	2,246,225	975,122	15,876,253	19,097,600
VaR 99.5%	5,820,638	3,229,190	24,561,222	28,166,869

Table 4.7: Estimated means and VaRs for IBNR, RBNP and RBNS claims as well as for the entire portfolio.

moments. A case study demonstrates the usefulness of this approach.

Acknowledgements

The authors wish to thank the Editor and two anonymous Referees for many constructive comments which greatly helped to improve the contents of this paper compared to the initial submission. Also, we thank Cindy Courtois for interesting exchanges about individual loss reserving methods.

References

- Allison, P. D. (1982). Discrete-time methods for the analysis of event histories. *Sociological Methodology* 13, 61-98.
- Antonio, K., Godecharle, E., Van Oirbeek, R. (2016). A multi-state approach and flexible payment distributions for micro-level reserving in general insurance. Available at SSRN: <https://ssrn.com/abstract=2777467>
- Arjas, E. (1989). The claims reserving problem in non-life insurance: some structural ideas. *ASTIN Bulletin* 19, 140–152.
- Bettonville, C., Denuit, M., Trufin, J., Van Oirbeek, R. (2020). Semi-Markov multi-state individual loss reserving model in general insurance. *Detra Note 2020-1* available at <http://detralytics.com>.
- Cebrian, A., Denuit, M., Lambert, Ph. (2003). Generalized Pareto fit to the society of Actuaries' large claims database. *North American Actuarial Journal* 7, 18-36.
- Cooray, K., Ananda, M.M.A. (2005). Modeling actuarial data with a composite Lognormal-Pareto model. *Scandinavian Actuarial Journal* 2005, 321–334.
- Charpentier, A., Pigeon, M. (2016). Macro vs. micro methods in non-life claims reserving (an econometric perspective). *Risks* 4, 1-18.
- Denuit, M., Trufin, J. (2017). Beyond the Tweedie reserving model: The collective approach to loss development. *North American Actuarial Journal* 21, 611-619.
- Denuit, M., Trufin, J. (2018). Collective loss reserving with two types of claims in motor third party liability insurance. *Journal of Computational and Applied Mathematics* 335, 168-184.

- Gabrielli, A., Wuthrich, M.V. (2018). An individual claims history simulation machine. *Risks* 6, 29.
- Hachemeister, C. A. (1980). A stochastic model for loss reserving. *Transactions of the 21st International Congress of Actuaries*, Vol. 1, pp. 185-194.
- Hesselager, O. (1994). A Markov model for loss reserving. *ASTIN Bulletin* 24, 183-193.
- Hurlimann, W. (2015). A simple multi-state Gamma claims reserving model. *International Journal of Contemporary Mathematical Sciences* 10, 65-77.
- Klein, N., Denuit, M., Lang, S., Kneib, Th. (2014). Nonlife ratemaking and risk management with Bayesian additive models for location, scale and shape. *Insurance: Mathematics and Economics* 55, 225-249.
- Norberg, R. (1993). Prediction of outstanding liabilities in non-life insurance I. *ASTIN Bulletin* 23, 95–115.
- Norberg, R. (1999). Prediction of outstanding liabilities II. Model variations and extensions. *ASTIN Bulletin* 29, 5-25.
- Scollnik, D.P.M. (2007). On composite Lognormal-Pareto models. *Scandinavian Actuarial Journal* 2007, 20–33.
- Taylor, G. (2018). The Long Road to Enlightenment: Loss Reserving Models from the Past, with Some Speculation on the Future. *Expanding Horizons* 57.
- Taylor, G. (2019). Loss reserving models: Granular and machine learning forms. *Risks* 7, 1-18.
- Willekens, F. (2014). *Multistate Analysis of Life Histories with R*. Springer, New York.
- Wuthrich, M.V., Merz, M. (2008). *Stochastic Claims Reserving Methods in Insurance*. Wiley, New York.
- Wuthrich, M.V., Merz, M. (2015). *Stochastic Claims Reserving Manual: Advances in Dynamic Modeling*. Available at SSRN: <https://ssrn.com/abstract=2649057>