Title: A Simple Multi-State Reserving Model

Topic: 3: Liability Risk – Reserve Models

Name: Orr, James

Organisation: Towers Perrin – Tillinghast

Address: 71 High Holborn, London WC1V 6TH

Telephone: +44 (0)20 7170 2729

Fax: +44 (0)20 7170 2222

E-mail: james.orr@towerperrin.com

Abstract:

Motivation. To explore how a simple, common process may underlie the development of claims arising from a portfolio of insurance contracts.

Method. This paper presents a simplified claims number model, with a Poisson arrival process for losses occurring during each Accident Year and Exponential waiting times in two intermediate states, being State 0 for Incurred But Not Reported (IBNR) and State 1 for Reported But Not Settled (RBNS) claims; paid losses are absorbed into State 2. Aggregate claim number development data is simulated using this model and standard reserving methods are applied, using the 'reserving' package in R to derive estimates of the ultimate claim numbers. The model itself is then fitted to the simulated data, using least squares and Bayesian approaches. Finally, extensions of the model to fit different real world circumstance are presented.

Results. Simulation from the simplified claims number model is shown to generate plausible data, to which existing reserving techniques may be applied. Alternative estimation approaches using least squares and Bayesian approaches are shown to produce similar results to the existing techniques. The model is also shown to be readily extendable to encompass a number of different circumstances that arise in practice, including inception (i.e. year of account) based accounting, catastrophes and the "negative" development of incurred claims.

Conclusions. Whilst not proven in a real world context, this model shows potential as an alternative basis for the study and estimation of insurance claims development. Its strengths include the incorporation of both paid and outstanding claims data within the estimation process and its ready expression in a Bayesian framework.

Keywords. Insurance claims reserving, claims number model, multi-state model, Markov Chain, Bayesian.

Acknowledgements:

This paper has been long in the gestation and I have spoken to a great many people about the ideas within it. I would like to record my particular thanks to Richard Verrall for his guidance at the start of my research, to Trevor Maynard for his encouragement, and to Markus Gesmann for first creating and then introducing me to the 'reserving' package in R.

1. Introduction - Motivation

Things should be made as simple as possible, but not any simpler.

Albert Einstein

The emerging development of claims from a portfolio of insurance contracts is complex. A portfolio may contain a range of different contracts, even within the same business line, with different terms and conditions, and periods of cover. Furthermore, individual contracts may generate a range of different claims, in terms of their size, occurrence, emergence and settlement characteristics. A "simple" motor book, for instance, may generate a large number of property damage claims, which can be processed in relative good order, but also a small number of expensive and potentially drawn-out liability claims.

Established reserving techniques seek to address the challenge of predicting future claims development through a top-down approach. These methods typically consider the overall development of the portfolio in previous years and assume that a similar development profile will apply in the future. In practice, adjustments will be made for known changes over time or special features of the historical data, such as different levels of exposure, claims inflation, premium rating strength, catastrophes and large losses, to name but a few.

It is not the purpose of this paper to challenge the hegemony of the established reserving techniques. They are in widespread, near universal, use within the general insurance industry, representing a useful standard of practice. The simplest technique, the "chain ladder", was developed in a pre-computer age and can be applied without specialist software. Furthermore, the latest stochastic reserving methods have built upon the basic concepts of development factor modelling, expressing anticipated future development as a proportion of observed development to date.

The ubiquity of development factor methods reflects their practicality, which can make them appear simple. However, many assumptions are implicit in their application. Moreover, particular knowledge of the underlying business and the application of judgement based on experience are needed before a "best estimate" can be arrived at that an actuary is prepared to put their name to. These deceptively simple tools must be used with skill.

This paper seeks to explore the idea that the claims development of an insurance portfolio, often presented as a "flattened s-curve" of incurred (reported) or paid claims over time, results from a common underlying process. The aim is to express that process in its simplest form, derive results for the expected development of future claims, simulate data from the resulting model and explore its use in predicting future claims development. The paper will focus on modelling claims numbers only, although a brief discussion of extending the model to claims amounts has been included.

If successful, this paper will provide some building blocks that contribute towards an alternative, claims level, approach to reserve estimation. Otherwise, it is simply a side-piece that explores how actuarial techniques developed in the fields of mortality and morbidity might be applied to an established and challenging problem in general insurance!

2. Formulation

The idea of modelling claims development using a multi-state approach has been discussed previously by Hachemeister (1980), Norberg (1993), and Hesselager (1994), and was recommended as an approach in the report of the GIRO "Cycle Survival Kit Working Party" by Line, et al (2003).

All of the above authors have presented or envisaged a general approach that would encompass the real world circumstance of claims development. However, our aim here is to look at a multi-state claims model in its purest (or simplest) form, in the hope of capturing the essence of the claims development process. For this purpose we have taken a "stripped down" version of the 3-state model presented by Hesselager, thus:

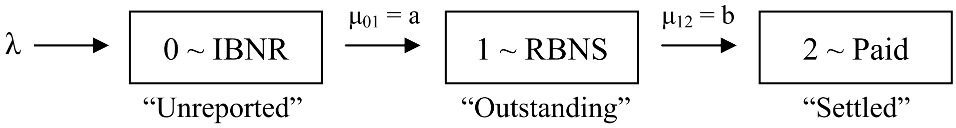

Here, losses from an insurance portfolio are assumed to arrive in State 0 as IBNR or "Unreported" losses. These losses are then reported in State 1 as RBNS or "Outstanding" claims, which are subsequently "Settled" in State 2 as "Paid" claims. It is common to refer to the sum of Outstanding and Paid claims as "Incurred" claims, although strictly this should be "Incurred and Reported" claims.

In this paper, we shall restrict the assumptions in the original model presented by Hesselager. First, we are only going to consider numbers of losses and claims within each of the three states over time. Secondly, losses will be assumed to arrive as a Poisson process with a constant arrival rate λ during each Accident Year. Thirdly, the instantaneous probabilities of movement between states, of individual claims, will be assumed constant, paying no regard to the numbers of claims in any state. This final assumption ignores the effects of resource constraints, as might be experienced during a major catastrophe where claims adjustors are in limited supply or unable to reach affected areas.

In the examples and illustrations that follow, we shall assume that $\lambda = 100$, as illustrated below.

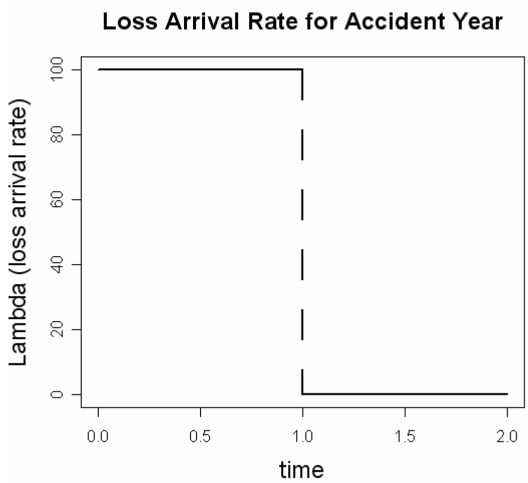

Therefore, the total number of losses in an Accident Year will be distributed as a Poisson random variable, with an expected value of 100 and a standard deviation of 10.

The instantaneous transition rates can, borrowing from mortality and morbidity theory, be defined as:

$$\mu_{01} = \lim_{h \to 0} \frac{\Pr(Claim\ in\ State\ 1\ at\ t+h \mid Loss\ in\ State\ 0\ at\ t)}{h}$$

$$= \text{“Force of Reporting”} = a$$

and

$$\mu_{12} = \lim_{h \to 0} \frac{\Pr(Claim\ in\ State\ 2\ at\ t+h \mid Claim\ in\ State\ 1\ at\ t)}{h}$$

$$= \text{“Force of Settlement”} = b$$

Again, from mortality theory, the waiting time in State 0, under “Force of Reporting” a will follow an Exponential distribution, with a parameter value of a. The expected waiting time is therefore $1/a$ and a higher force will lead to a shorter stay. In particular, a value of a that is greater than 1 will lead to an average waiting time of less than one year. Similar results follow for waiting times in State 1 and parameter b.

In what follows we shall assume that (a, b) takes one of two possible values:

for “short tail”: $(a = 4.00, b = 2.00) \Rightarrow$ average waiting times (3mths, 6mths)

for “long tail” : $(a = 0.40, b = 0.25) \Rightarrow$ average waiting times (2.5yrs, 4yrs)

A potential advantage of this formulation is that the transition rate assumptions can be presented with a real world meaning, in this case the average reporting and settlement delays. Indeed, we might expect to see relatively straightforward “short tail” business exhibit Exponential waiting times, for a homogeneous group of claims, and be able to monitor their average length as an input to the reserve estimation process.

However, more involved “long tail” claims would be expected to pass through a number of stages (for example, including assessment, discovery and trial) and a more plausible total waiting time might be the sum of waiting periods in different states. However, for the purpose of this paper, much simpler assumptions have been used and the “long tail” example is presented to illustrate the different methods discussed, as its development covers a number of years, as is common for many reserve estimation problems.

Finally, it is necessary to define the following notation:

$N(s, y, t)$ = Number of losses (or claims) in State “*s*” arising from Accident Year “*y*” as at Reporting Time “*t*”.

In practice, only data on the number of claims in States 1 and 2 will be observed, and then only up to the latest reporting period. The challenge will be to estimate the number of unreported (and therefore unobserved) losses in State 0, corresponding to the IBNR reserve on which a general insurance actuary is often called upon to opine.

3. Derivation

Our objective in this section will be to derive closed-form expressions for the expected numbers of losses and claims in each state at each point in time, *t*, which describes the expected development profile for the model.

To do this, it is first necessary to recognise that the claims development process differs between the Accident Year itself and subsequent "run off" years. During the Accident Year ($t \le 1$), losses can arrive in State 0 and transition (as claims) to States 1 and 2. After the end of the Accident Year ($t > 1$), no further losses can arrive, as they will attach to future Accident Years, but transitions will continue until all the claims have reached State 2 (i.e. have been paid). This "run-off" process is as discussed by Hachemeister (1980).

In what follows, the Accident Year index, *y*, has been dropped for brevity.

<u>During the Accident Year</u>

For $t \le 1$, the expected number of losses/claims in each state is given by:

$$E[N(s,t)] \quad = \quad \lambda t * [\alpha(0,t), \alpha(1,t), \alpha(2,t)]$$

The expression for the expected number of losses in State 0 at time t is determined by considering the instantaneous loss arrival rate λ at time s and the probability of then staying in State 0 (unreported) for the remainder of the period to time t, under the force of reporting a, then integrating across all values of s between 0 and t, thus:

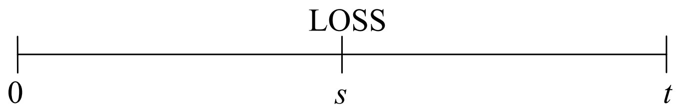

$$\lambda t * \alpha(0,t) \quad = \quad \int_0^t \lambda \exp[-a(t-s)]ds \text{ i.e.}$$

$$= \quad \frac{\lambda}{a}[1-\exp(-at)]$$

Similar arguments can be applied to determine the expected number of claims that have reached State 1 by time t, with losses arriving at time s and being reported at time r:

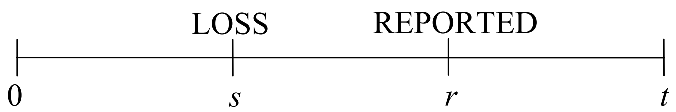

$$\lambda t * \alpha(1,t) \quad = \quad \int_0^t \lambda \int_s^t a \exp[-a(r-s)]\exp[-b(t-r)]drds$$

$$= \quad \frac{\lambda a}{(a-b)b}[1-\exp(-bt)] - \frac{\lambda}{(a-b)}[1-\exp(-at)]$$

Finally, the expected number of claims to have reached State 2 by time *t* is determined by considering losses at time *s*, which are then reported at time *r* and settled at time *q*, thus:

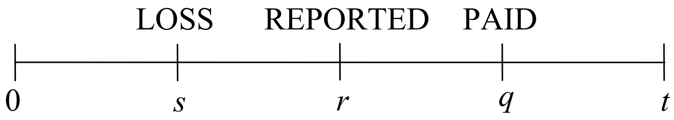

$$\lambda t * \alpha(2,t) = \int_0^t \lambda \int_s^t a \exp[-a(r-s)] \int_r^t b \exp[-b(q-r)] dq dr ds$$

$$= \lambda a \left\{ \frac{t}{a} - \frac{[1-\exp(-at)]}{a^2} - \frac{[1-\exp(-bt)]}{(a-b)b} + \frac{[1-\exp(-at)]}{a(a-b)} \right\}$$

It can be seen that the α 's sum to 1 and that the total expected number of losses/claims at time *t* is λt, as expected given the assumed Poisson arrival process during the Accident Year.

Where $a = b$, the equations for States 1 and 2 include division by zero and alternative expressions, based on the same integrals, can be derived for these special cases. However, in practice, it has been found more convenient to simply add a small adjustment (+0.1%) to *a* to avoid equality.

<u>After the Accident Year</u>

For $t > 1$, no further losses will occur (that would attach to the Accident Year) and the model operates purely as a multi-state Markov Chain, with the following transition intensity matrix.

$$Q = \begin{pmatrix} -a & a & 0 \\ 0 & -b & b \\ 0 & 0 & 0 \end{pmatrix}$$

Applying Kolmogorov's forward equations, the expected claim numbers in each state at time *t* (>1) can be determined as:

$$E[N(s,t)] = \lambda \pi(t)$$

$$= \lambda \pi(1) \exp[Q(t-1)]$$

where,

$$\pi(1) = [\ \alpha(0,1)\ ,\ \alpha(1,1)\ ,\ \alpha(2,1)\ ]$$

can be thought of as the state probability vector at time 1, representing the expected proportion of claims in each state at the end of the Accident Year, and

$$\exp[Q(t-1)] = P(t-1)$$

is the transition probability matrix whose (i,j)th element is the probability of a single loss/claim being in state *i* at time 1 and state *j* at time *t*.

Although it is possible to calculate the exponential of a matrix to a high degree of accuracy, by applying similar arguments to those used above it is possible to show that $P(t-1)$ can be expressed as:

$$\mathrm{P(t\text{-}1)} = \begin{pmatrix} p_{00}(t-1) & p_{01}(t-1) & p_{02}(t-1) \\ 0 & p_{11}(t-1) & p_{12}(t-1) \\ 0 & 0 & 1 \end{pmatrix}$$

where,

$$p_{00}(t-1) = \exp[-a(t-1)]$$

$$p_{01}(t-1) = \frac{a}{a-b}\{\exp[-b(t-1)] - \exp[-a(t-1)]\}$$

$$p_{02}(t-1) = 1 - \exp[-a(t-1)] - \frac{a}{a-b}\{\exp[-b(t-1)] - \exp[-a(t-1)]\}$$

$$p_{11}(t-1) = \exp[-b(t-1)]$$

$$p_{12}(t-1) = 1 - \exp[-b(t-1)]$$

The above formulae can be applied to derive the expected claims development profiles (i.e. expected numbers of claims in each state) for all value of *t*. The expected development profiles for the two example models ("short tail" and "long tail") are illustrated in the following graphs, with the average reporting and settlement delays shown in brackets.

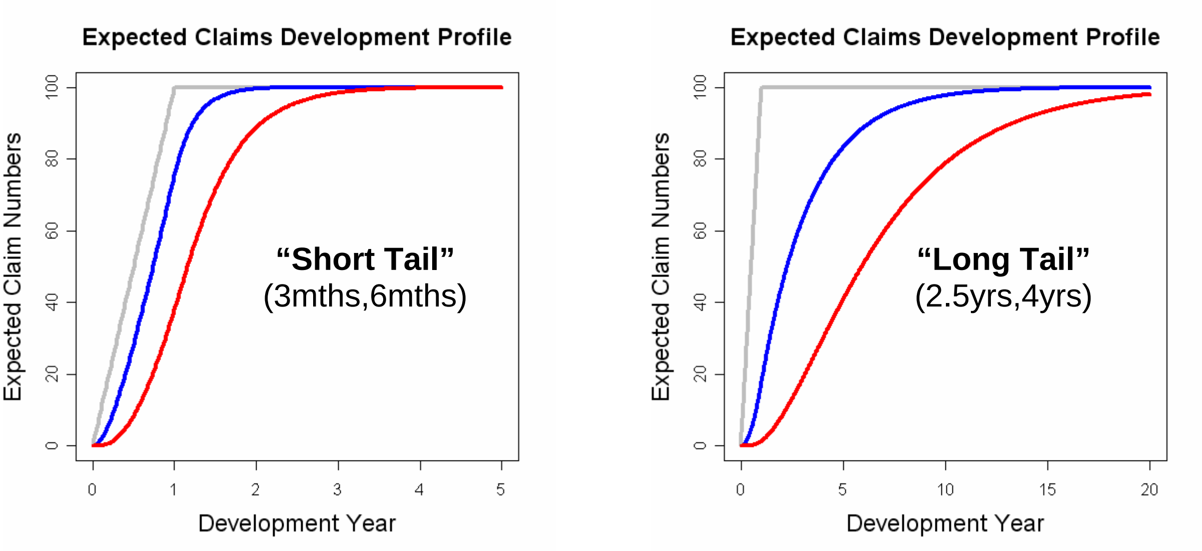

The above show the familiar claims development graphs for paid (i.e. State 2, shown in red) and incurred claims, which comprise the sum of paid and outstanding claims (i.e. States 1 plus 2, shown in blue). However, the graphs are unusual in also showing the underlying expected loss count (i.e. States 0 plus 1 plus 2, shown in grey), rising from 0 to 100 during the Accident Year; a quantity that can only be known in retrospect through examining loss occurrence dates from claims records, once they have all been reported.

Distribution of the Claims Counts

From the above, it is possible to consider the development of claims over time as a "partitioning" of the total loss count between the three states (i.e. some losses will not yet be reported and some claims will not yet be settled, but the remainder will), with different probabilities of being in each state at different points in time, *t*.

Under these conditions, it is a standard result that the unconditional number of losses/claims in each of the three states will follow independent Poisson distributions, with parameters:

$\lambda t * [\alpha(0,t), \alpha(1,t), \alpha(2,t)]$ for $t \leq 1$, and

$\lambda \pi(t)$ for $t > 1$.

Furthermore, *conditional* upon the total number of losses at time *t*, say N_y for Accident Year *y*, the numbers in each state will follow a Multinomial distribution with N_y trials and probabilities of being in each state given by the above expressions divided by λt and λ respectively. That is:

$[\alpha(0,t), \alpha(1,t), \alpha(2,t)]$ for $t \leq 1$, and

$\pi(t)$ for $t > 1$.

This result is central to the Bayesian approach to estimating N_y, which is presented in Section 6.

4. Simulation

Given our simple claims development model, it is a relatively straightforward matter to simulate data from that model, to which we might apply standard reserving methods and other estimation techniques based upon our model.

To do this, we need to simulate a series of loss, claim and settlement times, as described by $\{T_{s,i}\}$ where:

$T_{s,i}$ = arrival time of ith loss/claim in State s

so

$T_{0,i}$ = occurrence time for ith loss

$$= \sum_{j=1}^{i} LossInterval_j \quad \text{where } LossInterval_j \sim iid\ Exponential(\lambda)$$

$T_{1,i}$ = reporting time for ith claim

$$= T_{0,i} + R_i \quad \text{where } R_i \sim iid\ Exponential(a)$$

$T_{2,i}$ = settlement time for ith claim

$$= T_{1,i} + S_i \quad \text{where } S_i \sim iid\ Exponential(b)$$

Only losses that occurred during a particular Accident Year y will attach to it, and this can be determined by selecting the losses such that $T_{0,i} \in (y, y+1]$. The graph below shows a simulated incurred (blue) and paid (red) claims development profile up to Development Year 20 for Accident Year 1, using the above algorithm with the "long tail" model assumptions.

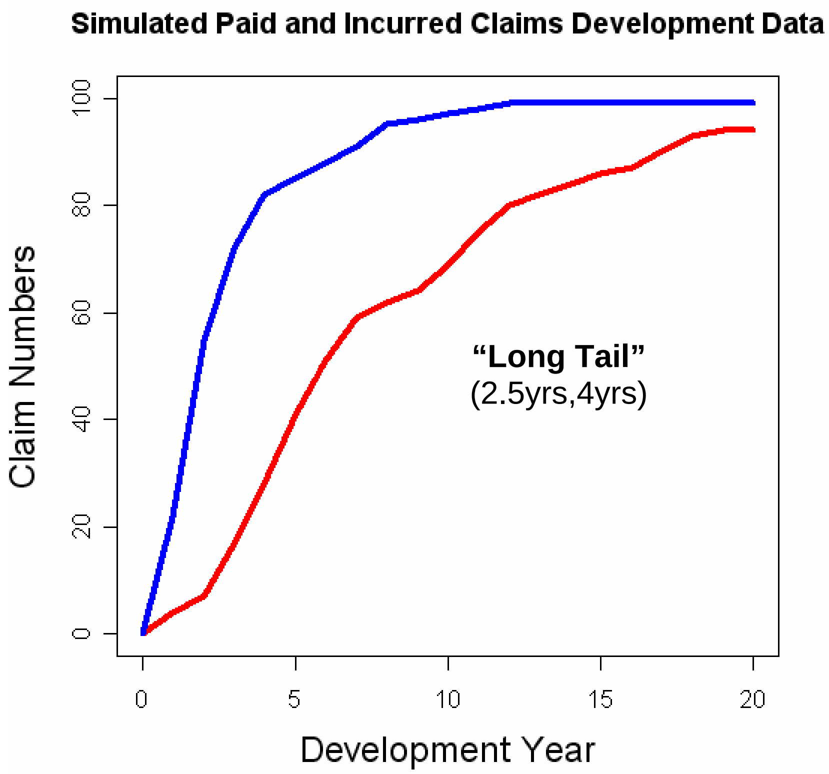

We have simulated full development data for 10 example Accident Years, to ultimate. The data triangles for paid and incurred claim counts that would be available at the end of Accident Year 10 are shown below, along with the ultimate claims counts.

Simulated Cumulative Incurred Claims Counts (limited to 10 years' data) plus Ultimates:

Accident Year	Development Year 1	2	3	4	5	6	7	8	9	10	Ult
1	22	55	72	82	85	88	91	95	96	97	99
2	22	57	80	91	103	108	115	118	120		122
3	23	57	79	91	96	100	103	104			107
4	21	48	70	81	90	94	98				107
5	19	48	67	85	95	102					109
6	15	40	58	66	78						96
7	18	48	66	78							97
8	16	46	65								115
9	21	42									100
10	12										88

Simulated Cumulative Paid Claims Counts (limited to 10 years' data) plus Ultimates:

Accident Year	Development Year 1	2	3	4	5	6	7	8	9	10	Ult
1	4	7	17	28	41	51	59	62	64	69	99
2	2	14	30	43	56	66	80	89	97		122
3	1	10	22	39	52	63	72	80			107
4	2	7	20	33	44	53	67				107
5	0	7	25	36	51	56					109
6	1	7	16	30	37						96
7	2	8	20	31							97
8	3	13	20								115
9	2	5									100
10	2										88

The above triangle data will be used in the sections that follow, with the aim of estimating the ultimate figures shown in the final columns.

5. Estimation Using Established Models

In this section, we shall apply established reserving techniques (i.e. Chain Ladder, Mack Method, Bootstrap Method and Munich Chain Ladder) to the simulated data from Section 4. The purpose of doing this is to set a benchmark for comparison with the techniques derived from the multi-state reserving model, which are presented in the Section 6.

In applying the established techniques, we have used the "reserving" package in R, developed by Markus Gesmann. [I am grateful to Markus for his time and patience in helping me with this. As ever, any mistakes are my own.]

5.1. Chain Ladder with Mack Method

The following output from R shows the results table for applying the "ChainLadder" procedure to the simulated incurred claims triangle, with the ultimate results added for comparison.

	Latest:	CL-Dev.%:	CL-Ultimate:	CL-Reserve:	Mack S.E.:	S.E.Res.%:	Ultimates
1	97	100.0	97.0	0.00	0.000	NaN	99
2	120	99.0	121.3	1.25	0.212	17.0	122
3	104	97.6	106.6	2.56	0.616	24.0	107
4	98	95.1	103.0	5.02	2.050	40.9	107
5	102	91.2	111.9	9.89	2.868	29.0	109
6	78	86.9	89.8	11.76	2.915	24.8	96
7	78	78.8	99.0	20.99	5.950	28.3	97
8	65	67.5	96.2	31.24	7.225	23.1	115
9	42	48.4	86.8	44.81	7.584	16.9	100
10	12	19.4	61.8	49.80	10.249	20.6	88

	Totals:	
Sum of Latest:	796	
Sum of CL-Ultimate:	973	1,040
Sum of CL-Reserve:	177	
Total Mack S.E.:	19	
Total S.E.% of Reserve:	11	

The above and the plot of "Actual versus Expected" ultimate claims counts overleaf shows that the method has performed well for the earlier (well-developed) years, but has not done so well on Accident Years 8 and 10 (represented by character "T" on the plot).

Inspection of the incurred data in comparison with the known ultimate figures shows that Accident Years 8 and 10 were tracking well below the norm for the latest available data. However, Accident Year 8 went from holding the second lowest value at the end of Development Year 3 to the second highest ultimate value. Accident Year 10 did have the lowest value at the end of Development Year 1 and the lowest ultimate value, but "caught up" more with the other years than would be expected from an inspection of the earlier years' development.

It can also be seen that the predicted ultimate values (excluding Accident Years 1 and 2, which are assumed to be 100% and 99% developed respectively) are within 3 standard errors of the actual ultimate.

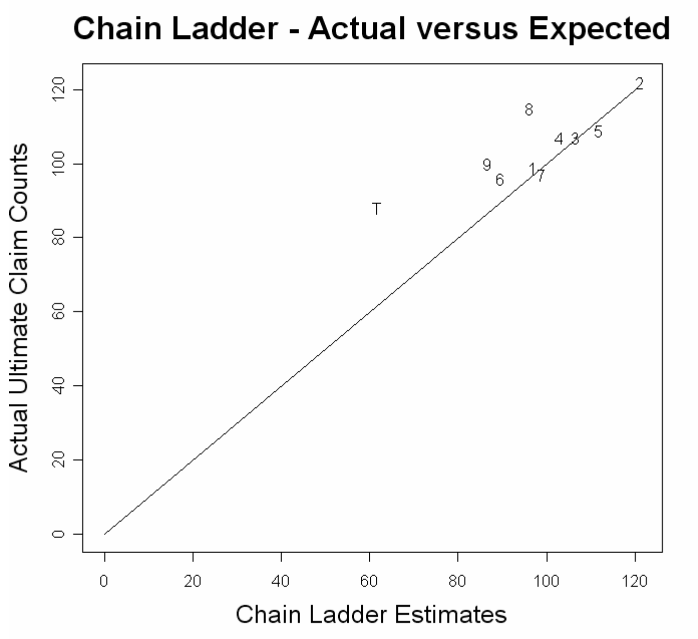

The exhibit below, which on the left summarises the absolute results, shows a wide range of estimated ultimate claim counts, where the ranges are represented by ± one standard error, with later years holding lower values. On seeing this, an analyst might seek to understand what could be driving such a "trend". However, in this case, with no trends in the underlying model, the figures simply reflect random fluctuations that are amplified for later years by the model. In practice, prior estimates might be used, as in the Bornhuetter-Ferguson method, to arrive at a more stable result.

Also, the right hand graph shows the relative size of the "reserves" (i.e. estimated unreported claim counts), illustrating the assumed development profile.

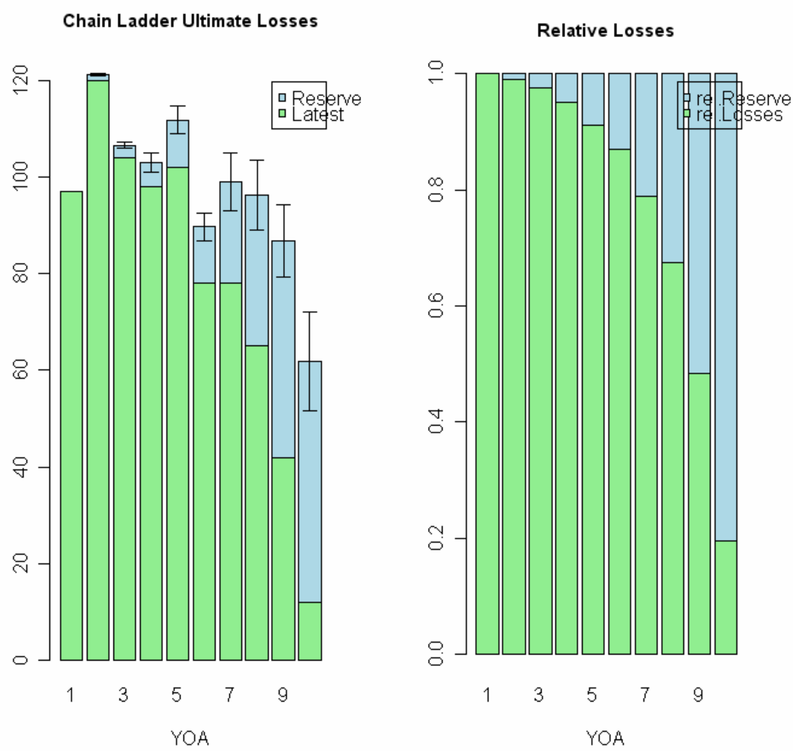

The following Residual Analysis shows no particular calendar year or development year effects, as expected given the formulation of the model.

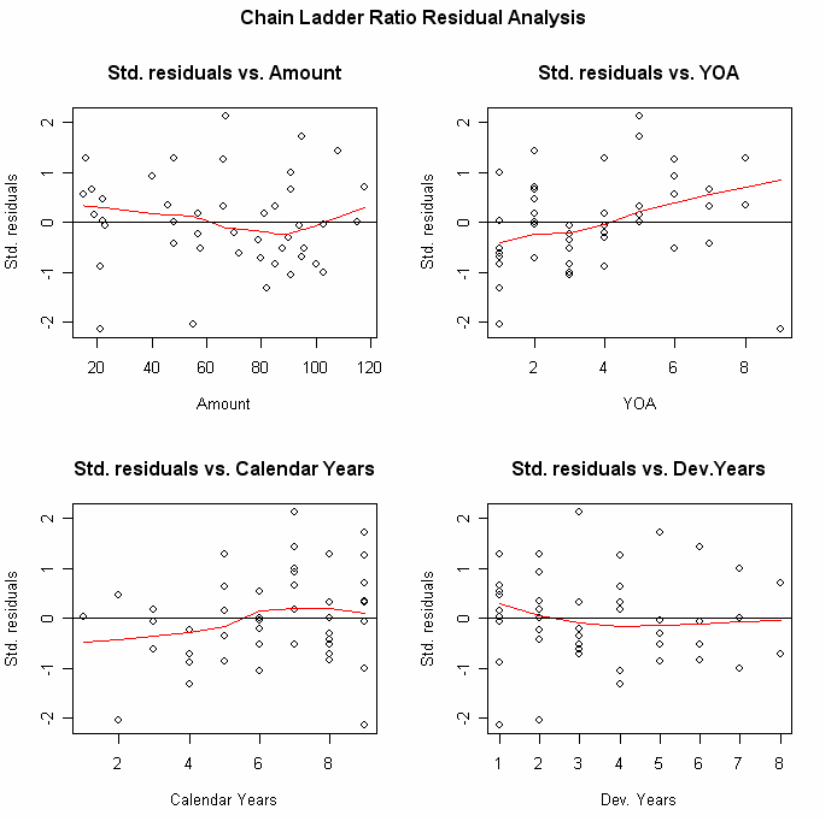

The following exhibit also shows the standard error "margin" estimated using the Mack method for each of the accident years.

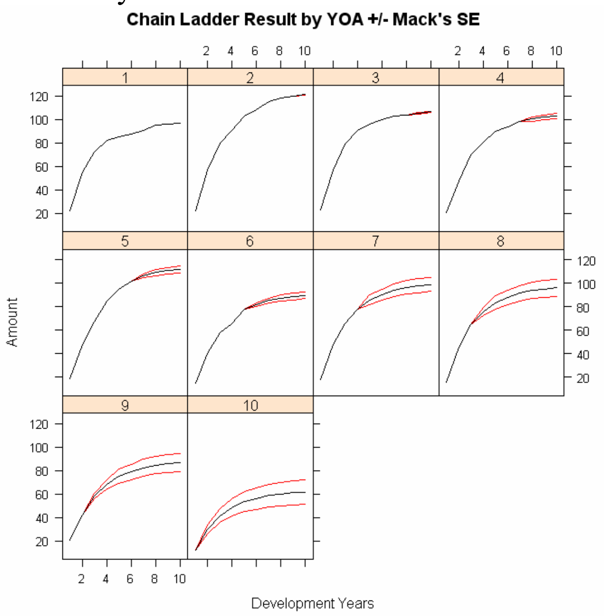

The above analysis was repeated with the "usetail=TRUE" option selected, which estimates a tail factor using log-linear regression through the chain ladder ratios. In this case an additional tail factor of 1.54% was estimated, which improved the estimates a little, increasing the total estimated ultimate claim count from 973 to 988, compared with the actual total of 1,040.

5.2. Chain Ladder with Bootstrap Method

In this example, broadly similar results were obtained after application to the simulated incurred claims triangle, but with the reserves and standard errors estimated using a re-sampling method applied to the residuals from the original fit, which are then "added back" on to the model values to arrive at alternative projections. It was notable that the procedure did not work if there were zero increments in the triangle.

	Latest:	Boot-Ultimate:	Boot-Reserve:	Boot-S.E.:	S.E.Res.:	S.E.Res.%:	Ultimates
1	97	97.0	0.00	0.000	0.00	NaN	99
2	120	121.2	1.24	0.811	1.35	108.7	122
3	104	106.6	2.57	0.932	1.74	67.7	107
4	98	103.0	5.00	1.142	2.30	46.0	107
5	102	111.9	9.89	1.533	3.18	32.1	109
6	78	89.7	11.74	1.495	3.34	28.4	96
7	78	98.9	20.90	2.347	4.71	22.5	97
8	65	96.3	31.32	3.262	6.07	19.4	115
9	42	87.0	45.01	5.102	8.33	18.5	100
10	12	62.3	50.27	10.151	13.80	27.5	88

	Totals:	
Sum of Latest:	796	
Sum of Boot-Ultimate:	974	1,040
Sum of Boot-Reserve:	178	
Boot total S.E.:	16	
Reserve total S.E.:	22	
Total S. E.% of Reserve:	13	

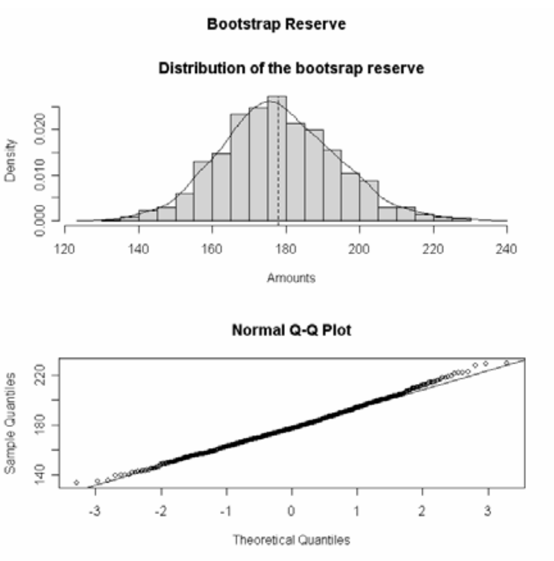

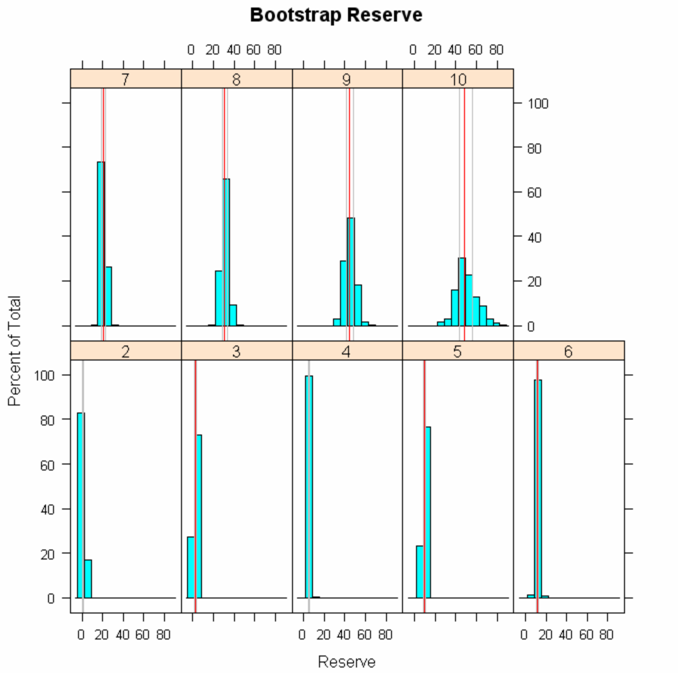

The left hand exhibit shows the estimated distribution of the overall reserve (in this case, the total count of IBNR claims) and that this is broadly normal, but with a slight positive skewness, as expected. The estimated distributions for the individual accident year reserves are shown on the right.

5.3. Munich Chain Ladder Method

This method, presented by Quarg and Mack (2004), seeks to adjust the chain ladder method applied to paid and incurred development data so that the paid-to-incurred (P/I) ratio is targeted at a consistent level. To quote from the paper:

> *Depending on whether the momentary (P/I) ratio is below or above average, one should use an above-average or below-average paid development factor and/or a below-average or above-average incurred development factor, respectively.*

The actual adjustment applied is based upon a regression of the residuals of the paid and incurred development factors against the residuals of the preceding (I/P) and (P/I) ratios development factors. The result of this is to use the paid development data in projecting the incurred and vice-versa, thereby making better use of the available data.

	Latest Paid:	Latest Incurred:	MCL-Paid:	MCL-Incurred:	MCL(P/I)%:	Ultimates
1	69	97	77.7	98.7	78.7	99
2	97	120	115.1	123.4	93.3	122
3	80	104	100.5	108.5	92.6	107
4	67	98	94.1	104.8	89.8	107
5	56	102	99.7	112.5	88.6	109
6	37	78	79.8	89.9	88.7	96
7	31	78	88.2	99.1	89.0	97
8	20	65	87.1	97.6	89.2	115
9	5	42	73.8	83.3	88.6	100
10	2	12	58.7	65.7	89.3	88

	Totals:	
Sum of Latest Paid:	464	
Sum of Latest Incurred:	796	
Sum of MCL-Ultimate Paid:	875	1,040
Sum of MCL-Ultimate Incurred:	983	1,040
Total MCL(P/I)%:	89	

The incurred ultimate estimates can be seen to have increased slightly against the Chain Ladder and Bootstrap methods, better estimating the ultimate claim counts.

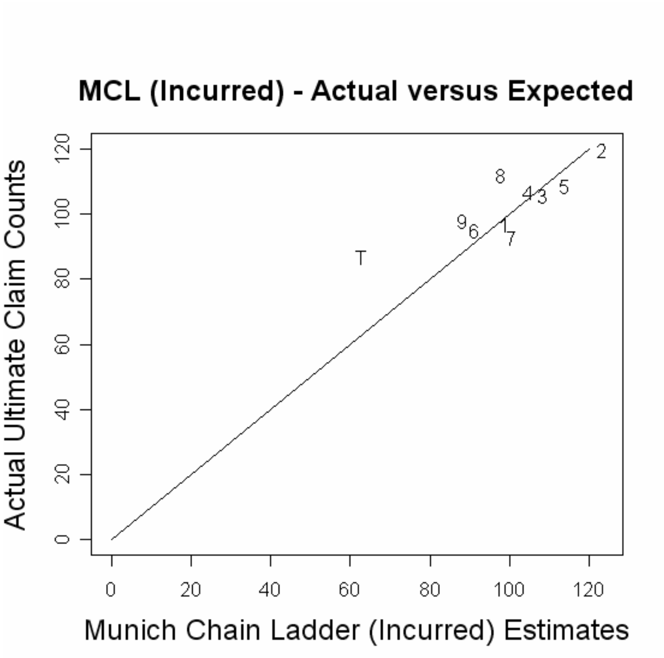

The following graph shows the paid and incurred data (in green, at the bottom) and estimated reserves to ultimate (in blue, on top of the observed data) for each accident year, with paid data marked “P”. It can be seen that a wide range of outcomes is predicted for later years.

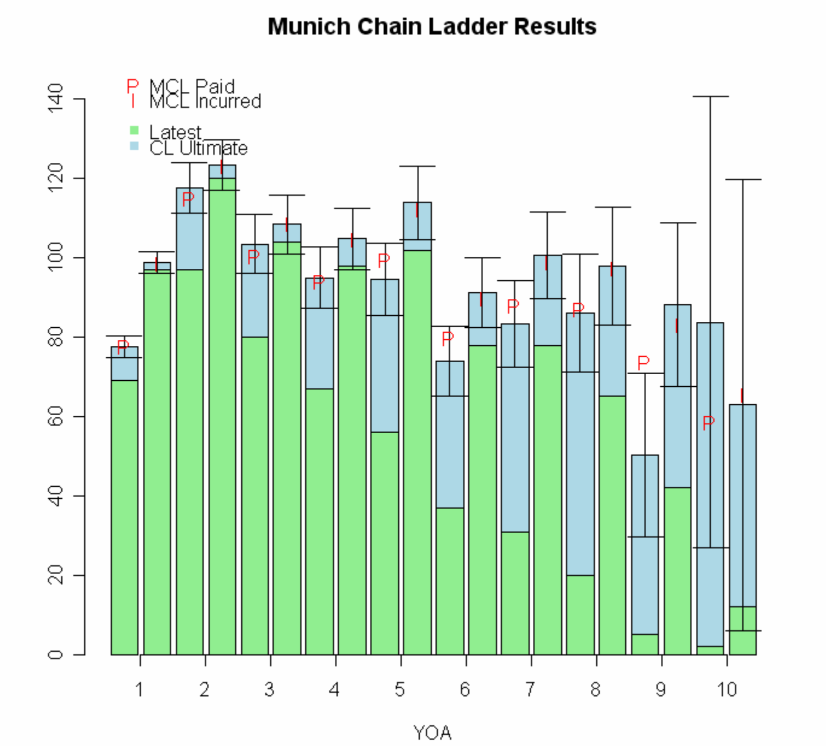

The next graph shows the regression of the residuals of the paid and incurred development factors against the residuals of the preceding (I/P) and (P/I) ratio development factors respectively. These demonstrate positive correlations, estimated at 0.475 for paid and 0.278 for the incurred development factors, which are used in the Munich Chain Ladder method.

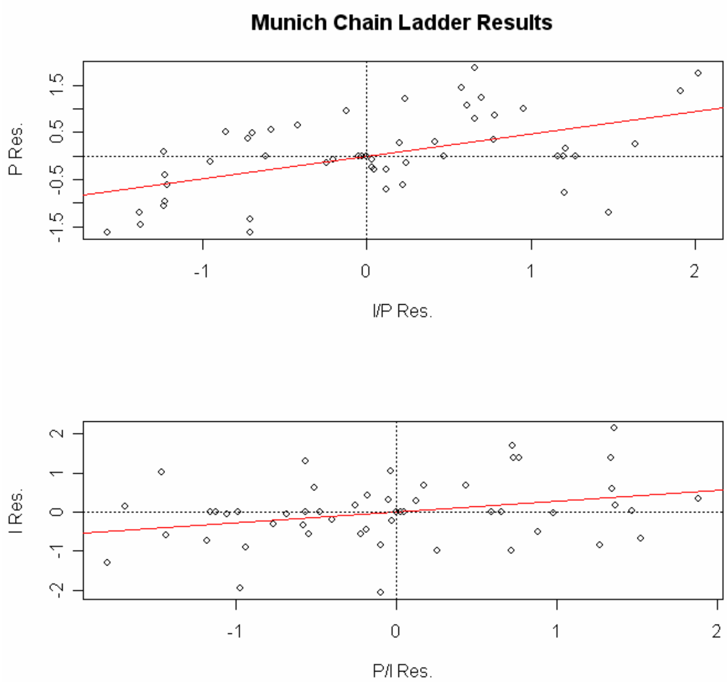

6. Estimation Using the Multi-State Model

We now consider how the multi-state model might be used to predict the ultimate number of claims for a particular underwriting year. Our approach in this section will be to consider how initial estimates of the two development profile parameters *a* and *b* might be determined and then look in more detail at a least squares approach and a Bayesian approach.

6.1. Obtaining Initial Estimates for the Development Profile Parameters

Under the multi-state model assumptions we know that, once all of the claims from an accident year have been reported and paid, we would expect the observed average waiting time in State 0 to be $1/a$ and the average waiting time in State 1 to be $1/b$. We can also assume that losses will arrive uniformly across the year, on average at time $t = 1/2$.

We can take each accident year's incremental incurred data to calculate a weighted average arrival time in State 1, here taking the value for Accident Year 2 as our initial estimate:

Accident Year	Development Year 1	2	3	4	5	6	7	8	9	10	Wtd Ave
1	22	33	17	10	3	3	3	4	1	1	2.93
2	22	35	23	11	12	5	7	3	2		3.22
3	23	34	22	12	5	4	3	1			2.72
4	21	27	22	11	9	4	4				2.88
5	19	29	19	18	10	7					2.92
6	15	25	18	8	12						2.71
7	18	30	18	12							2.31
8	16	30	19								2.05
9	21	21									1.50
10	12										1.00
										Selected Average =	3.22

We can apply a similar approach for the incremental paid data, but include all known outstanding claims (the cumulative incurred claims minus the cumulative paid claims) as at the latest evaluation date, as arriving in the next reporting period, this time taking the value for Accident Year 1 as our initial estimate of the average State 2 arrival time:

Accident Year	Development Year 1	2	3	4	5	6	7	8	9	10	11	Wtd Ave
1	4	3	10	11	13	10	8	3	2	5	28	6.86
2	2	12	16	13	13	10	14	9	8	23		6.03
3	1	9	12	17	13	11	9	8	24			5.74
4	2	5	13	13	11	9	14	31				5.69
5	0	7	18	11	15	5	46					5.28
6	1	6	9	14	7	41						4.83
7	2	6	12	11	47							4.22
8	3	10	7	45								3.45
9	2	3	37									2.83
10	2	10										1.83
											Selected Average =	6.86

From the above, we have an estimate for the average waiting time in State 0 of 2.72 (i.e. 3.22 minus the assumed average loss arrival time of 0.5) and an estimate for the average waiting time in State 1 of 3.64 (i.e. 6.86 minus 3.22). Taking the inverse of these two results gives an initial estimate for a of 0.37 and for b of 0.27.

In arriving at these rough "starting values" we have implicitly assumed that claims arrive at the end of each development year, which will overstate the result. Offsetting this, we have ignored the underestimation that results from using limited data and assuming that all outstanding claims will be paid in the next development year.

6.2. Using Least Squares to Estimate Ultimate Claim Counts and Improve the Model

With assumptions for a and b our development model is fully specified, allowing us to calculate expected development proportions at all durations. Naturally, with data simulated from the multi-state model itself, there is no "model" risk associated with this process, only parameter risk and stochastic variation. Consequently, what follows would not necessarily hold for a real claims number estimation exercise, but only for particularly simple cases!

For this exercise, we are going to look, for each accident year, at the claims count at each reporting point in time (say, $t-1$) and then estimate the next data points for the outstanding (State 1) and paid (State 2) claim counts, which can then be compared to the actual observed figures to calculate residual values that can be squared and summed to provide a goodness of fit measure.

Given our initial estimates for a and b all that we need specify to arrive at these "next value" estimates is an initial figure for the total claim count. Here, we shall use the highest observed incurred claims count figure (120 for our simulated data) as our initial estimate to populate $N(s,y,t)$ for $s \in \{0,1,2\}$, $y \in \{1,\cdots,10\}$ and $t \in \{1,\cdots,11-y\}$. That is, we will set:

$$N(0,y,t) \quad = \quad 120 - N(1,y,t) - N(2,y,t)$$

The first expected values $E[N(s,y,1)]$ can be found by applying the state probability vector $\pi(1)$ to 120, and then subsequent expected values, for $s \in \{0,1,2\}$, $y \in \{1,\cdots,10\}$ and $t \in \{2,\cdots,11-y\}$ can be found from:

$$E[N(s,y,t) \mid N(s,y,t-1)] \quad = \quad N(s,y,t-1)P(1)$$

where $P(1)$ is the probability transition matrix for a time step of one.

It is then possible to define a sum of squared residuals goodness of fit measure:

$$\text{SoS} \quad = \quad \sum_{y=1}^{10}\sum_{t=1}^{11-y}\{N(s,y,t) - E[N(s,y,t) \mid N(s,y,t-1)]\}^2$$

which we can then minimise using the "optim" procedure in R (applied repeatedly to ensure convergence), to find optimal values for the total claim count in each accident year and better values of a and b. The Actual versus Expected plots for the outstanding (State 1) and paid (State 2) claim counts are shown overleaf, indicating that a good fit has been achieved.

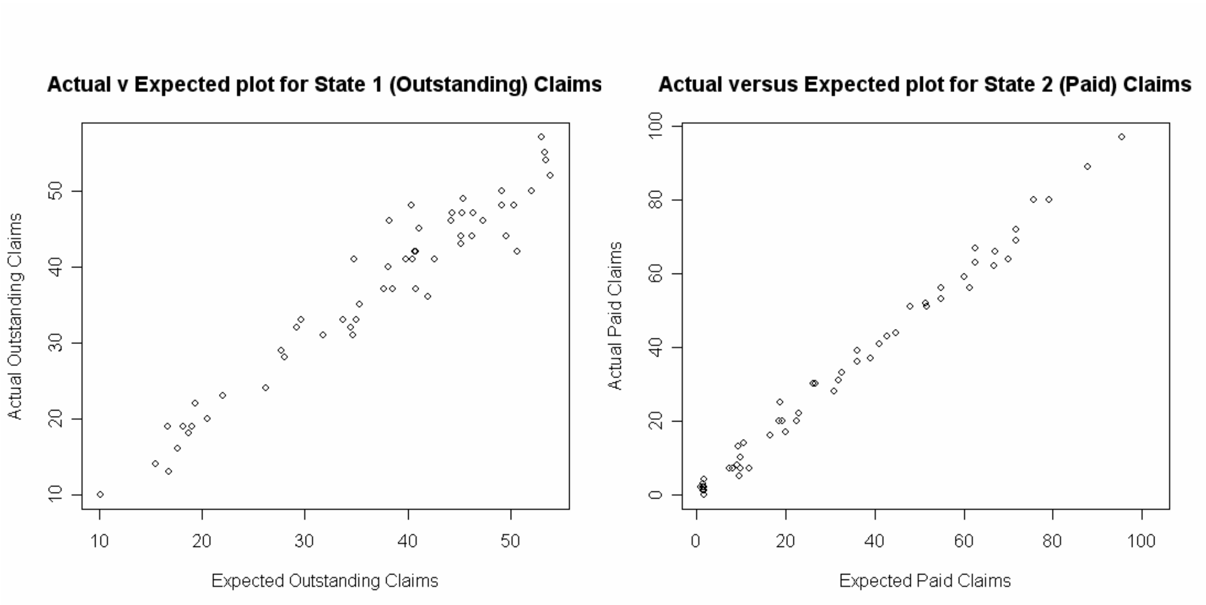

The above modelled values correspond to the following Least Squares (LS) estimates for the total number of claims in each year:

	1	2	3	4	5	6	7	8	9	10
Ultimate	99	122	107	107	109	96	97	115	100	88
LS Est	106.3	116.6	109.9	103.2	108.0	87.9	100.0	95.4	94.3	57.3

These ultimate claim number estimates sum to 978.9, compared with the known simulated total of 1,040. Also, the optimisation process arrived at updated values for a and b of 0.443 and 0.253 respectively.

Finally, it can be seen from the following plots that the results from this approach are comparable with those of the standard chain ladder method.

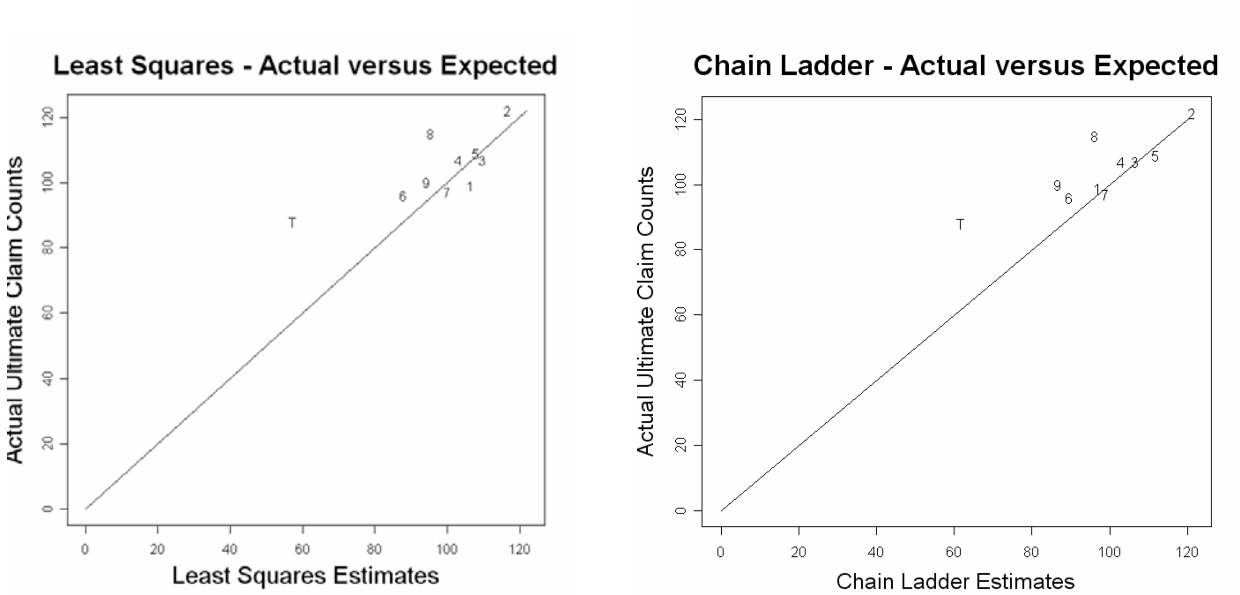

Although one might claim that this is a remarkable result for a model with only two parameters, it should be remembered that we know that the model is "correct" for the data and that the only risks of miss-estimation are from stochastic variation and parameter error.

6.3. Applying a Bayesian Approach to Estimating the Ultimate Claim Counts

With an increased focus on uncertainty within the reserving process, we now consider a Bayesian approach to the estimation of the ultimate claims count. For this purpose, we shall assume that our development model is fully specified with fixed value for a and b of 0.443 and 0.253, respectively, taken from the above exercise. Although it is possible to apply a Bayesian approach to dealing with a and b also, this is not covered in this paper.

Bayes Theorem can be applied in the context of our multi-state model, seeking to estimate the ultimate claims count N_y for year y, thus:

$$\Pr(N_y \mid N(s \in \{1,2\}, y, t)) = \frac{\Pr(N(s \in \{1,2\}, y, t) \mid N_y)\Pr(N_y)}{\sum_{\forall N_y} \Pr(N(s \in \{1,2\}, y, t) \mid N_y)\Pr(N_y)}$$

We have used the result that, conditional upon knowing the ultimate claims count, the number of claims in each state is distributed as a Multinomial distribution with probabilities specified by the state probability vector, $\pi(t)$, from which we can then calculate the probability of observing the known claim counts in States 1 and 2.

In this approach, we take each accident year separately and consider how our belief regarding the ultimate claims count (i.e. the Bayesian posterior distribution for N_y) would change for the claims count data at each reporting time, *t*. Our prior assumption for the distribution of N_y is an uninformative (i.e. uniform) distribution holding integer values between 50 and 200 (although similar results follow from a much wider prior range, the resulting graphs were less clear).

The graph below shows how the uninformative prior distribution becomes more precise for the more developed data.

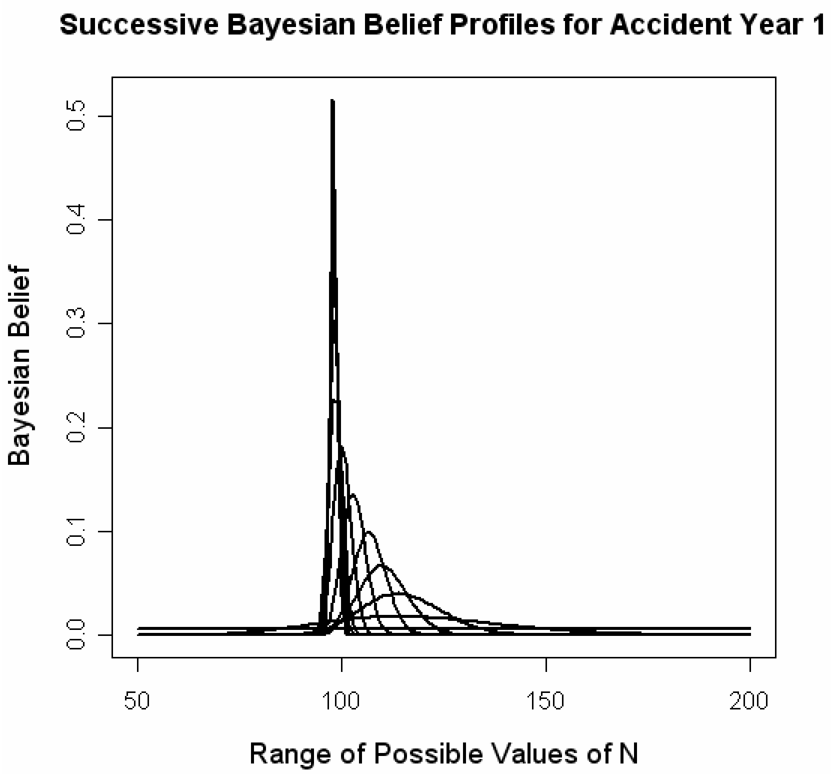

As the Bayesian approach produces a distribution of our beliefs regarding the values that N_y might take for each accident year, it is possible to interrogate (i.e. summarise) that distribution using different measures. For instance, the "best estimate" required for many actuarial opinions might be defined as the mean of the distribution based on the latest data, leading to the following estimates:

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Ultimate	99	122	107	107	109	96	97	115	100	88
Bayes Mean	98.2	121.5	109.6	103.3	109.1	87.3	98.8	95.9	92.6	70.8

With a total value of 987.1, against the known total of 1,040, this is a creditable performance, as demonstrated in the following "Actual versus Expected" graph.

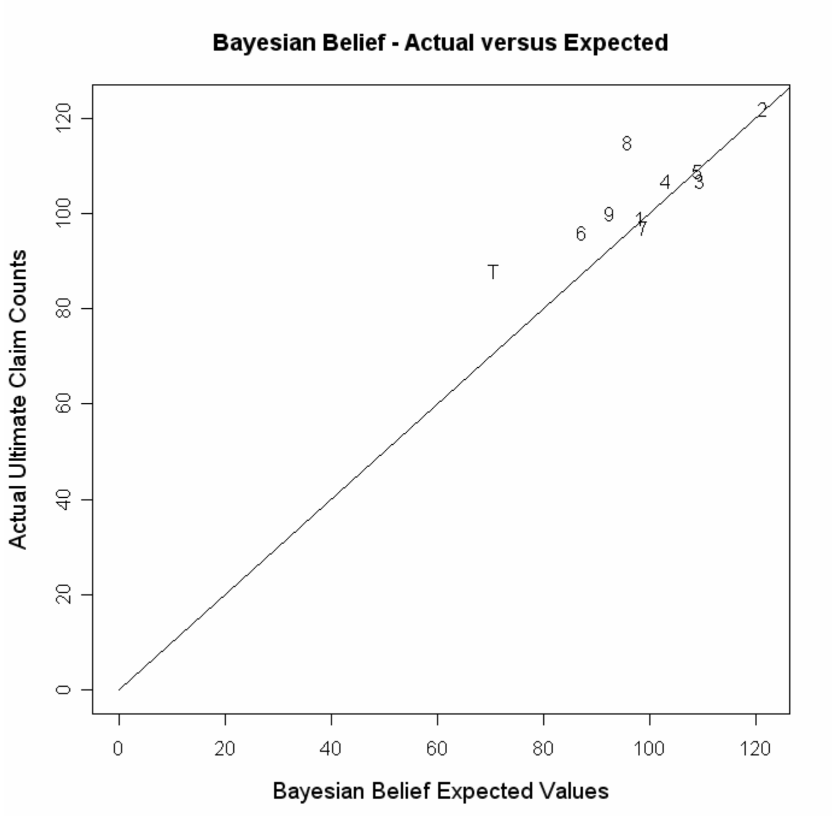

Once again, the data for Accident Years 8 and 10 (i.e. "T") has led our method astray! The exercise was also carried out with the "correct" values of a and b (i.e. 0.40 and 0.25 respectively), which improved the estimates significantly, to total 1,029.2, illustrating the contribution from parameter error.

As well as using the most up to date data to produce Bayesian estimates for a number of different measures, we can also evaluate those measures at each observed stage of the development data, which might be expressed as triangles. Examples of these are shown overleaf.

Expected Values of N_y :

Accident Year	Development Year 0	1	2	3	4	5	6	7	8	9	10
1	125.0	118.5	115.0	110.4	107.0	103.0	100.2	98.6	98.4	98.2	98.2
2	125.0	118.5	118.3	119.1	117.1	118.1	118.1	119.7	120.7	121.5	
3	125.0	123.7	119.3	118.5	116.9	114.1	112.1	110.8	109.6		
4	125.0	113.4	102.4	103.8	103.2	103.6	103.2	103.3			
5	125.0	103.1	100.2	100.1	104.4	107.1	109.1				
6	125.0	83.1	82.7	85.2	84.4	87.3					
7	125.0	97.9	99.2	98.8	98.8						
8	125.0	87.8	93.8	95.9							
9	125.0	113.4	92.6								
10	125.0	70.8									

Standard Deviation (i.e square-root of Variance) for N_y :

Accident Year	Development Year 0	1	2	3	4	5	6	7	8	9	10
1	43.6	22.2	10.0	6.0	4.0	2.9	2.2	1.7	1.3	1.0	0.8
2	43.6	22.2	10.1	6.1	4.2	3.0	2.2	1.6	1.2	0.9	
3	43.6	22.6	10.2	6.1	4.2	3.0	2.3	1.8	1.4		
4	43.6	21.8	9.4	5.7	3.9	2.8	2.1	1.6			
5	43.6	20.8	9.3	5.6	3.8	2.7	2.0				
6	43.6	18.0	8.4	5.2	3.5	2.5					
7	43.6	20.3	9.2	5.6	3.8						
8	43.6	18.9	8.9	5.5							
9	43.6	21.8	9.1								
10	43.6	14.7									

75% Percentile (actually, the first integer greater than the 75% percentile) Values for N_y :

Accident Year	Development Year 0	1	2	3	4	5	6	7	8	9	10
1	163	133	121	114	110	105	102	100	99	99	99
2	163	133	125	123	120	120	120	121	121	122	
3	163	138	126	123	120	116	114	112	110		
4	163	127	108	108	106	105	105	104			
5	163	116	106	104	107	109	110				
6	163	94	88	89	87	89					
7	163	111	105	102	101						
8	163	100	100	99							
9	163	127	98								
10	163	79									

In seeking to quantify and communicate uncertainty, the Bayesian approach has much to offer and, in the opinion of the author, is a natural framework within which to frame consideration of the appropriate reserves that should be held by an insurer.

7. Extensions to the Multi-State Model

In this section, a number of extensions to the multi-state model are presented, which have the potential to better adapt the model to real world circumstances. Due to the constraints on the length of this paper and the stamina of the reader (and author), we have not included full details of the derivation of the underlying formulae; the intention is simply to show what is possible. For illustrative purposes, the "short tail" model assumptions have been used in this section.

7.1. Year of Account "Inception Based Accounting" Development Patterns

Within Lloyd's and for some companies' internal underwriting performance measures, reserving is carried out at a "Year of Account" level, where claims are linked to the inception date of the underlying policies. This has implications for the distribution of exposure over time and the resulting claims development profile.

For a particular year of account, policies may attach throughout the year and will come off risk (assuming that they are all annual policies) during the next year. As a result, the exposures can be considered to build up and reduce over a two year period. An assumption of uniform inception dates and a constant loss arrival rate for each policy can be shown to result in a loss arrival rate that rises uniformly to a peak at the end of the first year, when all policies are on risk, and returns to zero at the end of the second year, as shown below.

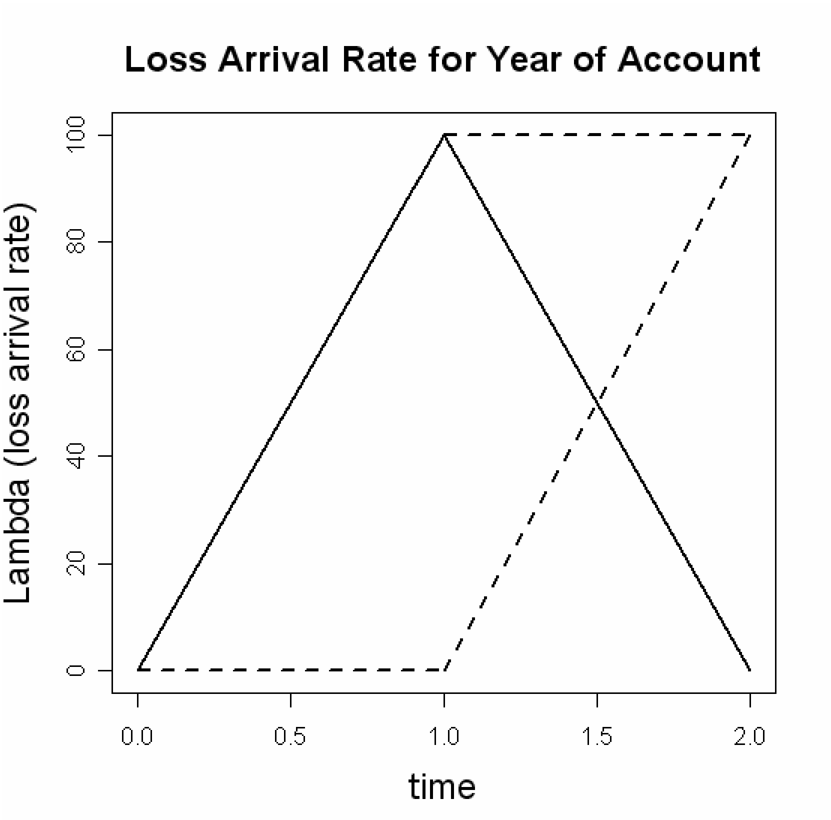

The dotted lines represent policies incepting throughout the year of account (starting at the bottom of the diagram) come off risk throughout the following year.

The resulting claims development profile is shown with the dotted lines on the following graph, in comparison with the standard accident year development profile, which is shown with the solid lines. The year of account development is more stretched out and curved, reflecting the longer period of exposure and its uneven distribution over time.

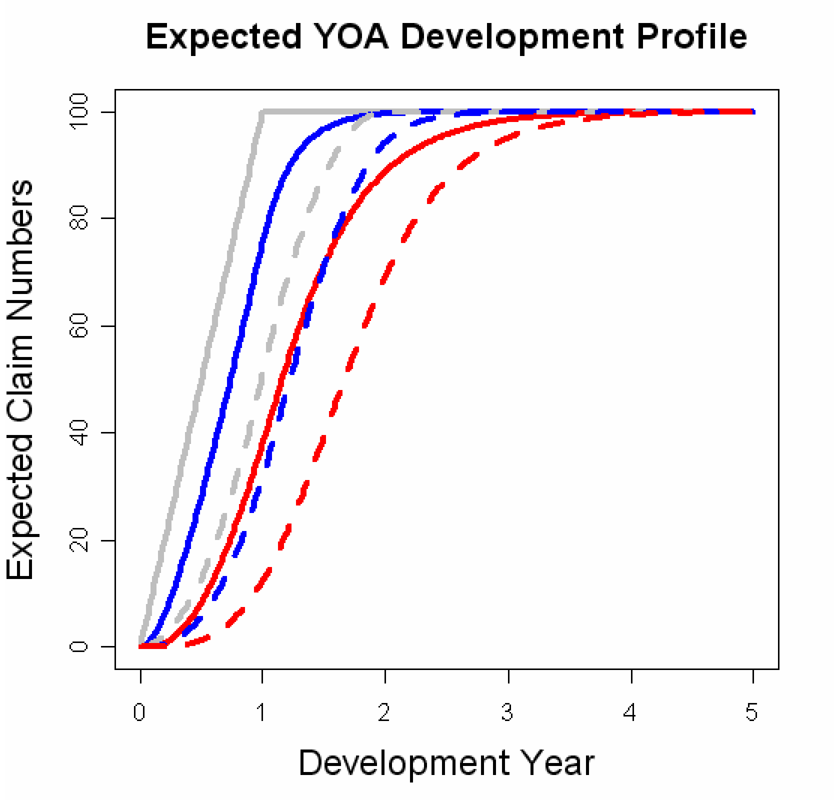

The above results have been derived through a numerical procedure, where multiple accident year development profiles have been combined with starting times distributed uniformly over a year. Although closed form expressions for the year of account development profiles may be achievable, a numerical approach offers greater flexibility to allow for circumstances where the distribution of the inception dates may be uneven, say due to more popular renewals dates such as the first day of each quarter.

7.2. Catastrophes

A particular challenge in reserving arises from the incidence of a great many claims on or around the same date as a result of a major catastrophe such as an earthquake, hurricane or terrorist attack. In these circumstances it is common to extract losses from catastrophes and treat these separately. However, with the assumption of a common run-off pattern, expressed using the transition intensity (or, equivalently, the probability) matrix, it is a relatively straightforward matter to incorporate allowance for catastrophic losses within the expected development profile.

The example shown overleaf is for the standard accident year development profile with a catastrophe loss of 40 claims occurring two-thirds of the way through the year (i.e. around 1st September). The dotted lines show how the overall loss count (in grey, at the top) increases instantly and how the incurred (in blue, in the middle) and paid (red, at the bottom) claims counts would follow.

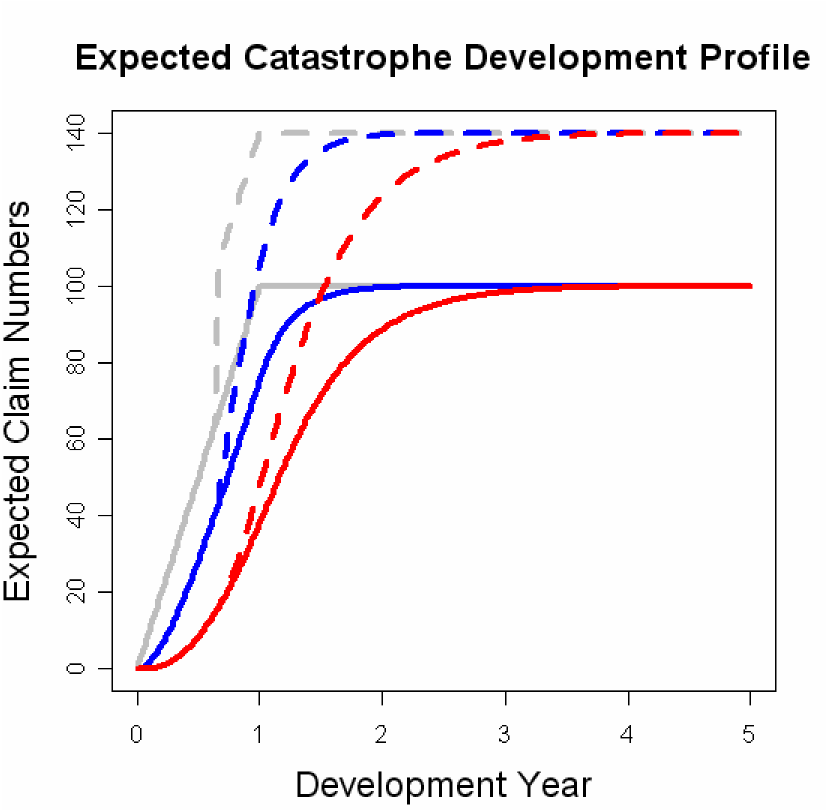

As previously, the adjusted development profile has been determined through a numerical approach and it is possible to allow for particular features of the catastrophe losses such as different rates of reporting and settlement. These different development characteristics might arise from resource constraints issues or delays from litigation relating to, say, the number of events at issue, or the proximate cause of the losses and related coverage issues.

7.3. Void Claims and "Negative" Incurred Development

It is common in claims reserving to find that one is unable to reconcile incurred and paid claims development. This might be simply as a result of insufficient paid development experience having emerged to date, as illustrated in the Munich Chain Ladder example earlier in this paper. Alternatively, systematic over-reserving at the individual claim level (e.g. as a conservative business practice) might lead to the incurred claims total reducing as individual claims are settled for less than their case reserve, resulting in "negative development".

Under the multi-state model formulation, we are considering claims counts only at this time. However, an analogy to claims settling for less than their reserved amount could be that a proportion of reported claims were not paid at all, being deemed "void" on settlement. This can be allowed for by including a fourth state in the model for void claims and assuming that a fixed proportion of all reported claims will be settled without payment, shown overleaf.

Four-state model with allowance for "void" claims:

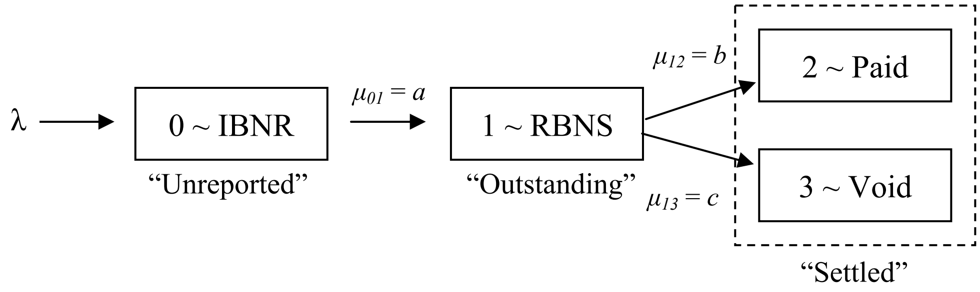

In this case, the total "Force of Settlement" equals the sum of the two transition intensities, $(b+c)$, and the proportion that will be settled as a void claim is $c/(b+c)$. The corresponding expected development profile equations can be derived in a similar fashion to those presented in Section 3 for the 3-state model.

The following graph shows the effect of including a 20% void rate in the "short" model (where the effect is more visible), requiring that $a = 4.00$, $b = 1.60$ and $c = 0.40$. The result is that 20 out of the 100 expected original losses are not expected to result in a claim payment. It is also interesting to see that the total loss count (in grey, at the top) reduces before one observes any negative development in the incurred claims count (in blue, in the middle) itself. In practice the claims department might be asked to monitor the void proportion, to provide a suitable input to the reserving process.

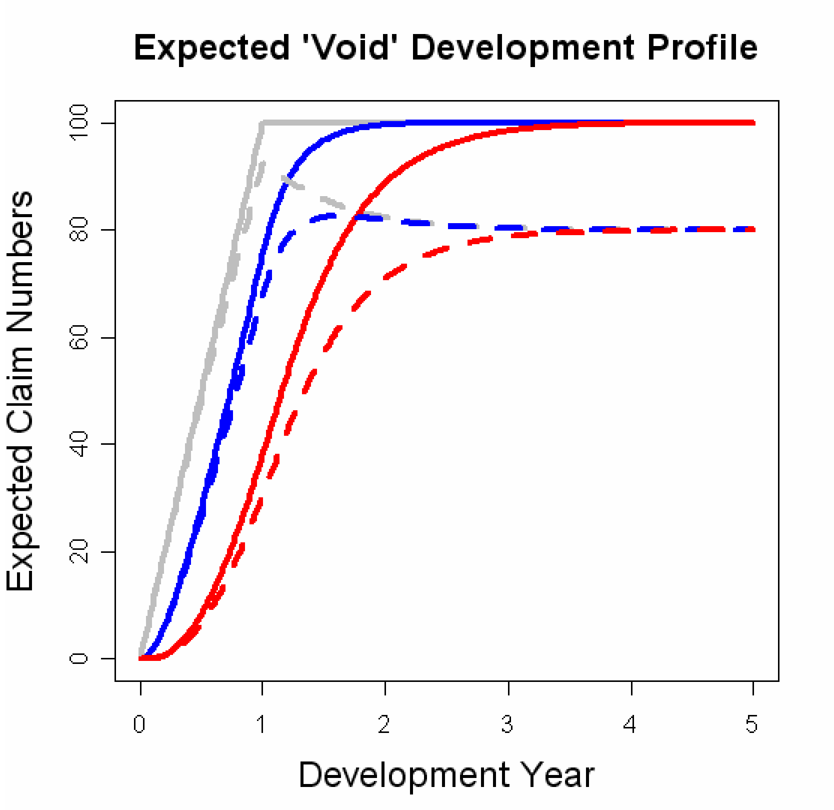

7.4. Modelling Claims Amounts

To be of practical use, a reserving model has to be able to explain and estimate future claims amounts and the model presented in this paper does not achieve this. However, it is apparent that steps could be taken to develop the model in this direction and these are described briefly below.

An Average Claims Model

A starting point, for a portfolio that generated individual claims of very similar amounts, might be to multiply the claims counts from the model by a parameter, ϕ, equivalent to the parameter in the Over-Dispersed Poisson model described by England and Verrall (2002).

As the underlying claims process is assumed to be Poisson in any case, this would be entirely consistent, and the loss and claims amounts in the different states would be expected to follow a distribution with the following mean and variance:

$$E[N(s,y,t)] = \phi\lambda t\alpha(t) \text{ for } t \leq 1;$$
$$\phi\lambda t\pi(t) \text{ for } t > 1.$$

and

$$V[N(s,y,t)] = \phi^2\lambda t\alpha(t) \text{ for } t \leq 1;$$
$$\phi^2\lambda t\pi(t) \text{ for } t > 1.$$

We do not discuss how the over-dispersion parameter might be estimated in this paper.

Incorporating Claim Amounts Distributions Assumptions

Panjer (1981) introduces a method to determine the distribution of an aggregate claims amount where the claims numbers follow a Poisson distribution and the individual claim amount distribution is discrete and defined on the positive integers. Again, the multi-state claims number model is well suited to this purpose and extension to an amounts basis should, in theory, be readily achieved.

Having said this, it is to be expected that this approach would be valid in a limited number of cases. This is because the reporting and processing times, and even the number of stages, may depend strongly on the size of claim involved. This would particularly apply where some kind of streamlined administration was in place for claims below a certain threshold.

For the multi-state model to be valid it needs to be applied to a homogeneous group of claims. It will therefore be necessary to carry out an analysis of the different claims types that might arise, their anticipated reporting, processing and settlement stages, the expected times within each stage and their likely amounts. From this, it is easy to see how a claims-level approach to reserving might quickly become too complex and unwieldy for practical purposes.

However, with modern computing power and procedures, greater data storage capacity and the automation of many processes, it should be possible to overcome these obstacles, particularly given the value to the business in better predicting the development of future claims.

8. Next Steps

As a next step, the Bayesian implementation of the model may be extended to encompass the parameter values a, b and c, and applied to real world claims count data. The extension of the model to include claims amounts may also be developed, through an average claims and an individual claim amount distribution approach.

To develop this as a real world model, detailed information on individual claims from an example portfolio, their size and processing times, will be required. The analysis of this data will help test the validity of the basic assumptions and point to where further development is required before a multi-state model approach can provide a viable basis for reserve estimation.

9. Conclusions

We believe the results in this paper demonstrate that there is a simple underlying process for the emergence of claims from a portfolio of insurance contracts and that real life circumstances are a variation upon this.

With such a simple model, simulated claims counts are readily produced, which can be modelled effectively using existing reserving techniques. Indeed, this facility may have value as a training resource for actuaries and analysts seeking to develop their reserving skills.

We have shown how the model can be applied to estimate the ultimate claims count for different years of account, producing comparable results to existing reserving techniques. We have also shown how a Bayesian approach might be applied, with the potential to produce results that better describe the uncertainty underlying reserve estimates.

References

England, P. D. and R. J. Verrall. 2002. Stochastic Claims Reserving in General Insurance. *British Actuarial Journal*. 8:443-544.

Gesmann, Markus. 2006. reserving: Actuarial Tools for Reserving Analysis. Version 0.1.3 edition, 2003-2006. *http://toolkit.pbwiki.com/ReservingPackage*

Hachemeister, C. 1980. A Stochastic Model for Loss Reserving. *Proceedings of the ICA*: 185-94

Hesselager, O. 1994. A Markov Model for Loss Reserving. *ASTIN Bulletin* 24, 2:183-93

Line, Nick, and Philip Archer-Lock, Steven Fisher, Ian Hilder, Shreyas Shah, Kevin Wenzel and Martin White. 2003. The Cycle Survival Kit: An Investigation into the Reserving Cycle and Other Issues. *Proceeding of the 2003 General Insurance Convention*.

Norberg R. 1993. Prediction of Outstanding Liabilities in Non-life Insurance. *ASTIN Bulletin 23*, no. 1:95-115.

Panjer H. H. 1981. Recursive Evaluation of a Family of Compound Distributions. *ASTIN Bulletin* 12, 22-26

Quarg, Gerhard and Thomas Mack. 2004. Munich Chain Ladder: A Reserving Method That Reduces the Gap Between IBNR Projections Based on Paid and IBNR Projections Based on Incurred Losses. Non-Refereed Paper/Article. *Munich Re Group.*

R Project for Statistical Computing. *http://www.r-project.org/*

APPENDIX

Comparison of Results from Different Methods

Simulation Basis: $\lambda = 100; a = 0.40; b = 0.25$

Accident Year	Latest Data: Paid	Latest Data: Incurred	Known Ultimate Claims	Estimated Ultimate Claim Counts: Chain Ladder	Bootstrap	MCL (Paid)	MCL (Inc'd)	Least Squares	Bayesian Mean	Estimated Standard Errors: Mack	Bootstrap	Bayes (stdev)
1	69	97	99	97.0	97.0	77.7	98.7	106.3	98.2	0.000	0.00	0.8
2	97	120	122	121.3	121.2	115.1	123.4	116.6	121.5	0.212	1.35	0.9
3	80	104	107	106.6	106.6	100.5	108.5	109.9	109.6	0.616	1.74	1.4
4	67	98	107	103.0	103.0	94.1	104.8	103.2	103.3	2.050	2.30	1.6
5	56	102	109	111.9	111.9	99.7	112.5	108.0	109.1	2.868	3.18	2.0
6	37	78	96	89.8	89.7	79.8	89.9	87.9	87.3	2.915	3.34	2.5
7	31	78	97	99.0	98.9	88.2	99.1	100.0	98.8	5.950	4.71	3.8
8	20	65	115	96.2	96.3	87.1	97.6	95.4	95.9	7.225	6.07	5.5
9	5	42	100	86.8	87.0	73.8	83.3	94.3	92.6	7.584	8.33	9.1
10	2	12	88	61.8	62.3	58.7	65.7	57.3	70.8	10.249	13.80	14.7
	464	**796**	**1040**	**973.4**	**973.9**	**874.7**	**983.5**	**978.9**	**987.1**			