

A Practical Guide To Measuring
Reserve Variability Using:
Bootstrapping, Operational Time
And A Distribution-Free Approach

Prepared by Julian Lowe

Contents

Section

1. Executive Summary
2. Bootstrapping
3. Operational Time
4. Distribution-Free Approach
5. Comparison Of The Results
6. Bibliography

Appendices

- I Example Of The GLIM Code For An
Operational Time Model
- II Comparison Of Some Operational Time Models
- III Breakdown Of The Steps In The Distribution-Free Approach

1. Executive Summary

The purpose of this paper is to describe some of the practical steps that are involved in various ways of measuring reserve variability. It is aimed at those who have been tempted by the theory but have been put off by the lack of explicit details of what on earth you do in practice!

The three methods described are Bootstrapping, Operational Time and Thomas Mack's Distribution-free approach. The first two of these methods were described in a 1993 Working Party paper on Variance in Claim Reserving. Following that Paper, various people expressed an interest in trying out some of the techniques, but stumbled at translating theory into practice. This stumbling was the prompt for this paper to be written.

It is not the intention to go into the theoretical considerations of the various methods in any great detail. For further information regarding Bootstrapping and Operational Time, the reader is referred to the 1993 Working Party Paper on "Variance in Claim Reserving". Thomas Mack's Distribution-free approach is described in his prize-winning CAS paper "Measuring The Variability Of Chain Ladder Reserve Estimates". Details of all three methods may be obtained from the various sources listed in the Bibliography.

The 1993 Working Party Paper compared the results from a variety of stochastic reserving methods as applied to three sets of real data. This work has been extended in this paper to include Thomas Mack's Distribution-free approach. The results of various different measures of variability are then compared (Bootstrapping, Operational Time, Distribution-free approach and Log-Linear Regression), to see how consistent they are in producing variability measures.

2. Bootstrapping

In order to put section 2.2 into context, section 2.1 is to a large extent a re-hash of the introductory section on Bootstrapping from the 1993 Working Party Paper.

2.1 So What's It All About?

The essence of Bootstrapping is to take a sample of data A, from an unknown distribution B, and then obtain information about a random variable C(A,B) by re-sampling the observed data A in an appropriate way. In a reserving context, a triangle of paid claims is the data sample "A". The unknown claims distribution is "B". The future claim payments, or Reserve, is the random variable "C(A,B)".

In a reserving context then, Bootstrapping lets us produce information about the Reserve, C(A,B), such as an estimate of its variance. A basic reserving method, such as a chain-ladder, only gives us a point estimate of the Reserve. Bootstrapping gives, in addition, an indication of the extent to which we expect the Reserve to vary either side of this expected value.

Bootstrapping can also shed light on more sophisticated reserving models. Take for example the Regression model based on Log-Incremental payments from the IOA claims reserving manual. The model of the claim process is described as:

$$\text{Log}(P_{ij}) = a(i) + b(j) + E_{ij}$$

where P_{ij} are the claim payments in Accident Year i at development period j , $a(i)$ and $b(j)$ are the parameters fitted by the model and E_{ij} is an Error term.


The method produces Maximum Likelihood Estimates for the future claims payments, $\text{MLE}(P_{ij})$. However, the MLEs are biased, that is:

$$E(\text{MLE}(P_{ij})) > E(P_{ij})$$

The MLEs are asymptotically unbiased, that is, as the sample size gets larger, $E(\text{MLE}(P_{ij}))$ gets nearer to $E(P_{ij})$. However, for "small" sample sizes, as is usually the case with reserving data, $E(\text{MLE}(P_{ij}))$ may be considerably different from $E(P_{ij})$. Bootstrapping enables this bias to be estimated and may provide a better estimate than traditional asymptotics with only a small sample size. The examination of bias was the original impetus for looking at Bootstrapping (or more generally, the Jackknife).

Some more sophisticated reserving models produce estimates of the variance of the projected reserve. Bootstrapping can give the modeller an indication of the extent to which the model variance is a result of the underlying "noise" in the data (statistical error) or due to uncertainty in the modelling process itself - such as mis-specifying the model, or the fact that the estimates of the parameters of a model are themselves random variables and contribute a degree of uncertainty to the predicted reserve.

The main steps in Bootstrapping reserve estimates are illustrated below:



The basic data triangle ((1) in the diagram above) is taken and a reserving model fitted to it. In the case of the chain-ladder, this may assume, for example, that each accident year has its own "level" (of ultimate claims) and that there is then a development pattern that is constant across all accident years. This model not only projects future payments, and hence allows one to make reserve estimates, it also produces a fitted model for the past data too ((2) in the diagram above), as will be spelt out in a bit more detail in section 2.2.

The difference between (1), the actual data, and (2), the fitted data, gives one a set of Residuals (shown as (3) in the diagram above). This is taken to be a typical representation of the extent to which the real data and the model may differ. In other words, if the data is just a realization of some random process, another realization of the process could lead to another set of data that differs at any point from the model by any one of the set of Residuals.

The trick is then to produce lots of these other possible realizations of the data by re-sampling the Residuals and adding them to the fitted model to produce lots of sets of possible data triangles, known as Pseudo-Data. For each triangle so produced, the reserving method is run, so that a series of Reserve estimates is produced - Pseudo-Reserves. If this is done many times, one will have a large collection of Pseudo-Reserves, which will have a certain distribution - that is, they will have an expected size, and will vary around that expected size by a certain amount which can be measured. It is the variation of these Pseudo-Reserves which gives us a measure of variability of the reserve estimate. These steps are described in a bit more detail in the following section.

2.2 The Steps In Practice

The following goes through the steps in arriving at the Residuals, one set of Pseudo-Data, and one Pseudo-Reserve, for a sample triangle of data. To produce an estimate of the reserve variability, many sets of Pseudo-Data and hence Pseudo-Reserves would be produced, and the distribution of those Pseudo-Reserves examined.

First of all, consider a triangle of data:

Bootstrapping Example Original Data Triangle

Accident Year	Actual Cumulative Paid Claims						
	Years of Development						
	0	1	2	3	4	5	6
1987	39,110	65,176	69,047	70,899	71,303	71,814	71,963
1988	35,877	58,094	61,884	63,330	63,980	64,254	
1989	39,907	66,009	72,310	74,273	76,390		
1990	51,296	89,666	95,878	98,097			
1991	64,109	107,279	118,753				
1992	73,944	122,541					
1993	76,050						

From this triangle we can calculate the Cumulative Paid Development factors necessary to perform a chain-ladder projection, as shown below:

**Bootstrapping Example
Original Data Triangle**

Accident Year	Cumulative Paid Loss Development Factors						
	Years of Development	1/0	2/1	3/2	4/3	5/4	6/5
1987		1.666	1.059	1.027	1.006	1.007	1.002
1988		1.619	1.065	1.023	1.010	1.004	
1989		1.654	1.095	1.027	1.029		
1990		1.748	1.069	1.023			
1991		1.673	1.107				
1992		1.657					
1993							
Vol Wtd Avg:		1.672	1.082	1.025	1.015	1.006	1.002

The development factors are just the ratios of cumulative paid claims from one period to the next, calculated in the usual fashion.

The chosen average factors can then be used to project the cumulative paid claims to ultimate. This defines a model for the claims process, namely a level for each accident year (the ultimate amount of claims), and a payment pattern by which that ultimate is reached (constant for all accident years). This model can then be used to fill in the rectangle - both for the future claims payments, and the past claims triangle, showing what the model would have predicted the historic payments to be. The fitted past cumulative payments are shown below:

**Bootstrapping Example
Fitted Data Triangle**

Accident Year	Fitted Past Cumulative Values						
	Years of Development						
	0	1	2	3	4	5	6
1987	37,924	63,418	68,614	70,330	71,400	71,814	71,963
1988	33,931	56,741	61,391	62,926	63,883	64,254	
1989	40,574	67,850	73,410	75,245	76,390		
1990	52,897	88,456	95,704	98,097			
1991	65,636	109,760	118,753				
1992	73,280	122,541					
1993	76,050						

Each accident year follows the pattern chosen, so the ratio of cumulative claims in development year 1 to those in development year 0 should be 1.672 for each accident year and so on. The last diagonal matches our actual data, and will be projected to the same ultimate as our original data triangle. What we are interested in is the incremental version of the fitted data. We are going to Bootstrap based on the Residuals between the incremental fitted payments and the incremental actual payments. We could Bootstrap based on the difference between the cumulative actual and fitted data, but to justify inferring results from the Bootstrapped reserves, we need to assume that all the Residuals are independent: this is unlikely to be the case for cumulative data. The incremental version of the cumulative fitted data, ((2) in the diagram in section 2.1), is shown below:

**Bootstrapping Example
Fitted Data Triangle**

Accident Year	Fitted Past Incremental Values (2) Years of Development						
	0	1	2	3	4	5	6
1987	37,924	25,494	5,196	1,716	1,070	414	149
1988	33,931	22,810	4,649	1,535	957	371	
1989	40,574	27,276	5,560	1,836	1,145		
1990	52,897	35,559	7,248	2,393			
1991	65,636	44,123	8,994				
1992	73,280	49,261					
1993	76,050						

From the fitted incremental claims we need to deduct the actual incremental paid claims, which are shown below:

**Bootstrapping Example
Original Data Triangle**

Accident Year	Actual Incremental Paid Claims (1) Years of Development						
	0	1	2	3	4	5	6
1987	39,110	26,066	3,871	1,852	405	511	149
1988	35,877	22,217	3,790	1,446	650	274	
1989	39,907	26,103	6,301	1,963	2,117		
1990	51,296	38,370	6,212	2,219			
1991	64,109	43,171	11,474				
1992	73,944	48,597					
1993	76,050						

This is triangle (1) in the diagram explaining the Bootstrapping process in section 2.1.

When we deduct (1) from (2) we get a triangle of Residuals, also shown below:

Bootstrapping Example

Calculated set of Residuals (3)

Accident Year	Fitted Incremental Values(2) - Actual Incremental Values(1)						
	Years of Development						
	0	1	2	3	4	5	6
1987	(1,186)	(572)	1,326	(136)	665	(97)	0
1988	(1,946)	593	860	89	307	97	
1989	668	1,173	(741)	(128)	(972)		
1990	1,601	(2,810)	1,036	174			
1991	1,528	952	(2,480)				
1992	(664)	664					
1993	0						

We could equally well deduct the Fitted Incremental Values from the Actual Incremental Values. The expected value of the re-sampled Residuals is, in this case, zero, so it makes no difference whether we subsequently add or subtract them to our Fitted data when producing the Pseudo-Data.

We can now produce as many sets of Pseudo-data as we want. We just need to pick a triangle from any of the points in the triangle of Residuals, each time picking the Residuals at random, so we can use a given Residual more than once each time we produce a set of Pseudo-Data. A typical re-sampled set of Residuals is shown below:

Bootstrapping Example

Re-sampled Residuals (4)

Accident Year	Simulated Residuals picked from whole triangle of Residuals						
	Years of Development						
	0	1	2	3	4	5	6
1987	593	0	(97)	(572)	(1,946)	307	89
1988	(1,186)	0	97	1,036	1,036	1,326	
1989	1,173	668	(664)	952	(572)		
1990	(1,946)	(97)	668	668			
1991	(1,186)	1,601	1,601				
1992	665	1,173					
1993	1,036						

To the set of re-sampled Residuals (4), we add the original fitted data. This produces a set of Pseudo-Data. This is one of the triangles that we think could equally likely have been produced from the claims process that produced the original set of claims. An example using the re-sampled Residuals is shown below:

Bootstrapping Example
Pseudo-data triangle

Accident Year	Fitted Incremental Past Data (2) + Re-sampled Residuals (4)						
	Years of Development						
	0	1	2	3	4	5	6
1987	38,517	25,494	5,100	1,144	(876)	721	239
1988	32,745	22,810	4,746	2,571	1,993	1,696	
1989	41,747	27,943	4,896	2,788	572		
1990	50,951	35,463	7,916	3,061			
1991	64,450	45,724	10,594				
1992	73,945	50,434					
1993	77,086						

We can now take this triangle of Pseudo-Data and perform our reserving method on it. To do this we need to go back to the cumulative version of the Pseudo-Data, calculate the development factors in the usual fashion, and then calculate the reserve estimates in the usual fashion. These three stages are shown below:

Bootstrapping Example
Pseudo-data triangle

Accident Year	Cumulative version of Fitted Data (2) + re-sampled Residuals						
	Years of Development						
	0	1	2	3	4	5	6
1987	38,517	64,011	69,111	70,254	69,378	70,100	70,338
1988	32,745	55,555	60,301	62,872	64,865	66,561	
1989	41,747	69,691	74,587	77,375	77,947		
1990	50,951	86,413	94,329	97,390			
1991	64,450	110,174	120,768				
1992	73,945	124,380					
1993	77,086						

Bootstrapping Example
Pseudo-data triangle

Accident Year	Cumulative Development Factors						
	Years of Development	1/0	2/1	3/2	4/3	5/4	6/5
1987		1.662	1.080	1.017	0.988	1.010	1.003
1988		1.697	1.085	1.043	1.032	1.026	
1989		1.669	1.070	1.037	1.007		
1990		1.696	1.092	1.032			
1991		1.709	1.096				
1992		1.682					
1993							
Vol Wtd Avg:		1.687	1.086	1.032	1.008	1.018	1.003

Bootstrapping Example
Pseudo-data triangle


Accident Year	Predicted future Cumulative claims						Predicted Ultimate Claims	Pseudo- Reserve
	Years of Development	1	2	3	4	5	6	
1987								
1988							66,788	2,534
1989						79,351	79,621	3,231
1990					98,172	99,940	100,280	2,183
1991				124,640	125,640	127,903	128,338	9,585
1992			135,098	139,429	140,549	143,080	143,567	21,025
1993	130,082	141,293	145,822	146,993	149,640	150,149	150,149	74,099
Total							<u>668,744</u>	<u>112,658</u>

Notice that the development factors are different from those chosen when using the original model.

To produce an estimate of the variability of the Pseudo-Reserves, we simply need to re-sample the Residuals many times, each time producing a different set of Pseudo-Data and a different Pseudo-Reserve. We will then have a whole series of Pseudo-Reserves; for example, looking at the total reserve, we may have a collection of estimates:

112,500; 110,500; 113,750; 112,000; 116,250; 107,500; and so on

We can then look at the shape of these reserves. For example, one might find that 30% were lower than 110,000, 20% were in the range 110,000 to 114,000 and the remaining 50% were above 115,000. If we have thousands of Pseudo-Reserves to observe, we will get a very smooth shape for this distribution. In fact, here's an example:



The actual numbers in the graph above are taken from a different simulation and are just to illustrate the type of shape one might expect.

2.3 Some Details Of The Steps Taken So Far

There have been a couple of areas so far where there have been leaps in going from one stage to the next, without spelling out some of the alternative ways of proceeding. One such stage is when one chooses Residuals to re-sample when producing the Pseudo-Data. In the example in section 2.2, a new triangle of Residuals was picked from anywhere in the actual triangle of Residuals. This worked fairly well, since there were not drastic differences in the size of the Residuals across the triangle. If one is looking at triangles where this is not the case, for example because the total payments in the initial years are very much larger than later payments, one may try and refine the model by partitioning the Residuals into two or three sets, so that the re-sampled Residuals in the first couple of development periods are only chosen from the first couple of periods of the original Residuals triangle and so on. This will lead to the relative variation in the size of claim remaining more nearly constant across the triangle: large total payments initially varying by 25%, say, and smaller total payments later on varying by a similar percentage.

If one lets the re-sampled Residuals in later development periods be chosen from earlier Residuals, one may find that the variation in the later payments is that much higher: fitted payments of a few hundred having re-sampled Residuals of a few thousand added to them and so on. One could argue that this is not entirely unrealistic, however, as in the tail, although one will get lower total payments, one often expects a small number of very large claims which may vary greatly in size. In any event, the impact at this far end of the triangle is not great on the total, and we are usually more concerned with the reserves for the more recent accident years. Bootstrapping is generally far too crude to infer anything reliable for older accident years, after five or more years development say. In the comparison of the methods in section 5, the re-sampled Residuals have been chosen from the entire triangle, and where the results are sensible, they correspond fairly well to those from other methods.

The other step implicitly taken in the example is that when calculating the Pseudo-Reserve, the actual cumulative payments to date were deducted from the Ultimate claims, as calculated from the Pseudo-Data. The alternative would be to deduct the cumulative claims taken from the Pseudo-Data triangle. Deducting the Pseudo-Data cumulative claims leads to a very low degree of variability: whenever one has a low set of Pseudo-Data one has a low Pseudo-Ultimate and vice versa. The variability of the Pseudo-Reserve is therefore greatly reduced. One could argue that we are interested in the difference between the Pseudo-Ultimate and the actual cumulative paid claims to date. We are saying that the actual data is just one realization of some random process and all the other re-sampled sets are just as likely, so we could just as easily have ended up with all the other Pseudo-Reserve estimates as our actual reserve estimate. Certainly comparing the Bootstrapping SE's with other methods, the results look far too low if one deducts the Pseudo-Data cumulative payments, and broadly similar if one deducts the actual Cumulative payments to date, all of which tends to suggest sticking to Pseudo-Ultimate less Actual for the Pseudo-Reserve.

Finally, there is the mechanism by which information on the Bootstrapped reserves is collated. For the sample data looked at in section 5, the Add-In package @Risk was used. This can be added on to Lotus, Symphony or Excel spreadsheets, amongst others. It basically just re-calculates the spreadsheet many times, letting certain cells in the spreadsheet be realizations of a random variable, and collates the results in a friendly fashion. There is no reason why this cannot be done in a normal spreadsheet without the Add-in, however, and this has been easily done in practice. A simple macro re-calculates the spreadsheet, each time picking a new set of re-sampled Residuals, and storing the results of each simulation in a separate area of the spreadsheet. The re-sampling can be achieved by, for example, placing all the Residuals in a numbered table, and using the spreadsheet's random number generator to pick out a new set of Residuals. The results can then be analysed and measures of variability such as standard errors and so on calculated.

3. Operational Time

3.1 Recap Of The Model

The Operational Time model is described in Tom Wright's CAS paper "Stochastic Claims Reserving When Past Claim Numbers Are Known". The Operational Time model attempts to represent the underlying claims settlement process. The starting premise is that the cost of settling claims and the order in which they are settled are related - that is, typically, the longer the period to settlement, the greater the final settlement cost is likely to be. The method therefore develops a model of the claim settlement cost as a function of the relative proportion of claims settled (this time-frame is known as Operational Time). In fitting this model, data is required for the amount and number of claims settled, as is an estimate of the ultimate number of claims.

3.2 When Is It Likely To Be Useful?

The method is likely to be of most use where the greatest cause of uncertainty in predicting ultimate claims is due to individual claim costs - for example, Motor Bodily Injury (BI) and other BI classes. It should also be of particular use when it is believed that settlement rates are changing, as the model may be able to capture these changes more effectively than traditional link-ratio approaches. Because it is a statistical model, standard errors of the reserve estimates can also be calculated to enable a view on the variability of reserves to be taken. It is also then possible to identify the components of this error.

3.3 Some Limitations

The Operational Time model does not utilise the case reserves and relies purely on the amounts of paid claims and numbers of claims paid and reported. Clearly information on case reserves is often a significant element of the reserving process. The model is sensitive to the estimated future number of settled claims, and these estimates need careful scrutiny. Inflation is a parameter that may be modelled and this is also an area where close scrutiny is required.

3.4 Outline Of The Main Steps

Operational Time (τ) is the number of claims closed to date expressed as a proportion of the ultimate number of claims. Thus, for a given accident year, it starts at zero and increases to 1. Transforming into Operational Time eliminates the need to model settlement rates (although assumptions about settlement rates are made at various stages in the process). The transformation also makes estimating the standard errors of claims more straightforward: larger claims tend to take longer to settle and have different claim size distributions than claims settled earlier; when modelling in development time, the time of settlement of a given claim will also be uncertain, so the appropriate claim size distribution is uncertain too - modelling in Operational Time avoids this difficulty.

Broadly the steps involved in the Operational Time model are:

1. Estimate the Ultimate number of claims for a given accident year. This can be done by traditional methods. This leads to ...
2. Calculate $\tau(i,d)$ for each accident year (i) and each development period (d) to date. This is done simply by dividing the number of claims settled to date by the estimated Ultimate number of claims (this is the definition of Operational Time). $\tau(i,d)$ may be different at a given point of development time for two given accident years if the settlement rates have changed. For the claims settled in a given period, say a quarter, the Operational Time is taken to be the average Operational Time during that quarter.
3. Calculate the average payment, $A(i,d)$, for each accident year and development period to date.

4. We can now express all the average payments $A(i,d)$ as $A(\tau)$, since, for a given (i,d) we know the corresponding τ . So, if we have a data triangle with "n" data points, we now have "n" mean payments A_1, A_2, \dots, A_n corresponding to "n" Operational Times $\tau_1, \tau_2, \dots, \tau_n$. At this stage we can fit a model to the mean claim size as a function of τ , $m(\tau)$, based on the "n" sample values $A(\tau_1), \dots, A(\tau_n)$. This is done using the Generalized Linear Modelling package GLIM. This enables us to fit a model to the mean claim size data with explicit assumptions about, for example, the error structure used in the fitting process, and calculate not only the parameters of the model, but also the standard errors of the parameters of the model. With some manipulation, estimates of the standard errors of future payments can then be made as well. The GLIM stage of the process is the only stage which needs anything other than straightforward manipulation in a spreadsheet. It is described further in section 3.5.

5. Estimate the run-off of the future numbers of settled claims for each accident year from the current date to Ultimate. This just involves producing a pattern for the settlement of the claims which we have estimated are still to be settled. In turn we can then...

6. Calculate the Operational Time, τ , for these future periods between the current date and Ultimate. This is obtained just by dividing the number of claims settled at a given point in time by the estimated Ultimate number of claims, in the same fashion as that done in stage 2.

7. Finally, we can now combine our estimates of the run-off (in Operational Time) of future numbers of settled claims from 6., with our model of the mean claim size as a function of τ from 4., to arrive at a stream of future payments and hence a reserve estimate (and by some manipulation, standard errors of the reserve estimate).

3.5 Using GLIM To Fit The Model

GLIM usually assumes that models have one of a certain number of standard error structures, such as a Normal or Gamma error structure. The Operational Time model has a non-standard error distribution, however, and so the model has to be set up using the user-defined facilities of GLIM. GLIM allows the systematic part of the model to include a variety of different terms. Typically these are of the form:

$$m(\tau) = \exp(\beta_0 + \beta_1\tau + \dots + \beta_n\tau^n)$$

In the nomenclature of GLIM, $\sum \beta_i \tau^i$ is the Linear Predictor, and log is the function (the Link function) that links the mean to the linear predictor ($\log(m(\tau)) = \sum \beta_i \tau^i$). Other terms such as $\log(\tau)$ can be used in conjunction with the various polynomial terms. In practice, polynomials up to, say, degree eight seem to enable a curve for $m(\tau)$ to be fitted satisfactorily in most cases. A printout of a GLIM (version 3.77) command file necessary to fit an Operational Time model is given in Appendix I.

The command file is fairly heavily annotated to describe what each line does. The first few lines read in the data. The items required are the average costs (c), the Operational Times (t), the origin years (y), the development periods (d) and the numbers of claims (n). These five items of data are assumed, in this instance, to be in a single file (which GLIM prompts the user to name) in five consecutive columns, in a file not more than 100 characters wide.

Two macros are defined, sres and dres, to produce output necessary to look at the standardised residuals and the deviance residuals. The particular GLIM command file shown does not go on to use these macros. If the user did want to do so, the following line would have to be added:

```
\use sres\use dres
```

The deviance residuals are the signed square root of the contribution that each observation makes to the deviance, where the sign depends on the difference between the actual and fitted data. They may be a better guide than standardised residuals for models which are non-Normal, for which standardised residuals may be markedly skew.

The tricky bit of the command file is in defining the non-standard error distribution. GLIM uses a technique called Fisher's scoring method to iteratively estimate the parameters in a model. This requires four pieces of information:

- (i) the relationship between the fitted values and the linear predictor,
- (ii) the derivative of the linear predictor with respect to the fitted values,
- (iii) the variance function,
- (iv) the contribution each observation makes to the deviance.

Normally when using GLIM, the user does not have to fret about what these functions are. One would normally choose from standard combinations of error/link functions, for which the above four pieces of information are pre-defined. In this case, (i) and (ii) are simple enough, as we know that we have a Log Link function connecting the linear predictor (lp) and the fitted values (fv). We can therefore see that the $fv = \exp(lp)$, and that $d(lp)/d(fv) = 1/fv$. The other two items are defined by the error structure assumed. Fortunately, Tom Wright sets out precisely how to calculate these last two pieces of information in the paper referred to previously. The four pieces of information are defined in the command file as the four macros $m1$, $m2$, $m3$ and $m4$, which are printed in bold in Appendix I, to show how important they are, and are reproduced below:

```
\mac m1 \cal %fv=%exp(%lp) \endmac
\mac m2 \cal %drl=1/%fv \endmac
\mac m3 \cal %va=(%fv*2)/n \endmac
\mac m4 \cal %di=wt*2*n*(-%log(%yv/%fv))+(%yv-%fv)%fv \endmac
```

wt is the weight given to each point: in the command file, negatives are weighted out. n is the number of claims read in previously, $\%fv$ are the fitted values, $\%yv$ are the values we are modelling.

The final section of the command file then uses the `\fit` command to fit any combination of terms the user requires. In the example given, the model fitted is an accident year level (YR) for each accident year, a constant force of inflation (I), and a series of polynomial terms of τ (T1, T2, T3 and so on). The user can specify how much information about the fit of the model he wishes to see using the `\disp` command. In the example, `\disp l e` shows the components of the linear predictor (l) and the parameter errors and their standard errors (e).

If the user wishes to use the output from the GLIM model in other programs, he can add a line to extract some of the results of the model to a file. For example, the following line extracts the parameter estimates and the variance-covariance matrix as vectors to a file:

```
\text %pe %vc
```

Alternatively (and this is the way the author has done it), the user can output all the results to a GLIM.log file, which can then be imported into, say, a spreadsheet, and manipulated as required.

Once one has obtained the parameters of the model, you then have a model of average claim size, $m(\tau)$ as a function of Operational Time. As we can calculate Operational Time for all future periods of time (stage 6 in section 3.4), we therefore have a set of average claim sizes at all future development times. Combining this with our projection of the future numbers of claims, gives us figures for all future payments. In practice, this has been done at quarterly intervals, although one could do this stage more precisely - as set out in Tom Wright's paper. There seems to be very little difference in the results, however one makes the link between payments in Operational Time and payments in real time.

The 1993 Working Party Paper explained that in using Operational Time models, one makes various initial assumptions which can then be examined and possibly relaxed. One such assumption is that the coefficient of variation of individual claim amounts is the same for all Operational Times. Generally the assumption of constant variation is not unreasonable. One can review the constant variation assumption and, if desired, adjust the model, so that bigger claims have more or less variation than smaller claims. One does this by replacing the macros m3 and m4 above with the macros m5 and m6 below. Rather than have $\%va = \%fv^{**2}/n$ in the m3 macro, which implies a constant coefficient of variation, one can have the variance function being a different scalar power of $\%fv$. This allows the coefficient of variation of individual claims to depend on the mean claim size. The GLIM code to achieve this is illustrated below:

```
\mac m5 \cal %va=(%fv**%a)/n \endmac
\mac m6
\cal pt1=(%yv**(1-%a)-%fv**(1-%a))/(1-%a) \
\cal pt2=(%yv**(2-%a)-%fv**(2-%a))/(2-%a) \
\cal %di=wt*2*n*(%yv*pt1-pt2) \
\endmac
```

One can then define different values of %a within the command file by, for example, adding a line before the "fit" directive:

```
\cal %a=1.6\
```

If the Residual plots against Operational Time tend to show a decreasing variance, this suggests trying a smaller value of %a and vice versa.

So, we can now fit a variety of models for $m(\tau)$. The question arises as to how to home in on a "good" model. If we consider fitting polynomials up to degree eight, there are 255 possible polynomials. The author took the approach that, with the speed of modern PC's, it was feasible to look at all possible polynomial models in an automated fashion, rather than examine a series of models on an ad hoc basis. The approach taken was to fit all possible 8-degree polynomials in one big GLIM command file, and home in on a "good" model by looking at a combination of statistics and diagnostics, including the overall deviance of each model (a measure of how closely it fits the data), the significance of the model (as measured by a simple T-test on the parameter estimates and their standard errors as calculated by GLIM) and the number of terms used in the model.

This was done by using a Lotus macro to create a series of GLIM command files. The Lotus spreadsheet initially prints a .PRN file, containing a series of GLIM commands, which GLIM can then use as a series of instructions to look at the fit of all possible 8-degree polynomial models. A summary of the diagnostics of this initial run, just giving information on the goodness of fit of the model, is then read back into Lotus. This is done by importing the GLIM.log file - to which the output of the GLIM fitting is directed. The Lotus macro then dissects this information and uses it to print a new .PRN file, this time instructing GLIM to output a greater range of diagnostics just for the best hundred models. The output of this second run is then used to choose a final selection of ten models, and this time a GLIM command file is printed, as a .PRN file again, to produce a complete set of diagnostics. This final stage runs to several hundred lines of output, so it would not be practical to perform this level of analysis for all 255 models at the initial stage. Although all the above may sound a bit cumbersome, it only takes about five minutes on a 66Mhz 486 machine.

This sifting process was used to examine how sensitive the final result is to the actual model chosen. After all that effort, the answer is not very!! This is not that surprising when you think about it. The model is fitting an n degree polynomial (with n less than nine) to a set of points in the range (0,1). Apart from some ridiculous models, like a straight line or a quadratic, any model with a reasonable number of parameters will fit a fairly coherently grouped set of data pretty well.

To illustrate the sensitivity, or otherwise, of the results to the model chosen, Appendix II gives a summary of the reserve estimates produced by the "top 50" polynomial models for one of the 1993 Working Party sets of data. For those of you who are interested, this is Class 3 from the 1993 Working Party Paper. For this set of data, all the reserves from the "top 100" were within $\pm 2\%$ of one another and the "top 10" were within $\pm \frac{1}{2}\%$. The original model chosen for the Working Party came 36th according to the grading system adopted, but the implied reserve was within $\frac{1}{4}\%$ of all the more highly graded models.

The upshot of the lack of sensitivity of the results to the chosen model is, in the author's opinion, that the user of such a model need not be too worried about going wrong at the GLIM fitting stage of the exercise (assuming one can get GLIM to work in the first place!). Where the results seem to be more sensitive, is in arriving at the ultimate number of claims and the settlement pattern of those claims. For most classes, ultimate claim numbers are relatively easy to model, and basic chain-ladder techniques should yield fairly consistent answers for the numbers of ultimate claims and their development.

3.6 Calculating The Standard Errors

So far we've concentrated on the reserve estimates only. Having fitted the models using GLIM, it is possible to produce estimates of the standard errors, and break down the reserve variability into components from different sources. This can be done in a spreadsheet by manipulating the GLIM output. This involves a series of matrix manipulations, which are set out in Tom Wright's original paper. Some of the algebra looks daunting, but basically you just have to follow your nose.

Because one is making explicit assumptions about various elements of the claims process, and using a statistical model to fit the parameters, it is possible to produce variability estimates for various different components of reserve uncertainty. Tom Wright identifies four sources of variability due to uncertainties in:

- (i) severity (of claim payments),
- (ii) parameter uncertainty,
- (iii) future inflation,
- (iv) frequency (of claim payment).

The first component is relatively easy to estimate. The variance of individual claim payments is given by:

$$\text{Var}(P(w,d)) = N(w,d) \times \phi^2 \times m(\tau)^2$$

Where $P(w,d)$ and $N(w,d)$ are the amounts and numbers of payments respectively for origin year w in development period d , ϕ is the coefficient of variation referred to earlier, and is estimated as part of the fitting process as the deviance (as produced from the GLIM fitting stage) divided by the degrees of freedom of the fitted model. The sum of these variances give the variance for the reserve as a whole.

The second component is probably the trickiest - see Appendix D of Tom Wright's paper. From this, we know that

$$\text{Variance due to parameter uncertainty} = \delta^T \times V \times \delta$$

Where V is the variance-covariance matrix of the parameter estimates and δ is a (column) vector of first derivatives of the outstanding claim amounts.

For a model of the form $m(\tau) = \exp(\beta_0 + \beta_1\tau + \dots + \beta_n\tau^n)$, the derivatives are relatively easy to calculate:

$$dm/d\beta_i = \tau^i \times m(\tau)$$

The $dm/d\beta_i$ components of δ , for example, will be a triangle of terms consisting of the incremental paid claims in each period, multiplied by the Operational Time cubed. Each column vector element of δ will be the sum of the rows of this $dm/d\beta_i$ triangle. A separate triangle is needed for each parameter. Although the whole triangle is calculated, it is only the row totals we are after which form the elements of δ .

If we model inflation, i , as an additional term in the model, $\exp(it)$, then similarly:

$$dm/di = t \times m(\tau)$$

The variance-covariance matrix can be output directly from GLIM (using the `\ext` command mentioned previously), and armed with the above expressions for the derivatives of the outstanding claim amounts, we can arrive at the parameter uncertainty by some simple matrix manipulation to obtain $\delta^T \times V \times \delta$. The square roots of the leading diagonal of this matrix gives the parameter uncertainty estimates.

Appendix F of Tom Wright's paper explains how to arrive at the variability due to claims inflation. In fact a slightly different route, cutting a few corners, was used by the author. An overall parameter variability measure was calculated, as described above, including the inflation parameter. The process was repeated, but excluding the inflation parameter. We thus arrive at two components of "parameter uncertainty excluding the inflation parameter" and "parameter uncertainty including the inflation parameter". The inflation uncertainty is approximately the square root of the difference of the squares of the two items above.

Finally, Tom Wright also sets out how the uncertainty due to claim frequency can be calculated, namely:

$$\text{Frequency variation} = (\tau_0 \times m(\tau_0) + \mu/M) \times \upsilon$$

Where τ_0 is the latest Operational Time reached for a given accident year (so we have already calculated this); $m(\tau_0)$ is the average claim cost associated with that time, which we can calculate from our model; μ is the estimated outstanding claim amounts for that accident year; M is the estimated ultimate number of claims for that accident year; υ is the standard error of M . All the above components, bar υ , are already known. υ can be calculated using standard log-linear regression techniques, for example.

Assuming the four components of variability are mutually independent, we can obtain a measure of the total variability by taking the square root of the sum of the squares. The 1993 Working Party gave some examples of such components of variability.

4. Distribution-Free Approach

4.1 Introduction

Thomas Mack has written a series of papers on the subject of the variability of chain-ladder estimates, most notably the CAS prize-winning paper "Measuring The Variability Of Chain Ladder Reserve Estimates". In the CAS paper, he derives a formula for the standard error of chain-ladder reserve estimates without assuming any specific claim amount distribution function. For ease of reference, the techniques used by Thomas Mack are described as the Distribution-free approach.

This section gives a brief summary of the paper, sets out how the detailed formulae in Thomas Mack's paper(s) can easily be broken down into a series of simple calculations in a spreadsheet, and briefly describes some of the diagnostic checks that can be carried out to test the assumptions that form part of the model.

4.2 The Assumptions

The foundation of the Distribution-free approach is the observation of three main assumptions which are shown to underlie traditional chain-ladder techniques. These are:

$$(i) E(C_{i,k+1}|C_{i1},...,C_{ik}) = C_{ik}f_k, 1 \leq i \leq I, 1 \leq k \leq I-1,$$

$$(ii) \{C_{i1},...,C_{iI}\}, \{C_{j1},...,C_{jI}\}, i \neq j, \text{ are independent,}$$

$$(iii) \text{Var}(C_{i,k+1}|C_{i1},...,C_{ik}) = C_{ik}\sigma_k^2, 1 \leq i \leq I, 1 \leq k \leq I-1.$$

Where C_{ik} denotes the accumulated total claims amount of accident year i up to development year k , f_k is the development factor from k to $k+1$, and σ_k are parameters. The first two assumptions seem intuitively sensible, although these can be demonstrated to be the implicit assumptions of the chain-ladder more formally. The third assumption is induced from the fact that the estimator of f_k is the C_{ik} -weighted mean of the individual development factors.

Thomas Mack goes on to show that a corollary of assumption (iii) is that the development factors are not correlated. That is, if we have a particularly high development factor in one period, there is no tendency for the subsequent factor to be particularly low (or high).

4.3 The Main Results

The estimate of the standard error of the reserve estimate for accident year i , \hat{R}_i , is:

$$\text{SE}(\hat{R}_i) = \hat{C}_{i1} \sum_{k=1}^{I-1} \frac{\hat{\sigma}_k^2}{\hat{f}_k^2} \left(\frac{1}{\hat{C}_{ik}} + \frac{1}{\sum_{j=1}^{I-k} C_{jk}} \right)$$

The estimate of the standard error of the reserve estimate for all accident years combined, \hat{R} , is:

$$\text{SE}(\hat{R}) = \sum_{i=2}^I \left(\text{SE}(\hat{R}_i)^2 + \hat{C}_{i1} \left(\sum_{j=i+1}^I \hat{C}_{j1} \right) \sum_{k=i+1}^{I-1} \frac{2\hat{\sigma}_k^2}{\hat{f}_k^2 \sum_{n=1}^{I-k} C_{nk}} \right)$$

C_{ik} , f_k and σ_k are just as before; a hat indicates an estimator of the particular figure. The estimators are as follows:

$$\hat{f}_k = \frac{\sum_{i=1}^{I-k} C_{i,k+1}}{\sum_{i=1}^{I-k} C_{ik}}, 1 \leq k \leq I-1$$

This is the traditional volume-weighted chain-ladder estimate of the development factors. Different estimators for f_k can be used, in which case the algebra is altered slightly.

$$\hat{C}_{i1} = C_{i,i+1-i} \prod_{j=i+1}^{I-1} \hat{f}_j$$

This is just the traditional chain-ladder method of calculating the ultimate claim amounts by multiplying the latest diagonal by all future development factors.

$$\hat{\sigma}_k^2 = \frac{1}{1-k-1} \sum_{i=1}^{1-k} C_{ik} \left(\frac{C_{i,k+1}}{C_{ik}} - \hat{f}_k \right)^2, 1 \leq k \leq 1-2$$

This is shown to be an unbiased estimator of σ_k^2 . An estimate is needed for σ_{1-1} ; this can be obtained in several ways. One way is to extrapolate as follows:

$$\hat{\sigma}_{1-1}^2 = \text{Min} \left(\frac{\hat{\sigma}_{1-2}^2}{\hat{\sigma}_{1-3}^2}, \text{Min}(\hat{\sigma}_{1-3}^2, \hat{\sigma}_{1-2}^2) \right)$$

Although the previous formulae look quite daunting, they consist of nothing more than basic arithmetic - addition, multiplication and so on. With a clear head, they can easily be programmed into a spreadsheet. There are no matrices to manipulate, no regression to perform, no GLIM models to be fitted and no simulations to be collated. Once the formulae have been set up, a new set of data can be imported into a spreadsheet, say, and a simple "calc" of the spreadsheet will yield the estimates of the standard errors of the reserves for each accident year, and the reserve as a whole, for the new set of data.

The author replicated the results in Thomas Mack's 1993 ASTIN paper, "Distribution-Free Calculation Of The Standard Error Of Chain-Ladder Reserve Estimates", by breaking down the formulae above into a series of shorter, less daunting, functions. The intermediate steps in this process are shown in Appendix III. III.1 shows the original cumulative data triangle, projected to ultimate, and the estimated development factors. III.2 shows the calculation of the first two intermediate functions, D_{ik} and E_k . These are defined to be:

$$D_{ik} = C_{ik} \left(\frac{C_{i,k+1}}{C_{ik}} - \hat{f}_k \right)^2$$

$$E_k = \sum_{j=1}^{1-k} C_{jk}$$

In the Appendix, the subscripts are shown in brackets and σ_k is shown as $s(k)$ - formatting in Lotus not being quite as friendly as in other packages! The D_{ik} are then summed and scaled to give the σ_k estimates, according to the formula at the top of the page.

III.3 defines one other function, F_{ik} , which is all that is needed to calculate the variance of the reserve estimates for individual accident years, $SE(\hat{R}_i)^2$, as follows:

$$F_{ik} = \frac{\hat{\sigma}_k^2}{\hat{f}_k^2} \left(\frac{1}{C_{ik}} + \frac{1}{E_k} \right)$$

$$SE(\hat{R}_i)^2 = C_{ii} \sum_{j=1}^I F_{ij}$$

This last expression is the estimate of the variance of the individual accident year reserve estimates - quite painlessly achieved.

To obtain the expression for the overall reserve variance, a few more functions are defined:

$$G_i = \sum_{j=i+1}^I C_{ji}$$

$$H_k = 2 \frac{\hat{\sigma}_k^2}{\left(\frac{\hat{f}_k^2}{E_k} \right)}$$

$$J_i = \sum_{j=1, j \neq i}^{I-1} H_j$$

$$K_i = SE(\hat{R}_i)^2 + C_{ii} G_i J_i$$

The estimate of the overall reserve variance is then, finally:

$$SE(\hat{R})^2 = \sum_{i=1}^I K_i$$

This is the result shown on III.4 of the Appendix. The figures correspond to the illustration in Thomas Mack's ASTIN paper.

4.4 Some Diagnostic Checks Of The Assumptions

The assumptions set out in section 4.2 can be validated, or otherwise, by some simple diagnostic checks.

Assumption (i) is checked by plotting $C_{i,k+1}$ or $R_{i,k+1}$ against C_{ik} . This should show a random spread about the diagonal. If this is not the case, it is possible to make an alternative assumption (i)':

$$(i)' E(C_{i,k+1}|C_{i1}, \dots, C_{ik}) = C_{ik}f_k + g_k, \quad 1 \leq i \leq I, \quad 1 \leq k \leq I-1$$

Assumption (ii) can be checked by plotting $\frac{\hat{R}_{i,k+1}}{\hat{\sigma}_k}$ against $i+k$ (that is, calendar year), again looking for a random spread.

Assumption (iii) can be adjusted by inspection of the model. For example, typically the $\hat{\sigma}_k$ will show some sort of decreasing progression from one year to the next, and it is possible to assume that $\sigma_k^2 \approx e^{(d-ck)}$, which leads to a revised (iii)':

$$(iii)' \text{Var}(C_{i,k+1}|C_{i1}, \dots, C_{ik}) = C_{ik}\sigma_k^2 e^{-ck}, \quad 1 \leq i \leq I, \quad 1 \leq k \leq I-1.$$

This leads to a more stable set of reserve variance figures. Under the original assumption (iii), the variance figures tend to jump up and down from one year to the next.

It is also possible to test various other aspects of the model, or look for distorting effects, such as a calendar year effect. Tests for correlations between development factors and calendar year effects are set out in Appendices G and H of Thomas Mack's CAS paper.

Some examples of the various diagnostic checks that can be performed will be shown at the Workshop session.

5. Comparison Of Methods

5.1 Introduction

The purpose of comparing the measures of variability of the different methods is to some extent by way of validation. Firstly, it is reassuring if there is a broad consensus as to the extent to which reserves can be expected to vary. Secondly, such a consensus would mean that the potential user could be just as happy with the broad indication of variability given by the simpler methods as opposed to the more complicated methods. Measures of relative variability may, for example, be required when considering the dynamics of capital requirements or performance measurement, so they are of interest for non-reserving applications too.

The simpler methods, Bootstrapping and the Distribution-free approach, need only a spreadsheet and a very basic level of programming to implement and so are not beyond anyone who can use a spreadsheet. The more complicated methods do, however, tend to offer more by way of diagnostics and further information about the model being fitted.

5.2 Comparison Of Results

The results for the three sets of data looked at by the 1993 Variance in Claim Reserving Working Party are summarized in the tables that follow:

Comparison of Standard Errors
Class 1

Accident Year	Reserving Method			
	Log-Lin	Bootstrap	Operational Time	Distribution -free
1922	11%	-	29%	-
1923	9%	-	22%	1%
1924	8%	60%	22%	2%
1925	7%	27%	21%	5%
1926	6%	17%	20%	7%
1927	6%	17%	22%	8%
1928	5%	16%	19%	9%
1929	5%	12%	16%	8%
1930	6%	7%	12%	7%
1931	6%	6%	10%	6%
Total	3%	6%	7%	3%

Reserve
to year 13

Reserve
to year 9

Reserve
to Ultimate

Reserve
to year 10

Comparison of Standard Errors
Class 3

Accident Year	Reserving Method			
	Log-Lin	Bootstrap	Operational Time	Distribution -free
1922	11%		101%	
1923	10%		84%	47%
1924	10%		61%	27%
1925	9%		43%	12%
1926	9%		32%	8%
1927	8%	61%	26%	10%
1928	8%	31%	19%	6%
1929	8%	15%	13%	6%
1930	9%	8%	9%	6%
1931	8%	4%	7%	4%
Total	5%	10%	5%	3%

Reserve
to year 13

Reserve
to year 9

Reserve
to Ultimate

Reserve
to year 10

Comparison of Standard Errors
Class 5

Accident Year	Reserving Method			
	Log-In	Bootstrap	Operational Time	Distribution -free
1922	22%		68%	
1923	19%		49%	21%
1924	17%		38%	16%
1925	15%	31%	27%	13%
1926	13%	17%	23%	14%
1927	12%	16%	19%	14%
1928	11%	18%	21%	14%
1929	11%	16%	18%	14%
1930	11%	15%	16%	15%
1931	12%	18%	15%	19%
Total	6%	9%	9%	8%

Reserve
to year 13

Reserve
to year 9

Reserve
to Ultimate

Reserve
to year 10

The results are not immediately comparable as some of the methods project the reserve to different periods of time: nine, ten, thirteen years and to Ultimate. To all intents and purposes however, there is going to be very little difference in the reserves for all except the oldest years (which were disguised as being from the 1920's to help preserve their anonymity!). We are, in this instance, most concerned with the variability of the reserves as a whole, which will be very little affected by the slight differences in the period to which the reserves are projected.

To remind the reader, the standard error (SE) is the standard deviation of an estimate of a variable, allowing for the uncertainty inherent in making that estimate. It thus gives an indication of the extent to which one may expect reserves to vary from the expected value according to the model we are looking at. If the variability about the mean is broadly symmetric, we can say roughly that in two cases out of three, the reserve should fall within one standard error of the mean, and in nineteen cases out of twenty it should fall within two standard errors of the mean.

The results indicated in the previous tables for older accident years do not generally mean very much. The reserves for the older years will tend to be very small relative to the reserves as a whole, so even if they are very variable, this will not affect the variability of the total greatly.

Class 1 could be described as Employer's Liability type business; it probably has an element of latent type claims which distorts the picture. Class 3 is relatively shorter tail but contains an element of Bodily Injury type claims. Class 5 is longer-tail and has quite a high Bodily Injury content.

Comparing the four methods, looking principally at the total reserve figures, we can say that they produce broadly similar results. On the whole, the Bootstrapping numbers tend to be higher than the measures indicated by other methods. Of the four methods, Bootstrapping is the least "scientific" in terms of the credibility I would give to the numbers, but if one wanted a ball-park number for overall reserve variability, then it seems to provide a comparable number. The overall figure for the Distribution-free method is reassuringly similar to that produced by the other two "complicated" methods. The Distribution-free method tends to produce SE's that are less stable from one year to the next; this is because the Distribution-free approach uses nine variance parameters as opposed to the other methods' one. Even for individual accident years however, the Distribution-free SE's look to be of the same order of magnitude.

In summary then, all four methods seem to show a good consensus of results. Either of the very simple variability measures, Bootstrapping and the Distribution-free approach producing comparable overall numbers, in terms of the measures of variability, to more complicated methods.

6. Bibliography

"Bootstrap methods, another look at the jackknife", 1977 Reitz lecture, *Annals of statistics*, Vol. 7, No.1, pages 1-26, Efron B. (1977)

"Bootstrapping a regression equation: some empirical results", *Journal of the American statistical association*, Vol.79, No.385, pages 155-161, Freeman D.A. & Peters S.C. (1984)

"Computer intensive statistical methods: validation, model selection and bootstrap", Chapman & Hall, Urban-Hjorth J.S. (1994)

"Monographs on statistics and applied probability 57: An introduction to the bootstrap", Chapman & Hall, Efron B. & Tibshirani R.J. (1993)

"A stochastic method for claims reserving in general insurance", *JIA* 117, Part III, pages 677-733, Wright T.S. (1990)

"Stochastic claims reserving when past claim numbers are known", *PCAS*, Wright T.S. (1992)

"Measuring the variability of chain ladder reserve estimates", *CAS Prize paper competition*, Mack T. (1993)

"Which stochastic model is underlying the chain ladder method?", *ASTIN Colloquium*, Mack T. (1993)

"Distribution-free calculation of the standard error of chain ladder reserves estimates", *ASTIN Bulletin*, Vol. 23, No.2, Mack T. (1993)

```

!title v\                                     !limits info sent to GLIM.LOG
!units 364\                                   !number of values in each data vector
!date c t n y d\                             !label for average cost data
!print ::'Enter filename for data:\           !prompt user for file
!input 10 100\                               !assign a channel number for data, with max width
!warn\
!macro sres                                   !start macro for standardised residuals
!cal r=(%yv-%fv)/%sqrt(%va)\                 !calculate standardised residuals
!print ::'Histogram of Standardised Residuals'\ !create plots
!hist (s=1) r '*'
!print ::'Standardised Residuals against Linear Predictor'\
!plot (s=1) r %lp '+'
!print ::'Standardised Residuals against Fitted Values'\
!plot (s=1) r %fv '+'
!print ::'Standardised Residuals against Accident Year'\
!plot (s=1) r y '+'
!print ::'Standardised Residuals against Development Period (years)'\
!plot (s=1) r d '+'
!endmacro
!macro dres                                   !start macro for deviance residuals
!cal sn=%if((%yv-%fv)<0,-1,1)\               !calculate sign of residual
!cal r=sn*%sqrt(%di)\                       !calculate deviance residual
!print ::'Histogram of Deviance Residuals'\ !create plots
!hist (s=1) r '*'
!print ::'Deviance Residuals against Linear Predictor'\
!plot (s=1) r %lp '+'
!print ::'Deviance Residuals against Fitted Values'\
!plot (s=1) r %fv '+'
!print ::'Deviance Residuals against Accident Year'\
!plot (s=1) r y '+'
!print ::'Deviance Residuals against Development Period (years)'\
!plot (s=1) r d '+'
!endmacro
!cal t1=1 : t2=t**2 : t3=t**3 \             !calculate extra variables
!cal t4=t**4 : t5=t**5 : t6=t**6 \
!cal t7=t**7 : t8=t**8 \
!cal i=y-1932+d\                           !calculate time index for inflation
!cal yr=%if(y>=1931,10,y-1921)\            !calculate year groupings
!fac yr 10\                                 !declare yr as a factor with 10 levels
!cal wt=%if(e<0,0,1)\                      !weight out negative amounts
!cal c=%if(e<0,0.001,c)\                   !assign negatives the value 0.001
!cal wt=%if(n<=0,0,wt)\                   !weight out negative numbers
!cal a=%if(n<=0,0.001,a)\                 !assign negatives the value 0.001
!macro m1 !cal %fv=%exp(%lp) !endmacro      !set up macros for user defined models
!macro m2 !cal %dr=1/%fv !endmacro
!macro m3 !cal %va=(%fv**2)/n !endmacro
!macro m4 !cal %di=wt**2*n*(-%log(%yt/%fv)+(%yv-%fv)/%fv) !endmacro
!var clown m1 m2 m3 m4 \                  !define model as above
!weight wt\                               !declare weights
!eye 30 %prt 1 c-6 1 c-8\                !specify settings for iter_n control
!acc 10\                                   !set max accuracy for output
!cal %lp=%log(%if((%yv<0,0.5,%yv+0.5)))\    !initialise linear predictor
!fit 1 + YR + T1 + T2 + T4 + T5 + T6 + T7 + T8 !disp 1 e!print ::\
!fit 1 + YR + T1 + T2 + T3 + T5 + T6 + T7 + T8 !disp 1 e!print ::\
!fit 1 + YR + T1 + T2 + T5 + T6 + T7 + T8 !disp 1 e!print ::\
.....
.....
!stop\                                     !end batch mode and exit GLIM

```

Operational Time Reserving: Summary of models
Group 3

Appendix II

Dev. & Sig.	Ordered by Dev.	Original Index	Dev.	Model	Number of polynomial terms	Percentage of terms Sig.	Total Reserve
1	15	179	13.320	1 + YR + T1 + T3 + T4 + T7 + T8	5	100%	921,752
2	19	211	13.323	1 + YR + T1 + T2 + T4 + T7 + T8	5	100%	919,305
3	22	173	13.324	1 + YR + T1 + T3 + T5 + T6 + T8	5	100%	920,813
4	26	157	13.327	1 + YR + T1 + T4 + T5 + T6 + T8	5	100%	923,107
5	29	117	13.329	1 + YR + T2 + T3 + T4 + T6 + T8	5	100%	920,529
6	31	205	13.330	1 + YR + T1 + T2 + T5 + T6 + T8	5	100%	918,258
7	38	115	13.334	1 + YR + T2 + T3 + T4 + T7 + T8	5	100%	924,464
8	41	109	13.335	1 + YR + T2 + T3 + T5 + T6 + T8	5	100%	923,777
9	42	158	13.335	1 + YR + T1 + T4 + T5 + T6 + T7	5	100%	918,863
10	44	110	13.337	1 + YR + T2 + T3 + T5 + T6 + T7	5	100%	919,272
11	46	94	13.339	1 + YR + T2 + T4 + T5 + T6 + T7	5	100%	922,538
12	47	181	13.339	1 + YR + T1 + T3 + T4 + T6 + T8	5	100%	917,097
13	54	121	13.354	1 + YR + T2 + T3 + T4 + T5 + T8	5	100%	915,810
14	55	174	13.357	1 + YR + T1 + T3 + T5 + T6 + T7	5	100%	915,409
15	56	93	13.359	1 + YR + T2 + T4 + T5 + T6 + T8	5	100%	925,693
16	59	213	13.365	1 + YR + T1 + T2 + T4 + T6 + T8	5	100%	913,864
17	60	118	13.367	1 + YR + T2 + T3 + T4 + T6 + T7	5	100%	914,719
18	62	62	13.373	1 + YR + T3 + T4 + T5 + T6 + T7	5	100%	925,083
19	64	185	13.390	1 + YR + T1 + T3 + T4 + T5 + T8	5	100%	911,951
20	67	206	13.391	1 + YR + T1 + T2 + T5 + T6 + T7	5	100%	911,949
21	68	61	13.409	1 + YR + T3 + T4 + T5 + T6 + T8	5	100%	926,924
22	70	182	13.412	1 + YR + T1 + T3 + T4 + T6 + T7	5	100%	910,599
23	71	229	13.419	1 + YR + T1 + T2 + T3 + T6 + T8	5	100%	909,475
24	73	135	13.436	1 + YR + T1 + T6 + T7 + T8	4	100%	918,498
25	75	217	13.443	1 + YR + T1 + T2 + T4 + T5 + T8	5	100%	908,137
26	77	122	13.446	1 + YR + T2 + T3 + T4 + T5 + T7	5	100%	908,773
27	78	214	13.473	1 + YR + T1 + T2 + T4 + T6 + T7	5	100%	906,608
28	81	186	13.522	1 + YR + T1 + T3 + T4 + T5 + T7	5	100%	904,401
29	85	233	13.535	1 + YR + T1 + T2 + T3 + T5 + T8	5	100%	903,206
30	87	230	13.576	1 + YR + T1 + T2 + T3 + T6 + T7	5	100%	901,507
31	89	124	13.601	1 + YR + T2 + T3 + T4 + T5 + T6	5	100%	901,224
32	92	218	13.617	1 + YR + T1 + T2 + T4 + T5 + T7	5	100%	899,954
33	95	11	13.626	1 + YR + T5 + T7 + T8	3	100%	917,646
34	98	241	13.684	1 + YR + T1 + T2 + T3 + T4 + T8	5	100%	897,125
35	100	7	13.702	1 + YR + T6 + T7 + T8	3	100%	927,372
36	69	252	13.412	1 + YR + T1 + T2 + T3 + T4 + T5	6	83%	915,447
37	12	203	13.320	1 + YR + T1 + T2 + T5 + T7 + T8	5	80%	922,770
38	28	171	13.328	1 + YR + T1 + T3 + T5 + T7 + T8	5	80%	924,400
39	49	155	13.343	1 + YR + T1 + T4 + T5 + T7 + T8	5	80%	925,610
40	57	107	13.359	1 + YR + T2 + T3 + T5 + T7 + T8	5	80%	926,494
41	82	71	13.523	1 + YR + T2 + T6 + T7 + T8	4	75%	918,012
42	84	139	13.534	1 + YR + T1 + T5 + T7 + T8	4	75%	912,722
43	86	39	13.576	1 + YR + T3 + T6 + T7 + T8	4	75%	918,017
44	88	75	13.592	1 + YR + T2 + T5 + T7 + T8	4	75%	913,784
45	91	43	13.616	1 + YR + T3 + T5 + T7 + T8	4	75%	915,112
46	96	51	13.676	1 + YR + T3 + T4 + T7 + T8	4	75%	911,777
47	97	29	13.678	1 + YR + T4 + T5 + T6 + T8	4	75%	912,822
48	99	45	13.700	1 + YR + T3 + T5 + T6 + T8	4	75%	910,516
49	37	199	13.334	1 + YR + T1 + T2 + T6 + T7 + T8	5	60%	925,737
50	52	167	13.348	1 + YR + T1 + T3 + T6 + T7 + T8	5	60%	926,421

Appendix III.1

Distribution-free Chain-ladder Breakdown of the example in Thomas Mack's ASTIN Paper All figures in £000s

Cumulative Claims ($C(i,k)$) Triangle

Accident Year (i)	1	2	3	4	5	6	7	8	9	10	R(i)
1	357.8	1,124.8	1,735.3	2,218.3	2,745.6	3,320.0	3,466.3	3,606.3	3,833.5	3,901.5	0.0
2	352.1	1,236.1	2,170.0	3,353.3	3,799.1	4,120.1	4,647.9	4,914.0	5,339.1	5,433.7	94.6
3	290.5	1,292.3	2,218.5	3,235.2	3,986.0	4,132.9	4,628.9	4,909.3	5,285.1	5,378.8	469.5
4	310.6	1,418.9	2,195.0	3,757.4	4,029.9	4,382.0	4,588.3	4,835.5	5,205.6	5,297.9	709.6
5	443.2	1,136.4	2,128.3	2,897.8	3,402.7	3,873.3	4,207.5	4,434.1	4,773.6	4,838.2	984.9
6	396.1	1,333.2	2,180.7	2,985.8	3,691.7	4,075.0	4,426.5	4,665.0	5,022.2	5,111.2	1,419.5
7	440.8	1,288.5	2,419.9	3,483.1	4,088.7	4,513.2	4,902.5	5,166.6	5,562.2	5,660.8	2,177.6
8	359.5	1,421.1	2,864.5	4,174.8	4,900.5	5,409.3	5,876.0	6,192.6	6,666.6	6,784.8	3,920.3
9	376.7	1,363.3	2,382.1	3,471.7	4,075.3	4,498.4	4,886.5	5,149.8	5,544.0	5,642.3	4,279.0
10	344.0	1,200.8	2,098.2	3,058.0	3,589.6	3,962.3	4,304.1	4,536.0	4,883.3	4,969.8	4,625.8

k	1	2	3	4	5	6	7	8	9	Total
f(k) Estimate	3.491	1.747	1.457	1.174	1.104	1.086	1.054	1.077	1.018	
Cumulative factor	14.447	4.139	2.369	1.625	1.384	1.254	1.155	1.096	1.018	

10 Accident years so $l = 10$

Intermediate Calculations (1) $D(i,k) = C(i,k) \times (C(i,k+1)/C(i,k) - f(k))^2$

Accident Year (i)	1	2	3	4	5	6	7	8	9
1	43.2	47.1	55.7	9.0	30.5	5.9	0.6	0.7	
2	0.1	0.1	16.8	5.6	1.4	7.2	0.1	0.5	
3	266.5	1.2	0.0	11.0	17.9	4.7	0.2		
4	360.5	56.9	142.0	38.6	1.1	6.7			
5	380.3	17.9	19.6	0.0	4.0				
6	6.2	16.6	17.0	11.7					
7	142.1	22.0	0.8						
8	77.0	102.3							
9	6.2								
10									

Intermediate Calculations (2) $s(k)^2 = (D(1,k) + ... D(l,k)) / (l-k-1)$

	1	2	3	4	5	6	7	8	9
$s(k)^2$	160.3	37.7	42.0	15.2	13.7	8.2	0.4	1.1	0.4

Intermediate Calculations (3) $E(k) = C(1,k) + + C(l-k,k)$

	1	2	3	4	5	6	7	8	9
$E(k)$	3,327.4	10,251.2	15,047.8	18,447.8	17,963.3	15,955.0	12,743.1	8,520.3	3,833.5

Intermediate Calculations (4)

$$F(i,k) = (s(k)^2) / ((i(k)^2) \times (1/C(i,k) + 1/(E(k)))$$

Accident Year (i)	1	2	3	4	5	6	7	8	9	SE(R(i))	SE(R(i))/R(i)
1	-	-	-	-	-	-	-	-	-	-	-
2	-	-	-	-	-	-	-	-	0.00019	75.5	80%
3	-	-	-	-	-	-	-	0.00032	0.00019	121.7	26%
4	-	-	-	-	-	-	0.00012	0.00032	0.00020	133.5	19%
5	-	-	-	-	-	0.00223	0.00013	0.00034	0.00020	261.4	27%
6	-	-	-	-	0.00368	0.00214	0.00012	0.00033	0.00020	411.0	29%
7	-	-	-	0.00376	0.00338	0.00197	0.00011	0.00031	0.00019	538.3	26%
8	-	0.00821	0.00324	0.00293	0.00172	0.00010	0.00028	0.00018	0.00018	875.3	22%
9	0.01027	0.00961	0.00377	0.00339	0.00198	0.00011	0.00011	0.00031	0.00019	971.3	23%
10	0.04219	0.01150	0.01073	0.00420	0.00377	0.00219	0.00012	0.00033	0.00020	1,363.2	29%
Total										2,447.1	13%

Intermediate Calculations (5)

$$G(i) = C(i+1,i) + \dots + C(1,i)$$

Accident Year (i)	G(i)
1	-
2	43,703.8
3	38,324.9
4	33,027.0
5	28,168.8
6	23,057.7
7	17,396.9
8	10,612.1
9	4,969.8
10	-

$$MSE(R(i)) = (C(i,i)^2) \times (F(i,1) + \dots + F(i,b))$$

$$SE(R(i)) = MSE(R(i))^{(1/2)}$$

Intermediate Calculations (6)		$H(k) = 2 \times (s(k)^2) / (f(k)^2) / E(k)$								
		Development Year (k)								
$H(k)$		1	2	3	4	5	6	7	8	9
		0.00791	0.00241	0.00263	0.00119	0.00125	0.00087	0.00006	0.00023	0.00022
Intermediate Calculations (7), (8)		$J(i) = H(1 + i) + \dots + H(i-1)$								
		$K(i) = SE(R(i))^2 + C(i,1) \times G(i) \times J(i)$								
Accident Year (i)	$J(i)$	$K(i)$								
1	-	-								
2	0.00022	59,128								
3	0.00046	109,089								
4	0.00052	108,901								
5	0.00139	258,562								
6	0.00264	480,624								
7	0.00384	689,819								
8	0.00647	1,231,705								
9	0.00888	1,192,254								
10	0.01678	1,858,191								
		$MSE(R) = K(2) + \dots + K(1)$								
		5,988,273								
		$SE(R) = MSE(R)^{(1/2)}$								
		2,447								