Swiss Re

Late claims reserves in reinsurance

Contents

Foreword

Late reported claims, especially on liability insurance, continue to be a problem for insurers and reinsurers alike. The industry has suffered from having to pay huge sums in respect of claims due to asbestos, pollution and mass torts, as well as for professional indemnity, medical malpractice and workers compensation. The value of bodily injury claims arising from, in particular, road accidents has increased markedly. Disparate as they may seem, these claims have one thing in common: they can be notified years – possibly even many years – after the insurance contract was drawn up, and therefore, they may not be *settled* for even longer. The socio-economic conditions at the time of settlement may then differ widely from what was expected when the insurance was underwritten, and this can have a large impact on the cost of settlement.

Actuaries all over the world have been working for years to develop methods that can reliably predict the eventual cost of settlement of such claims. Former Swiss Re actuaries, including Professor Erwin Straub, and Professor Hans Bühlmann[1] of the Swiss Federal Institute of Technology (ETH), have made considerable contributions to this development. Many actuaries within Swiss Re are working on reserving issues today, and they derive benefit from the huge amount of claims information and global claims network expertise available to the Swiss Re Group. They have found that, while a firm understanding of mathematics is required, elaborate mathematical methods are of limited practical use in determining adequate reserves. Indeed, Swiss Re's specialists set a convincing example that there is no substitute for experience coupled with business knowledge and a thorough analysis of the data.

Anthony Boulter
Dawson Grubbs
Swiss Re, Zurich

1 Professor Bühlmann recently retired from Swiss Re's Board of Directors.

Introduction

This publication is intended to assist general management, claims managers, underwriters and accountants in understanding why reinsurers need to hold additional claims reserves over and above the claims estimates that the ceding insurer supplies. It may also be useful to insurers' clients who wish to know more about this aspect of insurance generally.

Given such a wide audience, it is necessary to assume minimal knowledge of the field. As such, the more experienced reader will probably wish to skim through the initial sections. Later, the problem of assessing late claims reserves in practice is approached in detail.

Since the original edition of this publication, there has been much research in this field. Because no single estimate of reserves can ever be expected to be wholly accurate, one area of this research has been to try to quantify the range of outcomes that may be expected. While Chapter 1 introduces some of the techniques that are commonly used, the Appendix gives more detailed information about one particularly popular technique. Since no paper on reserving would be complete without some reference to discounting, a brief section on discounting concludes the first chapter.

Chapter 2 includes examples of some common methods of reserving in practice. For readers who want to pursue the methodology in more detail, a short bibliography concludes the written text. Finally, a glossary of relevant terms may serve as a handy reference.

1 The claims reserving process

What are claims reserves?

The premiums paid for insurance policies have to fund the claims that will be paid as well as the expenses of the underwriting company and the cost of the capital employed. Although there are many variations on the way the premiums are paid, the company will generally receive the premiums some time, and perhaps even a long time, before the claims are actually paid. In the meantime, the company must properly account for these claims liabilities by setting up provisions within its balance sheet to reflect as accurately as possible the eventual claims cost. These provisions are generally referred to as reserves.

These reserves[2] are divided into several categories, the principal ones being:

Table 1.1

Unearned premium reserve (UPR)	Each customer pays a premium to cover the risk of claims occurring over a certain period of time, often one year. It is usually done on a time-weighted basis so that after, say, three months, the company has only taken credit for one-quarter of the premiums paid, and the balance is held in the *unearned premium reserve.*
Notified claims	When a claim is notified to the company, a claims handler who has experience of dealing with similar claims in the past will assess the facts of the claim and place an estimate for its eventual cost in the company's books[3]. The claims handler will monitor this value as the claim moves towards settlement to ensure that the value held continues to be appropriate.
Incurred but not reported (IBNR)	There will always be some delay between when the event that gives rise to a claim occurs and when it is reported to the insurance company. Prudent accounting practice requires that the insurance company keeps a reserve that, at any time, will cover all the claims that have occurred but that have not so far been reported.

The UPR, although highly necessary, is not referred to as a claims reserve because it relates to risks in the future. During an accounting period, the net movement to or from the UPR, together with the new premium actually booked by the company (the written premium), defines the earned premium. This is important because it is a measure of the premium that corresponds to the claims that occur during the same accounting period.

A company may hold other claims reserves, including reserves against future catastrophic events (such as windstorms and earthquakes), and reserves to cover the expenses of running-off to settlement all the claims arising from the business that has been written. This publication refers only to the notified claims and the IBNR reserves.

2 There are variations on this structure eg Lloyds reserving, that are not considered here.

3 Sometimes companies will also use statistical methods to assess the value of notified claims. This is often encountered where there are a large number of low-value claims, such as damage to motor vehicles, in order to save the expense of having to case estimate claims that will be settled quickly.

Claims reserves for a reinsurer vs an insurer

A reinsurance company holds the same types of reserve as an insurance company. Because the reinsurer is one stage further removed from the claim, however, the balance may change. Consider the following example.

A reinsurance company agrees to cover 75% of an insurer's losses, which exceed USD 5 million, up to a limit of USD 25 million on the original loss. If the insurance company has already paid USD 9 million on a claim and considers that a further USD 8 million will eventually be payable, then the reinsurer must:

- pay the insurance company 75% of the claims already paid above the USD 5 million retention limit; this amounts to USD 3 million, and
- set up a case estimate of 75% of USD 8 million (ie USD 6 million) to cover the expected future claims payments.

This process will work well as long as the cedant (the original insurer) has set the case estimate correctly and informed the reinsurer promptly. Otherwise the reinsurer may find that its case estimates are inadequate and it will need to hold an additional reserve to compensate. In some cases, where the reinsurer has independent knowledge of a particular event, it will be able to set up an additional case reserve (ACR) relating to one or more specific claims. More generally the reinsurer will need to set up a reserve sometimes referred to as IBNER (incurred but not enough reserved). This reserve, together with the "pure" IBNR that covers claims not yet notified to the insurer, form the late claims reserve, and the total is likely to be bigger, relative to the case estimates, than would be the case for the insurer.

Also worth mentioning is that a reinsurer will be less certain than an insurer as to the amount of premiums it expects to achieve. For example, the amount of reinstatement premiums will depend on the experience of the insurer. This uncertainty is outside the scope of this publication, except that it should be borne in mind when using methods that require amounts of premium (such as Cape Cod or Bornhuetter-Ferguson – see Chapter 2).

What is "pure" IBNR?

Although many types of claim, particularly those involving loss of or damage to property, are apparent to the customer within days (at most), other types of claim may take months or years to emerge. Examples are:

Table 1.2

Type of claim	Reason for late notification
Claims involving bodily injury	1 The injury is not immediately apparent (delayed effects, psychological trauma); 2 The claimant may not realise the connection between his condition and the accident; 3 The claimant may not at first have the legal resources to pursue the claim.
Industrial disease claims	1 The effects of exposure to certain processes or substances may not be apparent for many years; 2 The claimant may not realise the connection between his condition and exposure to the processes or substances; 3 The claimant may not at first have the legal resources to pursue the claim.
Professional indemnity	The results of negligence may take years to be discovered.
Product liability and latent defects	The fault, its effects and the link between them may take years to be proved.

In some circumstances, even property claims can be notified some time after the event, when damage is done gradually, for example. Also, the claimant may initially contact an intermediary – such as his broker – rather than his insurance company, resulting in further delays.

To date the most serious IBNR issues have been caused by claims arising from exposure to asbestos fibres, pollution and other health hazards. Exposure to certain types of asbestos can give rise to medical conditions ranging from the debilitating to the fatal. The particular problem is the extended time delay that normally occurs between exposure and the time when the diseases develop, possibly even forty years or more. Pollution affecting, say, water supplies may take many years to seep through the soil and be discovered, and then traced to its origin. Product liability arising from the use of certain chemicals, drugs or other medical products has also caused serious strain on IBNR reserves many years after the insurance was issued.

IBNR reserves can only be assessed using statistical methods and the way in which these are applied is discussed below under "Assessing late claims using statistical methods"(see p 9).

What are the problems with case estimates?

Expert claims handlers working for the cedants will assess claims as they are notified and keep those claims under review as information about each claim emerges. To all appearances, reinsurers need not worry about case estimates.

Yet, in actual fact, there are several reasons why they should:

- Insurers are not necessarily under an obligation to advise the reinsurer of all claims;
- Not all the information affecting the value of a claim is necessarily available at the time the claim is notified;
- Claims may not be settled straightaway and any delay gives an opportunity for the value of the claim to be affected by inflation.

The following examines each of these points in more detail.

Insurers are not necessarily under an obligation to advise all claims to the reinsurer.

This depends on the terms of the contract. It is particularly likely when the reinsurer is accepting excess-of-loss business and the insurer does not realise that the claim(s) will reach the point at which he must notify the reinsurer. Any situation where additional parties stand between the insured customer and the reinsurer will add to the delay in notification. An example would be when the reinsurer is accepting retrocession business or where the cedant is co-insuring the business and not dealing directly with the client or broker. Even for simple proportionate treaties, there is likely to be some administrative delay.

Not all the information affecting the value of a claim is necessarily available at the time the claim is notified.

The insurer's claims handlers can only act on the information supplied to them. The true facts of the case – affecting not only the amount of the claim but also the extent of the liability of the insurer – may require a lengthy process of gathering evidence, involving loss adjusters, salvage companies, medical, legal and other experts. In some cases, there will be the possibility of recovering part or all of the costs from negligent third parties, if they are insured or otherwise have the means to pay.

Some companies will instruct their claims handlers to put an allowance in their estimates, based on their experience, to try to project the final outcome. Others will expect their handlers to assess claims purely on the information in front of them, relying on statistical assessments (see p 9) to make up the difference.

It is important for the reinsurer to know what practice each of his cedants adopts!

Claims may not be settled straightaway, and any delay gives an opportunity for the value of the claim to be affected by inflation.

An insurer cannot pay claims instantly. The amount and validity of the claim need to be checked and this takes some time, even for simple claims. In some cases, particularly where legal disputes are involved, the delay can be considerable, possibly amounting to several years.

In cases involving bodily injury, it may be necessary to wait some time for the claimant's medical condition to stabilise. Where a child is involved, and for legal reasons, it may not be possible to settle the case until the child has reached majority.

What's more, from time to time, insurers may experience changes in their administrative efficiency. At times when large numbers of claims are being notified, such as after a major storm, the insurer's claims organisation may be temporarily unable to cope. Reorganisation, IT changes or mergers and acquisitions can also cause significant disruption in the claims organisation, leading to changes in the pattern of claims settlement. The reinsurer needs to be aware of any significant events of this nature affecting his cedants.

Claims inflation can take various forms. At a minimum, it can be expected to follow general price/wage inflation, and this will apply particularly to property claims. Claims involving bodily injury, as a rule, have been seen to inflate much more rapidly because of changing legislation and social attitudes. Each type of claim will have its own inflation rate.

Assessing late claims using statistical methods

The problems, both of setting the IBNR reserve and addressing any deficiencies in the case estimates, can be solved together by the use of statistical methods. There are, in fact, many methods that have been suggested over the years, and both insurance and reinsurance companies will use a number of these to assess the additional reserves that they need to hold. The company will usually ask an actuary, an expert in the application of statistical techniques to the financial problems of insurance companies, to assess the reserves using these techniques.

All the methods depend on finding some pattern in the way that claims have been settled in the past that can be applied to the future. If, in fact, the pattern is changing in a way that is not detected, then the method concerned cannot be relied upon to give even an approximation to the correct answer.

The actuary will select the method only after the most thorough analysis of the data at hand, which will involve:

- Understanding the nature of all the major agreements that are in force, involving extensive discussions with one's own underwriters and possibly with cedants themselves;
- Checking the accuracy of all the data used;
- Classifying the data in the most appropriate way for the analysis. At the very least this would involve the separation of long-tail lines (such as liability) from short-tail lines (such as property), but in practice, the analysis will be done at a detailed level that involves separating business by treaty type (eg proportional and non-proportional), line of business and territory. Special features, such as claims-made policies, as well as large, catastrophic or other exceptional claims, will need careful treatment;
- Selecting what type of data to use, whether paid claims, case estimates or ACRs. Note that some methods have particular data requirements;
- Evaluating the result in the light of his or her knowledge of the business and comparing with available external benchmarks.

Even if the method chosen is completely appropriate – reflecting the patterns in the data exactly and the extent, if any, to which they are changing – the eventual outcome is still subject to statistical variation. Although modern methods permit this variation to be quantified, the most that anyone can say about the result is that it will fall within a range with a high probability. Moreover, nobody can say for sure that the method chosen is in fact the correct one, and this gives further scope for the eventual outcome to be different from the one projected.

Data triangles

The methods most often encountered in practice include such names as Chain Ladder, Cape Cod and Bornhuetter-Ferguson (see Chapter 2). All these classic methods involve grouping the claims data in a triangle. The data is classified in a row of the triangle according to when it originated and into a column of the triangle according to when it emerged[4]. The most common definitions of "origin" for these purposes are: accident year (claims are grouped according to year in which they occurred), and underwriting year (claims are grouped according to the year when the period of insurance in which they occurred began). Other definitions, such as according to the period in which claims are notified, may also be encountered.

Since calculations on different origin bases have different implications, it is important to know the origin basis. Projecting on an accident year basis gives a forecast of the cost of all claims arising from events occurring in the period covered by the data, whether they have been notified or not. This corresponds exactly to what is covered by the case estimates, additional case reserves and IBNR. If, on the other hand, the data is on an underwriting year basis, then the results will be an estimate for the entire period of insurance for all policies issued,

4 Sometimes one comes across the opposite convention, with the origin periods running across the columns and the development periods down the rows. As long as one or the other convention is applied consistently, the results will not be affected in any way.

including the period yet to be exposed that is covered by the UPR. By contrast, if claims are classified by the year in which they are notified, then the forecast only covers claims already notified and will therefore not include any allowance for "pure" IBNR. This basis is therefore less useful[5] in practice.

To some extent, the choice of using an accident year or underwriting year classification will depend on the accounting convention to be applied. For instance, accident year is the natural choice under US GAAP, whereas Lloyds accounting requires underwriting year. Even if the accident year basis is required for normal accounting purposes, an underwriting year basis may be useful for other purposes, such as assessing the adequacy of premium rates or determining whether a *premium deficiency reserve*[6] is required.

An example of a data triangle follows. The data shown is fictitious, but let us assume it relates to claims incurred (ie claims paid plus any outstanding case estimates). Note that the data is shown as cumulative across the columns.

Table 1.3

Accident	Development years					
year	1	2	3	4	5	6
1995	90	210	310	420	500	500
1996	130	280	360	460	600	
1997	140	290	440	600		
1998	160	240	420			
1999	120	260				
2000	110					

As seen in the triangle, the total claims paid plus outstanding estimates was 90 by the end of the year for claims occurring in 1995. By the end of 1996, it was 210, rising every year until it had reached 500 by the end of 1999. Then, it stabilised, so that the same value applied at the end of year 2000. For claims occurring in 1996, the value stood at 130 at the end of 1996, rose to 280 by the end of 1997 and eventually to 600 by the end of 2000. At the bottom of the triangle, only one value, of 110, is visible for claims occurring in 2000.

For triangulation methods to work, the patterns have to depend on the features of the triangle, ie on the:
- origin periods,
- development periods,
- calendar periods (diagonals),

either alone or in combination.

5 For claims arising from disease or pollution, where the date of loss is often unclear, the date of notification may be the only date known for certain. Under these circumstances, the accident and underwriting year classifications will not be meaningful, but a notification year analysis may be used to assess the adequacy of case estimates. Other methods, beyond the scope of this publication, will then be required to assess the IBNR.

6 Normally, the UPR should be sufficient to cover the costs of the remaining exposure on contracts already entered into. This may not be the case if business has been under-priced, in which case a special additional reserve may be required to bolster the UPR.

The actuary will look at the ratios that the values in each column bear to their immediate predecessors (one column to the left). These values, known as link ratios, themselves form a triangle, as shown in Table 1.4, to three places of decimals. Note that this table has one less row and one less column than the original table, because nothing can be deduced in respect of the first development year, and therefore the entry for year 2000 drops out altogether.

Table 1.4

Accident	Development years				
year	1 ▶ 2	2 ▶ 3	3 ▶ 4	4 ▶ 5	5 ▶ 6
1995	2.333	1.476	1.355	1.190	1.000
1996	2.154	1.286	1.278	1.304	
1997	2.071	1.517	1.364		
1998	1.500	1.750			
1999	2.167				

The values in each column of this second triangle are expected to be different. High values are expected in the early periods of development, as considerable volumes of claims are still coming in, but those values should tend to stabilise (towards 1.0) for the columns towards the right. Values failing to stabilise indicate that the triangle is not large enough, and that there is a "tail" of development even on the oldest years that cannot be found directly from the data in the triangle. Under these circumstances, the actuary may be able to infer the behaviour of this tail by fitting a curve to the ratios at hand. This needs to be corroborated, if possible, by reference to the behaviour of a similar portfolio of business for which a longer history is available.

Ideally, the link ratios in each row will not show a pattern down any particular column of the triangle, but simply vary randomly around a single value. Under these circumstances, the actuary can use a standard technique, such as the Chain Ladder method to exploit this regularity.

In practice, complete regularity is rarely observed. Looking down the column, the actuary may see a gradual change in the level (a trend) or a sudden shift in the values observed. These changes may affect just one column, or they may affect several columns, in which case the changes may well be correlated. In Table 1.4, for example, the first column of figures shows a generally declining pattern (except for 1999), with the opposite pattern in the second column. While the data here is too scanty to support such a conclusion directly, this may indicate that the development pattern is slowing down. Whatever the circumstances, the reasons for this will need to be investigated before the most appropriate method is selected.

High levels of past inflation will have an impact on the diagonals of a triangle and the effects of this are often difficult to identify. If a suitable inflation index can be found, it is possible to remove the effects of inflation by standardising the original triangle in constant money terms. This requires some manipulation of the underlying payments and estimates. The past inflation can then be factored back in when the analysis is complete. An explicit allowance for future inflation will then need to be included in the projections.

Alternatively, there are methods or tools that can estimate the diagonal effects directly. In any case, the actuary will have to estimate the extent to which future inflation will affect his results.

The actuary will look for stability in his results. One popular test is that, in the absence of any known reason to the contrary, the chosen method should produce approximately the same result even if the most recent data (ie the most recent diagonals) are excluded.

Looking at variability

Although a single estimate of the reserves is sufficient for the purposes of putting a figure in the company's accounts, no estimate can ever be wholly correct, and the result may fall within a range of possible values. For some purposes, such as assessing solvency margins or measuring the capital employed by the business, it is necessary to try to quantify that range.

Two basic methods have been used to do that. One is to embed the calculation method in a rigorous theoretical framework that enables the variability of reserves to be calculated directly. The Chain Ladder model in particular has been subject to much scrutiny (see Bibliography, particularly the work of T. Mack and A.E. Renshaw).

The other approach has been to use a simulation method known as the "Bootstrap" which is outlined in the Appendix. This requires a computer program, but will deal with any underlying method.

Whichever basic method is chosen, the insurer or reinsurer with business in many lines or territories must also assess the degree to which these lines or territories are independent. The greater the independence, the less variable the overall result will be. In practice, there may be significant correlation between the results for these components. Changes in the levels of award for bodily injury, for example, will affect many casualty lines, including motor liability, workers compensation and general liability. Extreme weather events may affect both motor and property lines and cut across neighbouring territories, the European storms just after Christmas 1999 being a case in point. As ever, careful judgement, after consideration of all the facts, will be needed in applying these techniques in practice.

Discounting of reserves

Reserves for life insurance have always been discounted for interest earned on the assets. The reserve would be the present value of the benefits expected to be paid from the contract, minus the present value of future premiums after allowing for expenses and future profit emergence[7].

For non-life insurance, where contracts are for short periods, there is usually little, if any, future premium income. An equivalent approach would then be to set claims reserves to the present value of future claims payments plus the expenses of settlement. Yet this is, in fact, uncommon, and most reserves are held undiscounted. The reason for this is mainly historical. In principle, there is no reason why any insurer or reinsurer should not hold discounted reserves.

In practice, there are three issues that need to be assessed carefully:

- Discounting the liabilities means that the treatment of income on the asset side of the balance sheet needs to be altered. Even if everything else remains the same and the claims run-off exactly as planned, an amount will need to be transferred each year from the asset income in order to compensate the claims reserves, the amount being determined by the rate of discount. This process is sometimes referred to as "unwinding" the discount.
- The rate of discount needs to be determined. If the rate is set too high, then there is a risk that future years' asset income will be insufficient to cover the unwinding of the discount, and reserves will need to be strengthened from other sources. This is another potential source of under-reserving. Even if the rate is set to a realistic level, there is still a risk that such under-reserving could occur because interest rates vary unpredictably. To avoid future shocks, the rate should be set to a prudent level that we would reasonably expect to be able to beat for as long as the reserves take to run-off.
- The amount of discount also depends on the assumed future payment pattern. This will need to be assessed from the historical run-off patterns of the business, and the techniques used to do this will be similar to those used to reserve the business (eg Chain Ladder). As with the reserving process itself, any changes observed in the rate of settlement will need to be carefully assessed: assuming too slow a pattern of future run-off will cause the reserves to be understated.

The risks of setting the discount rate too high may be offset by careful matching of assets to liabilities by term, although the degree of variability of non-life insurance reserves can make this more difficult to achieve than for life business. The derivatives market may also be of assistance here.

Practical solutions will depend on local considerations, including the reaction of regulators and the tax authorities.

7 This is, of course, oversimplifying the life reserving process greatly. For the present purpose, however, the key point is that the reserve is discounted.

This publication cannot possibly include a description of all the methods that are in use today; such a document would fill a large book. This chapter, therefore, simply offers a selection of the simpler and more common methods that are used.

The Chain Ladder

While probably the oldest, this is still the most popular method of projecting claims reserves. Easy to apply to any data that can be arranged in a triangular format, the Chain Ladder is also simple to explain.

The Chain Ladder method works by calculating an average factor for estimating the cumulative amount in each year from the cumulative amount in the previous year. This average can be formed just by averaging the link ratios, but it is often more satisfactory – particularly where the volume of business has changed significantly over the years – to calculate the average weighted by the size of the business. The usual way of doing this is to add up the values themselves in neighbouring columns and to divide one by the other. Let us use the same data as in Table 1.3:

Table 2.1

Accident	Development years					
year	1	2	3	4	5	6
1995	90	210	310	420	500	500
1996	130	280	360	460	600	
1997	140	290	440	600		
1998	160	240	420			
1999	120	260				
2000	110					

To get the development between years 1 to 2, we can add up all the values in the column for development year 2 and divide it by the total of all the values in the column for development year 1 *excluding the last one*. We exclude the last value to form a like-for-like comparison, and we have no such comparison for the last value in any column.

The result of this calculation is 1,280 ÷ 640 or exactly 2. Proceeding similarly for the development from years 2 to 3 gives 1,530 ÷ 1,020 or 1.5 and the later values are 1,480 ÷ 1,110 (or 1.3333), 1,100 ÷ 880 (or 1.25) and 500 ÷ 500 (or 1.0). These are called the age-to-age development factors, or simply the development factors.

We can use these development factors to fill in the empty cells in the triangle, starting from the left (estimated figures in italics):

Table 2.2

Accident	Development years					
year	1	2	3	4	5	6
1995	90	210	310	420	500	500
1996	130	280	360	460	600	*600*
1997	140	290	440	600	*750*	*750*
1998	160	240	420	*560*	*700*	*700*
1999	120	260	*390*	*520*	*650*	*650*
2000	110	*220*	*330*	*440*	*550*	*550*

Thus, the last row is formed by multiplying the latest actual value (110) successively by the development factors calculated as above. The previous row uses all the factors apart from the first one, which was for the development years 1 to 2, because we already have an actual value at development year 2.

If we add up the values in the last column (3,750) and take away the latest actual values in each row (these add up to 2,490) the difference of 1,260 gives an estimate of the reserve for late claims.

We can achieve the same result without having to calculate so many intermediate values by multiplying the development factors to form a series of factors to ultimate, the reciprocals of which are described as the *lag factors*. Starting from the right:

Table 2.3

Accident year (1)	Development factors (2)	Factor to ultimate (3)	Derivation of factor to ultimate (4)	Lag factor (=1 ÷ factor to ultimate) (5)
1996	1.0000	1.0000	(= Development factor)	1.0
1997	1.2500	1.2500	(= 1.0000 × 1.2500)	0.8
1998	1.3333	1.6667	(= 1.2500 × 1.3333)	0.6
1999	1.5000	2.5000	(= 1.6667 × 1.5000)	0.4
2000	2.0000	5.0000	(= 2.5000 × 2.0000)	0.2

Multiply column (3) by the latest actual values to give the same results. Note that it is unnecessary to include 1995 in this table because it is assumed to be fully settled.

The Chain Ladder method is intuitively appealing and simple to calculate, and these attributes ensure its continuing popularity. But it has the following problems:

(a) It is a purely multiplicative method, ie the estimate for each origin period is formed by multiplying the most recent value in each origin period by a development factor. If the most recent value is zero, (as it may well be for very recent periods on long-tail business, particularly when looking at paid claims), then the estimate will be zero! On the other hand, if the most recent value is unusually large, perhaps because of some large claim, the development factor may overstate the eventual losses for this period.

(b) The link ratios must be stable across the origin periods for the method to produce sensible results and such stability is rare. In particular, the method is vulnerable to changes in the pace of claims settlement, especially when applied to claims paid.

The following modifications to the Chain Ladder may help:

1 If the development pattern has changed over the years, then using only data from the most recent calendar periods will produce estimates that better reflect the current conditions. This will not help so much where a trend is continuing.

2 If the results are affected by some link ratios that appear to be highly unusual then it may be possible to down-weight or eliminate them when calculating the development factors. This should only be done after investigating the reasons, to the extent that they are due to past events that are now unlikely to reoccur.

Special adjustments to the data, using information on claims settled, can help to deal with any changes in the pace of settlement, provided of course that this information is available.

Mathematically, the Chain Ladder method works by calculating a series of linear regressions of the form (for a column N in the triangle)

$$y = A.x + \varepsilon$$

where x represents the values in column N, y represents the values in column N+1, A is the estimated parameter and ε is a random error.

There is, however, no particular reason for assuming that the regression takes this form and some research has been done into extending the formula. The simplest extension is to permit a constant term, so that the equation becomes:

$$y = A.x + B + \varepsilon$$

The values A and B are then estimated and tested for significance. If B proves to be not statistically significant, then the equation collapses to the first one and we return to the traditional Chain Ladder. If, on the other hand, A proves to be not significantly different from 1, then the actual values do not in fact help forecast the future development.

Because it is universally known, much research has been conducted into the theory underlying the Chain Ladder (see Bibliography). Research topics have included examining the distribution of ε and assessing confidence limits for the size of the reserves.

The underlying technique is also used in the course of applying other, more elaborate techniques. These include those shown below, as well as other techniques not discussed in detail here, such as Average Cost Per Claim. This can be used when suitable data on claim numbers is available in triangular form, and it involves deriving a triangle of average costs, applying the Chain Ladder to both the numbers and the average costs and multiplying the results together.

Cape Cod

This technique is named after the location of a conference where it was first devised. As well as the same underlying claims data as for the Chain Ladder, the method requires some additional information, namely the corresponding premiums for each origin period. For origin periods based on claim occurrence (eg accident years) the premiums chosen should be earned premiums, whereas for origin periods based on when the business was written (eg underwriting years), the premiums chosen should be written premiums.

For the sake of simplicity, we shall assume that we have earned premiums of 625 for each year.

The thrust of this method is, for each origin period, to balance the proportion of the eventual claims outgo we currently know about against a similar proportion of the premium. The reserve can then be calculated as the remaining proportion of the premium. How do we know what that proportion is? We don't, but we can estimate it using the Chain Ladder.

The first step, therefore, is to calculate the lag factors as for the Chain Ladder (see Table 2.3). We then form a table as follows:

Table 2.4

Accident year (1)	Earned premium (2)	Latest losses (3)	Lag factor (4)	(2) x (4) (5)	IBNR (6)
1995	625	500	1.0	625	0.0
1996	625	600	1.0	625	0.0
1997	625	600	0.8	500	124.5
1998	625	420	0.6	375	249.0
1999	625	260	0.4	250	373.5
2000	625	110	0.2	125	498.0
Total		2490		2500	1245.0

Note that 1995 has been included. This year is assumed to be fully run-off, so the lag factor must be 1.0.

The proportion that we need to bring into balance columns (3) and (5) in the table above is 0.996 (ie $0.996 \times 2{,}500 = 2{,}490$). We then form our estimate of the IBNR in column (6) by multiplying: $0.996 \times$ earned premium $\times$ (1 − lag factor) and adding up the result.

What is this factor of 0.996? It represents an estimated average loss ratio – losses divided by the premiums – over the period 1995–2000. The method assumes that late claims will emerge in accordance with this initial long-term loss ratio.

The Cape Cod method can be more robust than the Chain Ladder method under some circumstances. For instance, if the actual value in Year 2000 had been zero instead of 110, then Cape Cod would still give a plausible result for the year of 476 rather than zero. In other respects, however, the Cape Cod method inherits its advantages and disadvantages from the underlying Chain Ladder model.

The assumption of a level loss ratio over many years is also rarely justified, although the method can be modified, for example, by taking only the last few years into account when forming the proportion. This is often justifiable, because the bulk of the IBNR reserve will be derived from the more recent years. In the example in Table 2.4 above, if we were to take, say, the last three years into account, then our loss ratio would be (420 + 260 + 110) ÷ (375 + 250 + 125) = 1.053 and recalculating the reserve on this basis produces an answer of 1,316.

Bornhuetter-Ferguson

The Cape Cod method derives its result from fixed proportions of premiums, the proportions being determined by methods such as the Chain Ladder. Once those proportions are calculated, it is entirely automatic in its application, except for determining the period over which to average the loss ratio. But in any triangle covering several years of business, it is likely that the loss ratio will vary considerably from year to year, not least because of the operation of the underwriting cycle. Since the total amount of reserves usually depends mainly on the most recent years, getting the right pattern of initial loss ratios in those years is important.

The Bornhuetter-Ferguson (or B-F, for short) method, named after two US actuaries who devised this technique in the 1970s, works much like Cape Cod, but by selecting in advance a different initial loss ratio for *each* year. If these initial loss ratios can be estimated with sufficient accuracy, then it is likely that the B-F method will be more accurate than either Chain Ladder or Cape Cod.

This does, of course, beg the question of how to reliably estimate these loss ratios in advance. It will involve the actuary:

1 asking the underwriters and rate-makers about how they are pricing their business;
2 projecting the trend in loss ratios; and
3 reviewing the results of both these exercises in light of his or her general knowledge of the market and trends in claims.

With regard to point 2, it may be useful to record the trends in the loss ratio implied by the triangle, ie to divide the triangle of paid or incurred claims by the premium. For our data above, this would give:

Table 2.5

Accident	Development years					
year	1	2	3	4	5	6
1995	0.144	0.336	0.496	0.672	0.800	0.800
1996	0.208	0.448	0.576	0.736	0.960	
1997	0.224	0.464	0.704	0.960		
1998	0.256	0.384	0.672			
1999	0.192	0.416				
2000	0.176					

Each of the values above is the original value divided by the earned premium for the year. It is generally more sensible to look at these results incrementally, which gives:

Table 2.6

Accident	Development years					
year	1	2	3	4	5	6
1995	0.144	0.192	0.160	0.176	0.128	0.000
1996	0.208	0.240	0.128	0.160	0.224	
1997	0.224	0.240	0.240	0.256		
1998	0.256	0.128	0.288			
1999	0.192	0.224				
2000	0.176					

The task, as usual, is to fill in the bottom right area of the triangle. For this purpose, trends in the first column will not affect the result, and trends cannot be sensibly fitted where there are fewer than three data points. This means that we will look for trends in the current example for development years 2–4. *What follows here is inevitably subjective and many other views are possible.*

In development year 2, there is no evidence of any significant trend, and we can use the average value, which is 0.205. In year 3, there was apparently a marked shift between years 1995–96 on the one hand and 1997–98 on the other. In practice, we would investigate the reason for this, but here we will assume the average of the last two values (0.240 and 0.288) applies for each later accident year. In year 4, again, 1997 has a noticeably heavier development than 1995–96, and we will assume that the 1997 figure also applies to later years.

For years 5 and 6, we will simply assume the latest available figures continue. To assemble all this gives the following completed triangle of expected loss ratios (forecast values in italics):

Table 2.7

Accident	Development years						Total
year	1	2	3	4	5	6	
1995	0.144	0.192	0.160	0.176	0.128	0.000	0.800
1996	0.208	0.240	0.128	0.160	0.224	*0.000*	*0.960*
1997	0.224	0.240	0.240	0.256	*0.224*	*0.000*	*1.184*
1998	0.256	0.128	0.288	*0.256*	*0.224*	*0.000*	*1.152*
1999	0.192	0.224	*0.264*	*0.256*	*0.224*	*0.000*	*1.160*
2000	0.176	*0.205*	*0.264*	*0.256*	*0.224*	*0.000*	*1.125*

The Total column gives us our *initial loss ratios* and we now use these together with the actual data to give us an estimate of the reserves. First, we must select an underlying method to give us the lag factors, and we can use those already calculated for the Chain Ladder. The B-F estimate of the reserve is achieved by taking:

Premium x (1 − lag factor) × initial loss ratio, as shown in Table 2.8 in column (5), which is calculated by multiplying columns (2), (3) and (4) together.

Table 2.8

Accident year (1)	Earned premium (2)	Initial loss ratio (3)	(1 - lag factor) (4)	Reserve (5)	Latest value (6)	Final loss ratio (7)
1995	625	0.800	0.000	0.00	500	0.800
1996	625	0.960	0.000	0.00	600	0.960
1997	625	1.184	0.200	148.00	600	1.197
1998	625	1.152	0.400	288.00	420	1.133
1999	625	1.160	0.600	435.00	260	1.112
2000	625	1.125	0.800	562.50	110	1.076
Total				1433.50		

The B-F result is considerably larger than the earlier results from Chain Ladder and Cape Cod, because of the worsening initial loss ratio. The last column in Table 2.8 is calculated as the sum of columns (5) and (6), divided by column (2) and shows how the view of the loss ratio has been changed having gone through the B-F process. Compare columns (3) and (7).

Separation

As a final example, we will look at a technique that examines the diagonal (ie calendar period) effects. The most common such effect is inflation and this can be seen most easily if we concentrate on claims paid in each year (ie incremental payments) rather than cumulative claims paid or claims incurred. Under some circumstances, however, it will make sense to look at the changes in incurred losses, for example when there has been a change in estimating practice, or when legal changes have abruptly altered the likely cost of future settlements.

The key assumption here is that the origin period effects can be factored out of the data by scaling all the values in the triangle by a suitable measure of exposure. If we are looking at motor business, for example, and we know how many vehicles we have insured in each year in the past, then we can divide each row of our triangle by the corresponding number of vehicles to produce a triangle of average payments per vehicle insured. We would expect this to be more stable – except for inflation, which we will measure – than the raw figures, as we look down each column in the triangle.

In the absence of any specific exposure information, premiums or numbers of claims notified may be suitable for this purpose. Whatever measure is used, once this scaling has been carried out, the incremental payments can be modelled by:

(Incremental) lag factor $\times$ index

where the lag factor depends on the column in the triangle and the index depends on the calendar period.

The original method of calculating the parameters for this technique used an ingenious algorithm that calculated the values for the index and the lag factors in turn, starting with the latest index value and the most mature lag factor, and then working backwards. Later, it was noticed that the Separation and Chain Ladder models are very similar in form; both are multiplicative models involving two factors. This means that we can actually use the Chain Ladder logic after scaling by the exposure measure, rearranging the incremental payments by swapping rows and diagonals, and then accumulating them.

Taking the same data as for the previous examples, and using the earned premiums as the scaling factor, the basic triangle of data, after scaling, is the same as in Table 2.5 above. The incremental values are also shown above, in Table 2.6.

We swap the diagonals and rows in Table 2.6 to give a triangle of the same size and shape, as follows (Table 2.9):

Table 2.9

Calendar	Development years					
year	1	2	3	4	5	6
2000	0.176	0.224	0.288	0.256	0.224	0.000
1999	0.192	0.128	0.240	0.160	0.128	
1998	0.256	0.240	0.128	0.176		
1997	0.224	0.240	0.160			
1996	0.208	0.192				
1995	0.144					

and then accumulate the values across the rows (Table 2.10).

Table 2.10

Calendar	Development years					
year	1	2	3	4	5	6
2000	0.176	0.400	0.688	0.944	1.168	1.168
1999	0.192	0.320	0.560	0.720	0.848	
1998	0.256	0.496	0.624	0.800		
1997	0.224	0.464	0.624			
1996	0.208	0.400				
1995	0.144					

We can now calculate the Separation parameters by applying the Chain Ladder logic to Table 2.10, which gives results (to 3 places of decimals) as shown in Table 2.11.

Table 2.11

Calendar year (1)	Development factors (2)	Factors to ult. (3)	Lag factors (4)	*Lag factors fm Table 2.3 (5)*
1999	1.000	1.000	1.000	*1.000*
1998	1.212	1.212	0.825	*0.800*
1997	1.316	1.595	0.627	*0.600*
1996	1.486	2.370	0.422	*0.400*
1995	1.970	4.669	0.214	*0.200*

This table is very similar to Table 2.3; column (5) shows, for comparison, the lag factors from the earlier table. The new lag factors are generally somewhat larger than those for Chain Ladder. This is due to the removal of inflation in the Separation process. If inflation is generally positive, as it is here, then the effect of the later, more inflated, payments will be reduced and the separation lag factors should be larger than the Chain Ladder ones, as indeed they are. Now we can proceed to calculate the Separation parameters themselves.

Table 2.12

Calendar year (1)	Separation parameters Index (2)	Lag (3)	Residue to be paid (4)	Lag development (5)
2000	1.168	0.000	0.000	6
1999	0.848	0.175	0.175	5
1998	0.969	0.198	0.373	4
1997	0.995	0.205	0.578	3
1996	0.948	0.208	0.786	2
1995	0.672	0.214	1.000	1

An extra row has been added here for 2000; in what follows, we assume the lag factor for 2000 is 1.000, as there is no further development available. Column (2) of Table 2.12. is calculated as the latest value in Table 2.10 divided by the lag factor in column (4) of Table 2.11. For 1998, for example, the value is 0.800 ÷ 0.825 = 0.969. Column (3) in Table 2.12. is calculated as the difference between successive lag factors from Table 2.11, and this expresses the incremental amount paid in each year. Column (4) is calculated by summing column (3) from the top: these values will be used when we forecast the run-off. Column (5) is explained below.

How do we interpret the Separation parameters? Apart from being expressed incrementally, the lag factors are fundamentally the same as in the other methods we have examined. We also have to interpret them in relation to the development periods to which they apply, as shown in column (5) of Table 2.12. Thus, the separation lag factor shown in column (3) of Table 2.12 for 1995 corresponds to development period 1, that for 1996, to development period 2 and so on. This will be important when we make forecasts with specific allowance for future inflation.

The index values are just like an index of, say, prices of consumer goods. In fact, they may well be correlated with such a general index. In many cases, however, the index values we have calculated will not correlate very well with any general index, either because the factors driving the cost of claims are not measured in any general index (eg the cost of claims for bodily injury) or because the observed movements depend on internal procedural changes (eg a change in settlement practice).

To turn the results of these calculations into a forecast, we first have to decide what the index values will be for future calendar periods; in other words, what is the future rate of inflation. If our index values were correlated with a general index, we may be able to refer to economic forecasts of the general index to help forecast the future index values.

We will assume no such correlation here. Over the whole period 1995–2000, our payments have inflated by over 11% per annum (from 0.672 in 1995 to 1.168 in 2000), but that is heavily influenced by the sharp increase from 1995–96 and the sudden increase in 2000, following a reduction in 1999. The first is probably not statistically significant, as we only have one payment figure for the calendar year 1995. In practice, we would certainly want to look at the reasons why the payment levels in 2000 were much higher. In this case, of course, the numbers are artificial and we will calculate the result first assuming no inflation and then assuming 5% per annum.

With no inflation, the calculation is as shown in Table 2.13 below.

Table 2.13

Accident year (1)	Earned premium (2)	Latest index (3)	Residue (4)	Reserve (5)
1995	625	1.168	0.000	0.00
1996	625	1.168	0.000	0.00
1997	625	1.168	0.175	127.75
1998	625	1.168	0.373	272.29
1999	625	1.168	0.578	421.94
2000	625	1.168	0.786	573.78
Total				1395.76

The earned premium remains as before. The latest index value is taken from the year 2000 figure in column (2) of Table 2.12 and the residue comes from column (4) of Table 2.12 (note that the order of the years is now reversed). The reserve in column (5) is then calculated simply by multiplying together columns (2), (3) and (4).

If we include 5% annual inflation then the residues shown in column (4) of Table 2.12 can no longer be calculated just by adding up the numbers. Instead they must be put into a triangular form, where the lag factors in column (3) of Table 2.12 are multiplied by 1.05^n, n being the number of years from now until the payment is expected to be made (assuming, for simplicity, that all payments are made at the end of the year). This is shown in Table 2.14 below. Here, for instance, the value shown for development year 3 for year 2000 (0.226) is equal to 0.205 (Separation lag factor from Table 2.12 for development period 3) multiplied by 1.05^2.

Table 2.14

Accident year	Development years 1	2	3	4	5	6	Inflated residue	Reserve
1995						0.000	0.000	0.00
1996						0.000	0.000	0.00
1997					0.183	0.000	0.183	133.59
1998				0.208	0.192	0.000	0.400	292.00
1999			0.215	0.219	0.202	0.000	0.635	463.55
2000	–	0.218	0.226	0.230	0.212	0.000	0.886	646.78
Total								1535.92

The inflated residue is the total of the results for each development year. If we substitute the inflated residues from Table 2.14 into column (4) of Table 2.13, we generate the reserves shown in the last column of Table 2.14.

The difference between the reserves in Tables 2.13 and 2.14 shows the significant effect that future inflation can have. Even the reserves without inflation are much higher than for the Chain Ladder or Cape Cod. This is due to the large index value (1.168) for the year 2000, and this underlines the importance of investigating the diagonal effects that this method can reveal. The reserves from the Separation calculation are closer to the figures obtained under Bornhuetter-Ferguson, where the increased initial loss ratios (an origin period effect) led to the higher figures.

Stability tests

Let us compare our results for each of the above methods with those we would obtain if we strip out the most recent diagonal.

First, we can calculate the actual development during the year 2000 for years 1995–99. These are:

Table 2.15

Accident year	Position at 1999	Position at 2000	Movement
1995	500	500	Nil
1996	460	600	140
1997	440	600	160
1998	240	420	180
1999	120	260	140
Total	1760	2380	620

Repeating our Chain Ladder calculations without the latest diagonal gives lag factors of: 0.229, 0.449, 0.640, 0.840, 1.000 and an overall reserve calculation of 1,033 for these years. But if we look at Table 2.2, ignoring the line for year 2000, we see that we still need 820 (= 1,260 − 440) to cover 1995–99. Given that actual development of 620 has already occurred in year 2000, this indicates that our projection for these prior years has deteriorated from 1,033 to 1,440 (= 820 + 620). This is a large percentage increase and indicates that the Chain Ladder model does not fit the data very well.

Cape Cod fares just as badly. Recalculating without the latest diagonal gives an overall initial loss ratio of 0.892 (compared to 0.996 when year 2000 is included) and a reserve of 1,027. This compares to 747 for these years in the latest calculation (see Table 2.4) and again, when the actual development of 620 is added back, this indicates significant deterioration (from 1,027 to 1,367).

The main thrust of the Bornhuetter-Ferguson approach is to set the initial loss ratios; it is instructive to think what we might have selected without the latest diagonal. The triangle of incremental loss ratios would have been:

Table 2.16

Accident year	Development years 1	2	3	4	5
1995	0.144	0.192	0.160	0.176	0.128
1996	0.208	0.240	0.128	0.160	
1997	0.224	0.240	0.240		
1998	0.256	0.128			
1999	0.192				

The position for development year 2 is almost the same (the average is 0.200), but in year 3, much depends on whether we believe that the latest factor of 0.240 has set a new level or whether it is an one-off value that will not be repeated. This illustrates once again the importance of looking behind the data at the real reasons why the figures are what they are. We are also likely to under-estimate, compared to our year 2000 selections, the year 4 and year 5 factors.

We might finish up with something like the following, the forecast years in italics:

Table 2.17

Accident year	**Development years 1**	**2**	**3**	**4**	**5**	**Total**
1995	0.144	0.192	0.160	0.176	0.128	0.800
1996	0.208	0.240	0.128	0.160	*0.128*	*0.864*
1997	0.224	0.240	0.240	*0.176*	*0.128*	*1.008*
1998	0.256	0.128	*0.200*	*0.176*	*0.128*	*0.888*
1999	0.192	*0.200*	*0.200*	*0.176*	*0.128*	*0.896*

This still represents a considerable worsening of the initial loss ratio compared to the base year figure (and the peak in 1997 is still visible), but the overall level is much below what we were using for the year 2000 projections. Using both these initial loss ratios and the development factors listed above from the 1999 Chain Ladder calculation gives overall reserves of 1,051. The year 2000 B-F projection for these years is 871 (see Table 2.8) and adding back the year 2000 development of 620 indicates deterioration to 1,491.

Recalculating the Separation model gives a reserve of 941 (assuming 0% future inflation) or 1,036 (assuming 5% future inflation). The prior years' reserve at 5% inflation from Table 2.14 is 874, so adding back the year 2000 development indicates deterioration to 1,494, a similar position to the other methods.

None of the methods, then, gives stability as regards the prior years' development in year 2000. If this is due to the inflation in the year 2000 diagonal, this is not surprising, because this inflation was not a feature of the triangles as they stood at the end of 1999. Indeed these results, together with the parameters found for the Separation model in year 2000, tend to confirm that year 2000 inflation was the principal problem in this analysis. The causes of this would, in practice, need to be carefully investigated before coming to any conclusion about the level of reserves to be held.

For larger triangles, it is possible to perform this stability analysis over several years, removing successive diagonals and recalculating the results each time.

3 Conclusion

Setting the level of claims reserves for non-life insurance operations, and late claims reserves in particular, is undoubtedly difficult. It is neither simply a statistical exercise, nor can it ever be based purely on judgement and business understanding. Success in practice depends on achieving the right blend of rigorous analysis and judgement. These will be based on careful investigation of the facts and sound background knowledge, both of the business and the social, legal and economic conditions which prevail at any given time.

Given that no forecast of the future can ever be wholly accurate, it is important to understand both the magnitude of likely variations and their causes. Each year's new data provides information on the run-off of older years while also posing new problems in respect of the most recent year. Tracking the run-off of the older years against what was previously predicted is key to understanding the dynamics of the whole reserving process. Indeed, that tracking is part of a regular control cycle: analysis leads to evaluation, evaluation leads to conclusions that require testing and, once the testing has validated the conclusions, then what we have learnt helps improve the analysis when the next set of new data comes in. At Swiss Re, our Group reserve specialists carry out one circuit of this process every year.

The authors would be interested in hearing about readers' experience in using the common methods of reserving. More, at the address following, they welcome reactions to the methodology and materials presented in this publication.

Head of Group Reserves
Swiss Re
Mythenquai 50/60
8022 Zurich
Tel. +41 1 285 21 21
Fax +41 1 285 20 23

Measuring reserve variability using the "Bootstrap"

The Bootstrap idea comes from general statistics and was devised about 1980 by Bradley Efron. It can provide numerical solutions to a wide range of statistical problems. To use it in this context requires the following stages:

1 Calculate the normal method (ie for the point estimate) and record both the results and the parameters;
2 Use the parameters to fit the "expected" values, in other words, to produce a triangle that would exactly fit the parameters calculated in the first stage;
3 Subtract the expected values calculated in stage 2 from the actual values to produce a set of "residuals";
4 Produce a new set of pseudo-"actual" values by taking each expected value and adding to it one of the "residuals" *selected at random;*
5 Apply the original method to these pseudo-actual values and record the resulting estimate of the reserves;
6 Using a computer, repeat stages 4 and 5 a very large number, say, 1,000 times;
7 Sort the values recorded under each iteration of stage 5 and read off the points you are interested in. For example, if you want to know the range within which the result will fall 90% of the time then you can select (from 1,000 sorted values) the 50th and 950th values.

In practice, the Bootstrap is a little more complicated to apply than this. The residuals calculated in stage 3, in particular, need to be independently and identically distributed and this often requires some work.

To illustrate how we can calculate the residuals, consider our Chain Ladder example from Chapter 2. We derive our expected values by applying each of the lag factors to the projected ultimate values (ie the latest values plus the calculated reserves), giving the following results:

Table A 1

Accident	Development years						Projected
year	1	2	3	4	5	6	ultimate
Lag factors	*0.200*	*0.400*	*0.600*	*0.800*	*1.000*	*1.000*	
1995	100	200	300	400	500	500	500
1996	120	240	360	480	600		600
1997	150	300	420	600			750
1998	140	280	420				700
1999	130	260					650
2000	110						550

Each value in the triangle is calculated from the projected ultimate values on the right, multiplied by the lag factor in the same column. To compare this triangle with the original (Table 1.3) will show that all the values on the latest diagonal are the same, which is inherent in the Chain Ladder model, but the other values, in general, are not the same.

The residuals are therefore formed by taking each value in Table A1 away from the corresponding entry in Table 1.3, giving:

Table A2:

Accident	Development years					
year	1	2	3	4	5	6
1995	-10	10	10	20	0	0
1996	10	40	0	-20	0	
1997	-10	-10	20	0		
1998	20	-40	0			
1999	-10	0				
2000	0					

Because the residuals along the diagonal are constrained to be zero, they should be excluded from the rest of the process.

If the residuals appear to show a pattern, then they will need to be transformed. For example, if they appear to increase as the underlying values from Table A1 increase, which would often be the case, then it may help to express the residuals as proportions of the corresponding values in Table A1. The object is to get a set of residuals that do not show any pattern. In the present case, there is little evidence of any pattern and, for simplicity, this stage will be omitted.

The "Bootstrap" works by building new sets of pseudo-data, adding to each figure in Table A1 an element of Table A2 selected at random. It does not matter if a particular residual is selected two or more times; indeed it is essential that each selection can choose from any of the residuals (known as "sampling with replacement").

Each set of pseudo-data is projected using the chosen method – Chain Ladder here – and the result we are interested in is recorded. In this case, we are interested in the total reserves, but we may also look at any aspect of the projection we choose. If the triangle is of payments, for example, then we may be interested in the distribution of payments to be made in subsequent calendar years. This means recording the total value of payments in subsequent diagonals (the values in italics in Table 2.2, for example) converted to an incremental basis.

After the simulation, the results are sorted by size and an estimate of the range of likely results can then be read off from these sorted results as described above.

With any simulation method, the result can never be wholly precise, but the more simulations that can be done, the more precise the answers will be. How many simulations are required will depend on the degree of precision that is needed and also on the nature of the problem. If we are looking at features in the extreme "tail" of the distribution – trying to assess 99% confidence limits, for example, which will mainly depend on 0.5% of the values generated at either end of the distribution – then we may need to carry out thousands of simulations. With a computer program running on modern hardware, however, this should take only a few minutes at most.

Accident year: For any analysis of claims, it is necessary to classify them in some way. Here, claims are grouped according to the year in which they occurred (cf *notification year, underwriting year).*

Average cost per claim (ACPC): *Average cost per claim* refers to a claims projection method that uses two *triangles,* one of claims payments or *claims incurred* and the other of claim numbers. For a brief description, see the section on "Chain Ladder" in Chapter 2.

Acquisition costs: The part of the written premium relating to all costs of writing the contracts, including *commission, brokerage* and *profit commission.*

Additional case reserve (ACR): Provision held by a reinsurer in relation to a specific claim or claim event where there is reason to believe that the reserves reported by the cedant are likely to be inadequate.

Additional unexpired risk reserve (AURR): See *premium deficiency reserve.*

Allocated loss adjustment expenses (ALAE): *Allocated loss adjustment expenses* are the expenses of settling a claim that are particular to that claim, such as fees paid to loss adjusters or lawyers dealing with the claim. See also *ULAE.*

Bootstrap: Simulation technique to estimate (in this context) the monetary range within which, with a given probability, the eventual losses are expected to fall, according to the chosen reserving method. See the Appendix.

Bornhuetter-Ferguson: Claims projection technique that uses a claims *triangle,* information about premiums and initial *loss ratio* assumptions. See Chapter 2.

Brokerage: For reinsurance (where *commission* is paid to the *cedant* to cover expenses) this describes the fee paid to the intermediary.

Cape Cod: Claims projection technique that uses a claims *triangle* and information about premiums. See Chapter 2.

Case estimates: The sum of many individually assessed values, being the expected costs of settling each claim already notified to the (re)insurer.

Catastrophe provision: Reserve held to help cover the costs of large claims that occur infrequently and which are therefore unpredictable. Natural disasters, such as earthquakes, windstorms and floods, are typical. See also *equalisation provision.*

Cedant: Name to describe an insurance company that purchases reinsurance from a reinsurer.

Chain Ladder: Simple projection technique that can be applied to a data *triangle.* See Chapter 2.

Claims expenses: Internal and external costs incurred by the (re)insurer in administering claims.

Claims expense provision: Reserve held to cover the future *ULAE* of running off all claims arising from the (re)insurer's business, apart from claims already finally settled.

Claims incurred: In relation to a set of risks, this is the sum of claims paid to date plus the *claims reserves* held in respect of those risks.

Claims reserves: A term that covers provisions in respect of claims arising from risks that have already occurred. It includes the *notified claims reserve* and the *IBNR* provision.

Commission: An amount, usually a percentage of premium, paid (in the case of insurance) to the intermediary introducing the business. For reinsurance it is, where applicable, the amount paid by the reinsurer to the *cedant* to cover the latter's *acquisition costs.*

Deferred acquisition costs (DAC): When *written premium* is booked, the proportion that relates to *acquisition costs* will be set up as an asset (DAC). This DAC will be reduced over time as money is released from the *UPR* (ie as premiums are earned).

Discounting: The process of reducing the nominal value of a provision by taking credit for the investment income that is expected to be earned on the corresponding assets before the liability is due to be paid.

Earned premium: At any time, the portion of *written premium* that corresponds to the amount of risk already experienced by the (re)insurer. The balance of the *written premium* is retained in the UPR.

Equalisation provision: Reserve held to smooth the profit and loss over time by adding to the reserve in profitable years and releasing from the reserve when losses are made. Such reserves are not always permitted and, where permitted, may be subject to strictly defined conditions to prevent tax avoidance.

IBNER (reserve): Meaning *incurred but not enough reserved,* this describes a provision that is held to cover claims already notified, in addition to the *case estimates.* It will either form part of the *IBNR* provision or the *notified claims reserve.*

IBNR (reserve): Meaning *incurred but not reported,* this provision covers the estimated cost, including *ALAE,* of claims that are thought to have already occurred but which have not yet been notified to the (re)insurer. This defines a "pure" *IBNR;* sometimes the term also includes any *IBNER.*

Link ratio: For a data *triangle,* this is defined for each cell except for those in the first column, and is equal to the value in the cell divided by the value in the cell immediately to the left.

Loss ratio: The ratio of *claims incurred* to premium. The latter will be taken as *written premium* for business on an underwriting year basis and *earned premium* for business written on an accident year basis.

Notification year: For any analysis of claims, it is necessary to classify them in some way. Here, claims are grouped according to the year in which they were notified (cf *accident year, underwriting year).*

Notified claims reserve: Provision to cover the entire remaining costs of settling all claims that have been notified to the (re)insurer, including *ALAE.*

Pipeline premiums: An estimate of *written premiums* which have been contracted (for example by agents) but which have not been booked by the (re)insurer.

Premium deficiency reserve: Amount of reserve, if any, required in addition to the *UPR* to provide for future risks already contracted. Also known as *additional unexpired risk* reserve (AURR).

Profit commission: An additional *commission* paid by a reinsurer to an insurer that depends on the profitability of the business reinsured. It is often on a sliding scale that depends on the *loss ratio.*

Provision: In general accounting, a liability that can be reliably assessed. In (re)insurance the various provisions relating to aggregated claims costs and expenses of settlement are regarded as such, even though the cost of any particular claim might fail such a test. The term in this context is used interchangeably with *reserve,* although the strict accounting definition is quite different.

Reinstatement premium: On certain types of reinsurance (eg catastrophe reinsurance) the original premium only covers the first claim that might be made. In order to maintain cover after a claim has occurred, the *cedant* must then pay a further premium. This process may be repeated on subsequent claims; the reinsurance contract may, however, limit the number of reinstatement premiums that can be paid.

Reopened claims reserve: A (usually small) provision to cover any further costs that might arise from claims currently believed to be settled. Often just part of the *IBNR.*

Reserve: In (re)insurance accounting, this term is often used interchangeably with *provision.*

Retrocession: A second stage of reinsurance whereby a reinsurer protects his own account by purchasing reinsurance from another reinsurer.

Separation: Claims projection technique that makes explicit allowance for inflation. See Chapter 2.

Triangle: Common arrangement of aggregated claims data for reserving purposes. Claim movements (this term includes notifications, payments, changes in estimates and settlements) are aggregated in rows according to their origin grouping (eg *accident year, underwriting year*) and in columns according to how long it took from the beginning of the origin period before the movement was recorded.

Unallocated loss adjustment expenses (ULAE): Refers to the costs of settling claims that cannot reliably be attributed to particular claims, and is equal to the *claims expenses* minus the *ALAE.*

Underwriting year: For any analysis of claims it is necessary to classify them in some way. Here, claims are grouped according to the year when the period of insurance in which they occurred began (cf *accident year, notification year*).

Unearned premium reserve (UPR): Consists of the portion of premiums already booked relating to risk that has not yet occurred.

Written premium: Premium booked by the (re)insurer as soon as the contract is made, in respect of all risks assumed under the contract, including commission and expenses.

A short selection of publications which cover related subject matter follows. Some of the suggested works presuppose a considerable knowledge of mathematics.

S. Benjamin and J. M. Eagles. *Reserves in Lloyds and the London Market,* JIA 113, 1986.

R. L. Bornhuetter and R. E. Ferguson. *The Actuary and IBNR,* PCAS 59, 1972.

B. Efron and R. Tibshirani. *An Introduction to the Bootstrap,* Chapman & Hall, 1994.

T. Mack. *Measuring the Variability of Chain Ladder Reserve Estimates,* CAS Prize-winning paper, 1994.

G. S. Patrick. *Reinsurance,* Foundations of Casualty Actuarial Science (3rd Edition), CAS, 1996.

S. W. Philbrick. *Reserve Review of a Reinsurance Company,* Discussion Papers on Reinsurance, CAS, 1986.

A. E. Renshaw. *Chain Ladder and Interactive Modelling,* JIA 116, 1989.

G. C. Taylor. *Claims Reserving in General Insurance,* Elsevier Science Publishers, 1986.

Various authors. *Claims Reserving Manual,* Institute and Faculty of Actuaries, 1997.

Abbreviations:
JIA = Journal of the Institute of Actuaries
PCAS = Proceedings of the Casualty Actuarial Society (CAS)

Title: Late claims reserves in reinsurance

Authors:
Anthony Boulter, Head of Group Reserves
Dawson Grubbs, Group Reserves

Graphic design:
Markus Galizinski, Zurich

Editing and production:
RK, Corporate Communications
Reinsurance & Risk division

Additional copies and a brochure listing other Swiss Re publications (Swiss Re Publishing – Our expertise for your benefit) can be ordered from:

Swiss Re
Mythenquai 50/60
P.O. Box CH-8022 Zurich
Telephone +41 1 285 21 21
Fax +41 1 285 20 23
E-mail publications@swissre.com

Swiss Re publications can also be downloaded from our Website
www.swissre.com

Order no. 207_8955_en

R&R 12/00, 2000 en