

Triangle-free reserving

A non-traditional framework for estimating reserves and reserve uncertainty

Pietro Parodi*

[Presented to the Institute and Faculty of Actuaries in London, 4 February 2013]

Abstract

This paper argues that all reserving methods based on claims triangulations (the “triangle trick”), no matter how sophisticated the subsequent processing of the information contained in the triangle is, are inherently inadequate to accurately model the distribution of reserves, although they may be good enough to produce a point estimate of such reserves. The reason is that the triangle representation involves the compression (and ultimately the loss) of crucial information about the individual losses, which comes back to haunt us when we try to extract detailed information on the distribution of incurred but not reported (IBNR) and reported but not settled (RBNS) losses. This paper then argues that in order to avoid such loss of information it is necessary to adopt an approach which is similar to that used in pricing, where a separate frequency and severity model are developed and then combined by Monte Carlo simulation or other numerical techniques to produce the aggregate loss distribution.

A specific implementation of this approach is described, whose core feature is a method to produce a frequency model for the incurred but not reported claim count based on the empirical distribution of delays (delay = the time between loss date and reporting date), after adjustments to make up for the bias towards smaller delays. The method also produces a *kernel* severity model for the individual losses, from which the severity distribution of each year of occurrence can be derived. By combining the frequency and severity model in the usual way (e.g. through Monte Carlo simulation), an aggregate model for IBNR and UPR losses can be produced.

As for RBNS losses, we suggest using one of the many methods to analyse the distribution of IBNER (incurred but not enough reserved) factors to produce a possible distribution of outcomes.

A case study based on real-world liability claims is used to illustrate how the method works in practice.

Also, in a first step towards validating the method for calculating IBNR and comparing it with existing methods, a series of experiments with artificial data sets was undertaken, which show a drastic reduction in the prediction error of both the IBNR claim count and the IBNR total amount with respect to the standard chain ladder method. And what is perhaps most promising, the experiments show that the distribution of IBNR reserves is much closer (in terms of the Kolmogorov-Smirnov distance) to the “true” one than that based on Mack’s method in the way it is normally applied. The method promises therefore a more accurate assessment of the uncertainty around reserves.

*Correspondence to: Pietro Parodi, PhD, F.I.A., Willis Ltd, The Willis Building, 51 Lime Street, LONDON, EC3M 7DQ. E-mail: pietro.parodi@willis.com, pietro_parodi@yahoo.com

Keywords

Report delay; individual claims; frequency/severity modelling; IBNR; IBNER; RBNS; UPR; kernel severity distribution; chain ladder; reserving uncertainty; reserve distribution; granular reserving

1 Introduction

1.1 The “triangle trick”: glory and limitations

The idea behind most reserving techniques – that of compacting historical loss experience into a loss development triangle, and using these triangles to project the total claims amount to ultimate for every accident period (what one may call the “triangle trick”) – has proved to be a powerful tool for underwriters and actuaries, and has enabled them to have an at-a-glance view of the way that the total amount of claims develops over the years.

THE TRIANGLE TRICK – ILLUSTRATION
Development year

		0	1	2	3	4	5	6	7	8	9
Accident year	1	758,859	6,712,563	7,295,862	8,481,698	8,581,273	8,929,061	9,406,673	9,421,491	9,425,375	9,547,636
	2	588,009	1,786,021	2,187,149	2,365,737	2,474,465	2,842,739	2,842,739	2,882,701	3,398,944	
	3	514,089	1,532,487	2,331,175	8,377,877	8,954,659	9,117,566	9,138,301	9,147,275		
	4	419,422	2,882,030	4,009,785	4,413,923	4,468,089	4,616,335	4,823,964			
	5	261,482	2,089,735	3,050,709	3,684,369	4,130,221	5,036,548				
	6	893,053	2,121,944	4,368,448	4,546,849	6,942,262					
	7	481,366	954,766	2,026,609	2,481,851						
	8	696,678	1,505,950	2,283,808							
	9	4,336,497	5,355,547								
	10	433,625									

Figure 1. An example of a claims development triangle. E.g., there’s a total of £4,368,448 in claims that occurred in accident year 6 and were reported in development year 2 (i.e. by accident year 8).

This method has been used not only to obtain a point estimate of the projected loss amount for each policy year, but also to obtain a measure of the uncertainty around that point estimate, through what has been called *stochastic reserving*. Two popular methods are Mack’s method of calculating standard errors (Mack, 1993) and the bootstrap (England & Verrall, 2002). With some suitable distributional assumptions (e.g. that the distribution of outstanding reserves is a lognormal distribution), these methods can be used to estimate the full distribution of outstanding reserves.

The problem with triangulation methods: information compression

The triangle trick has been very successful and has the charm of simplicity. It is also, however, the “original sin” of reserving, and the most severe limitation of reserving techniques. By compacting the information into a small triangle, a large amount of information is lost and cannot be retrieved, no matter how clever one gets at manipulating the triangle itself.

It’s a bit like taking a picture with a high-resolution camera, then compressing it into a logo-like image of 16kB, and based on this compressed picture trying to reconstruct the content of the scene (see Fig. 2). Your 16kB compressed picture will give you some sense of the main features of the scene: in Fig. 2 (right) you can certainly see that this is a waterfront scene, with a number of houses, some small and some large, and with some cars parked in front of the water. However, if you want to recover the details – e.g., the name of the hotel on the waterfront, or the brands of the cars parked in front of it, or even the hour that the clock on the bell tower is showing – you have no alternative

A NON-ACTUARIAL EXAMPLE OF INFORMATION COMPRESSION




Figure 2. A waterfront scene from South Devon, at high (left) and low (right) resolution. Enlarge it to find out the name of the hotel, the brand of the cars on the waterfront, the time of the day as shown on the clock of the bell tower – then try to do the same with the image on the right!

but to go back to your original image and analyse that instead. This is exactly what this paper is proposing to do with your loss data set: give up on the heroic struggle of extracting information from your logo-like picture of past claims (the triangle), and go back to the original, information-rich image (the full loss data set), and start again from there with different methods.

Leaving metaphors aside, if one has for example 5,000 losses over a period of 10 years, and uses them to build a triangle such as that in Fig. 1, one is left with only $10 \times \frac{10+1}{2} = 55$ points to go by to project the losses to ultimate and to estimate the distribution of outstanding reserves. Is it any surprise that the infamous 99.5th percentile (Graham (2011)) is so difficult to capture with any accuracy¹?

The situation is not too dissimilar to the more traditional approach to pricing, where underwriters and actuaries rely solely on a burning cost exercise. In **burning cost analysis**, one looks at the total amount of losses incurred over a number of years, adjusts for claims inflation and possibly for IBNR (incurred but not reported) losses and other factors, then divides the exposure over the whole period to obtain the expected losses per unit of exposure. By analysing the volatility over the years, one can also estimate the standard deviation of the aggregate loss distribution and even venture to express an opinion on a few percentiles of the distribution. However, it is impossible to use burning cost to get an idea of the aggregate loss distribution with sufficient accuracy. For this reason, the modern approach to pricing, which is based on the collective risk model (see for example Klugman *et al.*, 2008), is to develop a frequency and a severity model *based on individual loss information*, and combine the two models (plus complications, such as a payment pattern model) with a Monte Carlo simulation or other numerical methods.

Pricing and reserving: two misaligned valuation frameworks

This leads us, by the way, to another dissatisfying feature of current triangle-based methods: whilst triangle-based reserving methods were “aligned” with the burning cost methodology for pricing, in the sense that the output of the claims triangle projection could be used as the input for an IBNR-adjusted burning cost exercise, the current state of the art is using stochastic methods based

¹ Admittedly, one can use a more granular loss triangle, e.g. intervals of three months instead of one year: however, that shifts the problem rather than solve it, and has some side effects, such as the fact that some periods may have no losses from one particular year and it may not be possible to use some of the projection techniques.

on a separate frequency and individual severity model for pricing and triangle-based methods based on the analysis of aggregates for reserving. This means that we have two very different valuation frameworks for what is essentially the same risk, but looked from two different points of view: prospectively (pricing) and retrospectively (reserving)!

1.2 A triangle-free reserving approach

This paper proposes a reserving method which, like in pricing, is based on building a separate model for the frequency and severity components of IBNR losses using individual loss information, and on using Monte Carlo simulation to produce the IBNR distribution. IBNER (incurred but not enough reserved/reported) is taken into account more traditionally by analysing the empirical distribution of the ratio between successive reserve estimates of open claims (and, where possible, studying the dependency on different variables such as development year, size, etc.). IBNER affects both the reserves for the reported but not settled (RBNS) claims *and* the estimated severity distribution of the IBNR claims. The method is not only “triangle free”, but also avoids aggregating individual losses in any other way when projecting the total claims to ultimate.

At the core of the method is (i) the use of the distribution of reporting delays (time between loss occurrence and loss reporting) to calculate the IBNR claim count, in the spirit of works such as that of Kaminsky (1987), Weissner (1978) and Guiahi (1986); (ii) the use of the distribution of the individual loss amounts, as in modern pricing techniques.

In simple terms, the method works as follows:

A. IBNR losses

- (1) Estimate the delay distribution, based on the empirical distribution of delays but *adjusting for the bias towards smaller delays* that is the inevitable consequence of observing delays through a limited time window.
- (2) Use the delay distribution in (1) to estimate the number of incurred but not reported (IBNR) claims based on the number of claims reported to date. Also determine the most suitable frequency model (e.g. Poisson, Negative Binomial) accordingly.
- (3) Model the severity distribution for the IBNR claims (this may be different for each loss year, or at least depend on claims inflation), also taking IBNER (incurred but not enough reserved/reported) claims into account
- (4) Combine the frequency and severity distributions via Monte Carlo simulation or another method (e.g. Fast Fourier Transform, Panjer recursion...) to produce an estimate of the aggregate distribution of IBNR losses

B. Future losses (UPR)

- (5) As a by-product of Steps (1) to (4) we have a frequency/severity model of future losses from those policies that have already been written but are not completely earned yet (what is normally referred to as UPR – unearned premium reserves). Producing a model for future losses is, by the way, a standard pricing exercise.

C. Reported but not settled (RBNS) losses

- (6) Separately analyse RBNS losses with one of the standard methods for IBNER analysis and produce a distribution of currently outstanding claims

TYPICAL OUTPUT OF A RESERVING EXERCISE WITH THE TRIANGLE-FREE APPROACH

Return Period (Years)	Percentile	Total Loss	Total Number
1 in 1.33	25%	14,205,625	230
1 in 2	50%	16,630,426	245
1 in 4	75%	19,631,593	260
1 in 5	80%	20,369,140	264
1 in 10	90%	22,627,203	274
1 in 20	95%	25,143,560	283
1 in 50	98.0%	29,441,708	292
1 in 100	99.0%	32,382,767	298
1 in 200	99.5%	35,057,817	304
1 in 500	99.8%	64,920,594	311
1 in 1000	99.9%	66,937,322	318
	<i>Mean</i>	17,363,917	245.1
	<i>Std Dev</i>	4,948,653	22.2

Figure 3. The typical output of a reserving exercise with the triangle-free approach is very much like the standard input of a gross loss model for pricing.

D. Overall reserve distribution

(7) Combine the distributions produce in A, B, C to produce the overall distribution of reserves.

The result of the reserving exercise looks therefore something like in Fig. 3

As can be seen from Fig. 3, the output is very much like that of a standard pricing exercise, giving the expected value of both the total losses and the total number of IBNR/UPR losses, the standard deviation of the same two variables, and various percentiles.

1.3 The comparison with the basic chain-ladder approach

The workings of this approach to a case study with real-world data are illustrated in Section 7. This is certainly important to understand how the method works in practice.

Anecdotal evidence based on a number of case studies is however not sufficient to validate the method – ultimately, a large-scale controlled experiment where different methods are allowed to compete is necessary.

In a first step towards model validation and comparison with competing methods, in this paper we have assessed the triangle-free method against the chain ladder by using controlled experiments where the underlying loss process was assumed to be known and the data was produced artificially based on that loss process. It is therefore possible to compare the estimated IBNR with the “true” IBNR.

Based on this, we have found that

- (i) The triangle-free approach normally leads to a more accurate estimate of the ultimate value of the claim count and of the total claim amount over a period. The difference in accuracy may be partly explained by the experimental set-up, but the main reason behind it lies probably in the fact that the triangle-free approach does not need to produce a different projection for every loss year, including loss years that are too immature and for which the ultimate outcome is

therefore very uncertain. However, this is not a very exciting result: why do we need an umpteenth method that claims to be more accurate at the price of extra complication?

TRIANGLE-FREE APPROACH v CHAIN LADDER
COMPARISON OF PREDICTION ACCURACY

	Chain ladder	Triangle-free (empirical)	Triangle-free (model)
Prediction error	7,392,262	4,423,237	3,842,561
Prediction error (%)	44.7%	26.8%	23.3%

Figure 4. Prediction error (in GBP) calculated as the mean squared error between the true value and the projected value for 100 different random data sets. The underlying model was in all cases a compound Poisson distribution ($\lambda = 100$) with a lognormal severity distribution ($\mu = 9.52$, $\sigma = 1.70$). The delay distribution was assumed to be exponential with an average delay $\tau = 3$ years. For simplicity's sake (although this is not a crucial assumption) the underlying loss model has been assumed to be inflation-free and the reserving process has been assumed to be IBNER-free.

(ii) What's more exciting, the experiments show that the reserving distribution produced by the chain ladder method is inadequate. It's not that it is more conservative or more optimistic – it's just not the right reserving distribution, and we should think twice about using it for deciding important matters, such as how much capital your company needs to satisfy Solvency II. The triangle-free approach produces a far more realistic result, or in other terms has more predictive power.

A larger-scale experiment on this has not been undertaken yet, but Fig. 5 shows an example in which the triangle-free approach and the chain ladder method have been used to calculate the IBNR distribution. Since we've already established that the triangle-free method performs better than the chain ladder on average, we have focused on a simulation where the projected ultimate was similar for the chain ladder and the triangle-free method

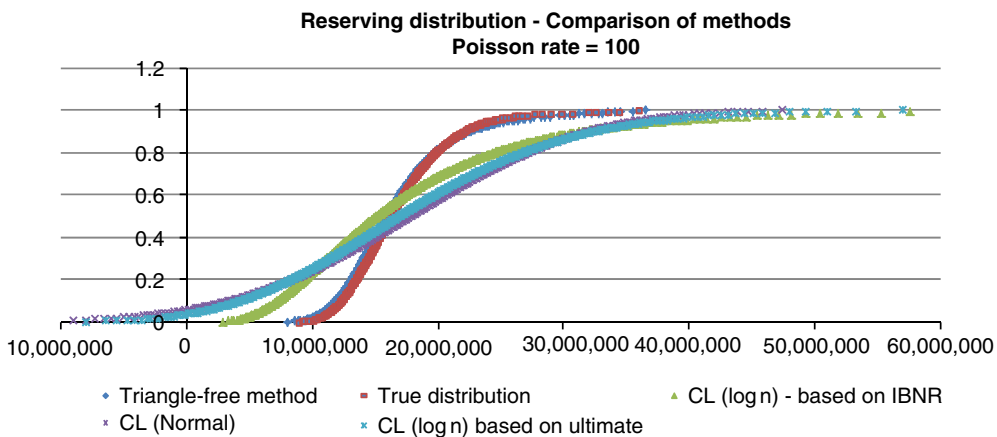


Figure 5. The cumulative distribution function (CDF) of the IBNR distribution as calculated by the triangle-free distribution method and the chain ladder method (with different assumptions on the resulting distribution), all of them compared with the true distribution, both graphically and (more formally) with the Kolmogorov-Smirnov statistic. The triangle-free distribution shows a far closer fit to the true distribution (the two distributions are almost indistinguishable on the chart). Note that *the role of parameter uncertainty has been taken into account* both in the triangle-free method and in the “true” distribution – as well as in the chain ladder method. The same assumptions for the tail were used in the two methods, to avoid introducing extraneous differences.

Apart from the question of accuracy and predictive power, the triangle-free approach has several advantages. Some of these are listed below:

- The approach can incorporate naturally any further information that we have on the risk: e.g. we can use a different model for losses above a given threshold, for which we only have market information and not client information;
- it doesn't break down as easily as the chain ladder when the number of claims per year becomes small and some of the accident years start with no claims at all;
- the calculation of the tail factor can be done more scientifically rather than in the heuristic fashion that is typical of triangle-based approaches;
- the methodology is completely aligned to the state-of-the-art methodology used in pricing, and therefore provides a consistent valuation framework of the risk under investigation, regardless of what use this valuation is put to.

1.4 This is a framework rather than a specific method

It should be stressed that although this paper often describes triangle-free reserving as “a method”, it is actually more of a **framework** for estimating reserves and their uncertainties. If we accept that this framework is useful, we still have to agree on a particular implementation of it. E.g., what distributions do we use to model the frequency and the severity of IBNR claims? How do we deal with IBNER? As in pricing, there is not a unique answer to these questions, but there is a framework that allows (and actually requires) to address them and, where necessary, make explicit assumptions.

If we embrace this different framework, this has consequences not only on the way we carry out the reserving analysis but also on the **data requirements**, since we need more detailed information on the losses and their characteristics: specifically, we need individual loss information including (for IBNR analysis) loss date and reporting date for each loss, and (for IBNER analysis) historical information on the reserve amount for each claim. The method can easily be used when some of this information is not available, but at the price of replacing empirical data with assumptions or at least approximations.

1.5 Limitations

The main limitation of this approach is that it is more complex than the approaches based on claims triangulations. This paper makes no apology for this – after all, the frequency/severity approach used in pricing to calculate the aggregate loss distribution describing a risk is more complex than a straightforward burning cost analysis. The extra granularity adds complexity but is necessary to a more in-depth understanding of the risk.

Another limitation is that the triangle-free approach doesn't provide an at-a-glance view of the inputs and outputs as the triangle-based methods do. No matter how sophisticated our tools are going to be in the future, triangles are here to stay at least as a representation tool and as a check and benchmark.

Finally, data requirements are more substantial in this approach, and good data capture and management is necessary (although nothing which is not already captured in insurance companies which follow good practice).

Despite all this, it should be noted that this approach is not a theoretical proposal but is currently routinely used in the author's company, often in a simplified fashion depending on the data available, and hybridised with triangulation methods as needed.

1.6 Further research

And then there's the question of the limitations of this research rather than the limitations of the approach itself. It is almost a cliché but further research is obviously needed.

What is needed the most is probably a large-scale experiment (we have limited ourselves to 100 different data sets here) with a number of different reserving methods (not only a simple chain ladder) competing with each other, possibly with data provided by an independent party or possibly produced with CAS's loss generator, in order to give more weight to the claim that the triangle-free approach actually gives better results (at least in estimating the reserving distribution if not the point estimate) than *any* method for which information has been compressed into a triangle.

Also, a systematic comparison of the approach presented here with other approaches that also advocate the use of more granular data for reserving should be undertaken. Since the first presentation of this approach at the GIRO conference in 2012 (Parodi, 2012b), several people have dug out references to works are similar in spirit to the present paper. Examples are the papers by Norberg (1993, 1999), by Taylor *et al.* (2008) and by Antonio & Plat (2012).

1.7 Outline

The core of the approach is the estimation of IBNR claim count based on the empirical distribution of delays, and therefore Section 2 ("Predicting IBNR claim counts") is the biggest chunk of the paper. Sections 3 to 6 illustrate how an aggregate loss model for the overall liabilities can be produced based on the triangle-free approach. Section 7 illustrates the application of triangle-free reserving to a real-world case study. Section 8 shows how the triangle-free approach can be validated and compared to the chain ladder methodology, and Section 9 explains what further research is needed.

2. Predicting the IBNR Claim Count

2.1 The problem

Suppose your company had a number of losses X_1, \dots, X_n over a period of k years. For each loss X_j we have the loss date LD_j and the reporting date RD_j , normally (but not always) expressed with the precision of one day. By aggregating the accident year we have n_1, \dots, n_k reported losses for years 1 to k . The objective is to calculate the projected number of losses N_1, \dots, N_k for each year:

The traditional approach to this problem is to aggregate losses in a frequency development triangle analogous to that in Fig. 1, which gives the number of losses for each accident year (rows) and for each development year (columns), then project to ultimate using a triangle projection technique, such as the chain ladder.

The main issue with this approach is that aggregating claim counts in a triangle throws away most of the information we have: e.g. if we have 1,000 losses over 10 years, we will compress the

information on the reporting delay (which is made of 1,000 different numbers) into a small triangle (which is made of 55 different numbers). Bizarrely, increasing the number of losses from, say, 1,000 to 10,000 does not increase at all the information that we can use to project to the final number of claims! This information compression is helpful for visualisation purposes and made sense in an age when calculations were performed by hand, and is surprisingly resilient in terms of finding an accurate projected number of claims for each year, but it ties our hands when trying to calculate the distribution of possible outcomes, making the uncertainty difficult to assess.

In Section 2.2 we propose an alternative approach. This approach is in the same spirit as the works by Kaminsky (1987), Weissner (1978) and Guiahi (1986).

2.2 A method for frequency projection

The method we propose is based on using the distribution $F(t)$ which gives the cumulative probability that a loss occurred at time 0 will be reported by time t .

2.2.1 The continuous case with known delay distribution F

Assume first that the delay distribution $F(t)$ is known. Assume t is the current date, and that we want to estimate the total number of claims that will ultimately be reported for the period $[0, t]$. Let $v(t)$ the frequency density of claims at time t – it can also be interpreted as a risk profile – where $v(t)$ is higher, the probability of having a claim is also higher. This may be due to seasonality (e.g. $v(t)$ may be higher in winter) or to non-seasonal systemic effects. The total (unknown) number of claims expected to occur in $[0, t]$ will be

$$E(\mu_t) = \int_0^t v(\tau) d\tau \tag{1}$$

The expected number of claims to be reported by time t is

$$E(r_t) = \int_0^t v(\tau) F(t - \tau) d\tau \tag{2}$$

Normally r_t , the number of reported claims, will be known, whereas μ_t is to be estimated: the estimated value of μ_t , $\hat{\mu}_t$, is given by

$$\hat{\mu}_t = \frac{\int_0^t v(\tau) d\tau}{\int_0^t v(\tau) F(t - \tau) d\tau} r_t \tag{3}$$

If the probability of having a claim can be assumed to be uniform, this simplifies to

$$\hat{\mu}_t = \frac{t}{\int_0^t F(t - \tau) d\tau} r_t \tag{4}$$

If we wish to project to ultimate the frequency for a period $(0, t')$ which is completely in the past ($t' < t$) then it is sufficient to replace the upper limit of the integral with t' (note that the point t inside the integral from which the CDF is calculated remains unaltered):

$$\hat{\mu}_{t'} = \frac{t'}{\int_0^{t'} F(t - \tau) d\tau} r_{t'} \tag{5}$$

Formulae (4) and (5) have a simple geometric interpretation, which is illustrated in Fig. 6.

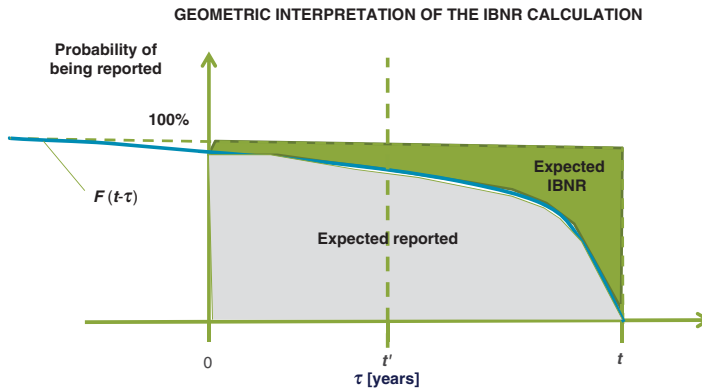


Figure 6. The denominator of the fraction in Equation (4) can be pictured as the gray area below the CDF and to the right of the y-axis (“Expected reported”) and the numerator as the rectangle with sides t and 1 (“Expected IBNR”+“Expected reported”). The factor to ultimate for the whole period $[0,t]$ is therefore the ratio of the two areas. Analogously, the denominator of the fraction in Equation (6) can be pictured as the area below the CDF, between the y-axis and the axis $\tau = t'$, and the numerator as the rectangle of sides t' and 1. The factor to ultimate for the period $[0,t']$ is again the ratio of the two areas.

2.2.2 The discrete case

In practice, of course, the number of reported claims in a period $[0,t]$ will only be known for discrete values of t , typically expressed in days: The estimated delay distribution $\widehat{F}(t)$ (more on it later) will also be a distribution known at discrete values. It is perhaps helpful, therefore, to give a discrete version of the formulae above.

In particular, if the time t is expressed in days, Equation (4) becomes

$$\widehat{\mu}_t = \frac{t}{\sum_{k=0}^t \widehat{F}(t-k)} r_t \quad (6)$$

Now if we define

$$\# \text{ of earned days in the period } [0, t] = \sum_{k=0}^t \widehat{F}(T-k) \quad (7)$$

We can rewrite the equation above in a much more intuitive way:

$$\widehat{\mu}_t = \frac{\# \text{ of days in the period } [0, t]}{\# \text{ of earned days in the period } [0, t]} r_t \quad (8)$$

Analogously, Equation (5) becomes, in discrete time:

$$\widehat{\mu}_{t'} = \frac{\# \text{ of days in the period } [0, t']}{\# \text{ of days in the period } [0, t'] \text{ earned by time } t} r_{t'} \quad (9)$$

Fig. 7 shows an example of calculation based on formulae (8) and (9) above, in which the number of earned days is calculated either for the whole period lumped together and for different accident years separately.

PROJECTION TO ULTIMATE -- ALL YEARS TOGETHER

Period	Days elapsed	Earned days	Factor to ultimate	Exposure	Latest reported	Ultimate losses	Standard error
2001-10	3,652.00	1,967.62	1.86	10,000	352	653.33	31.85

PROJECTION TO ULTIMATE -- YEAR BY YEAR

Year	Days elapsed	Earned days	Factor to ultimate	Exposure	Latest reported	Ultimate losses	Standard error
2001	365	347.94	1.05	1,000	56	58.75	3.16
2002	365	302.13	1.21	1,000	59	71.28	6.06
2003	365	275.29	1.33	1,000	43	57.01	7.24
2004	366	254.25	1.44	1,000	41	59.02	8.07
2005	365	220.78	1.65	1,000	42	69.44	9.18
2006	365	190.62	1.91	1,000	41	78.51	10.10
2007	365	158.90	2.30	1,000	23	52.83	10.98
2008	366	119.16	3.07	1,000	30	92.15	12.00
2009	365	71.35	5.12	1,000	14	71.62	13.10
2010	365	27.19	13.42	1,000	3	40.27	14.05
					<i>Mean</i>	65.09	
					<i>Variance-to-mean ratio</i>	3.28	

Figure 7. An example of calculation for the factors to ultimate, both when all years are lumped together (top) and when the calculations are carried out separately for each individual accident year (bottom). The exposure has been assumed constant in this specific case, although this is not a necessary assumption, as the case study in Section 7 shows.

The method that lumps all years together leads to more accurate results, since the projection of isolated years is affected by large errors for the more recent years. However, in many circumstances it will be useful to identify the IBNR in the claim count for each individual year. This will also allow us to estimate the year-on-year volatility.

The problem with calculating the IBNR separately for each year is of course that the more recent years may have some numerical instability. In the most extreme case, there might be no claims at all in the last year(s), and the projection would always yield zero.

This problem is common to triangulation methods such as the chain ladder. To get around this problem one can use a credibility approach, much in the same way as it is done with Bornhuetter-Ferguson (which uses a prior estimate of the loss ratio) or by Cape-Cod (which uses an average of the previous years). Section 7.1.3 shows in detail the workings of a possible credibility approach for triangle-free reserving in a practical case.

2.2.3 Calculating the uncertainty

The standard error on the projected number of claims can be easily calculated if we can make a reasonable assumption of what the underlying claim count process is – e.g. Poisson, negative binomial, etc.

If underlying claim count process is Poisson with rate μ_t , it can be easily proved that the number of IBNR claims also follows a Poisson distribution with rate $\mu_t - E(r_t)$. Since the variance of a Poisson process is equal to its mean, the variance of the process is also $\mu_t - E(r_t)$, which can be estimated from the data as $\hat{\mu}_t - r_t$. The square root of this can be used as an estimate of the standard error by which we know the ultimate claim count for the period $[0, t]$:

$$\text{s.e. (IBNR)} = \sqrt{\hat{\mu}_t - r_t} \quad (10)$$

(Note that we are not speaking here of the standard error on the estimated Poisson rate, which would be much smaller, but of the standard error of the projected claim count. This is why we do not need to divide by the number of observations.)

In case of an overdispersed Poisson process, e.g. a negative binomial process, this method can be easily modified by setting the standard error equal to

$$\text{s.e. (IBNR)} = \sqrt{(1 + r)(\hat{\mu}_t - r_t)} \quad (11)$$

where $1 + r$ is the variance-to-mean ratio.

2.2.4 A special case: exponential delays

It is often the case that the delay distribution can be well approximated by something simple,

e.g. an exponential distribution with average delay τ :

$$F(t) = 1 - \exp(-t/\tau) \quad (12)$$

In this case Equation (4) simplifies to:

$$\hat{\mu}_T(T) = \frac{T}{T - \tau(1 - \exp(-\frac{T}{\tau}))} r_T \quad (13)$$

and Equation (5) simplifies to:

$$\hat{\mu}_{T'}(T') = \frac{T'}{T' - \tau \exp(-\frac{T'}{\tau})(\exp(\frac{T'}{\tau}) - 1)} r_{T'} \quad (14)$$

(note that $T < T'$ in the equation above).

Although the exponential approximation may be viewed as simplistic, it becomes handy when the delay data is very scarce. In this case it might make sense – rather than relying on a patchy delay distribution – to assume an exponential distribution and just calculate the observed mean. Also, it provides a first rough approximation of the development factors without embarking on any complex calculation: e.g., assume we have a general liability risk with an average reporting delay of $\tau = 3$ years. We look at the reported claims as at 31/8/2012, and we want to project to ultimate the experience of the policy year started on 1/1/2010 and finished on 31/12/2010. The number of reported claims for the 2010 policy year is 10 claims. If you assume that the distribution of delays is exponential with an average of three years and we arbitrarily use the start date of the policy (1/1/2010) as $t = 0$, we can plug in the numbers $T' = 1$ y, $T = 2.67$ y, $\tau = 3$ y, $r_T = 10$ in Equation (14), obtaining

$$\hat{\mu}_{2.67}(1) = \frac{1}{1 - 3 \exp(-\frac{2.67}{3})(\exp(\frac{1}{3}) - 1)} 10 = 19.5$$

2.3 Estimating the reporting delay distribution

In Section 2.2 we have assumed that we knew the reporting delay distribution. In practice, all we typically have is a collection of claims for each of which the delay can be calculated as the difference between the reporting date and the occurrence date. We can then sort these delays from the smallest to the largest obtaining the empirical delay distribution.

One may be tempted to use this empirical distribution as a proxy for the true underlying report delay distribution. However, such empirical distribution will inevitably be biased towards shorter delays, as no reporting delay bigger than the observation period can of course be observed, and delays of length near the length of the observation period will also be very rare.

In this section we show how to correct the empirical distribution for this bias.

2.3.1 Calculating the bias-corrected empirical distribution

As we said above, we do not observe the true PDF $f(t)$ of the delays but an empirical, biased version of it. More accurately, we observe an empirical version of the distribution $f_a(t)$, which gives the probability that a delay of length t is observed in the observation window $[0, a]$. In order to reconstruct the unbiased $f(t)$, we need to be able to determine the relationship between $f_a(t)$ and $f(t)$.

First of all, let's introduce some notation:

- Let T be a random variable representing the delay between occurrence and reporting
- Let T_0 be the time at which the loss occurs

As a consequence, $T + T_0$ represents the time at which the loss is reported.

$f_a(t)$ is by definition the probability that the delay T is equal to t given that the time $T + T_0$ at which the loss is reported falls within $[0, a]$. Using Bayes' theorem,

$$\begin{aligned}
 f_a(t) &= \frac{\Pr(T = t | T + T_0 \leq a)}{\Pr(T + T_0 \leq a | T = t) \Pr(T = t)} \\
 &= \frac{\Pr(T = t)}{\Pr(T + T_0 \leq a)} \\
 &= \frac{\Pr(T_0 \leq a - t) \Pr(T = t)}{\Pr(T + T_0 \leq a)}
 \end{aligned}
 \tag{15}$$

The uniform case. Assuming for the moment that the loss is equally likely to happen at any time within the interval² $[0, a]$, we can write:

$$\Pr(T_0 \leq a - t) = 1 - \frac{t}{a}
 \tag{16}$$

and (by definition)

$$\Pr(T = t) = f(t)
 \tag{17}$$

² This is not a crucial hypothesis and can be easily modified if we have more information on, e.g., the different exposure or risk profile of different years.

As usual, calculating the denominator of the Bayes rule (the probability $\Pr(T + T_0 \leq a)$) is not crucial as this is only a normalising factor³.

We can therefore write

$$f_a(t) = \begin{cases} \frac{(1-\frac{t}{a})f(t)}{G(a)} & \text{for } t < a \\ 0 & \text{elsewhere} \end{cases} \quad (18)$$

Which we can invert to derive an expression for $f(t)$.

$$f(t) = \begin{cases} \frac{G(a)}{(1-\frac{t}{a})} f_a(t) & \text{for } t < a \\ \text{undefined} & \text{elsewhere} \end{cases} \quad (19)$$

There is no strictly empirical way of deriving $f(t)$ for $t > a$. However, if we make some assumption on the way the empirical distribution tails off based on the behaviour for $t < a$, this derivation becomes possible. This is explored further in Section 2.5.

The non-uniform case. As mentioned above, Equation (16) and all the equations derived from it only hold if it is equally likely that a loss comes from any point inside $[0, a]$. In actual fact, this will almost never be the case, because of changes in the exposure and in the risk profile. All of this can be captured in the variable $v(t)$ introduced in Section 2.2.1, which we then called the “frequency density”.

Equation (16) can now be rewritten as

$$\Pr(T_0 \leq a - t) = \frac{\int_0^{a-t} v(t_0) dt_0}{\int_0^a v(t_0) dt_0} = \int_0^{a-t} \frac{v(t_0)}{\int_0^a v(t') dt'} dt_0 = \int_0^{a-t} \tilde{v}(t_0) dt_0 \quad (20)$$

where $\tilde{v}(t)$ is the normalised version of $v(t)$. Equations (18) and (19) can now be amended to take (20) into account.

As a simple but very typical example, consider the case in which the risk profile from one year to the other only changes because of the different exposure (e.g. number of vehicle years for motor insurance, or wageroll/number of employees for employers’ liability). Following the standard (and rough) practice, we’ll assume that the exposure undergoes step changes at the beginning of every accident year.

Under these assumptions, it is easy to see that

$$\Pr(T_0 \leq a - t) = \int_0^{a-t} \tilde{v}(t_0) dt_0 = \frac{\sum_{i=0}^{\lfloor a-t \rfloor - 1} \varepsilon_i + \varepsilon_{\lfloor a-t \rfloor} \times (a - t - \lfloor a-t \rfloor)}{\sum_{i=0}^{\lfloor a \rfloor - 1} \varepsilon_i + \varepsilon_{\lfloor a \rfloor} \times (a - \lfloor a \rfloor)} \quad (21)$$

where ε_i is the (on-levelled) exposure in year i and $\lfloor x \rfloor$ is the integer part of x .

³ For completeness, however, the calculation goes as follow: define

$$G(a) = \Pr(T + T_0 \leq a)$$

where $G(a)$ is the cumulative distribution of $T + T_0$, which can be obtained as the convolution of the distributions of T and T_0 , which in turn can be obtained using Laplace transforms:

$$g(t) = \mathcal{L}^{-1} \left(\mathcal{L}(f(t)) \mathcal{L} \left(\frac{\theta(t) - \theta(t-a)}{a} \right) \right)$$

In the expression above, \mathcal{L} is the Laplace transform, and $\frac{\theta(t) - \theta(t-a)}{a}$ is the uniform distribution between 0 and a .

2.3.2 Geometric interpretation

The calculation in Section 3.1 has a simple interpretation in terms of stochastic geometry, as illustrated in Fig. 8. Imagine you have a machine that generates random segments with of length t , where the length of t follows a given distribution $f(t)$, and the segment starts at a point $t_0 \in [0, a)$, with an equal probability of starting at *any* point inside that interval.

The probability of a segment of length t being fully contained inside the slot $[0, a)$ is then

$$\Pr(\text{segment is inside } [0, a)) = \frac{a-t}{a} = 1 - \frac{t}{a} \tag{22}$$

and the probability that a segment of random length t is fully contained in the slot is $f_a(t)$ as calculated in Equation (20).




Figure 8. A geometric illustration of the probability of observing a delay of length t . Assume we are generating segments of random length t (according to the probability $f(t)$) with initial point between 0 and a and “letting them fall” – what is the likelihood that they’ll make it through the slot? If the length t is given, this probability is simply $1-t/a$ (the length of the segment as a proportion of the length of the slot), which decreases as the length t of the segment increases – and becomes zero when the length of the segment exceeds the length of the slot, because it is impossible for the segment to go through a slot which is smaller than the segment itself. The probability of a segment of random length making it through the slot is instead $f_a(t)$.

2.4 The case of exponential delays

We have seen in Section 2.2.4 that the expression for the factor to ultimate somewhat simplifies when the underlying distribution $f(t)$ is an exponential distribution with average delay τ :

$$f(t) = \frac{1}{\tau} \exp(-t/\tau) \tag{22-i}$$

Using this approximation, the observed distribution in a time window $[0, a)$ becomes

$$f_a(t) = \frac{\frac{1}{\tau} (1 - \frac{t}{a}) \exp(-t/\tau)}{1 - \frac{\tau}{a} (1 - \exp(-\frac{a}{\tau}))} \tag{22-ii}$$

and zero elsewhere.

As a consequence, the observed average delay is

$$\begin{aligned} \tau_{obs} &= \mathbb{E}(T) \\ &= \frac{1}{1 - \frac{\tau}{a} (1 - \exp(-\frac{a}{\tau}))} \int_0^a \frac{t}{\tau} (1 - \frac{t}{a}) \exp(-\frac{t}{\tau}) dt \\ &= \tau \left(1 + \frac{e^{-\frac{a}{\tau}} - \frac{\tau}{a} (1 - e^{-\frac{a}{\tau}})}{1 - \frac{\tau}{a} (1 - e^{-\frac{a}{\tau}})} \right) \end{aligned} \tag{22-iii}$$

The true underlying average delay can be estimated from Equation (22-iii) by finding τ numerically.

2.5 Calculating the tail factor

Equation (21) shows how to calculate the delay distribution $f(t)$ from the distribution observed through the slot $[0, a]$, $f_a(t)$. However, this prescription works by definition only for delays that are smaller than a . In order to estimate the “tail factor” – *i.e.*, the additional factor by which we have to multiply our projected ultimates in order to take into account delays larger than a – we need to have a prior model for $f(t)$.

If $f(t)$ can be assumed to be an exponential distribution with average mean delay equal to τ , then of course the PDF above a can be modelled as $f(t) = \frac{1}{\tau} \exp(-t/\tau)$. In general, even if $f(t)$ is not an exponential distribution, we can make a broad approximation and do as follows:

- (i) Calculate the average observed delay τ_{obs} based on our delay data set
- (ii) Assume that the delay distribution is exponential and invert Equation (22-iii) to calculate the implied average delay τ
- (iii) Approximate the delay distribution as

$$f(t) = \begin{cases} \frac{G(a)}{(1 - \frac{t}{a})} f_a(t) & \text{for } t < a \\ \frac{1}{\tau} \exp(-t/\tau) & \text{for } t \geq a \end{cases} \quad (22\text{-iv})$$

Note that in practice this corresponds to a **tail factor** φ_{tail} given by

$$\varphi_{\text{tail}} = \frac{1}{1 - \exp(-a/\tau)} \quad (22\text{-v})$$

This is of course a rough approximation, but it gives reasonable results in most cases, unless the tail of the delay distribution has a behaviour which is distinctly not exponential. In any case, there’s no reason to get married to this first approximation and if the data are there one can always make a more thorough study of the tail of the distribution by using extreme value theory and finding the parameters of a Generalised Pareto distribution for the tail.

The main risk, as is often the case, is to become a bit too clever and quibble about the exact behaviour of the tail of the delay distribution when there’s actually not enough data to support anything more than a simple, one-parameter distribution such as an exponential.

2.6 Other considerations

Non-constant delay distribution. The results above depend on the assumption that the delay distribution $f(t)$ remains constant over time. The method, however, can be easily modified to reflect changes of the delay distribution, *e.g.* by using a different delay distribution depending on the loss occurrence date. The exact calculations will depend on whether one can consider that the delay depends only on the loss occurrence date (*i.e.*, all losses occurred around a given time have the same average delay) or whether the remaining time to reporting is “accelerated” for all claims from a certain date forwards for all accident years. The key thing is to make sure that one has enough data to support any time-dependency of the delay distribution in order not to introduce spurious accuracy.

Note that the fact that the delay distribution is constant over time is itself a hypothesis that can be easily tested statistically if enough data is available, *e.g.* by comparing the delay distribution

(after the bias correction) before and after a certain time t^* for different values of t^* with the two-sample Kolmogorov-Smirnov statistic.

Size-dependent delay distribution. In the same way, we can test whether the hypothesis that the delay distribution is the same for claims of all sizes or if, on the contrary, we need a different delay distribution for, e.g., attritional and large losses. We can then apply different IBNR corrections for claims of different size class.

One simple approach in the case that all we need is a distinction between attritional and large losses is as follows:

- i. Revalue all claims with some measure of claims inflation, to bring them all to current terms
- ii. Identify the threshold for which the difference in the average mean delay between attritional and large losses is highest
- iii. Produce a separate delay distribution for attritional and large losses
- iv. Estimate the number of IBNR attritional losses and IBNR large losses
- v. Produce a separate severity model for attritional and large losses
- vi. Produce two different aggregate loss models for attritional and large losses, based on Steps iii to v
- vii. Combine the two models with a Monte Carlo simulation

The details of how to produce the severity model and how to set up a Monte Carlo simulation will be addressed later on in the paper – at this stage, we just think that it's important to show that a possible dependence of the delay on the size of the loss is not a showstopper. However, we caution against using too many fine distinctions in the case one can't really support them with sound experimental results.

Relationship with claims triangles. As we mentioned before, accident years can be developed individually to ultimate as well as all together. Considering IBNR as a whole leads to a more accurate result, since the development of the more recent years individually to ultimate introduces a large error.

Interestingly, the method can also be used as a complement to the usual claims triangle projection, when the last diagonal is not complete, as a means for grossing up the last diagonal. However, this is perhaps not worth the effort!

3 An Aggregate Loss Model for IBNR Losses

In order to estimate the IBNR distribution, we need to have not only a point estimate of the expected number of IBNR losses but a full frequency model (e.g. Poisson, Negative Binomial) with an estimate of its parameters and of the uncertainty (standard error) on the parameters themselves, as well as a severity model for IBNR losses (e.g. lognormal, GPD...). Finally, we'll need a protocol to combine the two models in order to produce an aggregate loss model for IBNR losses.

3.1 A frequency model for IBNR losses

In this paper we will assume that the underlying frequency distribution is either a Poisson distribution or a negative binomial distribution.

The Poisson distribution has its theoretical justification as the natural model for rare independent events arising from a stationary process. In the presence of systemic volatility, clustering of claims or simply parameter uncertainty, the dispersion of the Poisson distribution (for which the variance-to-mean ratio is 1) may not be sufficient to reproduce reality. The negative binomial distribution is simply one way to account for this additional dispersion, and there is nothing fundamental about this distribution – it’s just that it can be easily represented as a Poisson distribution whose Poisson rate is itself a random variable drawn from a Gamma distribution. It is very seldom the case that one has such a wealth of experimental data that the use of the negative binomial can be challenged!

The Poisson case

Under the Poisson assumption, the IBNR claim count can be represented as a Poisson distribution for which:

- The **Poisson rate** is the point estimate of the IBNR claim count, as calculated by Equation (5) (or Equation (6) for the discrete case);
- The **standard error on the Poisson rate** can be approximated as the square root of the Poisson rate itself (as this is the best estimate of how much variation we expect from one simulation to the other).

Note that although the process is assumed to be Poisson, when we use this frequency distribution later on for Monte Carlo simulation, we will need to take the parameter uncertainty into account and this will be equivalent to having an overdispersed Poisson distribution: specifically, we will use the negative binomial approximation by assuming a variance-to-mean ratio of 2 (as in this case the process variance is equal to the parameter uncertainty).

The negative binomial case

Under the negative binomial assumption, the IBNR claim count can be represented as a NB distribution for which:

- The **rate** is the point estimate of the IBNR claim count, as calculated by Equation (4) (or Equation (6) for the discrete case);
- The **variance-to-mean ratio** can be estimated by calculating the projected number of claims for different years, as in the “Ultimate losses” column of Fig. 7, calculating the variance based on those ultimate losses⁴, and dividing it by the average of the ultimate losses;
- The **effective variance-to-mean ratio** can then be derived from the variance-to-mean ratio by adding an allowance for the parameter uncertainty as in the Poisson case. As a simple rule of thumb (using a rough application of the central limit theorem), the effective variance-to-mean ratio can be written as $(V/M)_{\text{eff}} = (V/M) + 1/n$ where n is the number of years over which the variance was calculated. (It can be seen that in the case of Poisson this gives back the “variance-to-mean ratio=2” rule stated for the Poisson case.)

3.2 A severity model for IBNR losses

We now need to produce a severity model for IBNR claims, with the ultimate goal of producing an aggregate loss distribution for IBNR claims via Monte Carlo simulation. Since one of the things that

⁴ This, at least, is true in the simple case where the exposure is uniform over the years. If this is not the case, the formula for the calculation of the variance gets a bit more complicated. We won’t get into the details here.

we'll need to simulate will be the actual occurrence date for the claims, the distribution from which we sample the random loss amounts will in general be different depending on the occurrence date. This dependency is analysed in Section 4.1.

Another thing we need to take into account is IBNER, which is dealt with in Section 4.2.

Once we take these two elements into account, we'll have a way of sampling the severity of claims. The algorithm is illustrated in Section 4.3.

3.2.1 Dependency of the loss distribution of IBNR claims on the occurrence date

As in experience rating, the main difficulty in obtaining a severity model is that IBNR claims will come from different years of experience, and each accident year may well have a different severity distribution, depending on the specific business mix that was written in that particular year⁵.

There are at least two different elements that drive the dependency of the severity distribution *on the accident year*.

- The proportion of different types of claims may change, as a result of changing business conditions or simply the changing nature of risk due to environmental/technological changes (e.g., the relative number of “slips and trips” public liability claims might be increasing due to a more litigious environment, or decreasing due to better risk control mechanisms such as anti-slippery carpets)
- Claims inflation – losses occurring now will generally have a different cost (usually higher) than losses occurred a number of years ago. Furthermore, the severity distribution of losses that have not been reported yet might also be conditional on the reporting delay: e.g., losses that have occurred five years ago and were reported then might not follow the same distribution as losses that have occurred five years ago and are reported today. (It is not obvious whether they will be bigger or smaller – one can argue that a menial loss will either be reported immediately or not be reported at all, but might also argue that if a loss is really big it will be more evident and will certainly be reported shortly. Only empirical evidence can help us make an informed stance.)

In the following, we will focus on claims inflation and ignore changes coming from changes in business mix or in the type of claims: if this is not the case, we always have the option of splitting the loss experience into claims of different type and analyse them independently, assuming that the severity distribution within each segment remains the same except for claims inflation.

So how do we take claims inflation into account? In experience rating what one normally does is to analyse the past experience over a period and **revalue to current terms (or rather, to the mid-term of the policy one wishes to rate) all past claims by a certain index** (e.g. a wage inflation index, perhaps amended to include extra inflation) or a flat inflation rate which is assumed to incorporate all necessary elements of inflation.

In determining the severity model for IBNR losses, we need to go through a similar process, with the complication explained above and repeated here more formally: we are interested in the probability

⁵ Actually the subdivision in accident years (or accident periods) is rather artificial and simply corresponds to the frequency by which the information on the business relevant to the claims experience is updated, and we can say that the severity distribution of losses changes continuously through time.

$F_X^{t_0, \delta}(x)$ that a loss X be $\leq x$ conditional to X having occurred at time t_0 (assume the time is given in years) and being reported after δ years.

$$F_X^{t_0, \delta}(x) = \Pr(X \leq x, X \text{ occurred at } t_0, \text{ reported at } t_0 + \delta) \quad (22\text{-vi})$$

There are therefore two (possibly conflicting) effects: that of claims inflation and that of being reported after a number of years.

Ultimately, only empirical evidence helps tell us what this distribution looks like. One approximation that we can often use is to assume that $F_X^{t_0, \delta}(x)$ is simply a scaled version of the kernel distribution $F(x)$, where the scaling depends on both the year of occurrence and the reporting delay:

$$F_X^{t_0, \delta}(x) = F_X\left(\frac{(1+r)^{t-t_0}x}{(1+s)^\delta}\right) \quad (22\text{-vii})$$

where t is the current time, r is the claims inflation (assumed constant here for simplicity), and s is an inflation that captures the effect of the reporting delay: e.g., if claims tend to become larger when δ increases, then $s > 0$. Vice versa, it is < 0 .

This is just a rough working assumption, but it is often a good approximation of reality, as it is shown for example in Fig. 9, which shows the empirical distribution of employers' liability losses for a large UK company for the same accident year (2006) and different reporting years (2006, 2007, 2008 and later). Even without any delay-related inflation (i.e., assuming $s = 0$ in Equation (22-vii)), we see that the empirical distributions are quite close, e.g. there is no obvious effect on the severity distribution due to reporting delay.

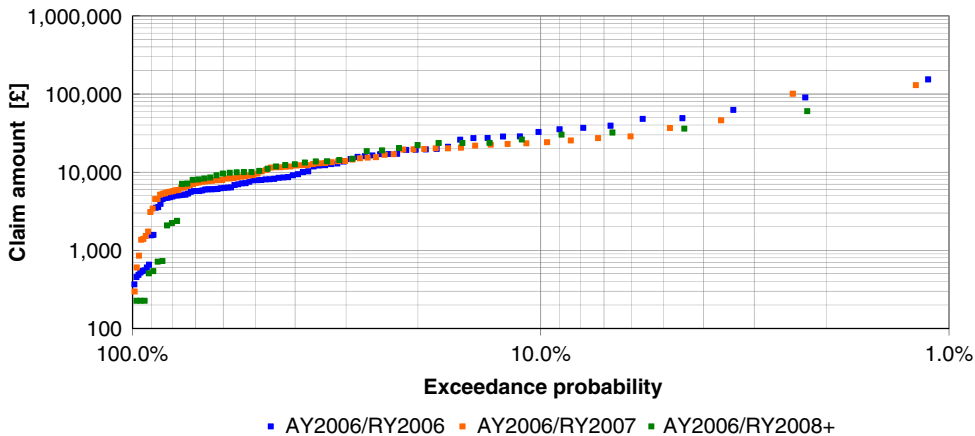


Figure 9. The graph shows the empirical distribution of employer's liability losses for a large UK company in the food services industry, after some scaling to make the losses unrecognisable. The empirical distributions for 2006 losses that have been reported in 2006, 2007 and 2008 or later have been compared, without adding any delay-related inflation. The empirical distributions are visually very similar, and this is confirmed by the Kolmogorov-Smirnov statistical test: the normalised KS distance between the differen pairs is $KS(RY = 2006, RY = 2007) = 1.22$ ($p = 10\%$), $KS(RY = 2006, RY = 2008+) = 1.36$ ($p = 5\%$), $KS(RY = 2007, RY = 2008+) = 0.74$ ($p \gg 10\%$).

A rough approximation may therefore be to simply ignore delay inflation and set r equal to a standard claims inflation rate for the line of business under consideration: e.g. $r = 5\%$ $s = 0$.

The situation needs of course to be assessed on a case-by-case basis and ultimately only empirical evidence can lead us one way or another – the important thing to remember as usual is that models that are more sophisticated than Equation (22-vii) need to be supported by a proportionate wealth of empirical evidence.

3.2.2 Determining IBNER in order to derive the severity distribution

Many of the losses that make up the loss data set and that we need to use to derive the severity distribution are still outstanding (not paid up) and as such are still liable to change in value. This is a type of *data uncertainty* which we denote generically as “IBNER” (incurred but not enough reserved/reported). IBNER has an impact on both the selection and the parameterisation of the severity distribution.

The idea of IBNER analysis is to identify systematic underestimation or overestimation of claims and to put it right by adding the relevant (possibly negative) amount to all open claims.

For large losses (and especially in reinsurance), this is normally done by using an individual claims development table which is made of triplets of rows like that in Fig. 10.

The calculation of the IBNER factors between two successive years of development of the claims is usually performed by an average (weighted or unweighted) over all open claims (settled claims will of course have no IBNER, unless they are re-opened) as in Fig. 11.

If the data is abundant, it is also possible to estimate the volatility of the IBNER factors around the mean value and actually the whole distribution of the IBNER factors, as in Fig. 12.

A more refined analysis

The simple method described above focuses on the dependence of the IBNER factor from the year of development, be it from the loss year or from the reporting year. However, it is obvious that the IBNER factor will depend in general on many factors, such as:

- development year d (the younger the claim, the more scope there is for it to deviate from the incurred estimate), as already mentioned;

CALCULATION OF IBNER FACTORS									
	2000	2001	2002	2003	2004	2005	2006	2007	2008
<i>Paid</i>	19,792	363,306	487,648	1,735,328	1,922,504	1,922,504	1,922,504	1,922,504	1,922,504
<i>O/S</i>	967,500	877,200	753,360	147,060	0	0	0	0	0
<i>Incurred</i>	987,292	1,240,506	1,241,008	1,882,388	1,922,504	1,922,504	1,922,504	1,922,504	1,922,504
<i>O/S ratio</i>	98.0%	70.7%	60.7%	7.8%	0.0%	0.0%	0.0%	0.0%	0.0%
<i>IBNER factor</i>		1.256	1.00	1.517	1.021	1.000	1.000	1.000	1.000

Figure 10. The history of a claim occurred in 2000 and reported in the same year. For each year of development, the paid, outstanding and incurred (estimated) amounts are given. The outstanding ratio (outstanding amount divided by incurred amount) is also given for each year. The last row shows the IBNER factors for this claim, calculated as the ratio between the incurred amounts of two successive years. The claim shown here is a real claim after scaling by a random amount and the removal of any identifiable information (Figure borrowed from Parodi (2012a)).

EXAMPLE IBNER CALCULATION

			Year 1 → 2		Year 2 → 3		Year 3 → 4			
IBNER factors			1.219		1.019		0.839			
Year			Year 1		Year 2		Year 3		Year 4	
Date of Loss	Year	Loss reported as non-zero	Estimate	O/S %	Estimate	O/S %	Estimate	O/S %	Estimate	O/S %
03/07/2004	2004	2004	2,486,700	100.00%	1,535,000	100.00%	472,780	61.00%	438,321	57.93%
04/07/2004	2004	2004	116,660	100.00%	85,960	100.00%	85,960	100.00%	-	0.00%
10/07/2004	2004	2004	30,700	100.00%	52,190	100.00%	39,910	66.50%	31,278	0.00%
15/07/2004	2004	2004	162,710	100.00%	184,200	100.00%	184,200	100.00%	6,844	100.00%
16/07/2004	2004	2004	1,334	0.00%	1,334	0.00%	1,334	0.00%	1,334	0.00%
12/07/2004	2004	2004	921,000	100.00%	921,000	100.00%	1,019,133	0.00%	1,019,133	0.00%
02/07/2004	2004	2004	1,123	0.00%	1,123	0.00%	1,123	0.00%	1,123	0.00%
23/07/2004	2004	2004	68,400	100.00%	49,120	100.00%	49,120	83.31%	49,120	83.31%
03/07/2004	2004	2004	39,910	100.00%	-	0.00%	-	0.00%	-	0.00%
28/07/2004	2004	2004	151,658	100.00%	168,850	100.00%	116,342	0.00%	116,342	0.00%
24/07/2004	2004	2004	171,920	100.00%	196,480	100.00%	35,682	0.00%	35,682	0.00%
26/07/2004	2004	2004	214,900	100.00%	-	0.00%	-	0.00%	-	0.00%
26/07/2004	2004	2004	4,169	100.00%	2,541	0.00%	2,541	0.00%	2,541	0.00%
03/08/2004	2004	2004	493	0.00%	493	0.00%	493	0.00%	493	0.00%
07/08/2004	2004	2004	127,098	100.00%	61,400	100.00%	61,400	100.00%	-	0.00%
10/08/2004	2004	2004	1,245	0.00%	1,245	0.00%	1,245	0.00%	1,245	0.00%
27/07/2004	2004	2004	1,588	0.00%	1,588	0.00%	1,588	0.00%	1,588	0.00%
07/08/2004	2004	2004	844,250	100.00%	905,650	100.00%	905,650	100.00%	905,650	100.00%
15/08/2004	2004	2004	33,770	100.00%	36,840	100.00%	36,840	100.00%	-	0.00%

Figure 11. Claims listing in a format suitable for IBNER calculation. The claims are actual anonymised general liability claims from a single large company (only the first three years of development of only the first few claims are shown). The IBNER factor, e.g., from year 1 to year 2 is calculated as a weighted average of the individual IBNER factors of non-zero claims that are still outstanding (settled claims obviously have no IBNER) from year 1 to 2: $IBNER(1 \rightarrow 2) = (1,535,000 + 85,960 + \dots) / (2,486,700 + 116,660 + \dots)$. If the data is abundant it is also possible to measure the volatility of this IBNER factor and the distribution around the mean value.

- size of claim x (larger claims may have more uncertainty about them);
- outstanding ratio r (the larger the outstanding amount, the more conjectural the estimate of the incurred value will be);
- type of claim t (for example, bodily injury claims will be more difficult than property claims to estimate reliably as the victim's health may worsen unexpectedly);
- ...

As in all situations in which one needs to understand how different factors impact a given variable, this is a problem of statistical learning (see for example Parodi (2012a), where the problem of IBNER analysis was specifically addressed as an example). A typical technique used in insurance to solve statistical learning problems is generalised linear modelling. According to this approach, we can write IBNER factors as a linear combination of different variables, possibly further transformed by a function h :

$$IBNER(d, x, r, t \dots) = h(a_1 f_1(d, x, r, t \dots) + \dots + a_n f_n(d, x, r, t \dots)) \quad (23)$$

and identify the relevant functional form of h, f_1, \dots, f_n and the parameters a_1, \dots, a_n that are suitable to describe the IBNER factors. For example, one might consider a very simple model such as

$$IBNER(d, x, r) = e^{-(a_1 + a_2 d + a_3 x + a_4 r)} \quad (24)$$

which is a special case of the more general Eq. (24), where $f_1(d, x, r, t, s \dots) = 1$, $f_2(d, x, r, t, s \dots) = d$, $f_3(d, x, r, t, s \dots) = x$, $f_4(d, x, r, t, s \dots) = r$, $h(x) = \exp(x)$.




Figure 12. The distribution of the IBNER factors from development year 1 to development year 2, based on the data set whose first few rows are shown in Fig. 11. Note that the largest factors (>2) have been removed for legibility’s sake. Also note the peak at 1, showing that many claims did not change their value from the first to the second year of development.

Output of the IBNER analysis

As a result of this GLM analysis – or even if we use the simple approach outlined at the beginning of this section – we will have:

- a. The IBNER factors to project each existing claim to ultimate. We can then use the projected values of the outstanding claims along with the settled claims to derive the severity distribution
- b. A distribution of values around each IBNER factor, which will allow us to consider each outstanding claim not as a point estimate but as a distribution of possible ultimate values

Note that the fact that there is uncertainty around the ultimate value of each claim, as captured by the distribution in Point (b) above, has the consequence of increasing the parameter uncertainty on the parameters of the severity distribution, as discussed for example in Parodi (2012a).

Note, finally, that there is no theoretical reason why we need to separate the claims estimates into accident years, as if claim reserves only could be revised at regular intervals, so the variable *d* in the equations above may be a continuous as well as a discrete variable.

Practicalities

In practice, the historical development of each claim may not be available, especially for small claims, in which case other solutions must be used:

- If we have at least an initial estimate for the claim and then the final settlement value, we can use those to estimate IBNER;

- If we only have the information on whether a claim is open or closed, we may want to check whether the distribution of the open and closed claims are close enough (in terms, e.g., of the Kolmogorov-Smirnov distance) and if not, we may want to only consider closed claims to derive the severity distribution. However, this is often dangerous and risks underestimating the average loss severity a great deal as large losses often take much longer to settle and therefore they may be hugely underrepresented. This method should be used only if most claims are settled quickly, and making sure that no significant bias is introduced;
- In some cases we may have both a database of individual losses with little or no information on the history of claim reserves, but we may be given a pre-cooked claims triangle (which is sometimes the only input used for reserving studies) with both the non-zero claim count and the aggregate losses. In this case we can derive a triangle for the average claim amount and from that derive a rough measure of IBNER, which we can then use to modify the individual claim amounts. This is a hybrid method which is unlikely to be satisfactory but it may be better than ignoring IBNER altogether.
- If we have no information at all which is relevant to IBNER, that is, we have no average claim triangles and we do not even know which claims are open and which are closed, then as a first approximation we'll simply need to assume that there is no IBNER, or use industry IBNER factors. E.g., in the US market IBNER factors are provided by institutions such as the NCCI and the ISO for different classes of business (e.g. workers' compensation, auto liability), by territory and by type of risk.

3.2.3 Overall algorithm to create a severity model for IBNR claims

Now that we have methods for revaluing IBNR claims (Section 4.1) and to adjust for IBNER (Section 4.2) we can put the pieces together and produce a severity model for IBNR claims. This is a bit different from what one does in pricing.

In the case of pricing, all we need is the severity distribution for the policy we wish to price. The severity model is simply a statistical distribution $F_X(x)$ that gives the probability that a given loss X has a value less than or equal to x : $F_X(x) = \Pr(X \leq x)$. This can be estimated empirically by fitting a suitable statistical distribution to the past losses after they're revalued and adjusted for IBNER.

The severity model for IBNR claims will need to be a slightly more complicated object, as we will need to take into account the fact that the size of the claim will also depend on when the claim occurred. What we need is therefore the probability that a loss X is less than or equal to x conditional on the fact that the loss has occurred at time t_{occ} and has not been reported by the as-at date (t_{as-at}):

$$F_X^{IBNR}(x, t_{occ}, t_{as-at}) = \Pr(X \leq x | T_{occ} = t_{occ}, T_{rep} \geq t_{as-at}) \quad (25)$$

A possible algorithm to create a severity model for losses X occurring at time t_{occ} and not yet reported as at t_{as-at} is as follows:

Input. The inputs to the severity analysis are the individual claim amounts (both paid and estimated)

1. Adjust individual past losses for IBNER, as in Section 3.2.2.
2. Revalue the individual past losses to current terms with some agreed index/inflation rate

- This we need to do so as to determine the joint underlying severity distribution from all years, as if all losses were to occur today
3. Model the severity distribution in current terms, according to standard methods (e.g. a log-normal distribution plus separate tail modelling), obtaining what we will refer to as the **kernel severity distribution** (since this is the nucleus from which the distribution for the different years will be constructed through transformations).
 4. Estimate the severity distribution for losses that occurred at time t_{occ} and that have not yet been reported by transforming the kernel severity distribution according to the method explained in Section 3.2.1 (perhaps using the approximation that there is no delay inflation), or any alternative valid method.

The algorithm outlined in Steps 1 to 4 (and which simply collects all observations dealt with in Sections 3.2.1 and 3.2.2) gives the theoretical background for the severity model of IBNR losses. In Section 3.3 we will see how we can actually sample from the IBNR distribution.

3.3 A protocol for simulating IBNR losses

Now that we have a frequency model and a severity model, we can combine them in the spirit of the collective risk model to produce an aggregate loss distribution. This can be done, e.g., by a Monte Carlo simulation.

We need a “simulation protocol” which takes all the subtleties of the IBNR distribution into account: we need to determine not only how many losses there will be in each simulated scenario, but also when each of these losses is assumed to have occurred, and from what severity distribution we need to sample.

There are many ways in which we can do this, but as an illustration we are going to use a property of the Poisson distribution – assuming, for now, that the Poisson rate is uniform over the years (i.e., the variable $\nu(t)$ in Section 2.2.1 is constant): an assumption we’ll drop shortly.

If we subdivide the period from which the losses are assumed to originate (normally there will be some constraint on how far back we can go) into small intervals $(t', t' + \Delta t)$ of length Δt , we can assume that the number of losses in each of these small intervals is either 1 or 0 (see, e.g., Ross, 2003), and that it is 1 with probability $\lambda(1 - F(t - t')) \Delta t$, where $F(t - t')$ is the probability that a loss from the interval $(t', t' + \Delta t)$ has been reported by time t (see Section 2).

This has a simple geometric interpretation: it is exactly like drawing a number (distributed as a Poisson variate with rate equal to the point estimate of the IBNR claim count over the years) of random dots uniformly from inside the dark green area of Fig. 6 (which we reproduce here for convenience), and select the x -coordinate as the simulated time of occurrence.

In the general case, the Poisson rate is not uniform, and the process may actually be a negative binomial process). Let’s deal with these two aspects in turn.

Non-uniformity. The geometric interpretation above is still valid if the Poisson rate is not uniform, but the geometric figure must be warped a bit: instead of the rectangle representing the expected ultimate claims as in Fig. 6 we have to consider the area below the function $\nu(t')$, and instead of the light-green area in Fig. 6 giving the expected reported we have the area below the function $\nu(t') F(t - t')$.

Negative binomial case. In case a negative binomial distribution process is used as a model for the IBNR loss count (even if only to represent the parameter uncertainty for the Poisson case), we can simply draw the number of IBNR losses from a negative binomial distribution by using the familiar Poisson/Gamma mixture process, and then sample it in proportion to $v(t') (1 - F(t - t'))$. The problem with this method is that it does not reproduce the extra year-on-year volatility which is typical of realistic loss processes. This is a known problem: “sample realizations of such processes look [like] Poisson ones – the variation is not *within*, but *between* processes” (Kozubowski & Podgorski, 2009). If we want to reproduce this extra volatility, we need to use another realisation of the NB process, e.g. the Poisson/Logarithmic mixture process, which creates clusters (“colonies”) of losses happening at the same time (*ibid.*). If we use this framework, it is these “colonies” of losses that are distributed at random in the dark-green area of Fig. 6.

Based on the considerations above, we outline the following protocol for simulating IBNR losses.

Simulation protocol

Input: Rate of IBNR losses, λ , effective variance-to-mean ratio (the one which includes the parameter uncertainty), $1 + R$, and the kernel severity distribution⁶, $F_X(x)$.

- i. For each scenario $j = 1$ to N_{sim} , simulate the number n_j of IBNR claims from a negative binomial distribution with rate $= \lambda$ and variance-to-mean ratio $= 1 + R$.
- ii. For each loss $i = 1$ to n_j , simulate the loss occurrence date as described above, in proportion to $v(t') (1 - F(t - t'))$, where t' is the occurrence date and t is the as-at date
- iii. For each loss $i = 1$ to n_j , if t_j is the loss occurrence date, sample a loss $x_i^{(j)}$ from the kernel loss distribution $F_X(x)$, then modify it to take into account the loss occurrence date, e.g. (using Alternative Assumption 1 in Section 3.2.1) $x_i^{(j)} = (1 + r)^{t_j - t} x_i^{(j)}$.
- iv. Loop over i
- v. Calculate the total losses for each scenario j , as $S_j = \sum_{i=1}^{n_j} x_i^{(j)}$
- vi. Sort all values of S_j in ascending order: $S_{(1)} < S_{(2)} < \dots < S_{(N_{\text{sim}})}$

Output: The empirical aggregate loss distribution, $\widehat{F}_S(s) = \max\{j | S_{(j)} \leq s\} / N_{\text{sim}}$.

4 A Model for Reported but not Settled (RBNS) Claims

In Section 3.2.2 we have seen how IBNER must be taken into account in order to derive the severity distribution of IBNR claims. Traditionally, however, IBNER is the amount of reserves that need to be put aside to deal with the possible adverse development of claims that have already been reported but not settled (RBNS) claims. This amount needs to be taken into account when estimating the total amount of reserves, and the uncertainty around that amount.

The calculations to estimate IBNER – which try to identify a systematic bias in the reserves, one way or the other – are exactly the same as in Section 3.2.2 but the output is used in a different way.

⁶ We ignore here the issue of parameter uncertainty for the severity distribution, which will need to be taken into account when comparing the results of this method with the chain ladder. We'll return to this issue in Section 7.

The IBNER analysis allows the estimation of an IBNER factor for each outstanding claim. As a result of this, the estimated ultimate value of the RBNS losses can be written like this:

$$\text{Ult_RBNS} = \sum_{k=1}^n \text{IBNER}(d_k, x_k, r_k, t_k \dots) \times X_k \tag{26}$$

In Equation (26), X_k is the k -th loss in the database. We have included here all losses, whether closed or not, with the convention that $\text{IBNER}(d_k, x_k, r_k, t_k \dots) = 1$ for settled losses, giving therefore no contribution to the estimated Ult_RBNS .

Since each IBNER factor $\text{IBNER}(d_k, x_k, r_k, t_k \dots)$ (except those for settled claims) is actually a random variate according to Equation (31), the estimated Ult_RBNS is also a random variate. If IBNER factors can be assumed to come from a simple statistical distribution such as a Normal distribution or a Gamma distribution (this is normally the case if IBNER is calculated as a result of a GLM distribution, or – even more simply – if it’s produced with a simpler, heuristic exercise), it is also easy to sample from the overall IBNER distribution:

The following protocol allows the estimation of the distribution of RBNS reserves.

Input: IBNER factors model

- (i) For each scenario $j = 1$ to N_{sim}
- (ii) For each open claim, sample a value of the IBNER factor from the empirical distribution or from an IBNER model (e.g. GLM), IBNER_k^j ;
- (iii) Calculate the contribution of each open claim to the estimated ultimate value of the RBNS, $\text{IBNER}_k^j \times X_k$
- (iv) Sum the IBNER contributions over all open claims, obtaining the simulated ultimate RBNS: $\text{Ult_RBNS}^j = \sum_{k=1}^n \text{IBNER}_k^j \times X_k$
- (v) Repeat for all scenarios

Output: Empirical ultimate distribution of outcomes for RBNS claims

A fleshed-out example of this algorithm is illustrated in Section 7.4, in the discussion of our case study.

5 An Aggregate Model for Future Losses

Reserves may also need to be put aside for losses that have not occurred yet but that are covered by the policies in force (or that are going to be in force). These are often called “unearned premium reserves” (UPR). E.g., the estimation of UPR is a requirement of Solvency 2.

Determining the distribution of future losses is a standard pricing exercise, by which one tries to determine the future loss distribution by building a frequency model and a severity model based on past losses (under the usual assumption that the past is to some extent a guide to the future). For this reason, we do not dwell on this, but we only sketch the approach we need to take.

In case the severity distribution is the same (apart from claim inflation) as that of previous years, one can use the kernel severity as calculated in Section 3.2.3 and revalue it to the mid-term of the policies in force in order to obtain the severity distribution of the future losses.

The frequency model can also be produced based on the results of Section 3: the Poisson rate for future losses will simply be given by the projected claim count over the period considered, divided by the total exposure and then multiplied by the future exposure.

By combining the frequency model and the severity model in the usual way, e.g. through Monte Carlo simulation, one obtains an estimate of the aggregate loss distribution of future losses.

In practice, this means that in many cases there is no need for a separate exercise for IBNR and UPR, but that they can be simulated based on the same model. An example of this is shown in the case study of Section 7.

6 A Model for the Overall Outstanding Liabilities

The overall outstanding liabilities are made of three components:

$$\text{Outstanding liabilities} = \text{IBNR} + (\text{RBNS} - \text{Paid}) + \text{Future losses} \quad (27)$$

We have dealt with (pure) IBNR losses in Section 3, with RBNS losses in Section 4, and with future losses in Section 5.

Since all the three components of the outstanding liabilities are random variates, the overall outstanding liabilities can also be considered as a random variate. The ultimate goal of a reserving exercise is to estimate the distribution of the overall outstanding liabilities, so that we know how much money will be needed at different levels of probability.

If we assume that IBNR, RBNS and future losses are statistically independent (see Section 6.1 below for a discussion), the distribution of the outstanding liabilities is the convolution of the three distributions for IBNR, RBNS and future losses (the “paid” amount can be subtracted at the end).

A three-variable convolution sounds a bit daunting from the mathematical point of view but with the help of numerical analysis is actually quite simple. Using Monte Carlo simulation, for example, the distribution of the overall reserves can be obtained by producing a large number of scenarios. For each scenario, a random value is drawn from the distribution of the IBNR losses, the IBNR losses and the future losses (in practice, we can use the outputs of the processes in Sections 3, 4 and 5), and the three values are summed together. We can then sort all simulations in ascending order and produce the empirical distribution of the overall reserves in the usual way.

6.1 Dependence among IBNR, RBNS and UPR

How good is the assumption that the IBNR, RBNS and UPR distributions are independent?

Conditional to the risk model which has been derived for each of these components (and involving, where applicable, a reporting delay distribution, a frequency distribution and a severity distribution) being correct, there is no residual dependency between the three components.

For example, the number of IBNR claims that will actually be reported in practice is – again if our frequency model is correct – independent of the number of claims that will be reported for the unearned portion of the premium (UPR), i.e. the future losses, much in the same way that the number of losses occurred in one year is independent of the number of losses occurred in another year, although the underlying frequency may be the same.

The same is true for our severity distribution and our reporting delay distribution – if we're sampling from the same severity distribution, subject to all the modifications discussed in Section 3.2.1, the actual samples for IBNR and UPR will be independent.

And that is also true for RBNS. You will have noticed that the modelling decisions one makes for RBNS have an impact on IBNR, RBNS and UPR, at least if open claims are used to produce the kernel severity distribution. However, if the model you have produced for RBNS is correct, the IBNER amount that you need to add to a particular claim to obtain the claim's ultimate true value is independent of the IBNER amount that you need for another claim, whether it's not occurred yet, incurred but not reported, or reported but yet outstanding.

However, that does not mean that we can ignore dependency under all circumstances. There are at least two types of dependency that we may want, or even need, to take into account.

1. *Modelling choices.* These are the correlations that arise out of the modelling process itself. The severity model that we use for UPR and IBNR is based on the same kernel severity model, and therefore any parameter and model choice that we make will affect both estimates. E.g., if we choose a lognormal model for both UPR and IBNR but the true model turns out to be a Gamma distribution, this wrong choice will affect both UPR and IBNR in the same way. So the correlation is not in the loss process but is an artifact of the calculation method.

How do we deal with this? Correlations arising out of the modelling process can be easily taken into account by making sure that the modelling choices (e.g. the parameters) are the same across the different components (RBNS, IBNR, UPR) of a given simulated scenario. The output of the simulation will therefore reproduce these correlations.

2. *Systemic shocks.* These are the correlation that are created by the fact that some of the assumptions behind the model need to be changed. An example is if the discount rate applied by the courts to produce an estimated lump sum for liability claims – if it were changed to reflect, say, lower returns on gilts, *all* claims that are not settled (whether reported but outstanding, incurred but not reported, or not incurred yet) might be affected and that would mean that the IBNER, IBNR and UPR distributions would change. On the frequency side, a change in legislation that encourages or discourages the reporting of claims will have an impact on both UPR and IBNR claims (not on RBNS, since the number of outstanding claims is fixed).

How do we deal with this? Correlations arising out of systemic shocks can be modelled by using a so-called correlation shock. Assume, for example, that we have convinced ourselves that the discount rate will go down by 0.5% with a probability of 1/50, and by 1% with probability of 1/100 (the probability of remaining the same being therefore 97/100). We then assess the effect of this discount rate change on the severity distribution, and as a consequence on all UPR, IBNR and RBNS. Finally, for each simulation scenario we choose at random one of the three situations (constante discount rate, 0.5% decrease, 1% decrease) with the relevant probabilities, and we apply the same assumption across all the three components of the simulation (UPR, IBNR, RBNS).

7 A Simple Case Study

In Sections 2 to 6 we have described a protocol for producing a model of the overall liabilities for a risk. Let's now see a specific implementation of this protocol in a simple (but not too simple) case study.

This case study is based on real-world data which has been rescaled and anonymised. To be more specific without being revealing, the data set is actually the combination of two data sets:

- a set of employers' liability claims from a large UK company, from which we have derived all the information except from IBNER. The losses span a period of 10.5 years, from 1/1/2001 to 31/7/2011. The EL policy period is 1/1 to 31/12 (unaltered through the years);
- a set of general liability claims from a large non-UK company, from which we have derived IBNER information.

We have applied a randomly chosen rescaling factor to all losses in (a) and (b), and a subset of (b) was excluded in order to make certain quantities (e.g. the probability that a 100% outstanding claim ends up as a nil claim) consistent between the two data sets.

The objective of this case study is to estimate the distribution of provisions needed for the period 1/1/2001–31/12/2011. Since we only have information up to 31/7/2011, that means that we have to consider:

- RBNS losses (losses that have occurred between 1/1/2001 and 31/7/2011 and have been reported, but have not been settled yet and could therefore develop further (IBNER));
- IBNR (losses occurred between 1/1/2001 and 31/7/2011 that have not been reported yet);
- UPR (losses that will occur between 1/8/2011 and 31/12/2011).

As we will see, IBNR and UPR can be modelled together.

7.1 Producing the frequency model

7.1.1 Estimating the report delay distribution

We have estimated the report delay distribution, based on the observed empirical delays. The adjustment for the bias introduced by the limited observation window has been carried out as explained in Section 2. We have *not* assumed a specific model (e.g. exponential) for the delay distribution. The bias-corrected empirical delay distribution is shown in Fig. 13.




Figure 13. The empirical reporting delay distribution, after correction for the bias introduced because of a limited observation window.

7.1.2 Counting the number of non-zero reported losses

We have then counted the number of *non-zero* reported losses for each accident year.

We had to take into account the fact that some of the losses reserved as non-zero may end up being zero (especially, and almost exclusively, if they're 100% outstanding). There are several ways to do this. When we have the historical loss amount for each loss, as in this case, we can calculate the percentage of non-zero, 100% outstanding claims that drop to zero for each year of development, and based on that we can estimate the number of currently reported claims that will end up being zero. This is illustrated in Fig. 14.

Year of development	Year-on-year probability of dropping to zero	Year-on-year survival probability	Cumulative survival probability
0	4.6%	95.4%	79.1%
1	9.4%	90.6%	83.0%
2	8.4%	91.6%	91.6%
3	0.0%	100.0%	100.0%
4	0.0%	100.0%	100.0%
5	0.0%	100.0%	100.0%
6	0.0%	100.0%	100.0%

Figure 14. Based on the historical individual loss information (that information that we have also used to analyse IBNER) we have calculated the probability that a claim drops to zero after being reported. E.g., there is a 4.6% that a claim which is non-zero in development year 0 will drop to zero in development year 1 – or, in other terms, there is a 95.4% probability that the claim will *not* drop to zero (what we call in the table the *survival probability*). Overall, the cumulative survival probability that a claim which is non-zero in development year 0 will eventually drop to zero is $100\% - 79.1\% = 20.9\%$.

Armed with the statistics shown in Fig. 14, we can count the number of claims that are fully outstanding for each year and from that we can determine the number of non-zero losses expected to survive, as illustrated in Fig. 15. This will be the input to calculate the projected number of claims for each year.

7.1.3 Estimating the IBNR/UPR claim count

We can now use the reporting delay distribution (Fig. 13) and the claim count (Fig. 15, column 6) to project the claim count to ultimate for each year, or on aggregate and produce a frequency model.

We have used a negative binomial to model the claim count. This is defined by two parameters: the mean claim count and the variance-to-mean ratio, which can be calculated by projecting each year independently and then calculating the year-on-year variance, as in Fig. 16.

The estimated parameters of our frequency model are a mean of 0.0586 claims per unit of exposure and a variance-to-mean ratio of 1.25 (which also takes into account parameter uncertainty).

Once we have an estimate of the frequency model per unit of exposure, we can estimate the projected IBNR amount for each year at the relevant exposure. The variance-to-mean ratio plays a role in the calculation of the standard error with which the projected IBNR amount is known.

Policy Year	No of losses	No of non-zero losses	No of losses with O/S=100%	No of losses with O/S=100% expected to survive	No of non-zero losses expected to survive
2001	65	51	0	0.0	51.0
2002	72	51	0	0.0	51.0
2003	52	35	0	0.0	35.0
2004	63	44	0	0.0	44.0
2005	45	35	0	0.0	35.0
2006	53	42	0	0.0	42.0
2007	50	36	1	1.0	36.0
2008	59	43	15	13.2	41.2
2009	28	23	12	10.3	21.3
2010	20	19	13	10.6	16.6
2011	12	11	11	8.7	8.7
Total	519	390	52	43.8	381.8

Figure 15. Each of the fully outstanding losses (column 4) has a survival probability that depends on its year of development. Summing all these survival probabilities for each claim occurred in a given year we can find the number of losses that are fully outstanding and that are expected to survive indefinitely (column 5). The number of non-zero losses expected to survive (column 6) is given by the number of non-zero reported losses (column 3)+the number of fully outstanding losses expected to survive (column 5) – the number of fully outstanding losses (column 4).




Figure 16. This chart shows the claim count (grey line) for each year after adjustment for exposure changes (light blue bars), the projected claim count for each year (blue line, with error bars of size equal to twice the standard error), and the estimated mean claim count (30.2) at an arbitrary level of exposure (dashed red line), calculated as a weighted average of years 2001–2010 (2011 is too immature). (Note: The arbitrary level of exposure was in this case chosen to be the estimated next year's exposure (515 arbitrary units): this makes sense in a pricing context but in this case all it matters is the frequency per unit of exposure.)

The IBNR calculations on an unadjusted basis (i.e., no corrections for changes in exposure) are shown in Fig. 17 (chart form) and Fig. 18 (tabular form).

Note that the 2011 projected figures include both IBNR and UPR.

The IBNR calculations illustrated in Fig. 17 and Fig. 18 are performed year by year. As a consequence, the projection of the last few years get increasingly more uncertain, as they're based on an increasingly smaller number of claims. This is especially true for 2011, which is strongly underdeveloped.

This problem is not specific to triangle-free reserving – exactly the same phenomenon happens with the chain ladder and with every attempt to project losses to ultimate based on last year's data alone.

A more accurate estimate of the overall number of IBNR claims can be obtained by lumping all years together, as shown in Fig. 19.

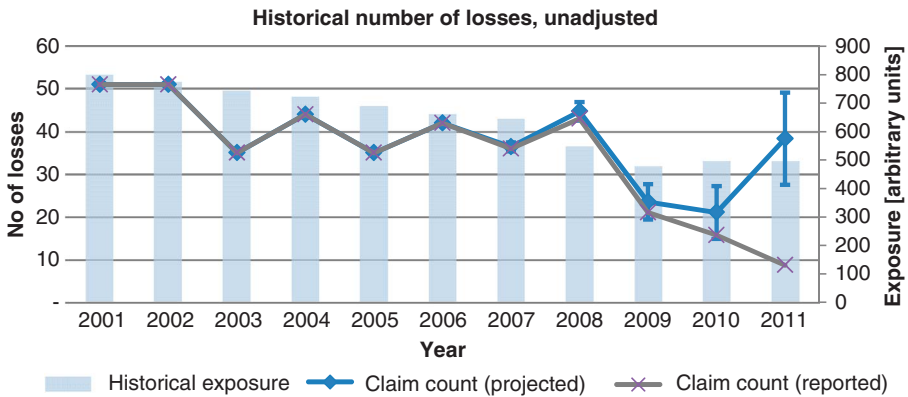


Figure 17. Projection of claim count (grey line) to ultimate (blue line) for each year, without adjustments for exposure changes (light blue bars).

PROJECTION TO ULTIMATE - YEAR BY YEAR													
	C1	C2	C3	C4	C5	C6	C7	C11	C12	C13	C15	C16	C17
	Year	Earned days	Factor to ultimate	Latest reported	Ultimate losses	Standard error	Exposure	IBNR only	IBNR percentage	"Prior" estimate of ultimate	Credibility factor	estimate of ultimate	Credibility estimate of IBNR
2001	365.0	1.00	51.0	51.0	0.0	800	0.0	0.0%	46.9	1.00	51.00	0.0	
2002	365.0	1.00	51.0	51.0	0.0	773	0.0	0.0%	45.3	1.00	51.00	0.0	
2003	365.0	1.00	35.0	35.0	0.0	742	0.0	0.0%	43.5	1.00	35.00	0.0	
2004	366.0	1.00	44.0	44.0	0.0	721	0.0	0.0%	42.2	1.00	44.00	0.0	
2005	365.0	1.00	35.0	35.0	0.0	688	0.0	0.0%	40.3	1.00	35.00	0.0	
2006	365.0	1.00	42.0	42.0	0.0	661	0.0	0.0%	38.7	1.00	42.00	0.0	
2007	360.6	1.01	36.0	36.4	0.6	645	0.4	1.2%	37.8	0.99	36.45	0.4	
2008	351.9	1.04	43.0	44.7	1.2	548	1.7	3.8%	32.1	0.97	44.31	1.7	
2009	326.4	1.12	21.0	23.5	1.9	478	2.5	10.6%	28.0	0.91	23.95	2.5	
2010	272.4	1.34	15.7	21.1	3.0	495	5.3	25.4%	29.0	0.80	22.62	5.7	
2011	82.7	4.42	8.7	38.3	5.2	495	29.6	77.4%	29.0	0.57	34.36	26.6	
Total	3,584.9		382.4	422.1		7,047	39.6		412.8		419.69	37.0	

Figure 18. The first nine columns show the same information as the chart in Fig. 17 but in tabular format. The input data (latest reported, exposure) is indicated in blue; the intermediate calculations are in black; the outputs we're interested in are in purple. The last four columns illustrate the credibility calculations that lead to a revised figure for IBNR, which is very similar to the one based on client data only for all years except for 2011.

PROJECTION TO ULTIMATE - ALL YEARS TOGETHER							
Year	Earned days (exposure-adjusted)	Factor to ultimate	Latest reported	Ultimate losses	Standard error	Exposure	IBNR only
2001-11	3,681.8	1.09	382.4	417.3	25.6	7,047	34.8

Figure 19. The same calculations as in Fig. 18, but without artificially splitting the estimate by year.

Another possible solution to the problem of having unstable results for the most recent year(s) is to adopt a credibility approach. This is not very different from what we do with Bornhuetter-Ferguson and Cape Cod for triangulation methods.

In order to produce a credibility-weighted estimate for all years, we need a benchmark against which to compare the raw projections. A natural benchmark is provided by the estimated number of claims for each year based on the frequency per unit of exposure derived above ($\hat{\lambda} = 0.0586$) and the exposure for each year, $\varepsilon(y)$:

$$\text{Projected}_b(y) = \hat{\lambda}\varepsilon(y) \quad (27a)$$

The volatility of such a figure will be given by the projected amount times the variance-to-mean ratio, yielding:

$$\text{Var_Projected}_b(y) = (1+r)\hat{\lambda}\varepsilon(y) \quad (27b)$$

This credibility factor can then be calculated by balancing this prior volatility against the standard error (squared) on the projected ultimate for each year (see, e.g., Parodi & Bonche (2010)). This standard error is calculated according to Equation (11) and leads to the numbers in Fig. 18 (C6). The credibility factor for year y is then given by:

$$Z(y) = \frac{\text{Var_Projected}_b(y)}{(\text{s.e.}(y))^2 + \text{Var_Projected}_b(y)} \quad (27c)$$

(Note that since the variance-to-mean ratio is included in all terms, in practice it plays no role in Equation (27c).)

The credibility estimate is thus given by

$$\text{Cred_Estim}(y) = Z(y)\text{Ult}(y) + (1-Z(y))\hat{\lambda}\varepsilon(y) \quad (27d)$$

Which in terms of the columns of Fig. 18 can be written as $C16 = C15 \times C5 + (1 - C15) \times C13$.

The credibility estimate for IBNR can then be calculated by applying the IBNR percentage (column C12) to the credibility estimate for the ultimate of Eq. (27d), leading to the results of column C17.

7.2 Producing a severity model

7.2.1 Adjusting for IBNER

We do a simple IBNER calculation based on the average of past IBNER factors by development year. This was done as explained in Section 3.2.2. The results are shown in Fig. 20.

Year of development	Average year-on-year IBNER ratio	Average cumulative IBNER ratio
0	1.091	1.133
1	1.061	1.038
2	0.974	0.979
3	1.014	1.005
4	0.991	0.991
5	1.000	1.000
6	1.000	1.000

Figure 20. Average IBNER factors for our data set. IBNER factors are assumed to be 1 after year 4 because the dwindling sample size makes it difficult to have confidence in the estimate.

The cumulative IBNER factors of Fig. 20 (column 3) are then applied to all open claims (open = more than 5% outstanding) of the relevant year of development in order to estimate the final settled amount.

7.2.2 Is there a common severity model to all years?

In its simplest form, our approach assumes that the severity distribution is the same for all years, except for claims inflation and possibly for the dependency of the average size on reporting delay. In this section we are going to test this assumption.

First of all, we revalue all claims by 5% p.a. (a common assumption for liability claims) to 2011, and we adjust all RBNS claims by the appropriate IBNER factor as derived in Section 7.2.1.

We then compare the empirical severity distribution of different years. To have samples that are statistically more significant, we have grouped years in sets of three: 2001–03, 2004–06, 2007–09, 2010–11 (this last set only has roughly 1.5 years).

Fig. 21 shows that the assumption that there is a common severity distribution to all years (except for claims inflation) is a good one for this risk. This can be appreciated visually but can

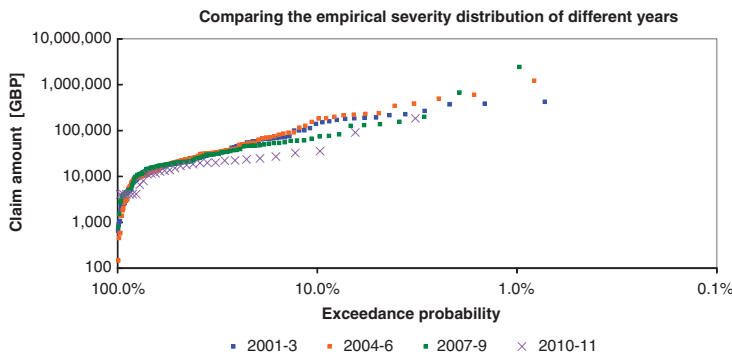


Figure 21. A comparison of the empirical severity between different years. Note that the comparison has been performed after adjusting for IBNER. Since the IBNER factors are spurious in this case (IBNER factors of data set (b) have been applied to data set (a)) it is likely that the differences between the years have been exaggerated.

Set 1	Set 2	Normalised KS distance
2001-03	2004-06	0.424
2001-03	2007-09	0.747
2001-03	2010-11	1.197
2004-06	2007-09	0.722
2004-06	2010-11	1.178
2007-09	2010-11	1.046

Figure 22. A statistical comparison among the data sets. All sets can be considered as coming from the same distribution with a significance level of 10% (remember that the normalised KS distance at a confidence level of 10% is 1.22).

also be confirmed statistically by looking at the two-sample KS distance among the different sets (Fig. 22).

7.2.3 The kernel severity model

We then bring all the revalued losses together and build a kernel severity model from which the model for the different years can be derived. We have modelled the losses with an empirical distribution (i.e., resampling of existing losses) up to £10,000 and then a GPD with parameters $\xi = 0.736$, $\sigma = 18,263$, $\mu = £10,000$ above that (Fig. 23).

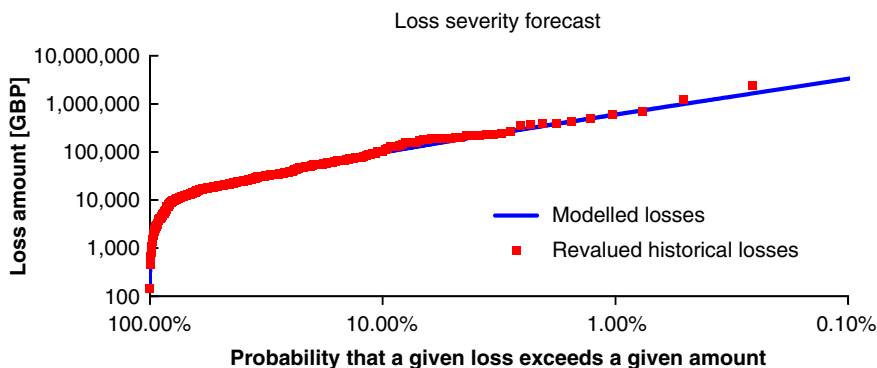


Figure 23. Severity model (blue line) vs revalued and IBNER-adjusted losses (red dots). Note that the empirical distribution was used below £10,000 (hence the perfect fit below that threshold).

7.2.4 The severity model for each year

Based on the kernel severity model produced in Section 7.2.3, we can now derive the severity model for each year.

In doing this, we assume that the loss severity does not depend on the reporting delay (this is often a good assumption, as shown in Fig. 9).

We proceed as follows:

- i. we sample 20,000 points of the kernel severity model (which is revalued to 2011), which yields a discretised version of the kernel severity model that is then easier to manipulate further;
- ii. The severity model for year $y = 2011 - x$ can be obtained by dividing each point of the sample above by 1.05^x .

Notice that we only need to produce a severity model for those years for which there is some IBNR! In our case, these are 2007–2011.

7.3 The aggregate model for IBNR/UPR losses

We can now combine the frequency model (Section 7.1) and the severity model (Section 7.2) for IBNR/UPR losses with Monte Carlo simulation. We can use a separate model for each year (which is the option we've adopted here) or we can combine all years in a single one by mixing the severity distributions in proportion to the frequency of IBNR claims.

(Note that we have used for the IBNR frequency the projection year by year based on column C11 of Fig. 18, not the credibility estimate – this is not because we believe this estimate is better but simply because we had to choose one of the cases for illustration purposes.)

Fig. 24 shows the output of this exercise.

Percentile	2007		2008		2009		2010		2011		All years combined	
	No of losses	Total amount	No of losses	Total amount	No of losses	Total amount	No of losses	Total amount	No of losses	Total amount	No of losses	Total amount
50%	-	-	1	28,506	2	54,494	5	165,329	29	1,301,057	39	1,767,666
75%	1	9,813	3	79,444	4	121,313	7	298,190	34	1,843,534	44	2,421,773
80%	1	15,157	3	98,607	4	149,049	7	348,429	35	2,020,449	45	2,615,037
90%	1	41,196	4	174,129	5	259,439	9	548,710	37	2,675,384	49	3,438,802
95%	2	82,308	4	273,492	6	410,515	10	842,833	40	3,733,321	51	4,646,643
98.0%	3	170,068	6	540,751	7	723,883	11	1,376,454	43	5,920,532	55	7,495,076
99.0%	3	285,339	6	879,810	7	1,109,637	12	2,200,778	45	8,269,363	57	10,083,991
99.5%	4	449,376	7	1,405,362	8	1,851,530	13	3,123,161	47	12,821,095	59	14,758,372
99.8%	4	915,845	8	2,499,818	9	4,132,847	14	5,866,062	49	19,249,445	61	20,596,802
99.9%	4	1,655,806	8	5,642,704	10	5,428,090	15	10,434,619	51	31,648,211	63	31,941,402
Mean	0.4	20,871	1.7	85,234	2.5	127,921	5.4	292,656	29.6	1,691,109	39.7	2,217,791
Std Dev	0.7	165,632	1.5	309,850	1.8	435,990	2.6	807,740	6.1	1,825,731	7.0	2,080,169

Figure 24. The output of the Monte Carlo simulation for each year of occurrence, and with all years combined.

7.4 The aggregate loss model for RBNS losses

In order to analyse the distribution of possible outcomes for RBNS losses, let us first look at the settlement pattern for open claims. Fig. 25 shows empirical statistics for our data set. These are obtained by looking at the percentage of claims of “age” (development year) d that have already been settled, and deriving from that the conditional probability that a claim which has already reached age d will be settled by the time it reaches age d' .

Given the crucial information of Fig. 25, we can now produce a distribution of possible outcomes for RBNS claims according to one of the methods mentioned in Section 4. In this specific case, we have adopted the following algorithm, which is a specific version of the general algorithm described in Section 4.

Year of development	Percentage settled	Percentage settled (incremental)	Conditional percentage settled (d>1)	Conditional percentage settled (d>2)	Conditional percentage settled (d>3)	Conditional percentage settled (d>4)
0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1	22.2%	22.2%	0.0%	0.0%	0.0%	0.0%
2	46.2%	23.9%	30.8%	0.0%	0.0%	0.0%
3	87.2%	41.0%	52.7%	76.2%	0.0%	0.0%
4	90.9%	3.7%	4.8%	6.9%	29.1%	0.0%
5	100.0%	9.1%	11.7%	16.9%	70.9%	100.0%
6+	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Figure 25. This table gives the probability that a claim settles at year s conditional to the fact that is still outstanding at development year d . For the purpose of this exercise, “settled” means that less than 5% of the overall estimated amount is still to be paid. (The reason why this is considered settled is that often after the indemnity part of a claim has been paid, the insurer might still incur relatively minor expenses.)

Loss ID	Development year (fixed)	O/S percentage (fixed)	Sampled settlement year	Sampled IBNER factor	IBNER-adjusted loss
Loss_1	2	81.9%	3	1.000	87,086
Loss_2	4	28.6%	5	1.070	32,207
Loss_3	4	20.0%	5	0.029	4,001
Loss_4	4	100.0%	5	1.314	5,904
Loss_5	1	49.1%	2	0.857	41,622
Loss_6	3	94.4%	5	1.000	52,843
Loss_7	3	42.3%	5	0.734	13,159
Loss_8	3	100.0%	5	0.032	1,452
Loss_9	1	95.6%	2	0.801	61,984
Loss_10	0	100.0%	3	2.732	50,376
Loss_11	2	83.8%	3	1.000	48,121
Loss_12	3	80.6%	5	0.807	48,455
Loss_13	2	100.0%	3	1.199	41,320
Loss_14	1	100.0%	2	0.983	13,422
Loss_15	3	100.0%	5	1.000	20,929
...
Loss_68	0	100.0%	3	0.232	6,960
<i>All losses</i>					<i>2,237,926</i>

Figure 26. Example of a simulation for RBNS losses. In our example, there are 68 open claims, each with their own development year and O/S percentage. The simulation works by sampling a random settlement year (column 4) and a random cumulative IBNER factor to go from year of development to the simulated settlement year. The sum of all IBNER-adjusted losses (£2,237,926 in the example) gives the result of this simulated scenario. This process has been repeated 10,000 times to yield the results of Fig. 27.

Input: All the open losses with their year of development, and all the IBNER factors model

- (i) For each scenario $j = 1$ to N_{sim}
 - a. For each open loss X_k
 - i. sample a value of the settlement year $s(j, k)$, based on the past stats for settlement time (see Fig. 25);

- ii. sample IBNER factors $IBNER_{d,d+1}, IBNER_{d+,d+2} \dots IBNER_{s(j,k)-1,s(j,k)}$ from the empirical distribution (i.e. the actual IBNER factors)
 - iii. calculate the cumulative IBNER factor for this loss by multiplying these factors together:
 $IBNER_k^j = IBNER_{d,d+1} \times IBNER_{d+,d+2} \times \dots \times IBNER_{s(j,k)-1,s(j,k)}$;
 - iv calculate the IBNER-adjusted value of the claim, $X'_k = IBNER_k^j \times X_k$
- b. Sum over all IBNER-adjusted claims, obtaining the simulated ultimate RBNS for simulation j : $S^j = \sum_{k=1}^n X'_k$
- (ii) Repeat for all scenarios
 - (iii) Sort all the values $S^j, j = 1, \dots, N_{sim}$ in ascending order

Output: Empirical distribution of possible outcomes for RBNS losses

In our case, we have used $N_{sim} = 10,000$ simulations, in each of which a settled value for each of the 68 open claims has been sampled. An illustration of this can be found in Fig. 26, which shows the results of the first simulated scenario (out of 10,000).

Percentile	Number of RBNS losses	RBNS total amount
50%	68	2,196,863
75%	68	2,533,844
80%	68	2,643,906
90%	68	3,052,900
95%	68	3,551,902
98.0%	68	4,717,209
99.0%	68	5,630,207
99.5%	68	6,013,556
99.8%	68	7,692,927
99.9%	68	9,315,854
<i>Mean</i>	68.0	2,377,816
<i>Std Dev</i>	68.0	734,858

Figure 27. The result of the RBNS simulation described in this section. Note that the number of losses remains constant as it is only the ultimate settlement value which can change here.

7.5 The overall distribution of liabilities

Finally, we can combine the contributions of IBNR/UPR and of RBNS claims to obtain the distribution of overall losses, as shown in Fig. 28. This can be done as usual by summing the simulations of RBNS and IBNR/UPR one by one if independence can be assumed.

Note, finally, that the distribution above also includes the paid component. For reserving purposes, one will then need to subtract the paid amount of the RBNS component from the overall losses.

7.6 Conclusions and limitations of this case study

The case study presented in this section illustrates how triangle-free reserving can be used in practice. We have chosen an example in which *most* of the features of the approach could be used,

Percentile	Number of IBNR/UPR losses	IBNR and UPR	Number of RBNS losses	RBNS total amount	Overall number of losses	Overall losses
50%	39	1,767,666	68	2,196,863	107	4,115,957
75%	44	2,421,773	68	2,533,844	112	4,895,375
80%	45	2,615,037	68	2,643,906	113	5,171,516
90%	49	3,438,802	68	3,052,900	117	6,220,342
95%	51	4,646,643	68	3,551,902	119	7,537,303
98.0%	55	7,495,076	68	4,717,209	123	10,329,807
99.0%	57	10,083,991	68	5,630,207	125	13,140,436
99.5%	59	14,758,372	68	6,013,556	127	16,884,020
99.8%	61	20,596,802	68	7,692,927	129	23,828,525
99.9%	63	31,941,402	68	9,315,854	131	34,267,323
<i>Mean</i>	39.7	2,217,791	68.0	2,377,816	107.7	4,595,607
<i>Std Dev</i>	7.0	2,080,169	68.0	734,858	7.0	2,218,176

Figure 28. The distribution of possible outcomes for the overall losses, obtained by combining “line-by-line” the simulation results for RBNS and for IBNR/UPR.

and we have used the simplest and most implementation possible for the various techniques used here (e.g. for IBNER analysis).

One limitation of this case study is that, indeed, not *all* features of the approach have been used. Specifically, we have not included a tail factor, because its effect (as calculated based on the current reporting delays) was negligible based on a loss period of 10 years and an exponentially decaying reporting delay. However, a more thorough analysis could have considered a fatter tail (e.g. Pareto) or could have included a separate allowance, e.g., for latent industrial disease claims.

8 Validation and Comparison with the Chain Ladder

It seems reasonable that a reserving method which uses all the information in the loss data set rather than a compressed version of it should provide a better assessment of the overall reserves to put aside for a given risk, and especially of the uncertainty of these reserves. However, proving that the triangle-free method is actually more accurate than the triangle-based methods is much more difficult.

Validating actuarial methods for reserving is famously difficult, especially because it ultimately involves waiting many years to compare the results of a model against what happens in reality, and do that for a large number of cases to make sure the result is statistically meaningful. Comparing different methods is equally difficult. The difficulty becomes perhaps insurmountable when we try not only to validate point estimates but full reserving distributions.

A different approach towards validation is to use an artificial data set of which we know (and actually decide) all the statistical properties. The artificial data set provides us by definition with the “true” answer, and therefore we can calculate the prediction error of any method we wish, and this method is fully objective.

The catch, of course, is that when the validation and comparison of methods is performed with an artificial data set, the stakeholders (in this case, the actuarial community and whoever else is

involved in reserving decisions) need to agree that the artificial data set used is an adequate replica of reality, and that it doesn't deliberately or unwittingly favour one method against the other.

As for the comparison between different methods, each method may come in different flavours and there is a risk that one compares one's favourite method with a straw man.

For these reasons, the artificial data set should probably be produced and agreed on by an independent party, and the methods should be provided by their advocates, a bit like in Axelrod's experiment on behavioural strategies (Axelrod, 1984).

Having all these limitations in mind, we have performed in this paper a few controlled experiments to validate the triangle-free method against a very standard version of the chain ladder. These experiments have some limitations (which we'll explain later) but are a first stepping stone towards a proper large-scale experiment.

In Section 8.1 we illustrate an experiment (Experiment #1) to assess the accuracy of the prediction of IBNR claim count of the triangle-free approach and compare it to that of the chain ladder.

In Section 8.2 we illustrate an experiment (Experiment #2) to assess the accuracy of the prediction of the IBNR total amount for the triangle-free approach and compare it to that of the chain ladder. Note that future losses and IBNER are ignored in this experiment, as the IBNR is the crucial element of the overall reserves (and that for which a development triangle is more useful).

In Section 8.3 we illustrate an experiment (Experiment #3) to compare *the whole IBNR distribution* of both the triangle-free approach and the chain ladder against the true distribution, which we know as we use an artificial data set.

Section 8.4 draws an overall comparison of the triangle-free approach and the chain ladder.

8.1 Experiment #1 – Predicting IBNR claim counts

In order to assess the prediction accuracy of the triangle-free approach and compare it to the chain ladder (CL) method we need to generate artificial data for which the true value of the projected claim counts is known. We need to run the experiment for many different values of the parameters (e.g. the mean delay, the underlying Poisson rate (in the case of a Poisson process), the number of years, etc.) and with a very large number of simulations. For the purpose of this paper, we have limited ourselves to looking at a Poisson rate with rate = 100 over 10 years, with different values of the mean reporting delay (1, 2, 3, 4, 5, 10, 20), and with just 100 different simulations.

Each accident year is assumed to start on 1/1 and, for simplicity, it is assumed that we have on record all losses reported by 31/12 of Year 10.

The experiment has been run as follows:

- (1) The underlying process has been chosen to be a Poisson process with rate 100
- (2) The true delay distribution has been chosen to be an exponential distribution with mean delay $\tau = 1, 2, 3, 4, 5, 10, 20$.
- (3) The number of years of loss experience (accident years) has been chosen to be 10.

- (4) A number $N_s = 100$ of simulations has been run, for each of which a random number of losses has been generated for each year based on the assumptions above. The output of each simulation is therefore the “true” number of claims for each accident year $i = 1, \dots, 10$. Fig. 29 shows the first ten simulations.
- (5) A random delay has been assigned to each of the losses generated in Step (4), drawn from an exponential distribution with mean delay τ .

Triangle-free method

- (6) Based on the delays calculated in Step (5), the number of claims reported by 31/12 of Year 10 was simulated, as in Fig. 30.
- (7) For each simulation, the average mean delay is calculated, and based on that and on Equation (22-iii) the projected number of claims is calculated⁷.

NUMBER OF LOSSES FOR EACH SIMULATION										
	Accident year									
	1	2	3	4	5	6	7	8	9	10
Sim 1	92	114	90	94	100	86	82	111	87	93
Sim 2	84	100	115	10	97	129	94	96	89	93
Sim 3	113	96	89	96	91	104	97	107	105	93
Sim 4	112	97	105	99	97	102	111	100	102	100
Sim 5	95	100	98	94	94	83	111	103	98	103
Sim 6	94	91	83	89	97	95	94	89	94	98
Sim 7	110	108	104	84	94	106	107	102	103	117
Sim 8	102	79	119	84	98	95	93	75	105	100
Sim 9	95	103	102	94	110	119	97	102	104	87
Sim 10	87	93	101	94	106	80	89	96	100	110

Figure 29. We generate the number of losses for each simulation and for each accident year, based on a Poisson process with uniform rate equal to 100. This table shows the first ten simulations.

NUMBER OF LOSSES FOR EACH SIMULATION										
	Accident year									
	1	2	3	4	5	6	7	8	9	10
Sim 1	9	6	8	14	5	4	6	4	3	0
Sim 2	3	5	9	8	6	5	5	4	4	2
Sim 3	8	5	9	4	5	3	2	3	1	1
Sim 4	7	8	11	3	5	5	4	5	4	0
Sim 5	6	4	5	6	6	7	1	6	2	1
Sim 6	15	8	7	8	5	4	2	3	1	1
Sim 7	11	4	11	8	6	10	6	6	4	1
Sim 8	1	10	7	8	4	6	4	1	2	2
Sim 9	11	11	9	9	9	4	7	6	2	0
Sim 10	3	5	8	12	10	8	6	5	4	2
...

Figure 30. We generate for each simulation and for each accident year the number of reported claims, based on the number of losses generated in Fig. 6 and on random delays drawn from an exponential distribution with mean delay $\tau = 3$.

⁷ Ideally, we should have used the empirical distribution instead of the exponential model assumption. By using a model for the delay distribution method and a model-free approach for the chain ladder method we create a disparity. However, this experimental set-up is much simpler and this is the reason why it was adopted in

- (8) For each simulation, the tail factor based on the exponential delays assumption is calculated, as in Equation (22-v), correcting the estimate produced in Step (7). As we'll see later, this can be based on a client's own delay or on market assumptions.
- (9) For each simulation j , the projected number of losses, $\hat{\mu}_a^{TF,j}$, is compared with the *true* number of losses μ_a^j (which we know because we have an artificial data set) and the prediction error is calculated as the square root of the mean squared error:

$$\text{Prediction error (TF)} := \sqrt{\text{MSE}} = \sqrt{\frac{1}{N_s} \sum_{j=1}^{N_s} (\hat{\mu}_a^{TF,j} - \mu_a^j)^2} \tag{28}$$

Chain ladder method

- (10) Based again on the random delays generated in Step (5), we create for each simulation a reported triangle, i.e. a triangular matrix which gives for each pair (AY, DY) the number of claims occurred in accident year AY and reported within DY years.
- (11) For each such triangle, we calculate the development factors from DY to DY + 1, for each acceptable value of DY
- (12) As for the tail factor, since there is no standard way of calculating the tail factor for the chain ladder method, we simply adopt the same tail factor as in Step (8). This is basically the same as not considering the tail factor at all in the comparison of the two methods, and should be to the advantage of the chain ladder method. (Alternatively, one could easily produce a set-up in which one pre-agreed particular method to calculate the tail factor for the chain ladder method is used.)
- (13) The projected number of claims is calculated for each simulation and for each accident year of each simulation, producing a new estimate $\hat{\mu}_a^{CL,j}$.
- (14) As for the delay distribution method, we can calculate the prediction error for the chain ladder method by comparing the projected values of claim counts using the mean squared error:

$$\text{Prediction error (CL)} := \sqrt{\text{MSE}} = \sqrt{\frac{1}{N_s} \sum_{j=1}^{N_s} (\hat{\mu}_a^{CL,j} - \mu_a^j)^2} \tag{29}$$

Comparison

- (15) The prediction error for the chain ladder method and the triangle-free approach has been compared for different values of τ both in the case where the tail factor has been estimated based on a single simulation or with a market assumption (in our case, using the average over all simulations), yielding the comparison tables below.

The results shown in Fig. 31 suggest that using the triangle-free method significantly increases the accuracy of the projection to ultimate claim count. Several elements drive this difference, two fundamental and the others accidental:

- a. The first fundamental element is that the triangle-free method doesn't need to segment the experience into accident years and can therefore project the full period at the same time, thus reducing the instability of the projection of the more recent years;

this first version of the paper. Hopefully it will be improved in the future, either by considering the empirical distribution for the delay distribution method or by using a chain ladder method enhanced with some prior knowledge on how the development factors decay towards 1.

USING CALCULATED TAIL FACTOR				USING MARKET ASSUMPTION FOR THE TAIL			
Average delay [y]	TF mean squared error	CL mean squared error	Error reduction	Average delay [y]	TF mean squared error	CL mean squared error	Error reduction
1	0.38	0.55	31%	1	3.83	5.48	30%
2	6.93	10.89	36%	2	5.65	9.58	41%
3	10.14	14.86	32%	3	9.82	13.56	28%
4	15.83	21.96	28%	4	11.15	17.33	36%
5	22.51	28.94	22%	5	14.81	19.24	23%
10	156.11	177.44	12%	10	22.14	32.86	33%
20	148.12	188.56	21%	20	26.99	58.34	54%

Figure 31. A comparison between the delay distribution (TF) method and the chain ladder (CL) method in terms of the error reduction achieved by the TF method, under different assumptions on the tail.

- b. The second fundamental element is that the triangle-free method uses more granular information, whilst the CL method condenses all information into a triangle, with significant loss of information
- c. The accidental element is that in this specific experimental set-up we are comparing a model-driven method (the delay distribution with an exponential assumption) with a fully empirical method (the chain ladder) – note that this is indeed accidental as the exponential model has been used only for the sake of simplicity of implementation – the triangle-free method can be used in a fully distribution-free fashion.

It is likely that elements (a) and (b) are the most important quantitatively, but until a different experimental set-up is used (based on the empirical version of the delay distribution method) this cannot be asserted conclusively.

8.2 Experiment #2 – Predicting IBNR total losses

Experiment #1 has demonstrated that the triangle-free approach is more accurate than the chain ladder in predicting the ultimate IBNR claim count. What we really want to know, however, is whether it is also better at predicting the IBNR total losses, which are a crucial part of the overall reserves.

In order to test this, we need to build up on the experimental set-up of Experiment #1 and introduce severities.

In order not to be distracted by useless complications, we have made the assumption that IBNER is negligible, and that we are not interested in future losses – this allows us to focus on the most important point, which is that of the IBNR losses. Also, we have assumed that there is no claims inflation.

This second experiment has been run as follows:

- (1) The underlying frequency and delay model have been chosen as in Experiment #1:
 - a. The claim count process is Poisson with rate 100
 - b. The delay distribution is exponential with mean delay $\tau = 3$ years.
- (2) As in Experiment #1, the number of years of loss experience (accident years) has been chosen to be 10, and the number of simulation is still $N_s = 100$. The random number of losses for each simulation is exactly the same as in Experiment #1 (see Fig. 6).

- (3) For each of the losses generated as in Step 2, a loss amount has been sampled from a lognormal distribution with parameters $\mu = 9.52$, $\sigma = 1.70$ (the choice of these parameters is purely incidental).
- (4) The random delay for each of the losses is the same as that in Experiment #1, and therefore the number of reported losses is also the same (see Fig. 7).

Triangle-free method

- (5) For each simulation, the projected number of losses of Experiment #1 is used, inclusive of the tail factor. The “market” tail factor has been used. Let $\hat{n}_{IBNR}^{TF,j}$ be the estimated number of IBNR losses.
- (6) For each simulation j , the projected total loss amount, $\hat{S}_a^{TF,j}$, is calculated with two different methods.
 - a. The first, “triangle-free (empirical)”, is by sampling $\hat{n}_{IBNR}^{TF,j}$ values from the set of losses generated for sample j (that is, no further attempt at modelling is made).
 - b. The second, “triangle-free (model)” first calculates the parameters μ_j, σ_j of the lognormal distribution based on the reported losses, and then samples $\hat{n}_{IBNR}^{TF,j}$ from that distribution.

The reason why we used two different methods, one implying modelling and the other not, is because we wanted to estimate the impact of using the knowledge on the true model on our prediction ability.

- (7) The projected total loss amount, $\hat{S}_a^{TF,j}$, is then compared with the *true* total loss amount S_a^j (which we know because we have an artificial data set) and the prediction error is calculated as the square root of the mean squared error:

$$\text{Prediction error (TF)} := \sqrt{MSE} = \sqrt{\frac{1}{N_s} \sum_{j=1}^{N_s} (\hat{S}_a^{TF,j} - S_a^j)^2} \tag{30}$$

Chain ladder method

- (8) Based again on the random delays of Experiment #1, we create for each simulation a reported triangle, much in the same way as we did in Experiment #1 for the number of claims, but this time for the total loss amount. Based on this, we can produce a new estimate $\hat{S}_a^{CL,j}$ of the total losses.
- (9) As for the delay distribution method, we can calculate the prediction error for the chain ladder method by comparing the projected values of claim counts using the mean squared error:

$$\text{Prediction error (CL)} := \sqrt{MSE} = \sqrt{\frac{1}{N_s} \sum_{j=1}^{N_s} (\hat{S}_a^{CL,j} - S_a^j)^2} \tag{31}$$

Comparison

- (10) The prediction error for the chain ladder method and the triangle-free approach has been compared, yielding the comparison table below.

As in Experiment #1, the results show that using the triangle-free method significantly increases the accuracy of the loss projection. The reasons that drive this difference are roughly the same as those listed in Experiment #1. However, in this case, notice that we have tried to quantify the impact of using our prior knowledge on what the correct severity distribution is by also producing an “empirical” estimate of the projected losses that doesn’t use any severity modelling.

8.3 Experiment #3 – Predicting the full IBNR distribution

The third experiment that we have carried out aims at calculating the accuracy not only of the mean projected IBNR for different methods, but the full distribution of IBNR losses. This allows us to determine how fit each reserving method is to assess the various percentiles of the reserving distribution (i.e., the distribution of past liabilities).

Ideally, we should compare the *true* distribution with the triangle-free approach distribution and the chain ladder distribution for a large number of simulations. However, since the process of obtaining a chain ladder estimate of the reserving distribution is currently not fully automated, a full experiment has not been performed yet.

What we've done as an interim solution was to choose one of the simulations produced in Experiment #2 (which has all the information we need) such that the projected ultimate for the triangle-free approach and the chain ladder approach are reasonably close, so that the difference between the two distributions does not simply depend on the fact that the projected values are quite different. After all, Experiment #2 has already taken care of the difference in the point estimates between the two methods.

In order for this experiment to make sense, we have to make sure that we define carefully what the “true distribution”, the “chain ladder distribution”, and the “triangle-free distribution” are in this experiment.

- a. For the **chain ladder distribution**, we used the lognormal approximation and the normal approximation of the reserving distribution obtained by the application of Mack's method, for which the distribution of the ultimate is the lognormal (or the normal) distribution such that the mean is equal to the mean projected liabilities, and the standard deviation is equal to the prediction error⁸, which takes into account both the process variance *and* the parameter uncertainty (see, e.g., Mack, 1993 and England & Verrall, 2002).
- b. For the **triangle-free distribution**, the IBNR distribution was simulated by combining a negative binomial distribution with rate equal to the projected number of IBNR claims, and variance-to-mean ratio equal to 2 (in practice, a Poisson distribution with some extra volatility to take parameter uncertainty into account), and a lognormal distribution whose parameters are calculated based on the reported losses only.

The parameter uncertainty on the parameters of the lognormal distribution was calculated based on the standard formulae for the maximum likelihood estimation, which state that (asymptotically at least) the estimated parameters are distributed normally around the true parameters as (note that the correlation matrix is diagonal so we don't need to worry about the interaction between the two parameters):

$$\hat{\mu} \sim N\left(\mu, \frac{\sigma^2}{n}\right), \hat{\sigma} \sim N\left(\sigma, \frac{\sigma^2}{2n}\right) \quad (32)$$

where n is the number of reported losses. In order to take the parameter uncertainty for the severity into account, for each simulated loss we first simulate a pair of values $(\hat{\mu}, \hat{\sigma})$ using Equation (32)

⁸ In the lognormal case, we have calculated it on two bases: (i) take the distribution of the *total ultimate losses* to be a lognormal with mean equal to the mean projected liabilities, and standard deviation equal to the prediction error, then calculate IBNR as the total losses minus the total reported claims, which in this case are assumed to have no IBNR; (ii) take the distribution of the *IBNR losses* to be a lognormal with mean equal to the mean projected IBNR, and standard deviation equal to the prediction error.

(after replacing the true parameters with the estimated ones) and then we sample a loss from a lognormal with parameters $(\hat{\mu}, \hat{\sigma})$. This has the double effect of increasing the volatility of results (which we want, because it reflects the additional uncertainty) and also of adding a systematic bias on the mean of the distribution, which we don't necessarily want. This effect becomes significant when the number of losses used to calculate the parameters is small, and one way to compensate for that is to rescale all losses so that the mean is preserved, and then sampling from the new set.

- c. As for the **true distribution**, this is subtler than it looks. Since we know the theoretical Poisson rate and the expected number of losses reported by the end of the tenth year, we know the theoretical number of IBNR losses:

$$\begin{aligned}
 (\text{\#IBNR losses})_{\text{theor}} &= \frac{\tau}{T}(1 - \exp(-T/\tau))(\text{total \# of claims})_{\text{theor}} \\
 &= \frac{3}{10} \times (1 - \exp(-10/3)) \times 1000 \\
 &= 289.30
 \end{aligned}
 \tag{33}$$

Also, we know the theoretical value of the lognormal parameters: $\mu = 9.52, \sigma = 1.70$. However, to make a fair comparison with the distributions calculated in one of the two methods above, we need to take parameter uncertainty into account *as if we didn't know the correct values of the parameter*, exactly in the same way as we do for the triangle-free approach, assuming a negative binomial distribution with rate equal to 289.30 and variance-to-mean ratio equal to 2. The “fictional” parameter uncertainty on the parameters μ, σ can be calculated via Equation (32) using $n = 1,000 - 289.30 = 710.70$.

The results of the comparison are shown in Fig. 33. The comparison has been carried out:

- by visually comparing the CDF of the aggregate loss distributions for the different methods, which shows that the triangle-free approach produces a far closer fit to the true distribution – in the figure it is actually rather hard to tell them apart;
- more formally, by calculating the normalised Kolmogorov-Smirnov distance between each of the tested methods and the true distribution. This confirms that the triangle-free approach provides a much better fit to the true distribution. However, notice how the value of the normalised KS distance ($d = 2.3$) is still quite large (corresponding to a p -value below 1%).

8.4 A comparison of the triangle-free approach and the chain ladder

In this section we compare the triangle-free approach and the chain ladder in general and not only in terms of prediction accuracy but more in general.

Advantages of the triangle-free approach

1. Increased accuracy in the projected IBNR, as demonstrated by the tables in Fig. 31 and Fig. 32.
2. More accurate estimation of reserving uncertainty, as shown in Fig. 33.

This is probably the main reason why we should move away from loss triangulation methods.

3. You don't need a complete last diagonal in order to apply the method⁹.

⁹ This is often a problem in pricing because the loss data set includes all losses reporting up to a given date and this date is normally a few months away from renewal, and therefore the last year is never complete.

					Chain ladder	Triangle-free (empirical)	Triangle-free (model)
Prediction accuracy (MSE)					7,392,262	4,423,237	3,842,561
Prediction accuracy (MSE) as a percentage of average true value					44.7%	26.8%	23.3%
Simul	IBNR (true)	IBNR (Chain ladder)	IBNR (Triangle-free, empirical)	IBNR (Triangle-free, model)	Error (Chain ladder)	Error (Triangle-free, empirical)	Error (Triangle-free model)
1	10,773,575	19,477,811	15,443,006	16,104,318	8,704,236	4,669,431	5,330,743
2	17,839,076	11,163,722	13,864,553	14,578,474 -	6,675,354 -	3,974,522 -	3,260,601
3	20,073,695	17,343,084	15,841,802	15,691,210 -	2,730,611 -	4,231,893 -	4,382,485
4	16,519,312	17,671,032	22,341,455	22,659,106	1,151,719	5,822,143	6,139,794
5	15,622,367	19,756,994	15,493,358	15,027,101	4,134,627 -	129,009 -	595,265
6	16,807,437	9,557,564	14,331,857	12,685,124 -	7,249,873 -	2,475,580 -	4,122,313
7	16,309,016	26,602,474	19,518,622	19,587,276	10,293,459	3,209,606	3,278,260
8	21,486,134	15,839,159	20,258,625	22,039,740 -	5,646,975 -	1,227,509	553,606
9	11,863,136	16,661,276	15,250,191	16,031,506	4,798,140	3,387,055	4,168,371
10	14,938,406	11,782,679	13,048,958	13,355,460 -	3,155,727 -	1,889,448 -	1,582,946
...

Figure 32. Prediction error (in GBP) calculated as the means squared error between the true value and the projected value for 100 different random data sets. The underlying model was in all cases a compound Poisson distribution ($\lambda = 100$) with a lognormal severity distribution ($\mu = 9.52$, $\sigma = 1.70$). The delay distribution was assumed to be exponential with an average delay $\tau = 3$ years. For simplicity’s sake (although this is not a crucial assumption) the underlying loss model has been assumed to be inflation-free and the reserving process has been assumed to be IBNER-free. The first ten simulations are included for illustration purposes.

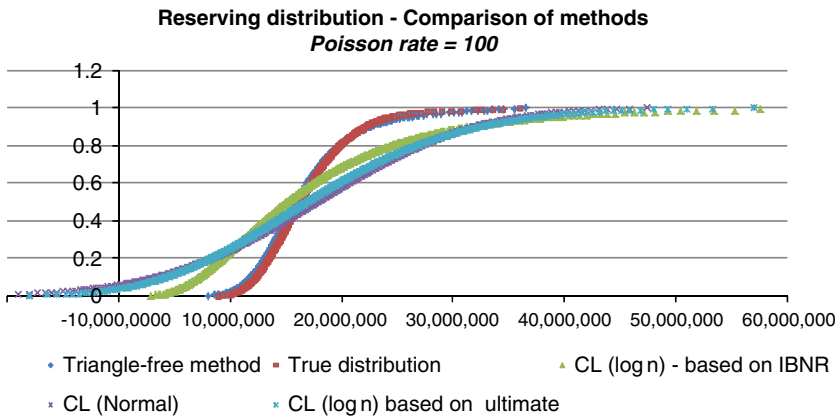


Figure 33. A comparison of the reserving distribution as calculated by the triangle-free method and the chain ladder (CL) method, both of which include a provision for parameter uncertainty. The results are compared to the “true” distribution – that in which an oracle has told us what the underlying parameters used to generate the loss data are, but also (paradoxically, it may seem) includes a provision for the parameter uncertainty. The triangle-free distribution shows a far closer fit to the true distribution (the two distributions are almost indistinguishable on the chart), and this is reflected in the lower value of the normalised Kolmogorov-Smirnov distance.

In the chain ladder method, if you have, e.g., only 7 months developed of the last accident year, you’ll have to discard the last diagonal, or find a way to gross it up, or change the discretisation window – e.g., use one month instead of one year (although this may cause other problems, such as numerical instability).

4. If a given year has no losses, this is a problem in the pure chain ladder method, because the projected losses will be zero (this is not necessarily the case with other methods based on the development of loss triangles), whereas it will affect less the performance of the triangle-free method, which relies on the number of claims reported over the *whole* period. There is therefore less need to resort to credibility-like methods such as Bornhuetter-Fergusson or Cape Cod.
5. The triangle-free approach allows the estimation of the tail factor in a more robust and scientific way, based again on a larger data set than in the case of chain ladder, where the only option is really to introduce a tail factor based on judgment or try to fit a function (e.g. exponential, power law, etc) through a very small number of development factors, as the best reserving packages do.

Disadvantages of the triangle-free approach

1. The main disadvantage is that the triangle-free approach is more complex than the chain ladder method...
2. ...and has more stringent data requirements: input data must be both more granular (individual loss claims) and accurate (e.g., the exact date of occurrence and of reporting for each loss).
3. Another disadvantage of the triangle-free approach is presentational: while all methods based on triangle development allow you to see at a glance in a simple matrix of size, say, 10 by 10, the pattern by which losses from each accident year develop, the triangle-free approach has no such snapshot view of total loss development.

9 Further Research

This initial research shows that what we called the “triangle-free approach to reserving”, which uses the full information in the loss data set rather than a compressed version of it, leads to a more accurate estimate of the projected losses, and especially of the uncertainty around this projection. This has been proved in a series of controlled experiments with an artificial data set.

As mentioned in the executive summary, some improvements to the experimental design are however needed to confirm these initial results. The main improvements are the following:

- The scale of the experiment must be enlarged: currently, we have built our considerations on 100 different data sets, but this number should be increased to 1,000 or even 10,000
- The experiment should be run with many different Poisson rates. Currently we have used a Poisson rate of 100, but a wide range of rates (say between 0.1 to 1,000) should be tried. Also, experiments with overdispersed Poisson distributions, e.g. with the negative binomial, should be carried out, to see if this leads to different results.
- The experiments should be run with many different mean delays. Currently we have run Experiment #1 with a good variety of mean delays, but Experiments #2 and #3 have been limited to an exponential distribution with a mean delay of 3 years. Also, other distributions other than the exponential distribution should be used to model reporting delays.
- The experiments should be run with different models for the severity distribution and with different parameters. The most crucial thing to try is probably to see the effect of having an increasingly heavy tail, by using a GPD with different values of the parameter ξ to generate losses above a certain threshold.
- Regardless of the way in which losses are generated, it is important to analyse the prediction accuracy of the version of the method in which no attempt is made at modelling the delay

distribution, but the empirical version is used (except for the tail, where we definitely need some model). This is the current implementation used for client data analysis, but it is slightly more difficult to automate it to the point of running a large-scale experiment, and so it hasn't been used here.

- The triangle-free approach should be compared not only to the chain ladder method but to a number of different, more sophisticated triangle-based approaches, such as BF, Cape Cod, fitting GLM to loss development triangles. This is needed to give more weight to the claim that the triangle-free approach actually gives better results (at least in estimating the reserving distribution if not the point estimate) than *any* method for which information has been compressed into a triangle.
- IBNER should also be considered in the comparison by introducing a model for producing artificial data with reserve errors.

Acknowledgements

I would like to thank my colleagues Andy Smyth, Arun Kurian, Chris Gingell, Claire Wilkinson, David Stebbing, Eamonn McMurrugh, James Coyle, Jen Ning Tan, Jeremy Brooks, Kyle Roper, Paolo Albini, Rajeev Aravind, Shubhanjali Gupta, for valuable help and feedback on this work, often through its use in their real-world projects. Also, I would like to thank Di Kuang, Frank Cuyper, James Orr, Jeremy Affolter, Joseph Lo, Kathryn Morgan, Maria Miranda, Michael Fackler, Robert Scarth and Xiaohan Fang for very helpful discussion and for their interest in this work.

References

- Antonio, K. & Plat, R. (2012). *Micro-Level Stochastic Loss Reserving*. KU Leuven, AFI-1270. Available at SSRN: <http://ssrn.com/abstract=2111134> or <http://dx.doi.org/10.2139/ssrn.2111134>
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.
- England, P.D. & Verrall, R.J. (2002). Stochastic claims reserving in general insurance. *British Actuarial Journal*, 8/3, 443–518, 2002.
- Graham, M. (2011). *The great 99.5th percentile swindle*, GIRO 2011.
- Guiahi, F. (1986). *A Probabilistic Model for IBNR Claims*, CAS Proceedings, Volume LXXIII, Number 139, 1986.
- Kaminsky, K.S. (1987). Prediction of IBNR claim counts by modelling the distribution of report lags. *Insurance: Mathematics and Economics*, 6, 151–159.
- Klugman, S.A., Panjer, H.H. & Willmot, G.E. (2008). *Loss Models: From Data to Decisions*, Third Edition. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2008.
- Kozubowski, T.J. & Podgorski, K. (2009). Distributional properties of the negative binomial Levy process. *Probability and Mathematical Statistics*, Vol. 29, Fasc. 1, pp. 43–71.
- Mack, T. (1993). Distribution-free calculation of the standard error of chain ladder reserves estimates. *ASTIN Bulletin*, 23/2, 213–225.
- Norberg, R. (1993). Prediction of outstanding liabilities in non-life insurance. *ASTIN Bulletin*, 23(1), 95–115.
- Norberg, R. (1999). Prediction of outstanding claims. II Model variations and extensions. *ASTIN Bulletin*, 29(1), 5–25, 1999.
- Parodi, P. & Bonche, S. (2010). Uncertainty-based credibility and its applications. *Variance*, 4(1), 18–29.
- Parodi, P. (2012a). Computational intelligence with applications to general insurance: a review. *Annals of Actuarial Science*, 6/2, 307–343.

- Parodi, P. (2012b). *Triangle-free reserving: A non-traditional protocol for estimating reserves and reserve uncertainty*. Proceedings of GIRO 2012. <http://www.actuaries.org.uk/research-and-resources/documents/brian-hey-prize-paper-submission-triangle-free-reserving>
- Ross, S.M. (2003). *Introduction to probability models*, 8th Edition. Academic Press, 2003.
- Taylor, G., McGuire, G. & Sullivan, J. (2008). Individual claim loss reserving conditioned by case estimates. *Annals of Actuarial Science*, 3, 215–256.
- Weissner, E.W. (1978). *Estimation of the Distribution of Report Lags by the Method of Maximum Likelihood*, PCAS LXV, 1978. p. I.