A Method For Projecting Individual Large Claims

Karl Murphy FIA and Andrew McLennan FIAA, FIA

Abstract

Motivation. The paper will address the issue of estimating the uncertainty in the run off of individual large claims in insurance portfolios, which is often the primary source of uncertainty in the reserving risk component of insurance risk.

Method. The paper begins by reviewing current methodologies for estimating the uncertainty in loss reserves. Methods until now have focused on aggregate modeling of gross or net of reinsurance loss reserves, and no direct connection between the distribution of gross and net reserves.

The paper develops a non-parametric framework to simulate the distribution of ultimate position of large claims, both reported and large IBNR claims. The method samples the development of individual claims based on the historic development of large claims, incorporating information at an aggregate level surrounding reserving strength. The model also predicts when claims will settle, and the timing of claim payments.

Results. The method developed is not intended to replace existing aggregate modeling, but is an improvement to traditional methods which estimate the variability of gross of reinsurance loss reserves, and is a useful tool to allow for reinsurance recoveries more accurately.

By individually projecting the ultimate position of large claims, we can explicitly allow for policy or contract limits. Further, we can apply any reinsurance program structure to the gross results, including allowance for aggregate deductibles, incomplete placements, retrocessions to captive reinsurers, indexation clauses, and different treaty attachment rules (ie Losses Occurring During vs Risks Attaching).

The paper then shows how the variability of attritional claims can be estimated using traditional stochastic methods, and the attritional and large results can be combined to estimate the variability of the aggregate portfolio of loss reserves.

Keywords. Reserving, Large Claims, Reinsurance, Stochastic Modeling, Simulation, Capital Modeling, IBNR.

1. INTRODUCTION

With an increased focus on understanding variability in claims reserves, a series of papers have been published which develop and add to existing literature on stochastic reserving, in particular England and Verrall[1]. However, almost universally, these papers consider aggregate claims triangles, and do not consider the range of possible outcomes of individual claims. We believe that for many classes of business, the primary source of uncertainty in reserve run-off stems from the uncertainty in large claims, and so a natural extension to the developments in stochastic claims reserving methods would be to produce stochastic outcomes of individual claims.

The paper develops a practical framework to simulate the distribution of ultimate position of large claims, both reported and large IBNR claims. The method samples the development of individual claims based on the historic development of large claims, and applies this development to the current position of claims. The model also predicts when claims will settle, and the timing of claim payments.

A practical by-product of having individually projected the ultimate position of large claims is that we can apply any policy contract limits to any claims, and any reinsurance program structure to the gross results in order to derive stochastic net results that are consistent with the gross without having to make simplifying approximations. For example, by having individual large claims information, excess of loss reinsurance can be properly allowed for. Other more complicated arrangements can also be considered, including allowance for aggregate deductibles, incomplete placements, retrocessions to captive reinsurers, indexation clauses, and different treaty attachment rules (i.e. "losses occurring during" treaties compared to "risks attaching" treaties). Reinsurance recoveries can then be allocated to specific contracts, enabling easier commutation and reinsurance bad debt calculations.

The paper then shows how the variability of attritional claims can be estimated using aggregate stochastic methods, and the attritional and large results can be combined to estimate the variability of the aggregate portfolio of loss reserves. By separating large and attritional claims in the estimation of the uncertainty in loss reserves, changes to the mix (by size and numbers) of large claims can be directly allowed for and modeled.

The structure of the paper will be as follows: first we are going to briefly discuss the main existing stochastic methods for estimating reserving risk. We will then look at a new method which we believe better identifies the main source of uncertainty in reserving risk. We will then show how the method can make exact explicit allowance for any historic reinsurance programs that protect the portfolio. By doing this, we show how to provide a very explicit link between gross and net reserving risk.

2. A BRIEF OUTLINE OF STOCHASTIC MODELLING TECHNIQUES

This section of the paper is intended to be a general review of existing techniques; hence we have kept existing theory to a minimum, quoting other papers or literature where a more theoretical explanation is required. In particular, readers are directed to the recent paper by England and Verrall [1] which sets out most techniques in theoretical detail.

Many stochastic techniques to date are based on some form of chain ladder technique.

Mack's model [2] was one of the first models used in practice to understand the variability in future claim amounts. Mack provided the first two moments of the future cumulative claim amounts, and assumed the model to be "distribution free". Ultimately, however, we are interested in the full predictive distribution of claims, rather than the first two moments. England and Verrall [1] provide a solution to this assuming the cumulative claims are normally distributed.

Renshaw and Verrall [3] introduced a statistical model assuming the incremental claim amount in each accident period and development period are independent random variables with an over-dispersed Poisson distribution.

Verrall [4] developed on the over-dispersed Poisson chain ladder model with the over-dispersed negative binomial model. A key difference between this and the over-dispersed Poisson model is the assumption that incremental claim payments are dependent on the cumulative claim amount at the previous period, similar to Mack's model.

In general, techniques to date have been designed for use on aggregate, portfolio level triangles of claim payment or incurred triangles. Making adequate, explicit allowance for reinsurance in practice has been, at best, an after-thought, often made using a deterministic gross to net ratio for each accident period, selected using information from aggregate modeling of the central estimate using traditional actuarial techniques. Techniques described above assume that all claims develop, on average, in a similar way, or that the mix of claims with different development patterns is constant throughout history. Due to the highly volatile occurrence and size of large claims, this may not be appropriate.

3. A METHOD FOR PROJECTING INDIVIDUAL LARGE CLAIMS

3.1 Introduction

One of the key assumptions in the aggregate stochastic methods described above is that the mix of claims with different development patterns over origin periods is stable. No allowance is made, for example, for increased variability for an accident year with "known" poor large claims experience. Also, no allowance is made for the status (i.e. open/settled) of large claims.

Perhaps more importantly, aggregate stochastic methods do not provide a process for linking the variability of gross and net of reinsurance reserves, where non-trivial treaties (such as quota shares) are in place.

We propose a model designed to cope with the problems described above, by separating the major source of uncertainty, large claims, from the remaining attritional losses, with a separate projection of individual large claims.

The remainder of this section will detail the specifics of our proposed method: uncertainty in known (reported) large claims, uncertainty in the numbers and amounts of unknown (un-reported, and reported, but not yet large) large claims, attritional claims and the aggregation of results.

3.2 Known Large Claims

We must first define by what we mean as "large". There are a number of practical considerations in choosing the threshold of large claims. The main concern is if we are going to use the results for calculating reinsurance recoveries under an Excess of Loss (XoL) program, we must choose a threshold below any historic excess of loss programs. Secondly, as we shall see, we need a significant pool of claims to sample from. To balance the above points, in the limit, we could apply this method to all claims in the portfolio, however computational and time limitations necessitate a cap on the size of the pool. It is important to frame question of choosing a threshold within context of the portfolio, for example, by considering the size of claims which are managed by the complex or large claims unit. In general, we have found this method produces reasonable results with as few as 200 individual large claims with the oldest years having had up to ten years of development.

We include all claims which were "ever" large in our method, that is to say, we include claims which could ultimately be small (or nil) but which were once estimated to be large.

We propose to adopt a stochastic chain ladder projection on individual large claims, where the simulated chain ladder factors are sampled from the observed chain ladder factors in historic large claims. Further, when simulating the development factor of the claim, we also sample the subsequent status of the claim. We therefore simulate chain ladder factors for open claims from historic claims which were open at the same point in development. Closed claims can be simulated at subsequent development periods from similar closed claims to allow for the possibility of re-opening; to the extent that they are present in the historic data.

Consider the following claims. For simplicity, assume all claims are settled by development year 3. To finalize the projection of large claims, we need to project claim D and E for one year and claim F for two years.

Table 1

Incurred Amounts			
	Development Year		
Claim	1	2	3
A	400,000	800,000	800,000
B	500,000	1,600,000	850,000
C	1,000,000	1,000,000	1,500,000
D	200,000	500,000	
E	300,000	200,000	
F	150,000		

Table 2

Development Factors		
Claim	Year 1 to Year 2	Year 2 to Year 3
A	2.00	1.00
B	3.20	0.53
C	1.00	1.50
D	2.50	
E	0.67	
F		

Table 3

Claim Status			
	Development Year		
Claim	1	2	3
A	Open	Closed	Closed
B	Open	Open	Closed
C	Open	Open	Closed
D	Open	Open	
E	Open	Closed	
F	Open		

To develop claim D to ultimate, we pick a claim that was open at development year 2. In this case, B and C were open at development year 2, and so we can either develop claim D by a chain ladder factor of 0.53 or 1.5.

To develop claim E to ultimate, we pick a claim that was closed at development year 2. In this simple example, only claim A was closed at the same point. Therefore, to simulate the ultimate position of claim E, we pick the chain ladder factor from claim A, that is 1.0.

To develop claim F, we must first project the position to development year 2 from open claims. Therefore, it can simulate chain ladder factors from any of claims A to E, with equal probability. If the claim follows the experience of claim B, C or D to development year 2, the claim remains open, and develops by a chain ladder factor of 3.2, 1.0 or 2.5 respectively. If the claim follows the experience of either claim A or claim E, then the claim closes and develops by a chain ladder factor of 2.0 or 0.67. Developing the position from year 2 to year 3 depends on whether the simulated claim closed in year 2 or remained open. If it remained open (i.e. was simulated from either B, C or D), then the development from years 2 to 3 is simulated from claims B or C (with equal probability) in a similar manner to claim D; if it closed (i.e. was simulated from A or E), then the development is simulated from claim A only (in a similar manner to claim E).

Based on this set of data, the possible range of outcomes for claim D is $265,625 to $750,000, for claim E is $200,000, and for claim F is $100,000 to $720,000 (the lower end of the range is attained if the simulation chooses claim E and then A, the upper end of it chooses claim B and then C). Note that the implied total ultimate chain ladder factor for the maximum simulated value of claim F is 4.8. This is more extreme than any ultimate chain ladder factors seen to date.

By explicitly identifying open and closed claims, we are adding extra information to the basic chain ladder model. The model will then capture the increased volatility of origin years which have a larger number or amount of large claims than average, and the reduced volatility of origin years with fewer large claims.

3.3 IBNR Large Claims

The above section deals with the uncertainty around claims which are already large. This is clearly only part of the picture. We must also deal with claims which become large at some point in the future. These claims can arise from genuinely new claims which have been incurred but not reported, and claims which have been reported, but which are not yet (or have never been) large.

Both the number and size of these claims need quantifying. The following sections detail how the method deals with these.

3.3.1 IBNR Large Claim Numbers

In dealing with the known large claims, we allow for the possibility that a currently large claim will ultimately settle below the large threshold. In our large number projection, we need a definition of large claim numbers that can cope with these outcomes. We deal with this by projecting a triangle of claim numbers, where a claim is counted once in the development year it became large. Claims which subsequently fall below the threshold are included in this triangle. We therefore are not making any assumption about how many of these claims will ultimately settle for less than the threshold in this step of the projection.

Standard stochastic chain ladder techniques can be applied to this data if desired, however we believe this may not be appropriate in this particular case. In particular, due to the generally small number of claims which are reported as large in development years one and two, the projected number of large claims for the most recent origin periods may be artificially unstable.

Further, we must ask ourselves if it is intuitive to suggest that if the most recent origin period has twice as many large claims per unit of exposure reported in development year one as the historical average, then it will have twice the number of large claims per unit of exposure ultimately. This does not seem to make sense in practice. Given that most aggregate stochastic methods are based on chain-ladder projections, in this instance the mean number of large claims may tend to be over-stated.

We suggest a more appropriate model for large claims numbers would be to assume the claim frequency per unit of exposure in each development period is independent of previous or subsequent development periods. The definition of exposure could include earned policy count, vehicle years, rate-adjusted earned premium or ultimate number of attritional claims.

Assuming the number of claims in a unique origin and development period follows a Poisson (or negative binomial) distribution, a number of claims that become large in each future time period can be simulated.

3.3.2 IBNR Large Claim Severity

A number of options are available to simulate the ultimate size of individual IBNR large claims.

The method we suggest is to sample from the (simulated) known large claims, where the claims are selected from the claims which became large in that development period. It may be necessary to group older development periods together to gain a significant pool of claims to sample from. By adopting this approach, we are allowing for any potential differences in average claim size by reporting development period, including the propensity for a claim to be ultimately small, and avoid the need to specify the claim size distribution. Appropriate adjustments for inflation are also required; a further refinement would allow the inflation factor selected to be stochastic.

A simplification to this method could be to sample from all simulated known claims, however if we are interested in the finalization date of claims, for example to calculate reinsurance recoveries under an excess of loss program with an indexation clause, we can run the risk of claims being finalized before they were reported as being large.

Instead of sampling from the simulated known claims, it is possible to parameterize the probability of a reported large claim finalizing as large, the finalization period of a large claim and the severity of ultimately large claims. These can be calculated from historic data, usually using a Bernoulli distribution for the probability of a reported large claim finalizing as large, a discrete distribution for the finalization period, and an appropriate distribution (perhaps Pareto or generalized Pareto) for the severity. These various distributions can then be reviewed against other market or portfolio benchmarks if available.

3.4 Combining Known and IBNR Large Claims

Now that we have separately generated the simulated ultimate position of known large claims and IBNR large claims, combining these results gives us the full picture of large claims in the run-off of reserves.

It is possible to apply a dependency structure to allow for correlations between the run-off of the known claims and the number and severity of large IBNR claims. Applying a positive correlation has an intuitive appeal; however it is very difficult to estimate the strength or shape of this relationship. We recommend at the very least scenario testing the results using various correlation strength and dependency shapes.

3.5 Non-Large Claims

To understand the variability of the aggregate reserve distribution, we need to allow for the variability of the non-large claims.

To do this we recommend using an aggregate triangle where each claim is "capped" at a certain value. For example, if a capping level of $100,000 is chosen, then the capped triangle contains all development up to the point where it reaches $100,000, and any amount in excess of this is omitted from the triangle. A claim which is reserved at $50,000 in year 1, $99,000 in year 2 and $150,000 in year 3 is included as $50,000, $99,000 and $100,000 for each respective development year. We prefer the use of a capped triangle as opposed to a triangle where large claims have been completely removed for a number of reasons, as we find it produces more stable results, and the historic triangle does not change when new diagonals of data are added (as large claims drop below the threshold and new large claims develop). In the example above, if large claims are removed from the triangle, then the development from the example claim is $50,000, $99,000, $0.

Once a capped triangle has been calculated, one of the traditional aggregate stochastic reserving methods described in Section 2 can be used to determine a range of outcomes for the "capped" reserve. This aggregate distribution can then be calculated as the sum of the capped claims and the excess of cap large claim amounts.

When selecting the capping level for the attritional claims, we recommend using a level above the "large" claim threshold. By selecting a cap above the large claim threshold, we are using information about the claims which are currently just below the cap and have a good chance of increasing above the cap at some stage in their development.

Again, it may be appropriate to introduce a dependency between the run-off of the capped and excess of cap claims.

4. REFINEMENTS TO THE METHOD AND KEY ASSUMPTIONS

4.1 Model Refinements

There are a number of refinements to the basic method that are worth outlining for completeness.

When simulating the known large claims, consideration should be given to measuring the development period as the time since the claim became large rather than as the time since accident (such that it is on a reporting period basis). This may be more appropriate for large claims due to the claim management and legal processes these claims are subject to, and generally these progress in a similar manner from the time a claim becomes large rather than from the time the accident occurs. Alternatively, a further split can be made by considering those reported "early" and "late", although this tends to reduce the sample from which to simulate from further.

We suggest splitting the large claims into at least two layers, to allow for different development patterns in the extremely large claims. For example, whereas a claim movement from $500,000 to $5 million is possible, it is perhaps less likely for a claim of $5million to increase to $50 million. Including the development factors from smaller large claims in the pool to project the extremely large claims may overstate the variability of possible outcomes for these claims. In determining which claims are in the upper layers (and indeed in the original large definition), it is important to standardize the historical claims for inflationary effects so as to not bias the claims towards more recent origin periods. It is also important to recognize claims can be in different layers at different development periods.

Selecting the very large threshold(s) is a difficult choice, and there is no one single correct method. We have found a threshold that varies by development period, such that between 10% and 20% of claims are in the top layer produces enough claims to sample from, and produces reasonably reliable results.

4.2 Key Assumptions

There are a series of assumptions underlying the model, which are worth pointing out so that their appropriateness or otherwise can be assessed.

We are assuming the historic observed chain ladder, and settlement patterns, contain the entire population of possible values. Clearly, over 1 period, this is not appropriate. However, as we are interested in the ultimate position of claims, often over a significant time period, the possible number of ultimate development factors (i.e. the product of the 1 period factors) even for a small number of possible factors (e.g. 50 at each period) becomes very large, and this assumption is not unreasonable.

We assume that chain ladder factors from one period to the next are independent, other than for changes in layer and claim status. This assumption is consistent with most other stochastic reserving methods. Further, we have assumed that individual claims develop independently within each period. This is potentially optimistic as there may be changes to internal case estimation procedures which affect all open claims, and there are external factors which also affect all open claims such as legal changes and economic factors. These global external effects can be allowed for within the model by overlaying these effects on the underlying process. By projecting claim status into the future, the effects can be applied only to open claims, as would happen in practice. If these effects are overlaid on the claims, it is important to remove any historic effects from the data to avoid double counting these shocks. Applying future inflation effects on top of the underlying projection is useful if this modeling is carried out as part of a wider capital modeling project, as it links in the reserving risk with the global economic scenarios.

As seen with the above simple example, for very new claims, the method can produce very wide ranging results. If the resulting range is thought to be too unstable, for example when considering the implied reinsurance recoveries at high layers, it may be appropriate to either adjust the range of possible results, or use a method similar to that developed for the IBNR large claims described above.

5. ALLOWANCE FOR HISTORIC REINSURANCE STRUCTURES

As we have now projected the ultimate position of all large claims, we can calculate any reinsurance recoveries exactly. For known large claims, we know all the reinsurance details which attach to the claim, and any quota share arrangements can be applied to the aggregate results.

It may be necessary to introduce a further refinement to the model if, say, the excess of loss treaty is placed on a risks attaching basis. For known large claims, we will know the underwriting year of the policy. For IBNR large claims, the underwriting year to which the claim attaches can be simulated. Typically the probability would be in proportion to the exposure that each underwriting year contributes to the accident year.

6. RECONCILIATION OF RESULTS WITH AGGREGATE MODELLING

Invariably, this work will form part of a larger piece of work; usually an outstanding claims review or part of a capital modeling project. The actuary may form a view of the reserves based on aggregate deterministic methods. This will not correspond with the results of the above method, or indeed any of the methods described in Section 2. This is less than ideal, as the practitioner would like to understand the variability around their central estimate, rather than some other result.

One way of ensuring consistency is to scale results by origin year so that the mean simulated result equals the actuary's best estimate of reserves, or try a different method. This can be done by either applying a multiplicative scaling factor for each accident year, or alternatively by adding on a fixed loading for each accident year. This can lead to undesirable results, either with negative reserves in some instances of additive scaling, or extreme results if the multiplicative scaling factor is large.

If the outstanding claims review uses consistent data (in terms of separately modeling capped and excess claims, and considers ultimate counts as well as amounts), then there are additional diagnostics available to the actuary, such as the following:

- Large Claim Frequency;
- Large Ultimate Claim Frequency;
- Large Excess Ultimate Claim Size;
- Large Excess Ultimate Burning Cost / Loss Ratio;
- Capped Claim Burning Cost / Loss Ratio;
- Large Excess Cost as a Percentage of the Total Claims Cost.

With these, the actuary has the ability to understand which piece of the projection is producing results inconsistent with the aggregate modeling.

7. CASE STUDY

7.1 Introduction

The concepts described above are more readily visualized as a case study. The data modeled is from a UK auto account, and contain 16 years of historic data. For individual large claims above £100,000, the data included the accident date, report date, and the year-end paid and incurred positions, as well as a history of the claim status.

The layers were chosen such that 80% of the claims in each development period were in the lower layer, and 20% in the upper layer. The actual layer limits can be seen in Appendix 1.

7.2 Analysis of the Gross Results

Figure 1 shows the simulated development of a claim which has just been reported as being large, with a current incurred position of £125,000. The lighter shades of gray represents the more extreme percentiles, with the dotted lines representing the 90th, 75th, 50th, 25th and 10th percentiles. The mean development is represented by the solid line. As can be seen, we expect the case reserve to be ultimately inadequate, with the expected ultimate amount being just above £300,000. However, using the method described in this paper, can see that 90% of the time, the claim will settle for £700,000 or less. Occasionally, however, the claim develops much more significantly. Figure 2 shows an individual simulation where the claim grows to more than £1,000,000.

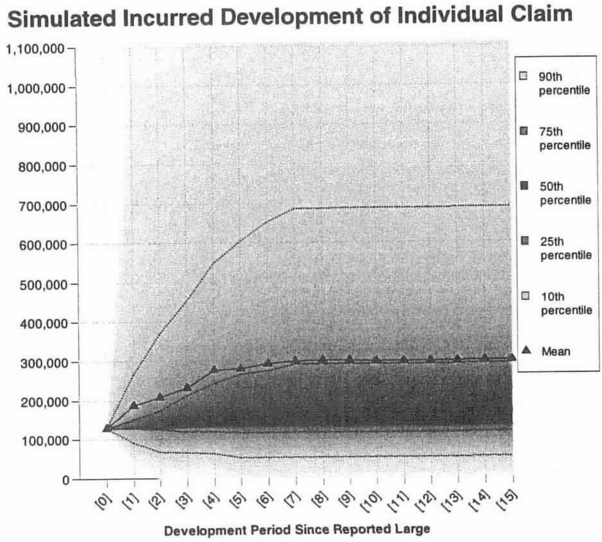

Figure 1

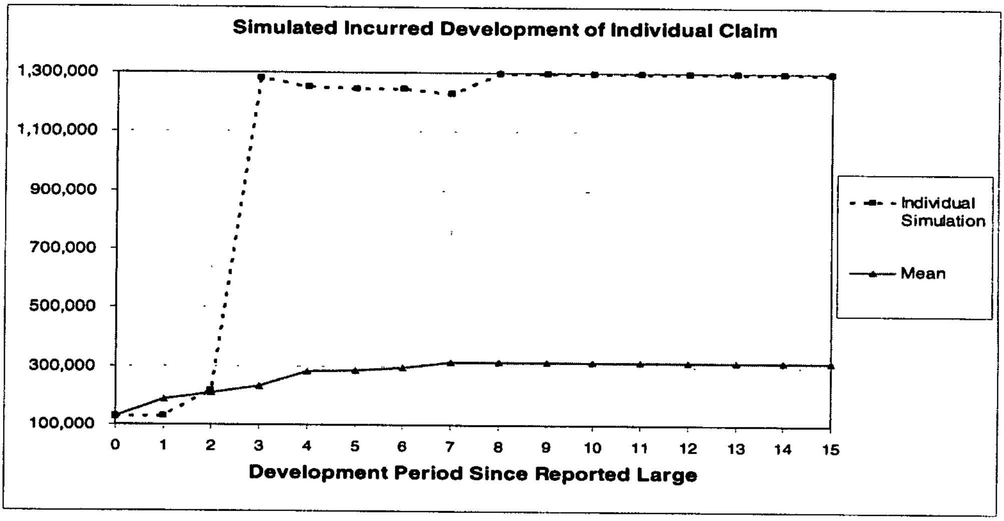

Figure 2

Even for claims that have been reported as large for several years, there is uncertainty over the development. Figure 3 shows the simulated development for a claim that has been reported large for four years, using the same percentile descriptions as for Figure 1. On average, the claim is expected to run off at an increase to the current incurred. Note that the variability around this is still quite significant.

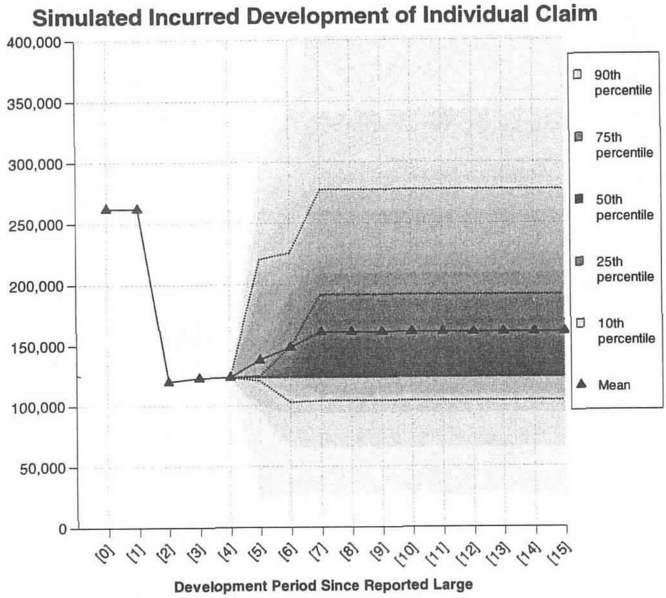

Figure 3

We mentioned in Section 4 that we would typically expect to see different loss development factors for individual "small" large claims than for "large" large claims. This is illustrated in Figure 4.

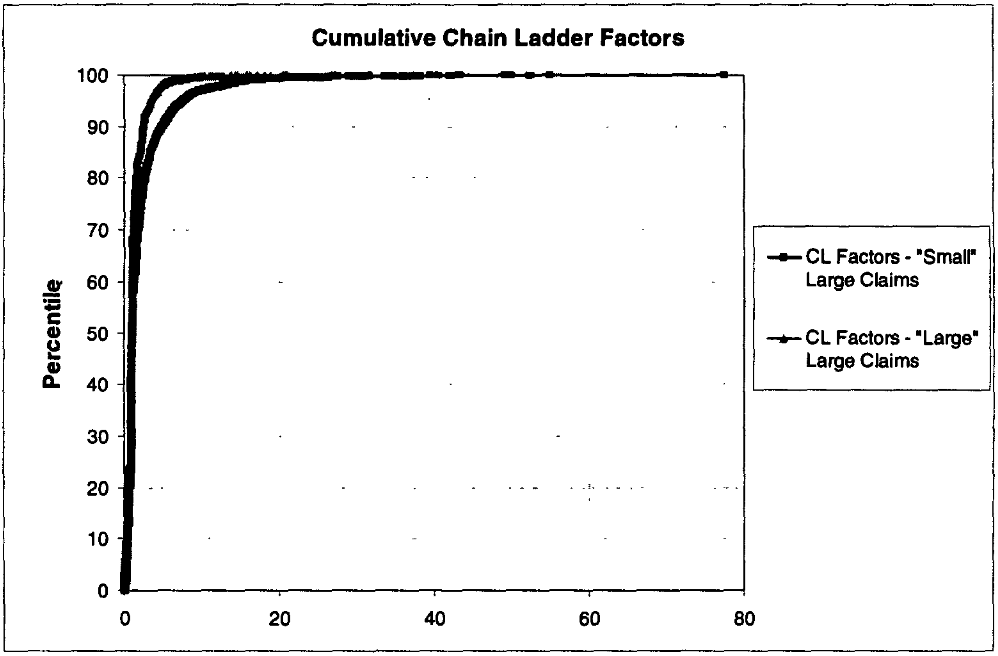

Figure 4

The darker line represents the distribution of cumulative loss development factors for "small" large claims in the first development period, the lighter line the distribution for "large" large claims. As expected, it is much more unlikely to have a large development factor for the "large" large claims, although it is quite possible.

To analyze full accident year results, we have estimated the uncertainty surrounding the attritional claims using Mack's method on a triangle based on a combination of incurred and paid data. Figure 5 shows the percentile plot of the total unpaid liabilities of the capped claims.

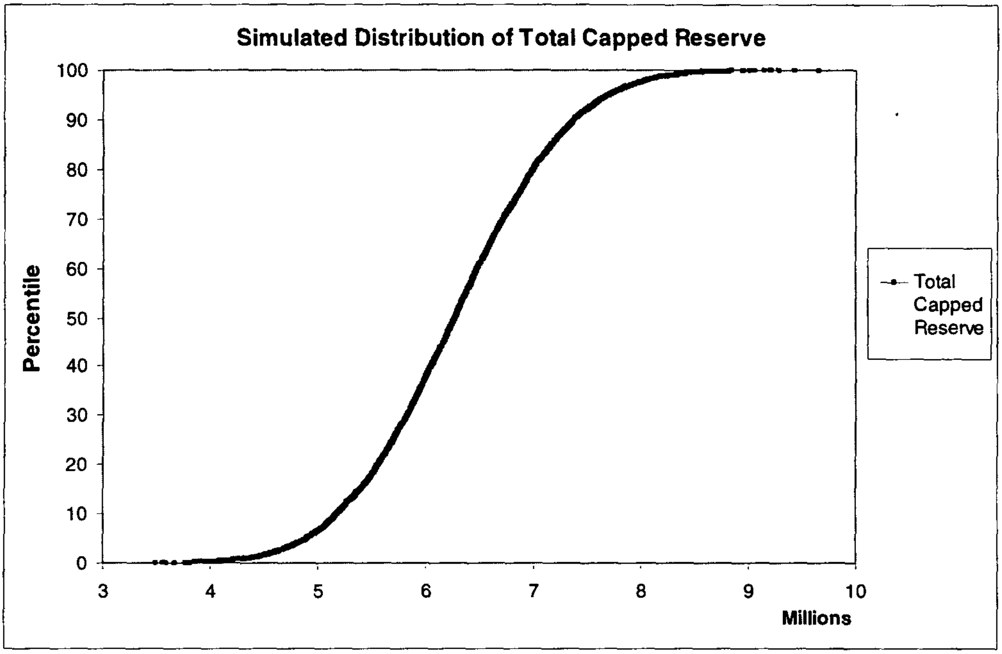

Figure 5

The table below shows the results of our projections and compares the results with those obtained by modeling the aggregate triangle using a Mack bootstrap. The 75th and 95th percentiles are given as percentages of the mean reserve. The coefficient of variation (C.o.V.) indicates the variability in the results.

Table 4

		Individual Claim Projection Method			Mack Bootstrap		
Accident Year	Mean Reserve	C.o.V.	75th percentile	95th percentile	C.o.V.	75th percentile	95th percentile
1998 (and Prior)	499,653	60.28%	120.51%	212.55%	15.60%	110.21%	126.61%
1999	2,836,912	16.52%	107.26%	135.45%	8.65%	105.73%	114.29%
2000	4,525,560	21.08%	109.64%	137.95%	25.06%	116.66%	141.34%
2001	6,582,895	24.46%	112.33%	144.18%	23.08%	115.10%	138.96%
2002	7,073,142	28.99%	114.39%	153.25%	27.41%	117.91%	146.37%
2003	12,608,970	32.81%	113.21%	161.58%	18.46%	112.19%	130.81%
2004	12,265,893	25.46%	113.75%	147.09%	29.15%	118.97%	149.76%
2005	15,134,996	30.66%	114.44%	154.19%	30.97%	120.17%	153.40%
Total	61,528,020	12.53%	107.29%	121.70%	13.56%	109.04%	123.04%

It can be seen that similar estimates are produced by the two methods for the C.o.V. of the total gross reserve. However the results for individual accident years can be significantly different. Figures 6, 7 and 8 show the gross reserve distributions for 2003, 2004 and 2005 respectively. In all three graphs, the individual claim projection method produces a distribution which is heavier in the upper tail than the aggregate modeling.

On investigation, the cohort of claims in 2003 contain a higher proportion of open large claims than average, including one claim of £6m, which results in the greater uncertainty than implied by the aggregate projection. The extra information provided by the individual claim projections arguably enables a more realistic projection of the true underlying uncertainty in the liabilities.

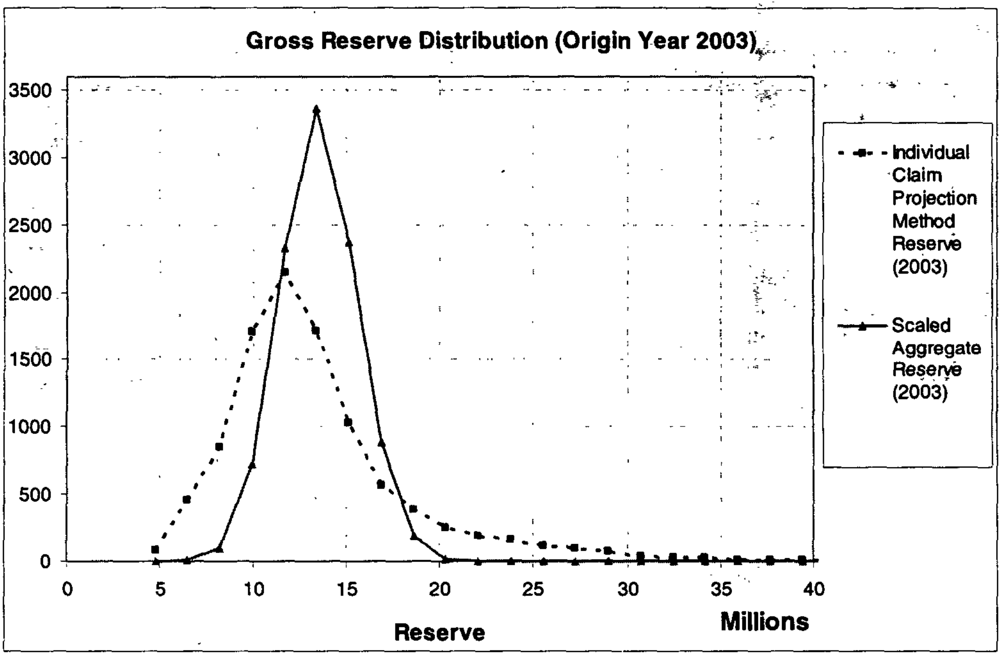

Gross Reserve Distribution (Origin Year 2003)
3500
3000
2500
2000
1500
1000
500
0
0
5
10
15
20
25
30
35
40
Reserve
Millions
Individual Claim Projection Method Reserve (2003)
Scaled Aggregate Reserve (2003)

Figure 6

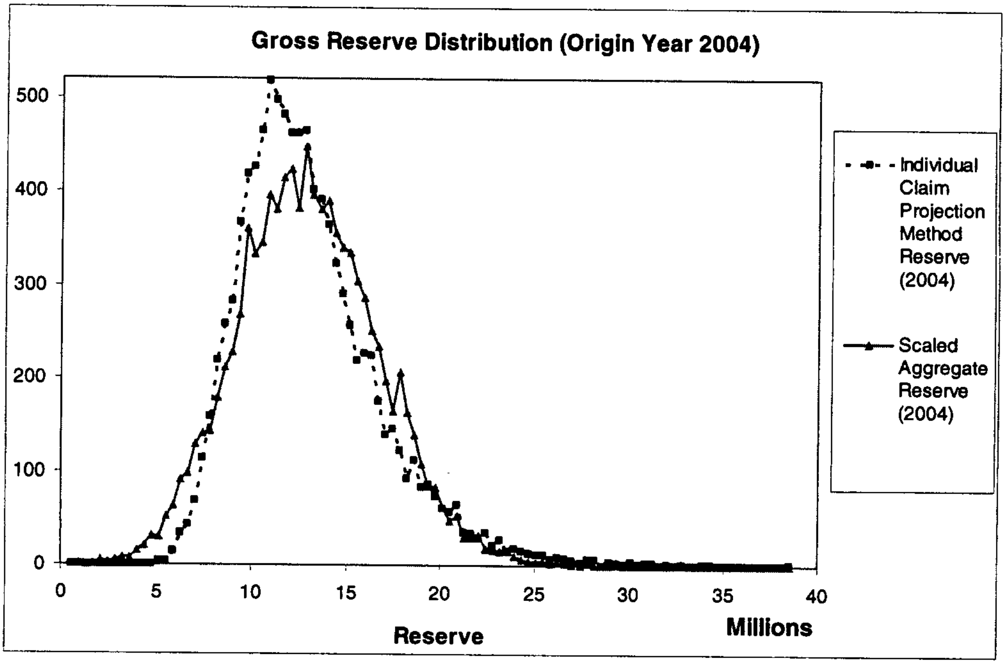

Gross Reserve Distribution (Origin Year 2004)
500
400
300
200
100
0
0
5
10
15
20
25
30
35
40
Reserve
Millions
Individual Claim Projection Method Reserve (2004)
Scaled Aggregate Reserve (2004)

Figure 7

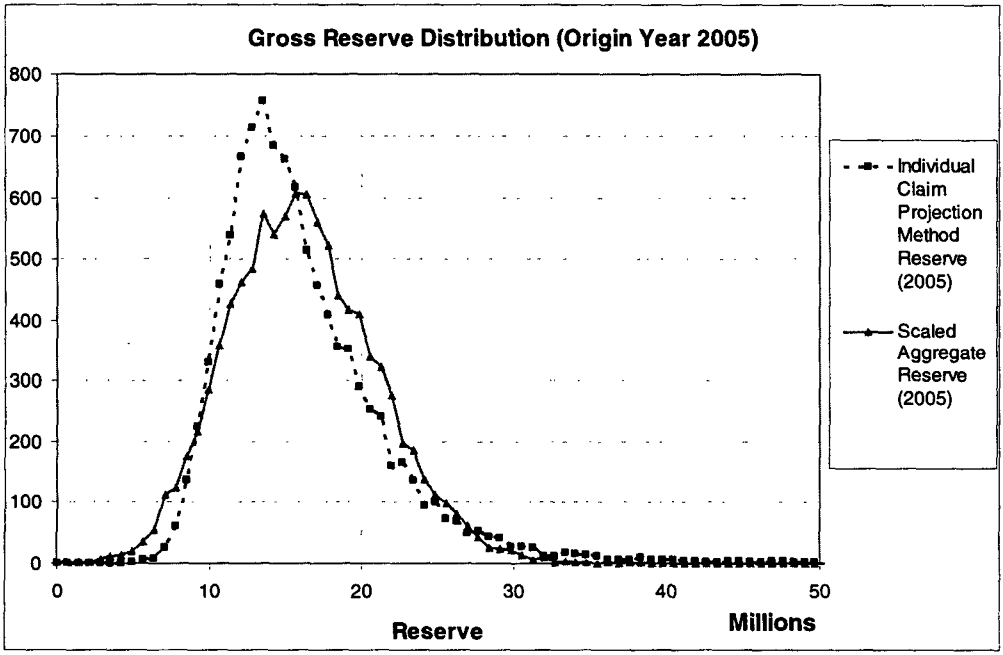

Figure 8

Figure 9 shows the cumulative distribution of the aggregate unpaid liabilities across all accident years based on the two methods. It can be seen that the two methods produce very similar results for the total gross reserve although the individual claim projection method produces a slightly heavier upper tail. This is highlighted in Figure 10, which compares the two distributions in the upper tail.

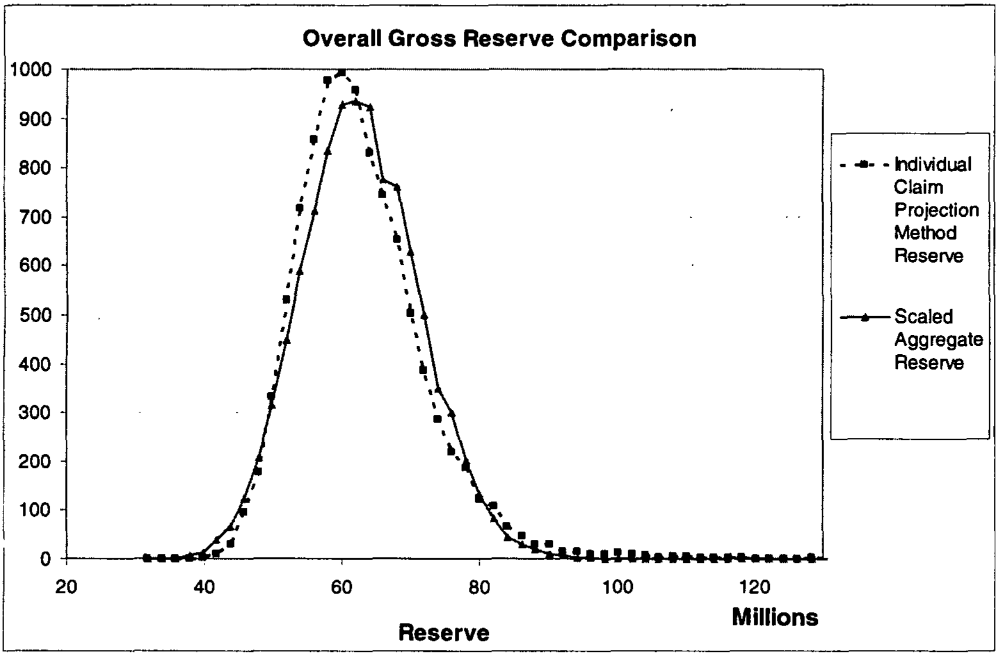
Overall Gross Reserve Comparison
1000
900
800
700
600
500
400
300
200
100
0
20
40
60
80
100
120
Millions
Reserve
Individual Claim Projection Method Reserve
Scaled Aggregate Reserve

Figure 9

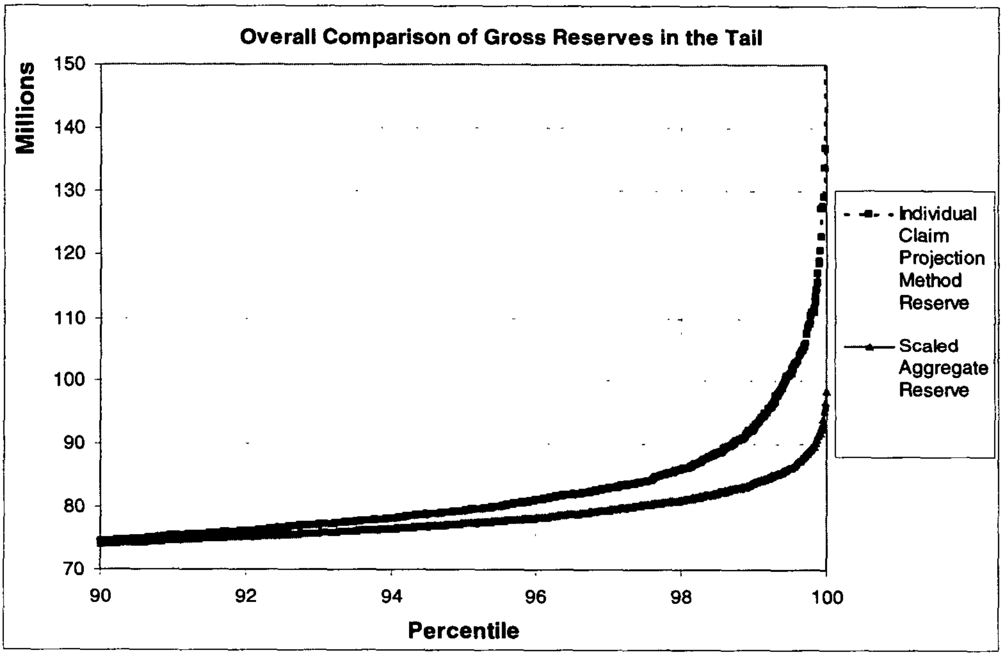

Overall Comparison of Gross Reserves in the Tail
Millions
150
140
130
120
110
100
90
80
70
90
92
94
96
98
100
Percentile
Individual Claim Projection Method Reserve
Scaled Aggregate Reserve

Figure 10

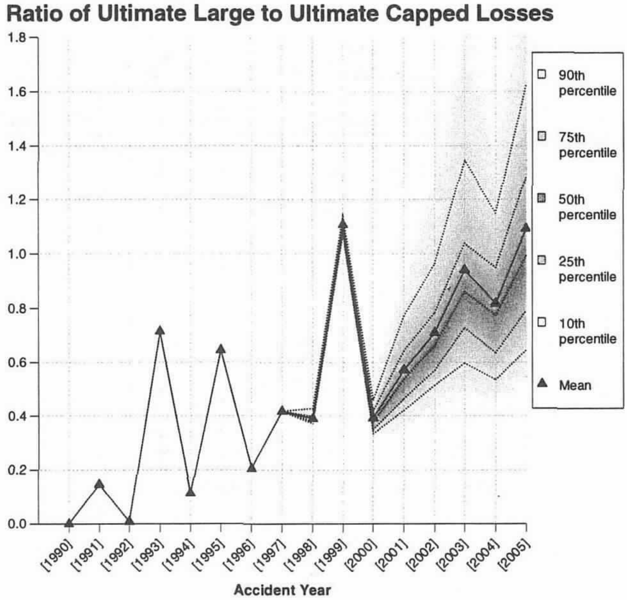

Figure 11

One reason for the heavier upper tail produced by projecting the individual large claims can be seen in Figure 11 (using the same percentile description as in Figure 1). The graph implies that the ultimate large claim proportion is increasing in recent years, the appropriateness of which can be tested in the aggregate modeling. This trend, if true, will not be allowed for adequately in the aggregate stochastic methods.

7.3 Analysis of the Net Results

Once we are comfortable with the gross results, we can calculate reinsurance recoveries on individual claims using the appropriate reinsurance terms. Table 5 shows the net results for the individual claim projection method.

Table 5

		Individual Claim Projection Method		
Accident Year	**Mean Reserve**	**C.o.V.**	**75th percentile**	**95th percentile**
1998 (and Prior)	465,519	46.96%	119.36%	200.79%
1999	977,976	26.68%	108.93%	147.16%
2000	2,983,380	13.90%	109.34%	124.26%
2001	3,496,184	18.09%	112.00%	131.59%
2002	3,457,150	23.91%	116.35%	141.11%
2003	5,755,890	17.85%	111.07%	131.46%
2004	8,518,805	16.93%	110.45%	129.51%
2005	9,999,956	16.62%	110.57%	128.90%
Total	35,654,860	8.19%	105.26%	114.12%

Figure 12 shows the net and gross reserves for the 2005 accident year. The gross reserves have been scaled to have the same mean as the net reserves. As would be expected, the netting down has resulted in a large reduction in variability.

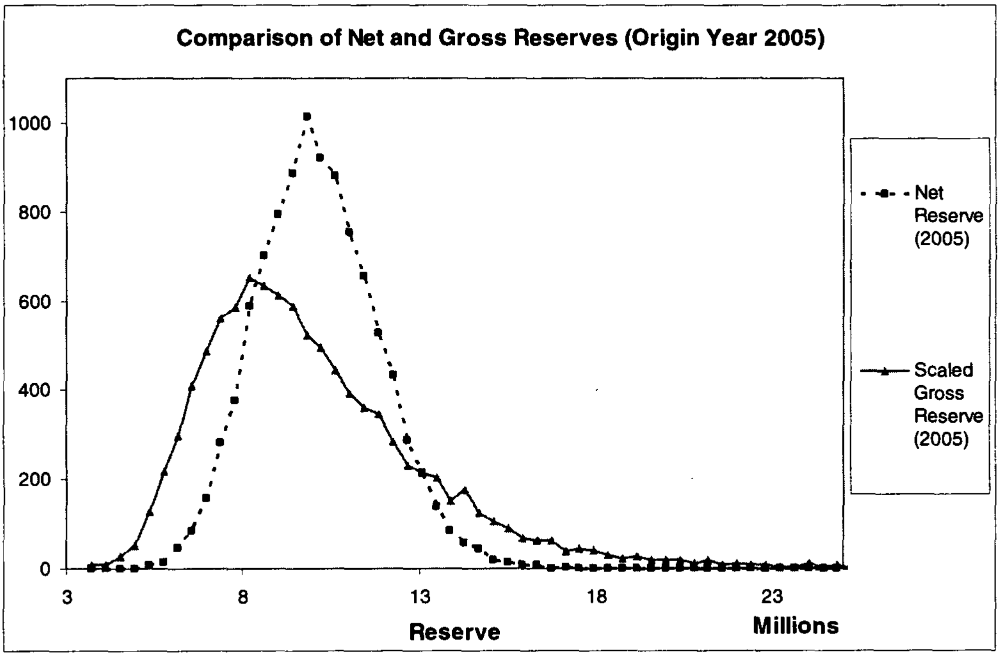

Figure 12

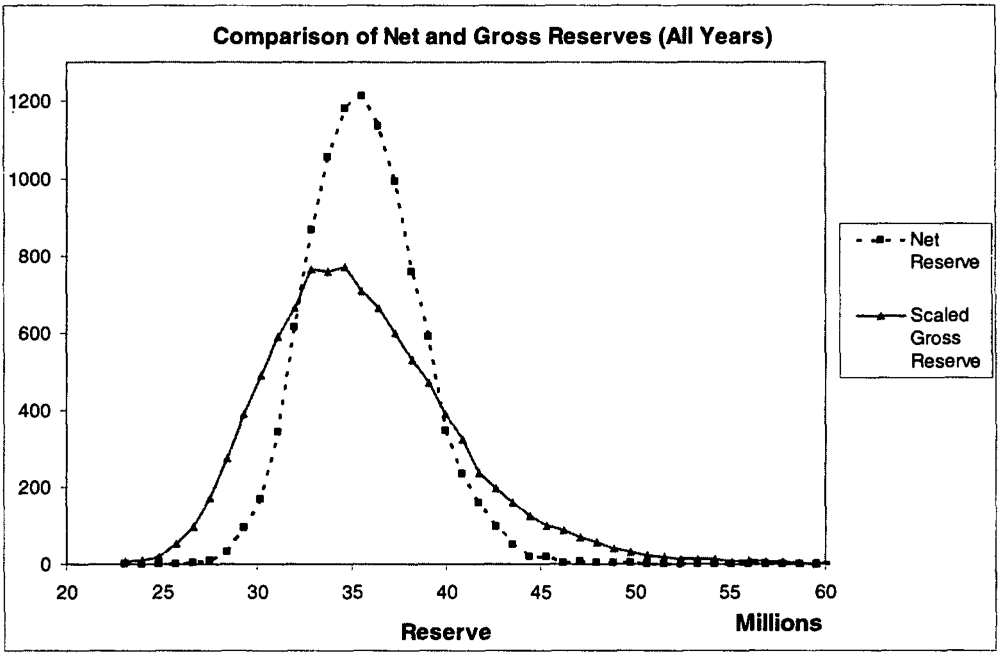

Figure 13

Figure 13 shows the overall net and gross reserves. The net reserve again shows a substantial reduction in variability.

It has already been noted that one of the additional benefits of the method described in this paper is the ability to accurately examine the performance of reinsurance cover. Figure 14 shows the distribution of recoveries associated with an aggregate deductible of £2.25m attaching to a layer of £400k in excess of £600k for the 2002 accident year. As can be seen, approximately 10% of the time the deductible is fully blown and losses pass through to the reinsurer.

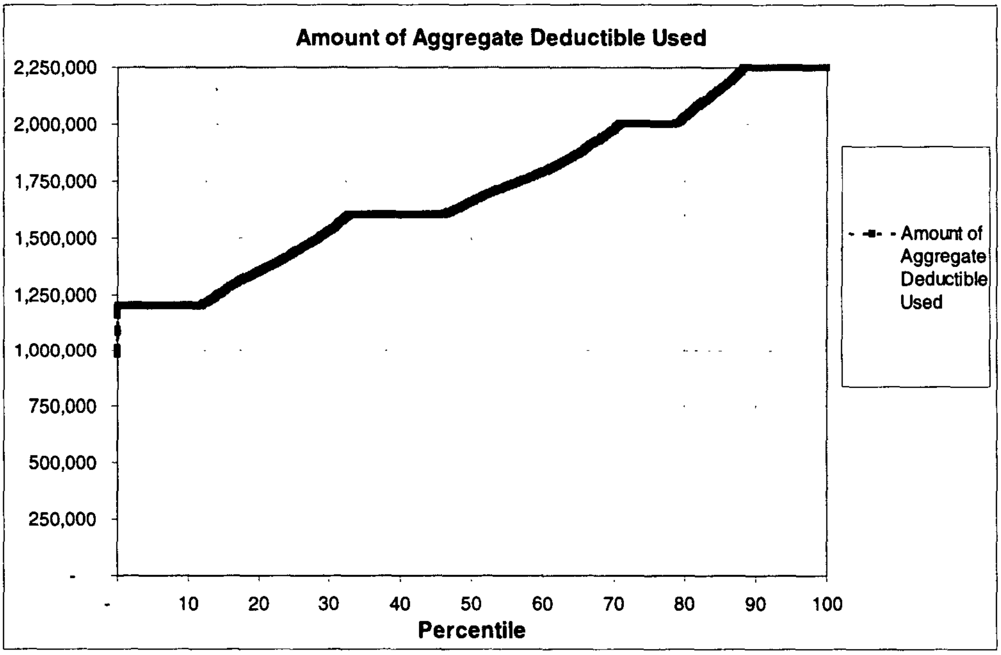

Figure 14

This allows us to consider the value of this contract and whether it represents value for money.

The method described in this paper also provides the complete predictive distributions of the gross ultimate position and ultimate reinsurance recoveries of individual large claims. Therefore the mean net ultimate position for each simulated claim can be correctly calculated. Some netting down methodologies we have seen used in practice implicitly assume that the mean of the reinsurance recoveries equals the mean of the gross claim less the retention. The one-sided nature of reinsurance means that this is flawed. The error associated with this assumption can be seen in Table 5, which shows the gross, reinsurance (RI) and net ultimate incurred position for two claims, with an excess of £1,000,000. The ultimate figures have been calculated on the our stochastic basis and also on a deterministic basis. The final two columns correspond to the stochastic calculations, where the mean net position takes into account the variability of the ultimate gross position.

Table 6

Netting Down - Comparison of Methodologies						
	Current Incurred	Mean Gross Ultimate	*Deterministic Ultimate Reinsurance Recoveries*	*Deterministic Net Ultimate*	Mean Ultimate Reinsurance Recoveries	Mean Net Ultimate
Claim 1	500,000	829,180	0	829,180	226,604	506,962
Claim 2	1,000,000	1,337,416	337,416	1,000,000	464,964	872,452

In this case, the deterministic basis is likely to lead to an overestimation of the net position, and is therefore a conservative basis. While in itself this is not a cause for concern, a desirable property of any reserving exercise would be to ensure a consistent basis for gross and net reserves.

8. INTEGRATION AND APPLICATION WITHIN CAPITAL MODELS

In recent years, there has been considerable time invested in the development of capital models to understand and quantify the risks faced by an insurance business. A significant piece of this work has been an analysis of reserving risk, which forms part of the wider insurance risk. In our experience of the UK market, there are two main methods used by practitioners to estimate net of reinsurance reserving risk. Both methods project gross aggregate triangles, with a different approach to netting down for reinsurance recoveries.

The first method arrives at net results using a deterministic net to gross ratio applied to the stochastic gross results. This method has the advantage of simplicity and transparency, however it in effect gives no credit for the expected reduction in volatility that non-proportional reinsurance should provide.

The second method projects both gross and net triangles, with some link between the projections in an attempt to ensure consistency and nonsensical simulations are avoided (for example, simulations where net reserves are higher than gross). While this should allow for the reduction in variability not captured by the first method, it is likely that the reinsurance has changed over the years (for example, reinsurance excess points have changed), and the observed historical figures may not be appropriate to apply to the newer accident years.

In forecasting the ultimate large claim severity, it is important to allow for parameter uncertainty. We would further recommend including development uncertainty. Currently, most ultimate loss generators are parameterized from some projected ultimate claim figures, allowing for IBNER, which is assumed to be known and fixed. Parameter error, using various techniques is included in the forecasting of the large claims. However the ultimate position of the claims used to parameterize the distribution are not known or fixed. To not make allowance for this will understate the true uncertainty of the underlying distribution.

9. CONCLUSIONS

Existing methods available to help gain understanding of the variability of insurance liabilities have focused on aggregate gross data, with no explicit allowance for changing mix of claims, and with no obvious adjustment to allow for non-trivial reinsurance. We have developed a method based on a small number of key assumptions to explicitly project the development of individual large claims. We show how various refinements can be made to the standard method and implement this method via a case study using actual data from a UK motor injury portfolio.

By explicitly projecting individual claims we show how to make appropriate allowance for policy limits and the reduction in variability arising from non-proportional reinsurance. By separately considering attritional and large claims, we can directly allow for changes in the mix of claims in our portfolios.

A range of diagnostics is available to the practitioner to aid understanding of the results, and to ensure it is not applied in a mechanical fashion.

Appendix 1 – Example Data

The following table shows the layer limits used in the Case Study. The lower layer lower bound is the threshold above which claims are individually simulated. The upper layer lower bound defines the boundary between 'small' large claims and 'large' large claims, in order to partition the development factors. Development Periods 11 and above have been grouped due to scarcity of data.

Table 7

Development Period	Lower Layer Lower Bound	Upper Layer Lower Bound
1	100,000	360,000
2	100,000	500,000
3	100,000	520,000
4	100,000	400,000
5	100,000	680,000
6	100,000	500,000
7	100,000	630,000
8	100,000	320,000
9	100,000	310,000
10	100,000	260,000
11+	100,000	120,000

ACKNOWLEDGEMENTS

The authors are very grateful for the work put in by John Berry to help complete this paper.

REFERENCES

[1] England, P.D. and Verrall, R.J., "Stochastic Claims Reserving in General Insurance," *B.A.J.*, **2002**, Vol 8, Part III, No. 37

[2] Mack, T., "Distribution-free calculation of the standard error of chain-ladder reserve estimates," *ASTIN Bulletin*, **1993**, Vol 30, 213-225.

[3] Renshaw, A.E. and Verrall, R.J., "A stochastic model underlying the chain-ladder technique," *B.A.J.*, **1998**, Vol 4, 903-923.

[4] Verrall, R.J., "An investigation into stochastic claims reserving models and the chain-ladder technique," *Insurance: Mathematics and Economics*, **2000**, Vol 26, 91-99

Biography(ies) of the Author(s)

Karl Murphy is Partner at EMB Consultancy LLP in the UK. Karl graduated from Trinity College, Dublin in 1990, with a first class honors degree in Mathematics and Economics. His main areas of study were in statistics and econometrics, and received the Gold Medal for his final exam results.

On graduation, Karl joined the non-life department of consulting actuaries Bacon & Woodrow, before joining EMB in 1993. Karl has considerable experience in many areas of non-life insurance, and clients include insurance companies, Lloyd's syndicates, captives and brokers. His main expertise is in personal lines insurance, and is considered to be one of the UK's leading personal lines pricing experts. Karl has played an important role in the development of EMBLEM, EMB's pricing tool that is now used by 100 insurers around the world. More recently Karl has being introducing the new capital modeling techniques developed by EMB to personal lines clients.

Karl has spoken at conferences throughout the world, including at several American Casualty Actuarial Society seminars, the South African Actuarial Society's General Insurance Convention, the UK's GIRO conference, and for the Actuarial Society of India. He co-authored a paper entitled "Using Generalized Linear Models to Build Dynamic Pricing Systems" that appeared in the CAS Forum, and has written several articles that have appeared in the insurance press on various topics.

He has served on the Institute of Actuaries' GIRO Committee and Research Steering Committee, and is currently on the Institute of Actuaries' GRIP (Pricing) committee.

Andrew McLennan is Senior Consultant at EMB Consultancy LLP in the UK. Andrew joined EMB in 2003, having previously worked in the non-life department of Trowbridge Consulting in Australia. Since moving to the UK, Andrew has been involved in a number of capital modeling projects for both personal lines clients and London Market clients. Andrew has also been involved in developing new actuarial techniques for application in the personal lending market. Andrew had considerable experience in the Australian general insurance market, including significant work with major government accident compensation schemes, including the design of the funding policy of one authority, and was involved in the due diligence review of the world's largest private equity floatation in 2002. Andrew qualified as an actuary in 2003.

The authors can be contacted by email at karl.murphy@emb.co.uk