



Micro-level stochastic loss reserving for general insurance

Katrien Antonio, Richard Plat

DEPARTMENT OF ACCOUNTANCY, FINANCE AND INSURANCE (AFI)

Micro-level stochastic loss reserving for general insurance

Katrien Antonio ^{*} Richard Plat [†]

May 29, 2012

Abstract

¹ This paper demonstrates the merits of loss reserving using detailed information on the development of individual claims. Apart from few exceptions, the vast literature on stochastic loss reserving is developed for data aggregated in run-off triangles. However, a triangle is a summary of an underlying, more detailed, data set. We refer to this data set at individual claim level as ‘micro-level’ data. A realistic micro-level data set on liability claims from a European insurance company is analyzed. We specify a stochastic model for each event in the development of a claim: the time of occurrence of the claim, the delay between occurrence and reporting of the claim to the insurance company, the occurrence of payments and their size and the final settlement of the claim. We calibrate the resulting model to historical data and use it to project the future development of open claims. Through an out-of-sample prediction exercise we show that the micro-level approach provides the actuary with detailed and valuable reserve calculations. For the case-study developed in this paper, the micro-level model outperforms the results obtained using traditional loss reserving methods based on aggregate data.

Key words: loss reserving, general insurance, poisson process, recurrent events, survival analysis, prediction.

^{*}University of Amsterdam and KU Leuven (Belgium), email: k.antonio@uva.nl. Katrien Antonio acknowledges financial support from the Casualty Actuarial Society, Actuarial Foundation and the Committee on Knowledge Extension Research of the Society of Actuaries, and from NWO through a Veni 2009 grant.

[†]University of Amsterdam and Richard Plat Consultancy, email: rplatconsultancy@gmail.com

¹The authors would like to thank Jan-Willem Vulto and Joris van Kempen for supplying and explaining the data. Please note that the original frequency and severity data have been transformed for reasons of confidentiality.


1 Introduction

We develop a stochastic model for the run-off of general insurance (also called ‘non-life’ or ‘property and casualty’) claims. Our paper demonstrates the merits of a loss reserving model designed for the development of individual claims (as depicted on a time line in Figure 1). This micro-level loss reserving model is a valuable alternative for traditional loss reserving techniques using data aggregated in run-off triangles. We motivate that our micro-level model is able to deal with many of the problems that arise when analyzing data aggregated in triangles. Moreover, an out-of-sample prediction exercise (see Section 6) shows that the predictive distributions obtained from the micro-level model reflect reality more accurately than models based on run-off triangles do.

Figure 1 illustrates the run-off (or development) process of a general insurance claim. The structure in Figure 1 is generic for the kind of information that is available and is labeled here as ‘micro-level’. A claim occurs at a certain point in time (t_1), consequently it is declared to the insurer (t_2) (possibly after a period of delay) and one or several payments follow until the settlement (or closing) of the claim. Depending on the nature of the business and claim, the claim can re-open and payments follow until the claim finally settles. At the present moment (say τ) the insurer has to set reserves aside to fulfill his liabilities in the future. This actuarial exercise is known as ‘loss’ or ‘claims reserving’. General insurers distinguish between reserves for RBNS and IBNR claims. ‘RBNS’ claims are claims that are Reported to the insurer But Not Settled, whereas ‘IBNR’ claims Incurred But are Not Reported to the company. For an RBNS claim occurrence and declaration take place before the present moment and settlement occurs afterwards (i.e. $\tau \geq t_2$ and $\tau < t_6$ (or $\tau < t_9$) in Figure 1). An IBNR claim has occurred before the present moment, but its declaration and settlement follow afterwards (i.e. $\tau \in [t_1, t_2]$ in Figure 1). The interval $[t_1, t_2]$ represents the so-called reporting delay. The interval $[t_2, t_6]$ (or $[t_2, t_9]$) represents the settlement delay.

With the introduction of Solvency 2 and IFRS 4 Phase 2 insurers face major challenges. IFRS 4 Phase 2 will define a new accounting framework for insurance contracts, based on the market value of liabilities. This framework requires the insurer to measure its liabilities (for reserving) using ‘best estimates’ of the contractual cash flows, discounted with current market discount rates. A margin is added to compensate market participants for bearing risk. Under Solvency 2 the so-called Solvency Capital Requirement (‘SCR’) will be risk-based, and market values of assets and liabilities will be the basis for these calculations. Therefore, in light of the introduction of these new guidelines, the measurement of future cash flows and their uncertainty is gaining importance. The question naturally rises whether current techniques can be

Figure 1: *Development of a general insurance claim*



improved. Reserving for general insurance is nowadays based on data aggregated in run-off triangles. A run-off triangle summarizes observable variables per arrival and development year combination. The *arrival year* ('AY') or *year of occurrence* is the year in which the accident takes place. For a claim from AY t the first *development year* is year t itself, the second development year is $t + 1$ and so on. An example of a run-off triangle is given in Table 3 and 4 below. An overview of techniques for loss reserving based on triangles is given in England and Verrall (2002), Wüthrich and Merz (2008) or Kaas et al. (2008).

England and Verrall (2002) and Taylor et al. (2008) question the use of aggregate loss data. With aggregate data, lots of useful information about the claims development process remain unused. Information from policy, policy holder or the past development process is not taken into account, because each cell in the run-off triangle is an aggregate figure. Quoting England and Verrall (2002) (page 507) "[...] it has to be borne in mind that traditional techniques were developed before the advent of desktop computers, using methods which could be evaluated using pencil and paper. With the continuing increase in computer power, it has to be questioned whether it would not be better to examine individual claims rather than use aggregate data". Many problems or issues may arise with triangular data. Kunkler (2004) discusses the problem of zero or negative cells in the triangle. Verdonck et al. (2009) put focus on robustness properties and the influence of outliers on triangular methods. The number of observations in a run-off triangle is typically small, and for recent accident years only very few observations are available. In that respect Wright (1990) and Renshaw (1994) discuss the over-parametrization of the chain-ladder method, which is a well known method for analyzing data in triangles. Bornhuetter and Ferguson (1972) discuss the instability in ultimate claims for recent arrival years. The separate assessment of true IBNR and RBNS claims in a run-off tri-

angle is not straightforward, see Schnieper (1991) and Liu and Verrall (2009). Neither is the combination of different sources of information, like paid and incurred losses (i.e. paid losses plus case estimates set by experts), see the different approaches presented in Quarg and Mack (2008), Posthuma et al. (2008) and Merz and Wüthrich (2010). All these references present useful adjustments to the chain-ladder method, but they have not been applied simultaneously. The existence of this substantial literature illustrates that a chain-ladder type analysis of aggregate data is not always adequately capturing the complexities of stochastic reserving for general insurance.

A small stream of literature has emerged with focus on stochastic loss reserving at individual claim level. Arjas (1989), Norberg (1993) and Norberg (1999) formulate a mathematical framework for the development of individual claims. Using ideas from martingale theory and point processes, these authors present a probabilistic, rather than statistical, framework for individual claims reserving. Haastrup and Arjas (1996) continue their work and present a detailed implementation of a micro-level stochastic model for loss reserving. Their use of non-parametric Bayesian statistics complicates the accessibility of their approach. Furthermore, their case study is based on a small data set with fixed claim amounts. Recently, Larsen (2007) revisits the work of Norberg, Haastrup and Arjas with a small case-study. Zhao et al. (2009) and Zhao and Zhou (2010) present a model for individual claims development using (semi-parametric) techniques from survival analysis and copula methods. However, a case study is lacking in their work.

Our work is an extensive case-study developed in the probabilistic framework of Norberg (1993) and Norberg (1999). We analyze a realistic data base from practice. We motivate our distributional assumptions and derive predictive distributions for quantities of interest. Finally, we check the performance of the micro-level model through an out-of-sample prediction exercise and compare our approach with triangular methods.

Our approach specifies stochastic models for the occurrence of claims, the reporting delay, the development process and the payments made during the development of a claim. This allows explicit quantification of the Incurred But Not Reported (IBNR) reserve as well as the Reported But Not Settled (RBNS) reserve. We show that the quality of reserves improves when using micro-level data. Moreover, the problems encountered when analyzing data in run-off triangles will no longer exist. The large data set available at micro-level allows flexible modeling of the claims development process. For example, covariate information (like deductibles, policy limits, calendar year) can be included in the projection of the cash flows. The use of individual data avoids robustness problems and over parametrization. Problems with negative or zero cells in the triangle are avoided, and small and large claims can be handled simultane-

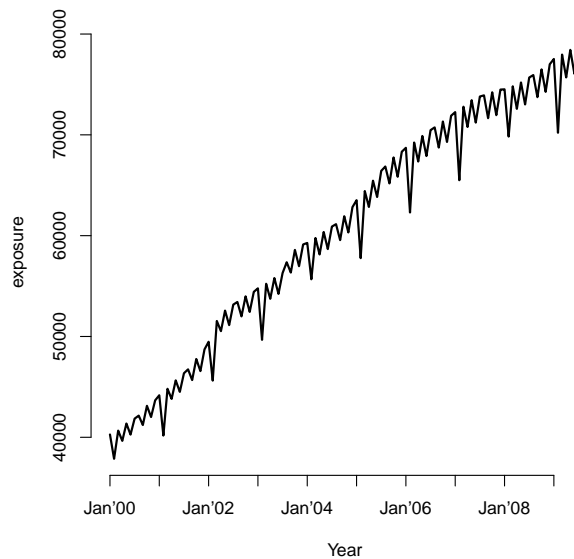
ously. Furthermore, individual claim modeling can provide a natural solution for the dilemma whether to use triangles with paid or incurred losses. We propose using the initial case reserve as a covariate in the projection process of future cash flows.

2 Data

The data come from a general liability insurance portfolio (for private individuals) of a European insurance company. The exposure per month is given from January 2000 till August 2009, as well as a detailed track record of each claim filed with the insurer between January 1997 and August 2009. We are missing exposure information between January 1997 and December 1999, but the impact of this lack on our reserve calculations is negligible. We present the data and focus on the different pieces of information available. Each of these will be used in the micro-level model.

Exposure. Exposure is expressed as ‘earned’ exposure. A policy covered during the whole month of January will contribute 31/365th to the exposure of that month, 10/365th if it’s only covered during 10 days, and so on. Figure 2 shows the exposure per month. Note that the downward spikes correspond to the month February.


Figure 2: *Available exposure per month from January 2000 till August 2009.*



Type and number of claims. We distinguish two types of claims: material damage (‘material’) and bodily injury (‘injury’). Figure 3 shows the total number of claims

per arrival year (solid line), as well as the number of open claims at the end of the observation period (i.e. end of August 2009). We analyze material damage and bodily injury claims separately, because their development pattern and corresponding loss distributions are very different (see further for some descriptive statistics).

Figure 3: *Number of open and closed claims of type material (left) and injury (right).*



Development processes (overview). The claim file consists of 1,525,376 records corresponding with 491,912 claims. Figure 4 shows the development of three claims, taken at random from our data set. It shows the timing of events (namely: occurrence, declaration, payments and settlement) as well as the size of payments (if any). Payments are indicated as jumps in the figure. Starting point of the development process is the accident date. This is indicated with a sub-title in each of the plots and corresponds with the point $x = 0$. The x -axis is in months since the accident date. The y -axis represents the cumulative amount paid for the claim.

Reporting and settlement delay. Important drivers of the IBNR and RBNS reserves are the reporting and settlement delays. Figure 5 (upper) shows the reporting delay registered for material and injury claims. This delay is measured in months since occurrence of the claim. Obviously, it is only available for claims that have been reported to the insurance company before the end of the observation period. Figure 5 (lower) shows the settlement delay registered for injury and material claims. It is measured in months since the reporting of the claim and only available for closed claims. These figures show that the empirical distribution of observed reporting delays is similar for material and injury claims. However, the distribution of observed settlement delays is far more skewed to the right for injury than for material claims.

Figure 4: *Development of 3 random claims from the data set. The x-axis is in months since the accident date. The y-axis represents the cumulative amount paid for the claim.*


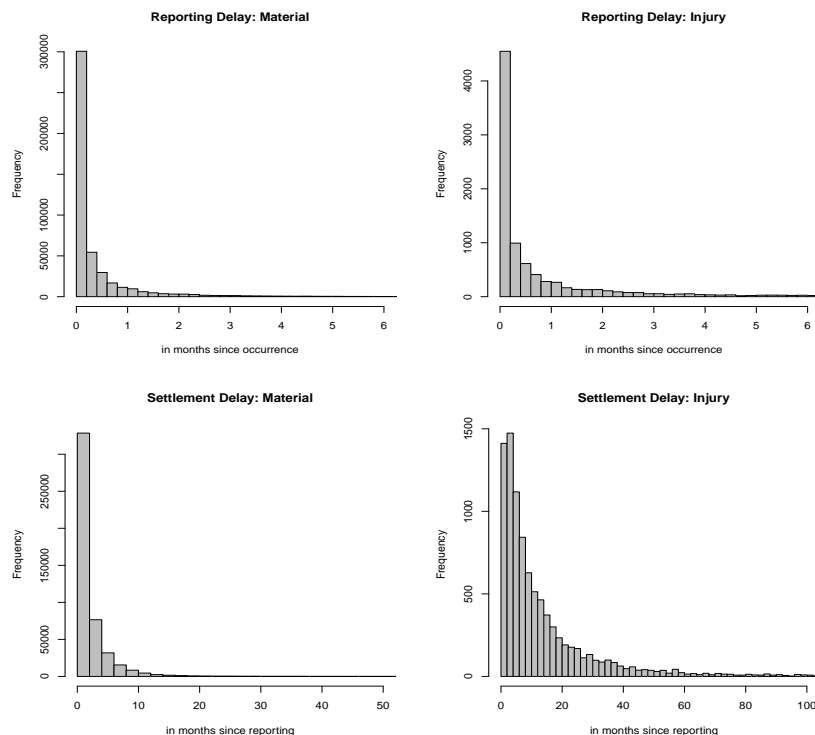


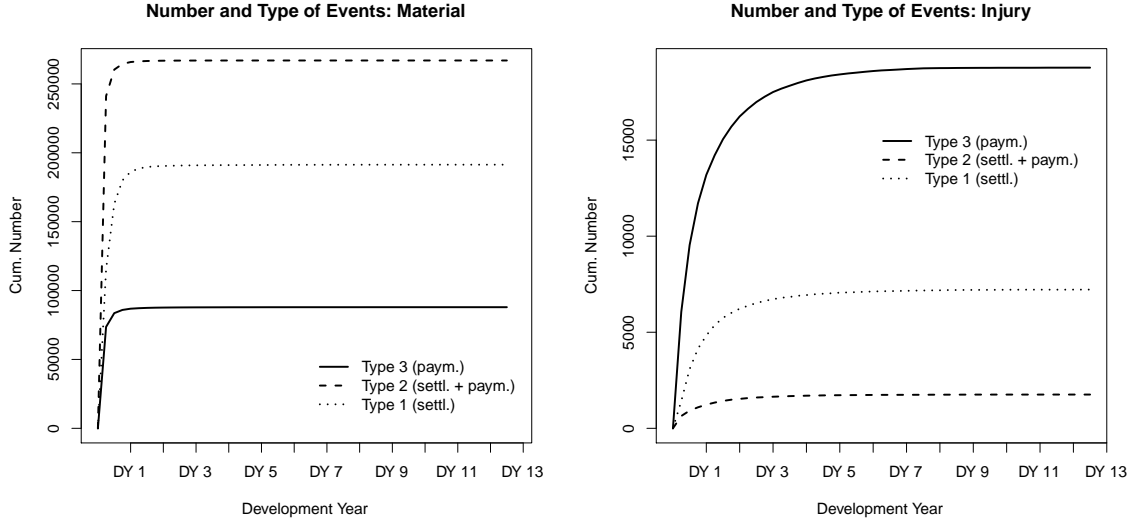
Figure 5: *Upper: reporting delay for material (left) and injury (right) claims. Lower: settlement delay for material (left) and injury (right) claims.*



Events in the development. We distinguish three types of events that can occur during the development of a claim. ‘Type 1’ events imply settlement of the claim without a payment. A ‘type 2’ event is a payment with settlement at the same time. Intermediate payments (without settlement) are ‘type 3’ events. Figure 6 gives the cumulative number of events observed over the development of individual claims. The graph shows that injury claims settle more slowly than material claims. Injury claims typically re-

quire more intermediate payments before settlement than material claims do.

Figure 6: *Cumulative number of events over development years: material (left) and injury (right).*



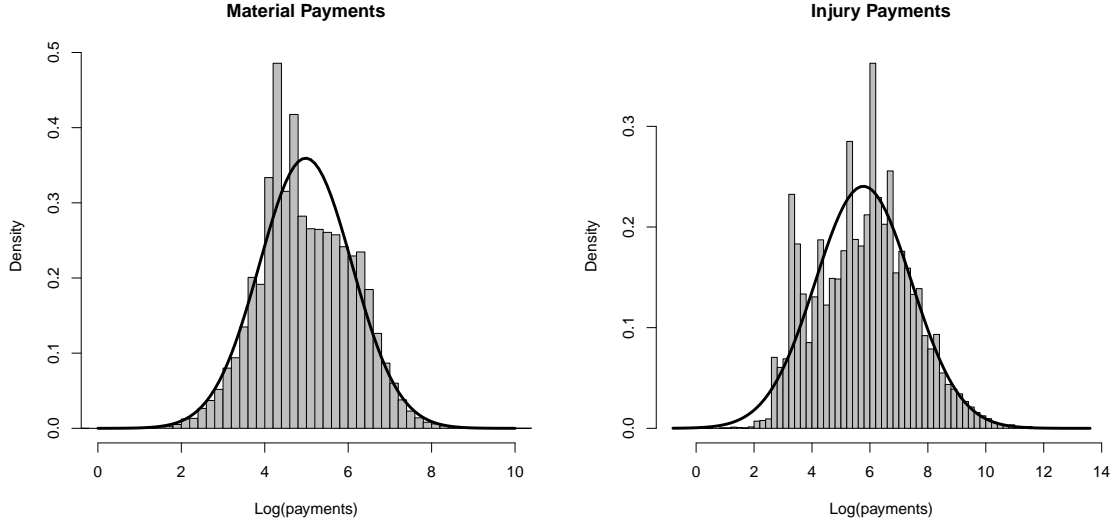
Payments. Events of type 2 and type 3 come with a payment. The distribution of these payments differs materially for the different types of claims. Figure 7 shows the distribution of payments with respect to material damage (left) and injury claims (right). The payments are discounted to 1-1-1997 with the Dutch consumer price inflation, to exclude the impact of inflation on the distribution of the payments. The figures suggest a lognormal distribution for the distribution of the payments (see Section 4 for further discussion). Table 1 gives characteristics of the observed payments for both material and injury losses. As expected, the distribution of injury payments has a heavier right tail than the distribution of material damage payments. The table reveals the presence of a very large injury payment (namely, 779,398 euro versus the empirical 99% quantile of 16,664 euro). We will discuss this further in Section 6.

Table 1: *Characteristics observed payments.*

	Mean	Median	Min.	Max.	1%	5%	25%	75%	95%	99%
Material	277	129	8×10^{-4}	198,931	12	25	69	334	890	1,768
Injury	1,395	361	0.4875	779,398	16	25	89	967	4,927	16,664

Initial case estimates. As explained in Section 1, the joint modelling of incurred and paid losses is documented in the literature on stochastic loss reserving. It is worthwhile

Figure 7: Distribution of payments for material (right) and injury (left) claims. Normal density reference line is included.



investigating the added value of case reserves when projecting future payments. We categorize the initial case reserve (for material damage and injury claims separately) and use it as an explanatory variable when modeling the distribution of payments. For each category, Table 2 shows the number of claims, the average settlement delay (in months) and the average cumulative paid amount for these categories. The table clearly reveals that claims initially set as ‘large’, have large settlement delays and lead to high ultimate cumulative amounts.

Table 2: Initial reserve categories: number of claims per category, average settlement delay and average cumulative paid amount.

Material				Injury			
Initial Case Reserve	# claims	Average settl. delay (months)	Average Cum. payments	Initial Case Reserve	# claims	Average settl. delay (months)	Average Cum. payments
$\leq 10,000$	465,015	1.87	252	$\leq 1,000$	3,709	9.87	2,570
$> 10,000$	385	10.88	7,950	(1,000 -15,000]	5,165	15.17	3,872
				$> 15,000$	360	35.2	33,840

3 The statistical model

By a claim i is understood a combination of an occurrence time T_i , a reporting delay U_i and a development process \mathbf{X}_i . The development process \mathbf{X}_i is a jump process. It is modeled here using two separate building blocks: the timing and type of events and

their corresponding severities. The different types of events are specified in Section 2. \mathbf{X}_i is short for $(E_i(v), P_i(v))_{v \in [0, V_i]}$. $E_i(v_{ij}) := E_{ij}$ is the type of the j th event in the development of claim i ; it occurs at time v_{ij} , expressed in time units after notification of the claim. V_i is the total waiting time from notification to settlement of claim i . If the event at time v_{ij} includes a payment, its corresponding severity is given by $P_i(v_{ij}) := P_{ij'}$ where j' runs over all payments in the development of this claim. Denote the present moment with τ . IBNR, RBNS and settled claims can be distinguished as follows:

- for an IBNR claim: $T_i + U_i > \tau$ and $T_i < \tau$;
- for an RBNS claim: $T_i + U_i \leq \tau$ and the development of the claim is censored at $(\tau - T_i - U_i)$, i.e. only $(E_i(v), P_i(v))_{v \in [0, \tau - T_i - U_i]}$ is observed;
- for a settled claim: $T_i + U_i \leq \tau$ and $(E_i(v), P_i(v))_{v \in [0, V_i]}$ is observed.

3.1 Position dependent marked Poisson process

Following the approach in Arjas (1989) and Norberg (1993) we treat the claims process as a Position Dependent Marked Poisson Process (PDMPP), see Karr (1991). In this application, a point is an occurrence time and the associated mark is the combined reporting delay and development of the claim. We denote the intensity measure of this Poisson process by λ and the associated mark distribution by $(P_{Z|t})_{t \geq 0}$ with t the occurrence time of the claim. The mark distribution $P_{Z|t}$ is specified by the distribution $P_{U|t}$ of the reporting delay (given occurrence time t) and the distribution $P_{X|t,u}$ of the development (given occurrence time t and reporting delay u). The complete development process is a Poisson process on claim space $\mathcal{C} = [0, \infty) \times [0, \infty) \times \chi$ with intensity measure:

$$\lambda(dt) \times P_{U|t}(du) \times P_{X|t,u}(dx) \text{ with } (t, u, x) \in \mathcal{C}. \quad (1)$$

The reported claims (which are not necessarily settled) belong to the set:

$$\mathcal{C}^r = \{(t, u, x) \in \mathcal{C} | t + u \leq \tau\}, \quad (2)$$

whereas the IBNR claims belong to:

$$\mathcal{C}^i = \{(t, u, x) \in \mathcal{C} | t \leq \tau, t + u > \tau\}. \quad (3)$$

Since both sets are disjoint, both processes are independent (see Karr (1991)). The pro-

cess of reported claims is a Poisson process on \mathcal{C} with measure

$$\begin{aligned} & \lambda(dt) \times P_{U|t}(du) \times P_{X|t,u}(d\mathbf{x}) \times 1_{[(t,u,\mathbf{x}) \in \mathcal{C}^r]} \\ = & \underbrace{\lambda(dt)P_{U|t}(\tau-t)1_{(t \in [0,\tau])}}_{(a)} \times \underbrace{\frac{P_{U|t}(du)1_{(u \leq \tau-t)}}{P_{U|t}(\tau-t)}}_{(b)} \times \underbrace{P_{X|t,u}(d\mathbf{x})}_{(c)}. \end{aligned} \quad (4)$$

Three parts can be identified in (4). The occurrence measure for reported claims is in (a). This is the intensity of the Poisson process driving the occurrence of claims, $\lambda(dt)$, adjusted for the fact that focus is only on reported claims. (b) refers to the distribution of the reporting delay, conditioned on the fact that reporting should take place. (c) represents the distribution of the development process, given occurrence time t and reporting delay u of the claim. Similarly, the process of IBNR claims is a Poisson process with measure:

$$\underbrace{\lambda(dt) \left(1 - P_{U|t}(\tau-t)\right) 1_{(t \in [0,\tau])}}_{(a)} \times \underbrace{\frac{P_{U|t}(du)1_{u > \tau-t}}{1 - P_{U|t}(\tau-t)}}_{(b)} \times \underbrace{P_{X|t,u}(d\mathbf{x})}_{(c)}, \quad (5)$$

where similar components can be identified as in (4), conditioned on the fact that claims have not been reported.

3.2 The likelihood

The observed part of the process consists of the development up to time τ of claims reported before τ . We denote these observed claims as follows:

$$(T_i^o, U_i^o, \mathbf{X}_i^o)_{i \geq 1}, \quad (6)$$

where the development of claim i is censored $\tau - T_i^o - U_i^o$ time units after notification. The likelihood of the observed claim development process can be written as (see Cook and Lawless (2007)):

$$\begin{aligned} \Lambda(obs) & \propto \left\{ \prod_{i \geq 1} \lambda(T_i^o) P_{U|t}(\tau - T_i^o) \right\} \exp \left(- \int_0^\tau w(t) \lambda(t) P_{U|t}(\tau - t) dt \right) \\ & \times \left\{ \prod_{i \geq 1} \frac{P_{U|t}(dU_i^o)}{P_{U|t}(\tau - T_i^o)} \right\} \times \prod_{i \geq 1} P_{X|t,u}^{\tau - T_i^o - U_i^o}(d\mathbf{X}_i^o). \end{aligned} \quad (7)$$

The superscript in the last term of this likelihood indicates the censoring of the development of this claim $\tau - T_i^o - U_i^o$ time units after notification. $w(t)$ gives the exposure

at time t .

We use techniques from survival analysis to model the reporting delay and the development process. The reporting of the claim is an event that only occurs once during the existence of the claim. Therefore, reporting delay is modeled using standard distributions from survival analysis. The statistical framework of recurrent events (see Cook and Lawless (2007)) is suitable to model the development of a claim. A hazard rate is specified for each type of event that can occur during the development. Denote these hazard rates with h_{se} , h_{sep} and h_p , corresponding to type 1 (settlement without payment), type 2 (settlement with a payment at the same time) and type 3 (payment without settlement) events, respectively.

Events of type 2 and 3 come with a payment. We denote the distribution of a severity payment with P_p . Using this notation the likelihood of the development process of claim i is given by:

$$\left\{ \prod_j \left(h_{se}^{\delta_{ij1}}(V_{ij}) \times h_{sep}^{\delta_{ij2}}(V_{ij}) \times h_p^{\delta_{ij3}}(V_{ij}) \right) \right\} \times \exp \left(- \int_0^{\tau_i} (h_{se}(u) + h_{sep}(u) + h_p(u)) du \right) \times \prod_{j'} P_p(dP_{ij'}). \quad (8)$$

Here δ_{ijk} is an indicator variable that is 1 if the j th event in the development of claim i is of type k . j runs over all events registered in the observation period for claim i . This observation period is $[0, \tau_i]$ with $\tau_i = \min(\tau - T_i - U_i, V_i)$. j' runs over all payments made during the development of this claim.

Combining (7) and (8) gives the likelihood for the observed data:

$$\begin{aligned} \Lambda(obs) &\propto \left\{ \prod_{i \geq 1} \lambda(T_i^o) P_{U|t}(\tau - T_i^o) \right\} \exp \left(- \int_0^\tau w(t) \lambda(t) P_{U|t}(\tau - t) dt \right) \\ &\times \left\{ \prod_{i \geq 1} \frac{P_{U|t}(dU_i^o)}{P_{U|t}(\tau - T_i^o)} \right\} \times \prod_{i \geq 1} \prod_j \left(h_{se}^{\delta_{ij1}}(V_{ij}) \times h_{sep}^{\delta_{ij2}}(V_{ij}) \times h_p^{\delta_{ij3}}(V_{ij}) \right) \\ &\times \exp \left(- \int_0^{\tau_i} (h_{se}(u) + h_{sep}(u) + h_p(u)) du \right) \times \prod_{i \geq 1} \prod_{j'} P_p(dP_{ij'}). \end{aligned} \quad (9)$$

3.3 Distributional assumptions

We discuss the likelihood in (9) in more detail and present distributional assumptions for its building blocks. Our final choices and estimation results will be covered in Section 4.

Reporting delay. A large number of claims is reported in the days immediately following occurrence of the claim, see Figure 5. We use a mixture of one particular standard distribution (f_U) with n degenerate distributions for notification during these first days:

$$\sum_{k=0}^{n-1} p_k I_{\{k\}}(u) + \left(1 - \sum_{k=0}^{n-1} p_k\right) f_{U|U>n-1}(u), \quad (10)$$

where $I_{\{k\}}(u) = 1$ if reporting takes place on the k th day after occurrence and $I_{\{k\}}(u) = 0$ otherwise.

Occurrence process. When optimizing the likelihood for the occurrence process, we use the reporting delay distribution and its parameters (as obtained in the previous step). The likelihood

$$\left\{ \prod_{i \geq 1} \lambda(T_i^o) P_{U|t}(\tau - T_i^o) \right\} \exp \left(- \int_0^\tau w(t) \lambda(t) P_{U|t}(\tau - t) dt \right), \quad (11)$$

needs to be optimized over $\lambda(t)$. We use a piecewise constant specification for the occurrence rate: $\lambda(t) = \lambda_l$ for $d_{l-1} \leq t < d_l$, $l = 1, \dots, m$ and $d_0 = 0$. Hereby $\tau \in [d_{m-1}, d_m)$ and $w(t) := w_l$ for $d_{l-1} \leq t < d_l$. Let the indicator variable $\delta_1(l, t_i)$ be 1 if $d_{l-1} \leq t_i < d_l$, with t_i the occurrence time of claim i . The number of claims in interval $[d_{l-1}, d_l)$ is given by

$$N_{oc}(l) := \sum_i \delta_1(l, t_i). \quad (12)$$

The likelihood corresponding with the occurrence times then becomes

$$\begin{aligned} & \lambda_1^{N_{oc}(1)} \lambda_2^{N_{oc}(2)} \dots \lambda_m^{N_{oc}(m)} \left\{ \prod_{i \geq 1} P_{U|t}(\tau - t_i) \right\} \\ & \times \exp \left(-\lambda_1 w_1 \int_0^{d_1} P_{U|t}(\tau - t) dt \right) \exp \left(-\lambda_2 w_2 \int_{d_1}^{d_2} P_{U|t}(\tau - t) dt \right) \\ & \times \dots \exp \left(-\lambda_m w_m \int_{d_{m-1}}^{d_m} P_{U|t}(\tau - t) dt \right). \end{aligned} \quad (13)$$

Optimizing over λ_l (with $l = 1, \dots, m$) leads to:

$$\hat{\lambda}_l = \frac{N_{oc}(l)}{w_l \int_{d_{l-1}}^{d_l} P_{U|t}(\tau - t) dt}. \quad (14)$$

Development process. A piecewise constant specification is used for the hazard rates. This piecewise specification can be integrated in a straightforward way in likelihood specification (9), although the resulting expression is complex in notation. The optimization of the likelihood expression can be done analytically (which results in very elegant and compact expressions) or numerically. It might be worthwhile to specify a separate hazard rate for ‘first events’ in the development and ‘later events’. This will be investigated in Section 4.

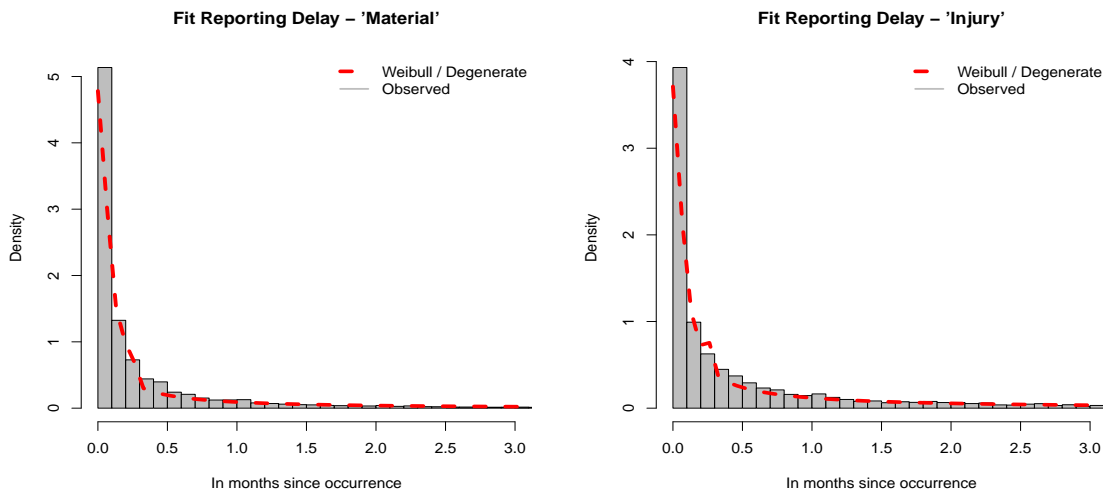
Payments Events of type 2 and type 3 come with a payment. We investigate a selection of univariate distributions suitable for modeling severities. Covariate information (e.g. the initial reserve category and the development year) is taken into account.

4 Estimation results

We motivate our distributional assumptions and show the outcomes of the calibration process. Optimization of all likelihood specifications was done with the Proc NLMixed routine in SAS.


Reporting delay. A mixture of a Weibull distribution and 9 degenerate components corresponding with settlement after 0, \dots , 8 days, is used. Figure 8 illustrates the fit of this mixture of distributions to the actually observed reporting delays.

Figure 8: *Observed reporting delays for material (left) and injury (right) claims and fit obtained with 9 degenerate components combined with a truncated Weibull distribution.*



Occurrence process. Using the distribution for the reporting delay, we optimize the likelihood (13) for the occurrence times. $\lambda(t)$ is constant on monthly intervals, ranging from January 2000 till August 2009. Figure 9 shows point estimates and corresponding 95% confidence intervals.

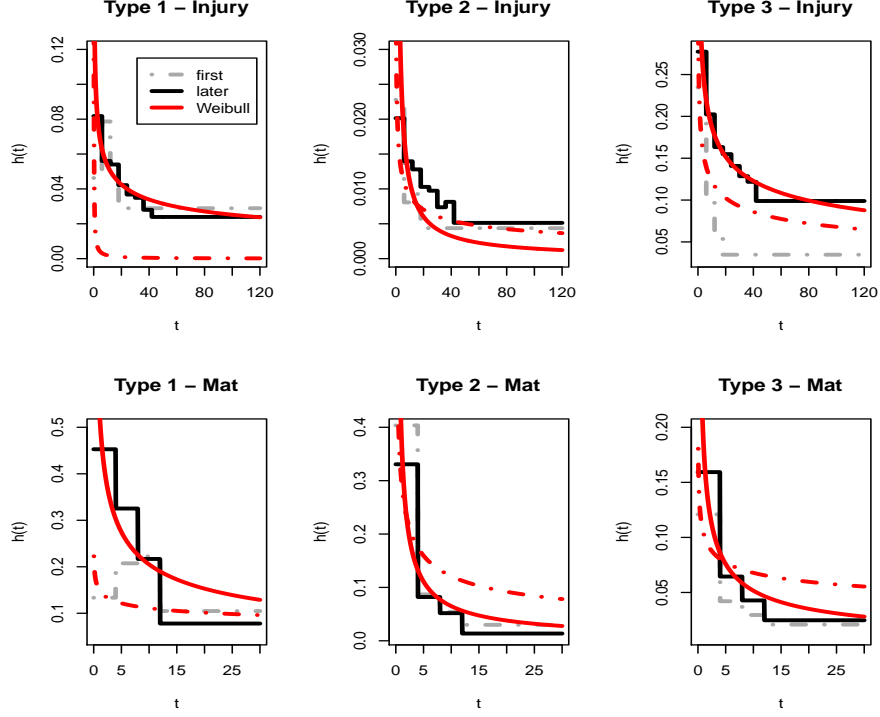
Figure 9: *Estimates of piecewise specification for $\lambda(t)$: (left) material and (right) injury claims.*



Development process. We compare the use of a Weibull as well as a piecewise constant specification for the hazard rates h_{se} , h_{sep} and h_p . For material claims, we assume the hazard rate is constant on four month intervals: $[0 - 4)$ months, $[4 - 8)$ months, \dots , $[8 - 12)$ months and ≥ 12 months. For injury claims, the hazard rate is constant on intervals of six months: $[0 - 6)$ months, $[6 - 12)$ months, \dots , $[36 - 42)$ months and ≥ 42 months. Figure 10 shows estimates for Weibull and piecewise constant hazard rates. Separate hazard rates are specified for ‘first events’ and ‘later events’. The piecewise constant specification reflects the actual data. The figure shows that the Weibull distribution is reasonably close to the piecewise constant specification. In the rest of this paper we use the piecewise constant specification. Section 5 includes brief discussion of the use of the Weibull hazard rates in the simulation routine.

Payments. We examine the fit of a Burr, gamma and lognormal distribution, including covariate information. Distributions for the payments are truncated at the coverage limit of 2.5 million euro per claim. A comparison based on BIC shows that the lognormal distribution achieves a better fit than the Burr and gamma distributions. When including the initial reserve category as covariate or both the initial reserve category and the development year, the fit further improves. The latter approach is used in the

Figure 10: Estimates for Weibull and piecewise constant hazard rates driving the occurrence of events in the development of claims: (upper) injury claims and (lower) material claims. ‘Type 1’ events represent settlement without payment, ‘type 2’ refers to settlement with payment and ‘type 3’ is for intermediate payments.



prediction routine. We include covariate information in both mean (μ_i) and standard deviation (σ_i) of the lognormal distribution for observation i :


$$\begin{aligned}\mu_i &= \sum_r \sum_s \mu_{r,s} I_{DY_i=s} I_{i \in r} \\ \sigma_i &= \sum_r \sum_s \sigma_{r,s} I_{DY_i=s} I_{i \in r}.\end{aligned}\tag{15}$$

Hereby r is the initial reserve category and DY_i is the development year corresponding with observation i . $I_{DY_i=s}$ and $I_{i \in r}$ are indicator variables referring to observation i belonging to DY s and reserve category r , respectively. Figure 11 shows corresponding qqplots.

5 Predicting future cash-flows

Distinguishing IBNR and RBNS claims is necessary in the prediction routine. The following step by step approach is implemented.

Figure 11: Normal qqplots for the fit of $\log(\text{payments})$ including initial reserve and development year as covariate information.



Predicting IBNR claims. As noted in Section 3, an IBNR claim occurred already but has not yet been reported to the insurer. Therefore, $T_i + U_i > \tau$ and $T_i < \tau$ with T_i the occurrence time of the claim and U_i its reporting delay. The T_i s are missing data: they are determined in the development process but unknown to the actuary at time τ . The prediction process for the IBNR claims requires the following steps:

- (a) **Simulate the number of IBNR claims in $[0, \tau]$ and their corresponding occurrence times.**

IBNR claims are driven by a Poisson process with intensity:

$$w(t)\lambda(t)(1 - P_{U|t}(\tau - t)), \quad (16)$$

where $\lambda(t)$ is piecewise constant. With $N_{\text{IBNR}}(l)$ being the number of IBNR claims occurring in time interval $[d_{l-1}, d_l)$, we know

$$N_{\text{IBNR}}(l) \sim \text{Poisson} \left(\lambda_l w_l \int_{d_{l-1}}^{d_l} (1 - P_{U|t}(\tau - t)) dt \right). \quad (17)$$

Given the simulated number of IBNR claims $n_{\text{IBNR}}(l)$ for each interval $[d_{l-1}, d_l)$, the occurrence times of the claims are uniformly distributed in $[d_{l-1}, d_l)$. These are simulated as well.

- (b) **Simulate the reporting delay for each IBNR claim.**

Given the simulated occurrence time t of an IBNR claim, its reporting delay is

simulated by inverting the distribution:

$$P(U \leq u | U > \tau - t) = \frac{P(\tau - t < U \leq u)}{1 - P(U \leq \tau - t)}. \quad (18)$$

Using a mixture of a Weibull distribution and 9 degenerate distributions requires numerical evaluation of this expression.

(c) Simulate the initial reserve category.

For each IBNR claim an initial reserve category has to be simulated for use in the development process. Given m initial reserve categories, we use the following discrete probability density function:

$$f(c) = \begin{cases} p_c & \text{for } c = 1, 2, \dots, m-1 \\ 1 - \sum_{k=1}^{m-1} p_k & \text{for } c = m. \end{cases} \quad (19)$$

The probabilities used in (19) are the empirically observed percentages of claims in a particular initial reserve category.

(d) Simulate the payment process for each IBNR claims.

This step is common with the procedure for RBNS claims and is explained in the next paragraph.

Predicting RBNS claims Given the RBNS claims and the simulated IBNR claims, the process proceeds as follows.

(e) Simulate the next event's exact time.

In case of RBNS claims, the time of censoring c of claim is known. For IBNR claims this censoring time is $c := 0$. The next event – at time v_{next} – can take place at any time $v_{\text{next}} > c$. To simulate its exact time we need to invert: (with p randomly drawn from a $\text{Unif}(0, 1)$ distribution)

$$P(V < v_{\text{next}} | V > c) = p. \quad (20)$$

From the relation between a hazard rate and cdf, we know

$$P(V \leq v_{\text{next}}) = 1 - \exp \left(- \int_0^{v_{\text{next}}} \sum_e h_e(t) dt \right), \quad (21)$$

with $e \in \{se, sep, p\}$. Numerical methods are required, for instance with a Weibull specification for the hazard rates. With a piecewise constant specification for the hazard rates numerical routines as well as closed-form solutions are available.

- (f) **Simulate the event type.** Given the exact time v of the next event, it is of type $e \in \{se, sep, p\}$ with probability:

$$\frac{h_e(v)}{\sum_e h_e(v)}. \quad (22)$$

- (g) **Simulate the corresponding payment (if any).** Given a claim's covariate information, payments are drawn from the appropriate lognormal distribution. The cumulative payment cannot exceed the coverage limit of 2.5 million per claim.
- (h) **Stop or continue.** Depending on the simulated event type in step (f), the prediction stops (in case of settlement) or continues.

In the next section, this prediction process will be applied separately for the material claims and the injury claims.

Comment on parameter uncertainty. With respect to the uncertainty of predictions, process as well as estimation or parameter uncertainty (see England and Verrall (2002)) should be taken into account. We cover process uncertainty by sampling from the distributions selected in Section 3 and 4. To include parameter uncertainty the bootstrap technique or concepts from Bayesian statistics can be used. While a formal Bayesian approach is very elegant, it generally leads to significantly more complexity, which is not contributing to the accessibility and transparency of the techniques towards practicing actuaries. Applying a bootstrap procedure is very computer intensive in this application, in light of the large sample size and the different stochastic processes used. We deal with parameter uncertainty through the asymptotic normal distribution of the maximum likelihood estimators. At every step in the routine, each parameter is sampled from its corresponding asymptotic normal distribution. Note that –due to our large sample size– confidence intervals for parameters are narrow. This is in contrast with run-off triangles where sample sizes are typically very small and parameter uncertainty is an important point of concern.

6 Numerical results

We present the results of an out-of-sample prediction exercise. The micro-level results are compared with the results of a triangular analysis (based on aggregate data). The out-of-sample test estimates reserves per 1-1-2005. The available data are summarized in run-off triangles (see Table 3 for material damage and Table 4 for bodily injury). The

actual observations registered between January 2005 and August 2009 are displayed in bold.

Table 3: *Run-off triangle material claims (displayed in thousands), arrival years 1997-2004.*

Arrival Year	Development Year							
	1	2	3	4	5	6	7	8
1997	4,380	972	82	9	36	27	34	11
1998	4,334	976	56	35	76	24	0.57	17
1999	5,225	1,218	59	108	108	12	0.39	0
2000	5,366	1,119	161	14	6	4	0.36	10
2001	5,535	1,620	118	119	13	3	0.35	2
2002	6,539	1,547	67	65	17	5	9	8.80
2003	6,535	1,601	90	21	31	7	1.7	
2004	7,109	1,347	99	76	20	13		

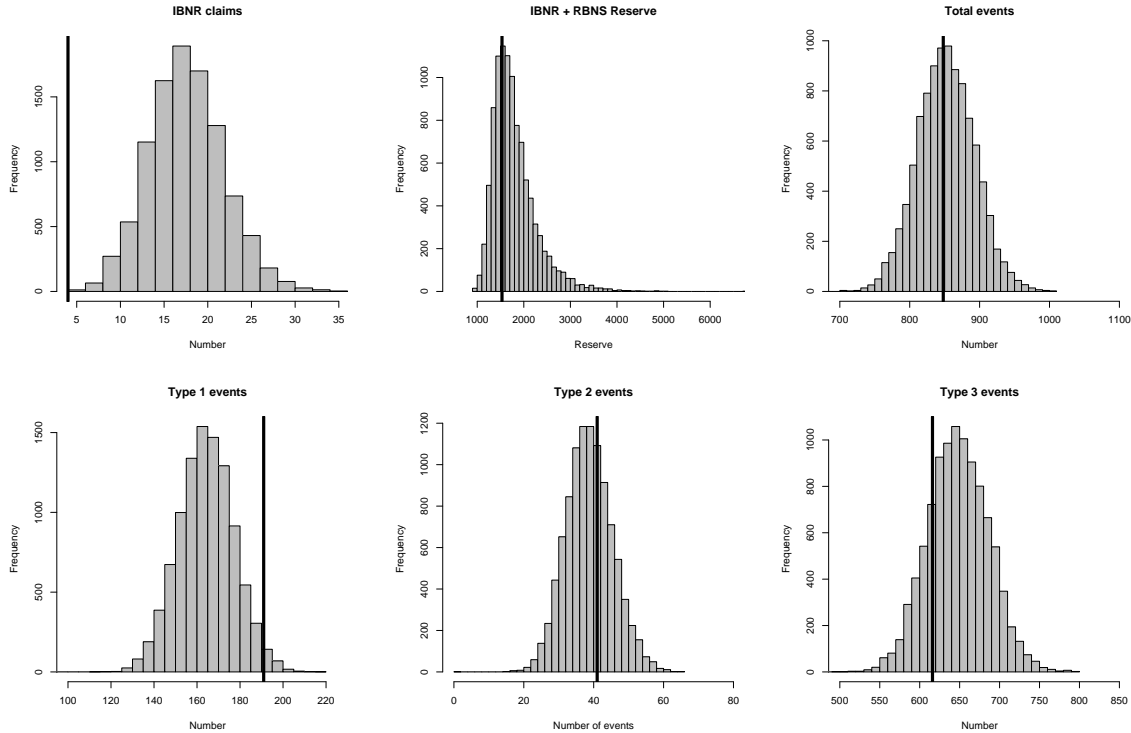
Table 4: *Run-off triangle injury claims (displayed in thousands), arrival years 1997-2004.*

Arrival Year	Development Year							
	1	2	3	4	5	6	7	8
1997	308	635	366	530	549	137	132	339
1998	257	482	312	336	269	56	179	78
1999	292	590	410	273	254	286	132	97
2000	317	601	439	498	407	371	247	275
2001	466	846	566	567	446	375	147	240
2002	314	615	540	449	133	131	332	1,082
2003	304	802	617	268	223	216	173	
2004	333	864	412	245	273	100		

Output from the micro-level model. We consider the different types of output coming from the micro-level model. Figure 12 shows the results for injury payments made in calendar year 2006, based on 10,000 simulations. In Table 4 calendar year 2006 corresponds with the diagonal going from 412, 268, ..., up to 97. The first row in Figure 12 shows (from left to right): the number of IBNR claims reported in 2006, the total amount paid in this calendar year and the total number of events occurring in 2006. The IBNR claims are claims that occurred before 1-1-2005, but were reported to the insurer during calendar year 2006. The total amount paid in 2006 is the sum of payments for RBNS and IBNR claims, which are separately available from the micro-model. In the second row of plots we take a closer look at the events registered in 2006 by looking separately at type 1, 2 and 3 events. In each of the plots the black solid line indicates

what was actually observed. This figure shows that the predictive distributions from the micro-level model are realistic. However, the actual number of IBNR claims is far in the tail of the distribution, but note that this corresponds with a small number of IBNR claims.

Figure 12: *Out-of-sample exercise per 1-1-2005, injury claims. Results are for calendar year 2006, based on 10,000 simulations from the micro-level model. Top row (from left to right): number of IBNR claims, total reserve (i.e. IBNR plus RBNS reserve), total number of events. Bottom row (from left to right): number of type 1, 2 and 3 events. The black solid line indicates actually observed quantities.*



Comparing reserves. We compare the results obtained with the micro-level model with reserve calculations based on aggregate data. More specifically, we consider Mack's chain-ladder model, a stochastic overdispersed Poisson and lognormal chain-ladder model. The model specifications for overdispersed Poisson (see (23)) and lognormal (see (24)) are given below, where Y_{ij} denotes cell (i, j) in a run-off triangle (and corresponds to arrival year i and development year j). The models specified in (23)

and (24) are implemented in a Bayesian framework.²

$$\begin{aligned} Y_{ij} &= \phi M_{ij} \\ M_{ij} &\sim \text{Poi}(\mu_{ij}/\phi) \\ \log(\mu_{ij}) &= \alpha_i + \beta_j; \end{aligned} \tag{23}$$

$$\begin{aligned} \log(Y_{ij}) &= \mu_{ij} + \epsilon_{ij} \\ \mu_{ij} &= \alpha_i + \beta_j \\ \epsilon_{ij} &\sim N(0, \sigma^2). \end{aligned} \tag{24}$$

Figure 13 shows the reserves (in thousands) for material claims, as obtained with the different methods (from left to right: micro-model, chain-ladder overdispersed Poisson and chain-ladder lognormal). The histograms are based on 10,000 simulations of the total reserve. Corresponding numerical results are in Table 5. The total reserve predicts the complete lower triangle (i.e. all bold numbers, plus the three missing cells in Table 3). The solid black line in each plot indicates what has really been observed, i.e. the sum of the numbers in bold in Table 3. We use the same scale on the x -axis of the histograms representing the micro-level and the overdispersed Poisson model. However, for the lognormal model a different scale on the x -axis is necessary, because of the presence of a long right tail. These unrealistically high reserves (see Table 5) are a disadvantage of the lognormal model for the portfolio of material claims. We conclude from Figure 13 and Table 5 that the overdispersed Poisson as well as the lognormal model overstate the reserve; the actually observed amount is in the left tail of the corresponding histogram. The predictive distribution obtained with the micro-model is more realistic. The corresponding best estimate is closer to the true realization than the best estimates from aggregate techniques.

Figure 14 shows the distribution of the reserve (in thousands of euro) obtained for bodily injury claims (based on 10,000 simulations). In contrast to the plots in Figure 13 the plots in Figure 14 use the same x -axis. Corresponding numerical results are in Table 5. The observed run-off triangle in Table 4 shows a very large total payment (779,383 euro) in occurrence year 2002, development year 8. This payment is much larger than all other payments in the data set (see the statistics in Table 1). The micro-level model reflects this appropriately, i.e. the observed total amount is rather in the right tail of the predictive distribution. The aggregate models again tend to overstate the reserve.

We conclude that, for the case-study under consideration, the micro-model outper-

²The implementation of the overdispersed Poisson model is in fact empirically Bayesian. ϕ is estimated beforehand and held fixed. We use vague normal priors for the regression parameters in both models and a gamma prior for σ^{-1} in the lognormal model.

Figure 13: *Out-of-sample exercise per 1-1-2005, material claims. Results are for the total reserve (i.e. IBNR + RBNS reserve), based on 10,000 simulations. From left to right: reserve calculations using the micro-level model, the aggregate overdispersed Poisson chain-ladder and the aggregate lognormal chain-ladder model. The black solid line indicates the amount actually paid.*





Figure 14: *Out-of-sample exercise per 1-1-2005, injury claims. Results are for the total reserve (i.e. IBNR + RBNS reserve), based on 10,000 simulations. From left to right: reserve calculations using the micro-level model, the aggregate overdispersed Poisson chain-ladder and the aggregate lognormal chain-ladder model. The black solid line indicates actually observed quantities.*



forms the aggregate models under consideration and reveals a more realistic predictive distribution of the reserve.

7 Conclusion

Continuing the work by Arjas (1989), Norberg (1993) and Norberg (1999) this paper demonstrates the usefulness of micro-level stochastic loss reserving as a way to quantify the best estimate of the reserve and its uncertainty. Stochastic models for the occurrence time, the reporting delay and the development process (including intermediate

Table 5: *Out-of-sample prediction per 1-1-2005: numerical results for material damage and injury claims (in thousands), as obtained with Mack’s chain-ladder model, an overdispersed and lognormal stochastic chain-ladder model and the micro-level model. Real observed outcomes are also displayed.*

Model	Type	Expected value	Median	s.e.	VaR _{0,95}	VaR _{0,99}
Total Reserve						
Chain-ladder Mack	MD	2,865		349		
	BI	9,562		1,154		
Chain-ladder ODP	MD	2,803	2,785	361	3,426	3,846
	BI	9,611	9,533	1,214	11,700	13,230
Chain-ladder LogN	MD	7,073	3,660	1,549	21,824	90,712
	BI	10,246	10,030	1,931	13,651	17,590
Micro-model	MD	2,208	2,054	596	3,305	5,074
	BI	7,386	7,209	1,259	9,721	11,725
Observed	MD	> 1,861				
	BI	> 7,923				

payments and settlement) of a claim are fit to a data set with the development of individual claims. A micro-level approach allows much closer modeling of the claims process. The method is not restricted by limitations that exist when using aggregate data.

We perform an out-of-sample test with respect to a general liability insurance portfolio from a European insurance company. The paper shows that micro-level stochastic modeling is feasible for real life portfolios. We compare prediction results from the micro-level model with results obtained by analyzing a run-off triangle. Conclusion of the out-of-sample test is that – at least for the comparisons made here – traditional techniques tend to overestimate the real payments. Predictive distributions obtained with the micro-model reflect reality in a more realistic way: ‘regular’ outcomes are close to the median of the predictive distribution whereas pessimistic outcomes are in the very right tail.

The results obtained in this paper make it worthwhile to further investigate the use of this technique for loss reserving. Several directions for future research can be mentioned. One could try to refine the performance of the individual model with respect to very pessimistic scenarios by using a combination of e.g. a lognormal distribution for

losses below and a generalized Pareto distribution for losses above a certain threshold. Connected to this suggestion, we intend to explore the possibilities of the micro-level approach in a reinsurance context. Analyzing the performance of both the micro-level model and techniques for aggregate data on simulated data sets and new case studies will bring more insight in their performance. More careful modeling of inflation effects and taking the ‘time value of money’ into account will be important in future research. Studying the micro-level approach in light of the new solvency guidelines, is another path to be explored.

References

- E. Arjas. The claims reserving problem in non-life insurance: some structural ideas. *ASTIN Bulletin*, 19(2):139–152, 1989.
- R.L. Bornhuetter and R.E. Ferguson. The actuary and IBNR. *Proceedings CAS*, 59:181–195, 1972.
- R. Cook and J. Lawless. *The statistical analysis of recurrent events*. Springer New York, 2007.
- P.D. England and R.J. Verrall. Stochastic claims reserving in general insurance. *British Actuarial Journal*, 8:443–544, 2002.
- S. Haastруп and E. Arjas. Claims reserving in continuous time: a nonparametric Bayesian approach. *ASTIN Bulletin*, 26(2):139–164, 1996.
- R. Kaas, M. Goovaerts, J. Dhaene, and M. Denuit. *Modern actuarial risk theory: using R*. Kluwer, The Netherlands, 2008.
- A.F. Karr. *Point processes and their statistical inference*. Marcel Dekker INC, 1991.
- M. Kunkler. Modelling zeros in stochastic reserving models. *Insurance: Mathematics and Economics*, 34(1):23–35, 2004.
- C.R. Larsen. An individual claims reserving model. *ASTIN Bulletin*, 37(1):113–132, 2007.
- H. Liu and R. Verrall. Predictive distributions for reserves which separate true IBNR and IBNER claims. *ASTIN Bulletin*, 39(1):35–60, 2009.
- M. Merz and M. Wüthrich. Paid-incurred chain claims reserving method. *Insurance: Mathematics and Economics*, 3:568–579, 2010.
- R. Norberg. Prediction of outstanding liabilities in non-life insurance. *ASTIN Bulletin*, 23(1):95–115, 1993.
- R. Norberg. Prediction of outstanding liabilities ii. model extensions variations and extensions. *ASTIN Bulletin*, 29(1):5–25, 1999.

- B. Posthuma, E.A. Cator, W. Veerkamp, and E.W. van Zwet. Combined analysis of paid and incurred losses. *CAS E-Forum Fall 2008*, 2008.
- G. Quarg and T. Mack. Munich chain ladder: a reserving method that reduces the gap between IBNR projections based on paid losses and IBNR projections based on incurred losses. *Variance*, 2:266–299, 2008.
- A. Renshaw. Modelling the claims process in the presence of covariates. *ASTIN Bulletin*, 24:265–285, 1994.
- R. Schnieper. Separating true IBNR and IBNER claims. *ASTIN Bulletin*, 21(1):111–127, 1991.
- G. Taylor, G. McGuire, and J. Sullivan. Individual claim loss reserving conditioned by case estimates. *Annals of Actuarial Science*, 3(1-2):215–256, 2008.
- T. Verdonck, M. Van Wouwe, and J. Dhaene. A robustification of the chain-ladder method. *North American Actuarial Journal*, 13(2):280–298, 2009.
- T.S. Wright. A stochastic method for claims reserving in general insurance. *J. Inst. Actuar.*, 117(1):677–731, 1990.
- Mario V. Wüthrich and Michael Merz. *Stochastic claims reserving methods in insurance*. Wiley Finance, 2008.
- X. B. Zhao and X. Zhou. Applying copula models to individual claim loss reserving methods. *Insurance: Mathematics and Economics*, 46(2):290–299, 2010.
- X. B. Zhao, X. Zhou, and J. L. Wang. Semiparametric model for prediction of individual claim loss reserving. *Insurance: Mathematics and Economics*, 45(1):1–8, 2009.