Section 2
INTRODUCTION TO RESERVING

[2A]
DESCRIPTION OF CLAIM PROCESS

The insurance background and the influences that need to be taken into account in the reserving process are described in some detail in later sections. In this Introduction we shall give a brief overview of what the claims process in insurance entails, and how the problem of reserving arises.

We may start by supposing that we have a "risk" situation, associated with an insurance "cover". The essential features are that there is a person or corporate body, whose financial condition is directly affected by the occurrence of certain "events" occurring over a defined period of time. An obvious example would be an individual with Private Car Damage insurance who would suffer repair costs following an accident causing damage to their vehicle.

If an insurance cover exists, an event occurring under the cover will give rise to an insured loss, which subsequently becomes a claim on the insurer.

Typically there will be some delay between the event giving rise to the claim, and the ultimate settlement of the claim with the insured. In the case of a Motor claim, this delay from event to settlement may stretch from a number of days to several years, depending upon factors such as complexity or severity of the claim.

Other significant dates are involved in this process. Following the occurrence of the event, a significant date would be the date upon which that event becomes known to the insurer. Whilst notification would normally occur quite soon after the event, there can be circumstances where a considerable period of time may elapse between the occurrence of an event and the notification of a claim to an insurer (for example, when a ship is damaged in harbour, but the damage becomes evident only when she is dry-docked at some later date).

Claim Reserves

The delay between event and settlement dates means that the insurer must set up "reserves" in respect of those claims still to be settled. The reserves required at any time are the resources needed to meet the costs, as they arise, of all claims not finally settled at that time. The insurer must be able to quantify this liability if it is to assess its financial position correctly, both for statutory and for internal purposes.

Throughout this Manual we are dealing with claims reserves in respect of events which have already occurred. This is distinct from future claims, arising from risks covered by the insurer over the remaining period of the policy, where the insured event has not yet occurred. Such liabilities are covered through the mechanism of an Unearned Premium or Unexpired Risk reserve, which are outside the scope of this Manual.

We are concerned then with reserves in respect of claims which have occurred as at a particular date — which we shall call the "valuation date". We can distinguish two categories of such claims:

- claims for which the event has occurred, and which are already known and reported to the insurer;

- claims for which the event has occurred, but which have not yet been reported to the insurer.

Reserves relating to the former category are normally referred to as "Known (or Reported) Claims" reserves; the latter as "Incurred But Not Reported (IBNR)" reserves.

This definition of IBNR applies usually in relation to Direct Insurance: in the context of the London Market, the usage of "IBNR" will normally include, in addition to reserves for unreported claims, allowance for any future deterioration in amounts outstanding on claims already reported. The latter element is normally called "IBNER" (incurred but not enough reserved).

A number of variations on the type of claims reserve are needed because of various features including:

- a claim may involve a number of partial payments separated over an extended period of time, culminating in its final settlement;

- claims may be settled prematurely by an insurer, and then require to be reopened for further payments or recoveries that subsequently come to light (this may occur on more than one occasion for a particular claim);

- the insurer may share the liability for the claim with other insurers, either by reinsurance or by co-insurance.

Thus the need may arise for reserves for re-opened claims, and for reserves net of reinsurance or co-insurance.

From this description, it is clear that reserves represent an attempt, at a point in time, to attribute a financial value to those payments still to be made in respect of a set of incurred losses, as yet unsettled. This cannot therefore be quantified with precision, but must be the subject of estimation. Varying assumptions about future influences on the outcome of those losses will lead to greater or smaller estimates in a given context, leading to the idea of strength of reserves. In consequence, information will often be needed as to the likely adequacy of a given estimate of reserves. This may in turn involve careful examination of the methods by which estimates are reached, and the assumptions on which they are based.

The description of claim events given above is somewhat idealised. In many instances, it will be possible to identify a finite event or occurrence which gave rise to a claim in the way described. However, in others, the question of defining the relevant event can be one of considerable difficulty and argument. We need only consider the situation with claims for compensation for employment-related disease, where the dates of the relevant event, or indeed its definition, can be difficult to determine. The need for modification and extension of the picture included here will be addressed in certain specific areas outlined below.

Need for Reserves

The need to calculate reserves arises in a number of different circumstances, of which the following are examples:

- assessing the financial condition of an insurer, since movements in reserves over a period are key to assessing its progress;

- pricing insurance business in the sense of estimating the future cost of claims on risks yet to be taken on (by extrapolation of past paid and reserved claim cost);

- assessing the solvency of an insurer, in terms of its ability to meet its liabilities (requiring assessment of likely upper limits of outstanding claim cost);

- putting a value on the net worth of an insurer, particularly for purposes of sales or acquisitions;

- commutations and reinsurance to close: that is, putting a financial value on the run-off of a portfolio of insurance business.

The strength that is appropriate for the reserves may vary from one of these circumstances to another.

Reserving Methods

Until the early 1970s, the approach to reserves commonly related solely to the area of reserves for known claims. The practice was for each claim to be individually assessed by a claims official at an early stage of its existence, and possibly at subsequent stages. These individual "case" estimates would then be aggregated to form a total reserve for outstanding claims. With the passage of time, and the increasing ability to subject the results of this process to statistical scrutiny, it was found that other methods or approaches to reserving might be more appropriate. Nevertheless, it is still the case that, in some areas of business, particularly where the numbers of claims are relatively small, or where they are particularly complex, case estimation is employed, possibly in conjunction with other methods

However, in many instances the volume of claims is such that it would be impractical or too expensive to assess each claim individually and, in such cases, an alternative approach is required. Furthermore, in the case of IBNR reserves, specific claim files are not available for examination. In such cases, an alternative method of reserving is required.

The case estimator — claims official — requires access to details of policy cover, claim event and subsequent history before making their assessment of each claim. Similarly, in using statistical estimation methods, the reserver requires access to relevant data relating to the group of claims that are required to be reserved.

The more obvious of such data would normally relate to claim cost levels and perhaps numbers of claims settled in the recent past — together with some information such as the numbers of claims in the group unsettled at the present time. It might however be that, in some instances, the claims reserving specialist has available not only what might be called "hard fact" relating to actual past settlement experience, but additional "softer" information relating to the claims concerned. This may include more generalised market and economic information.

As will be discussed in Section 2B below, reserving methods would normally involve the application of a series of assumptions to underlying data from which a reserve would be evaluated. It is crucial, if use is to be made of soft data, that this aspect is fully reflected in the methodology employed and reports written. Indeed, in some methodologies, it might be felt that the nature of factual data employed (for example, assuming ratios of claims paid and outstanding to premiums for past claim years), is too questionable to lead to valid results. This is when the "softer" data may be of some value, even if they relate to more general market conditions.

In reviewing methods of reserving to be employed, an important initial point of consideration will be the nature and quality of data upon which the method will be based, together with the extent to which the use of a particular model is likely to be valid for those data.

For this reason, it is crucial that the individual carrying out the reserving exercise should be as familiar as possible with the underlying business concerned. If that person is situated within an insurance operation, then they should maintain close contact with underwriters and claims management. If not, then it is still important to understand as closely as possible the origins of whatever data sources are to be used.

<>

[2B]
RESERVING METHODOLOGY — GENERAL

Introduction

This section gives a general overview of the methodology used in reserving. Most of the comments that follow could equally well apply to any situation where one is constructing a model, fitting it to past observations, and using it to infer results about future statistics of interest.

Whatever the purpose of reserving, in essence it involves the following steps:

1. Construct a model of the process, setting out the assumptions made.
2. Fit the model, using past observations.
3. Test the fit of the model and the assumptions, rejecting or adjusting it.
4. Use the model to make predictions about future statistics of interest.

By "model" we refer to an artificial creation whose function is to represent what is important in the process under consideration, but to omit aspects not considered relevant to the particular area of understanding. The design of an appropriate model thus involves a process of selection from among many possible models. Normally, as an aid to understanding, we would try to select as simple a model as possible — indeed, the model can be regarded as a deliberate simplification of the underlying process itself.

There is a very close analogy with the underlying scientific method, which involves the successive refinement and replacement of models to improve their validity and accuracy for the purpose in hand.

Unfortunately, in the real world of claims reserving, restricted availability of data often places severe limitations on the models that can be applied in practice. It is clearly inappropriate to use a model which cannot be supported by the data available.

We are thus led to a situation where, in some practical instances, although a particular model may be felt to have a number of disadvantages, the available data and knowledge of the claims environment do not support any further refinement of this model.

The alternative is to revert to an earlier stage of model development, and to produce an alternative approach based on somewhat differing assumptions as to the underlying nature of the claim process. This might be used to create an alternative version of the original reserving exercise. This process may in turn be repeated, each such development producing alternative estimates of reserves.

In practice, a consequence of the limited availability of information may be that we have no alternative but to introduce a fifth step to those described above, as follows:

5. Apply professional judgement and experience to choose a number.

In other contexts, particularly those where direct access to the underlying claims data files can be made, it may be possible to develop apparently differing approaches to a point of broadly similar estimate values. The more elaborate models to which this development gives rise are among those contained in Volume 2 of the Manual. Even here, however, judgement as to the choice of model and the value to be chosen to represent the reserve is still largely at the reserver's discretion.

Ultimately, our objective is to formalise, as far as possible. the use of numeric or other available information, together with the judgement of the reserver, through the use of appropriate models.

Section 2E discusses at some length the elements that make up professionalism and judgement, the need for which cannot be emphasised too strongly.

The following sections describe some of the general considerations that affect the steps outlined above, and some of the more specific points that arise when performing such a process in the context of reserving.

Constructing a Model

The type of model constructed will depend on the purpose of the exercise. Estimates of future claims for pricing purposes may be required at a very detailed level (by type of car or geographical location, for example), and to be produced quarterly or monthly. Reserves for high level management information or statutory purposes may be required only by broad classification or on an annual basis.

Similarly, the model must reflect the data that are available and their limitation. Some types of information may not be adequately recorded (details of claims settled at no cost, for example). In other instances, past data may not be available for a sufficient period of time historically, or be of too small a volume, to form a credible basis for fitting the model. When data of an adequate quantity or quality are not available, it may be possible to supplement the available company-specific data by reference to industry-wide data. Section E13 refers to some of the possible sources of these additional external data. This section also sets out techniques for projecting data beyond the available experience, by using regression or graphical techniques. Data availability and limitations are also discussed further in Volume 1 Section B.

The foundations of any model are the assumptions that underlie it. Ideally, any modelling process should start with a clear statement of the assumptions being made. The model can then be fitted to the data, and the assumptions tested. Some simple reserving methods do not set out explicitly the assumptions underlying the method. This may make it hard or impossible to gauge the appropriateness of that particular model.

There is no easy classification of the types of model that can be used for reserving, but there are various distinctions that can be made. In the first instance, one can make the distinction between models that do not appear to start from a formal set of assumptions. and those which rigorously set out the foundations on which the model is built.

Many of the methods in Volume 1 fall into the former category, and those in Volume 2 the latter. To produce valid reserve estimates using methods that fall into the first category, it is essential that the user should consider the implicit and explicit assumptions being made, and endeavour to test their validity or otherwise. Just because a method jumps straight to step 4, does not mean that steps 1, 2 and 3 can be forgotten! The apparent simplicity of the arithmetic involved in performing some of the simpler methods in Volume 1 should not distract the user from the complexity of the underlying assumptions being made when applying these techniques.

A second general distinction can be made in the approach taken in constructing a model. Many of the simpler models do not, on the surface in any event, start from a base point that considers the underlying process influencing the future claims payments (cars crashing, houses burning down, people receiving compensation for injury, and so on). Instead, at a crude level, they are simply taking one aggregate set of values (the past claims experience), and making some estimate from this as to how the progression of values will continue (the future claims experience). Step 3 of the reserving process (testing and adjusting the model), will, of course, bring information about the nature of the claims into the process, as will step 5.

By contrast with this initially arithmetic approach, other models begin with a recognition of the underlying nature of the risks, perhaps starting with assumptions about the frequency and severity of claims, and modelling how these may change over time.

A final distinction that can be made is between deterministic models and stochastic models. The future claims payments predicted by a reserving model are random events. We can never know with certainty what the future payments will be. The best that any model can do is to produce an estimated value of those payments. Deterministic models just make assumptions about the expected value of future payments. Stochastic models also model the variation of the future payments. By making assumptions about the random component of a model, stochastic models allow the validity of the assumptions to be tested statistically. They therefore provide estimates not just of the expected value of the future payments, but also of the size of the possible variation about that expected value.

All the models in Volume 1 are deterministic. Several of the models in Volume 2 are stochastic. The introduction to Volume 2 describes the general nature of stochastic models, and sets out some of their strengths and weaknesses.

Fitting and Testing the Model

Fitting a model may very often be a simple arithmetic process. Testing a model is not so straightforward. A model can be broken down into a number of parameters. For example, the basic chain-ladder model for an $n \times n$ triangle of data may be summarised as n different "levels" for each year of origin, and n different development factors for the n years of claims development, making $2n$

parameters in all. Other models may have more or fewer parameters for the same sized triangle of data.

The more parameters a model has, the better it may appear to "fit" the observed data, but the less stable it will become for predicting future values. As an extreme example, if one had, say, twenty data points, one could fit a model with twenty parameters that would "fit" the data perfectly. The model would, however, be completely unstable. A small change in any of the data points may result in a large change in the parameters fitted and the future values predicted. There is always a compromise between having enough parameters to fit the data adequately, but few enough to produce a model with a certain amount of stability and predictive power.

The amount of data available for reserving is often very limited. It is important therefore to appreciate the limitations of an over-parameterised model. This may be done, for example, by examining the effect on the model, and the predicted results, of small changes in the data to which the model is being fitted. We know the data are just one realisation of some random process, so a given data point could reasonably be a little larger or smaller than in fact was the case. For a model to be acceptable, ideally, the predicted future payments should not be greatly affected by small changes in the observed data. The user should therefore check carefully whether the results are heavily influenced by a few data points.

Testing the assumptions of a model can take many forms. Whilst there are a variety of statistical techniques that can be used for this purpose, the approach in practice may often be more pragmatic. This is particularly so for the methods in Volume 1, where "soft" information may be used as a guide in choosing development factors, for example. This is not easily amenable to explicit testing. Many of the simpler, chain-ladder based, models are based on broad assumptions about the stability of the types of business, and the speed of settlement of claims, for example. The user of a reserving model should attempt to validate that these assumptions apply to the particular class of business being considered. Ideally, this should be done in a quantitative fashion but, as a practical compromise, it may have to take the form of a qualitative assessment. For example, this may involve a review of the types of policy written, or a visual examination of whether settlement rates appear to have remained stable.

Throughout the text of Volume 1, the reader's attention is drawn to some of the assumptions underlying the different methods. An awareness and questioning of these and other implicit assumptions should be kept in mind whenever one applies the methods in practice. Section B of Volume 2 discusses in a little more detail the requirements of a "good" model, and some of the more general techniques for testing models in general and stochastic models in particular.

[2C]
RESERVING IN CONTEXT

So far, reserving has been presented in terms of points of principle, removed from the particular features of specific classes of general insurance business. As reserving must in the end be a practical application, it is worth setting out the main characteristics of some specific reserving situations, as they in turn will affect the choice of appropriate reserving method. The examples given are not intended to be exhaustive, but indicate the range of characteristics found in general insurance.

Reference should be made to a suitable general insurance primer for basic details of policy coverage found in different classes of business.

Private Car

The main classifications for analysing data are:

cover type
- third party
- comprehensive

peril
- fire
- theft
- windscreen
- third party bodily injury
- third party physical damage
- accidental damage

In general this business is amenable to statistical treatment, as there will usually be a large number of similar policies, generating a large volume of claims.

Claims are mainly in respect of property damage, with a smaller number of liability claims. Third Party bodily injury claims are subject to the greatest uncertainty, since

- the size of claims can vary enormously, and for seemingly similar situations
- claims can remain outstanding for 7–10 years in some cases
- claims inflation is difficult to extrapolate, as judicial inflation in particular tends to be fairly erratic in its development.

Creditor

The claims should be split between the main perils (accident, sickness and unemployment), as the frequency and duration of claims may differ markedly, and vary according to the economic cycle.

If sufficient information is available, it is possible to use an annuity-type approach to reserving. Each outstanding claim is reserved according to the number of months the claim has already been running.

Experience can differ markedly by scheme, reflecting different selling practices and customer bases.

Mortgage Indemnity Guarantee

Following a period of severe market losses, the standard insurance cover changed during 1993/4. The new cover required the insured (the mortgage lender) to retain part of the loss on a coinsurance basis, together with the imposition of a maximum indemnity per property. Whilst this will serve to spread the loss more evenly between the insured and the insurer, the basic claim characteristics remain unaltered, namely:

- claims arise from a combination of repayment default and loss of property value. Whilst there is a steady underlying level of default due to marital breakdown and individual financial problems, the main underlying cause is economic. Sudden sharp rises in interest rates, and increasing unemployment, are two such influences;

- the cost of claims is heavily geared to house price deflation;

- the delays in claim settlements reflect the degree of forbearance of the lender following the borrower's first falling into arrears, the attitude of the courts to granting repossession orders and, finally, the length of time it takes to sell the property in a depressed market.

New forms of cover continue to evolve, and there is currently considerable discussion about the possibility of substituting some form of aggregate excess of loss contract for the cover described above. Under this, the insurer would be able to cap its liability, both in terms of the aggregate losses retained by the insured before the contract comes into play, and in terms of the upper limit of aggregate losses then payable by the insurer.

Because, under existing arrangements, cover normally extends over the whole life of a loan, the main reserving problem concerns the Unexpired Risk Reserve — i.e. the reserve needed if the cost of claims yet to arise on existing loans seems likely to exceed the Unearned Premium Reserve. The determination of reserves for outstanding or IBNR claims for this class of business is relatively straightforward, compared to that of Unexpired Risk.

Catastrophes

Most catastrophes arise from natural causes (windstorms, earthquakes, floods, etc.). They will usually give rise to a large number of related individual losses occurring within a short period of time defining that event.

Other catastrophes may be due to man-made causes (explosions, air crashes, etc.). These may give rise to a relatively small number of claims, but ones of exceptionally high cost, both in terms of material damage and of liability.

If we confine ourselves to property insurance, the main thing to note is their atypical features:

- reporting patterns are a function of the exact nature of the catastrophe, as the speed with which the insured reports the claim will depend on the seriousness of the loss, whether it was possible or practical to report losses initially, and any emergency procedures set up by the insurer;

- settlement patterns are also likely to be different for each catastrophe, particularly in the first few days, being dependent on the insurer's ability to cope with the problem.

In view of the singular behaviour of each natural catastrophe, it is preferable to try to project each such event separately. This may take the form of a curve-fitting exercise. Alternatively, it may be possible to refer to the daily development pattern of previous catastrophes, subject to the problems in the immediate aftermath as the insurer gears up to tackling the problem. An assessment of the insurer's aggregate exposure may be utilised in this initial period. In the extreme case, it may be necessary to rely on any external market comment or assessment of total insured loss available at that time.

Reinsurance and the London Market

The London Market specialises in writing those risks that are too large for the smaller direct insurers to handle. A significant part of this business relates to overseas risks. In view of the size of the risks involved, much of the business is written on a co-insurance basis, spreading the risk amongst several insurers. The London Market is also a major centre for the acceptance of reinsurance business, again on an international basis.

A given piece of London Market business could therefore range from a very specific property insurance on one risk, to an excess of loss reinsurance contract covering an insurance company's entire world-wide general liability business. In practice, the common features of both Reinsurance and London Market business are:

- the data available to the insurer of this business are limited particularly for retrocession business;

- numbers of claims and individual claims information will often not be available, particularly for proportional business. Hence, any reserving methods requiring such data are not applicable;

- classification of the business is difficult, and the description and nature of the cover may need careful interpretation;

- sub-dividing the data into too small groups, to improve homogeneity, can give rise to excessive volatility as the statistical significance of the development model becomes increasingly suspect;

- the development of claims data is generally medium/long-tailed. Hence, IBNR often forms a significant element of the reserve amounts;

- the length of tail also means that projections are usually based on incurred claims data, rather than on paid claims, since the latter may require development factors to ultimate which are too large and sensitive to be reliable. This is particularly so in the early periods of development, when little or no claim payment may have been made for some classes;

- differing underlying currencies and inflation rates may distort aggregated data;

- There may be a "spiral" effect for catastrophe losses, due to the number of retrocession contracts written by all the companies making up the market. This means that such losses can cycle round the market, particularly for those years where there is an active retrocession market;

- Latent claims are even more difficult to deal with, as different cedants may use different triggers to determine reinsurance coverage. This will have a knock-on effect on the reinsurer's own reinsurance coverage.

Latent Claims

Most of the methods examined in the Claims Reserving Manual are dependent on having historic data which can be modelled statistically, in order to project the total liabilities still outstanding. However, there are circumstances where this is not possible, latent claims being a prime example. These are losses which may lie dormant for several years after the originating event before manifesting themselves. Industrial diseases and environmental pollution are typical sources of such claims.

Even when the latent hazard becomes a reality and claims are being settled and paid, as is now the case with asbestosis, standard projection methods are still not applicable. The normal course of events for latent claims is that, when the nature of the hazard becomes apparent. claims start to be reported at roughly the same time, irrespective of the underwriting year to which they relate. Amounts then increase rapidly.

This is completely different from the normal course of events, where regular development by "development year" might be expected. In the case of latent claims, however, large movements occur across each of the affected diagonals of the data when displayed in the familiar triangle format. Conventional "triangulation" methods of projection will not work in this situation.

Another problem with latent claims is that the historic data may not exist. Even when they do exist, they give little indication of what will happen in the future and, hence, how claims will develop. This is because the claims are dependent on the results of court actions, and these are highly unpredictable. The court actions determine whether coverage applies, if it is deemed to apply then who is liable, when coverage is deemed to apply, and the amount of liability (which may include considerable amounts for punitive damages). Also, expenses are very high covering, inter alia, lawyers' contingency fees, and are often incurred even when coverage is found not to apply.

Although the previous development of claims data may not be of much use in forecasting future development, there are alternative approaches that might be tried. These might include the following:

Share of total market: tracking losses

This involves initially determining what the total loss will be world-wide and the sources of the loss.

For example, for pollution in the USA, this may be done by site, first for known sites and secondly for not-yet-known sites. Assessment of the size of loss by site and the probability that insurers will have to pay ("win-factors") have to be determined, followed by how the losses are aggregated and spread by year and between insurers and reinsurers. This is dependent on court actions to date and will vary state by state.

The individual insurer or reinsurer then has to follow through the alternative insurance, reinsurance and retrocession linkages to obtain its own ultimate forecast figures. This might be done via a decision tree, based on tracking the total market loss through the consecutive layers of insurance, reinsurance and retrocession. The total market loss will have been estimated by splitting the exposure into four markets (US, Lloyd's, London Market and Others), considering retained/insured percentages and share of the insured amounts by market at each level.

Share of total market: relativities

Again, the total market loss is considered. However, for this simpler approach the percentage of the total market is then estimated, using suitable market parameters (for example, by considering the insurer's premium income for the relevant classes of business relative to the total premium income for those same classes/insurers over the whole market).

Exposure

The book of business written is examined to determine, for each contract written, what the possible exposure is to the various latent hazards, and somehow estimating numbers of losses to each contract. This must take into account numbers of reinstatements, aggregation issues, occurrence and aggregate excesses and limits, and so on.

Market Practice

This is based on the assumption that, as the results are so unpredictable, the best that can be done is to be in line with the market. Therefore, if the comparable market is in general using an IBNR of (say) 1.5 times reported outstanding claims for asbestos-related claims, then this would be used as a benchmark. A higher multiplier might be considered as placing an unnecessary strain on the resources of the company; a lower multiplier might be questioned by the company's auditors, unless it could be proved that the company's relative exposure differed from that of the whole market.

Further Research

Much work is currently being undertaken worldwide in researching approaches to reserving for latent claims, particularly in the USA. To date, however, there is no universally accepted approach to these problems. Each case must therefore be considered carefully on its own merits by the reserver, bearing the above points in mind.

◇

[2D]
COMMONLY USED SIMPLE RESERVING METHODS

Introduction

The following section lists all the reserving methods found in Volume 1. It briefly describes the model underlying each method, together with its advantages and disadvantages, and refers to where more detailed discussions and examples can be found.

The main purposes of this section are:

a) while reading through the Manual, to act as a simple summary to help recap and consolidate the knowledge gained to date,

b) having read the Manual, to help identify a shortlist of possible methods to apply to a given reserving situation, and to act as a quick reference point for those wanting to refresh their memory about a particular method.

Consequently some of the terminology used may not be familiar until subsequent sections have been read.

Methods of estimating ultimate liability which do not explicitly allow for inflation

In the following descriptions, unless otherwise stated, "ultimate liability" is taken to mean the ultimate loss from a particular claim year, however defined. The total incurred claims cost can then be found by adding across all claim years and deducting the paid to date.

Method and References	Description of the Model	General Comments
1. Grossing up (general) D6 El–4, 11–12 F3,7	Ultimate liability = paid* at delay *d* / grossing up factor for delay *d* [various ways in which the grossing up factor can be derived]	Simple. Assumes a stable run-off pattern, but is susceptible to error if this is not the case (e.g. if inflation is varying rapidly), particularly for the most recent years.

2. Grossing up (case reserves) paid at delay d F5–6	Ultimate liability = + (case reserves at delay d / grossing up factor for delay d) [various ways in which the grossing up factor can be derived]	As l. Tests the adequacy and consistency of case reserves if used with data grouped by report year.
3. Link ratio (general) D6 E5–12 F4–7	Ultimate liability = paid* at delay d × link ratio at delay d × link ratio at delay $d+1$ × ... × last link ratio [various ways in which the link ratios can be derived)	As 1.
4. Link ratio (basic chain ladder) E8	As 3. with the link ratio at delay d derived from the run-off triangle as the sum of column $d+1$ divided by the sum of column d excluding the last entry.	As 1. Trends and anomalies in the data are not allowed for if operated blindly. Produces a single rigid estimate, without any indication of how to look for possible variations.
5. Loss ratio (general) D6–7 G2, 10–12	Ultimate liability = premium × ultimate loss ratio [various ways in which the ultimate loss ratio can be derived, and in the way that premium and loss ratio can be defined]	Very simple. Requires little information. May be the only method if claims data are scanty or unreliable. Ignores the claim development pattern to date. Prejudges answer if used naively.
6. Loss ratio (step-by-step) G9	Ultimate liability = paid* at delay d + premium × (loss ratio for delay $d+1$ + loss ratio for delay $d+2$ + ... + loss ratio for last delay) [various ways in which the loss ratios can be derived, and in the way that premium and loss ratio can be defined]	Simple. Insensitive to data in the most recent years. Loss ratios are based on the observed claims development pattern to date, rather than chosen arbitrarily.

7. Bornhuetter Ferguson D6 G3–8	Ultimate liability = paid[1] at delay *d* + (premium × ultimate loss ratio × proportion of the ultimate liability which will emerge in future) [various ways in which the ultimate loss ratio can be derived, and in the way that premium and loss ratio can be defined]	Credibility approach between statistical estimate and predetermined figure. Insensitive to data in the most recent years. Prejudges answer to some extent.
8. Average cost per claim D8 H1–5	Ultimate liability = ultimate number of claims × ultimate average cost per claim [various ways in which the number of claims and average cost per claim can be derived]	Makes use of extra information. Easier to detect and allow for changes in the development pattern of numbers. Doesn't allow for changes in the number of zero claims or in the company's definition of settled. Difficult to apply if claims are relatively few in number.

[1]incurred claims (paid plus case reserves) may be used as an alternative to paid claims only, and can give a further perspective on the estimating of ultimate liability. However, whereas methods based on paid data require only a stable settlement pattern, those based on incurred data also require a stable reporting pattern and consistency in the setting of case reserves.

Methods of estimating ultimate liability which explicitly allow for inflation

These methods can be used when inflation is varying rapidly.

Method and References	Description of the Model	General Comments
9. Inflation adjusted methods J2	Most of the above methods can be based on historical payments inflated to current money terms. The resulting projected payments are increased in accordance with expected future inflation.	Need to choose appropriate past inflation rates.
10. Bennett & Taylor Method A J3	Ultimate (report year) liability = paid at delay d + number of claims reported × (inflation adjusted average payments in delay $d + 1$ + inflation adjusted average payments in delay $d + 2$ + ... + inflation adjusted average payments in last delay)	Need to choose appropriate past inflation rates.
11. Separation method J4	Ultimate liability = paid at delay d + number of claims × (proportion of payments in delay $d + 1$ × index of average payments in payment year $a + d + 1$ + proportion of payments in delay $d + 2$ × index of average payments in payment year $a + d + 2$ + ... + proportion of payments in last delay × index of average payments in payment year a + last delay)	Past inflation rates derived from data — no need to choose appropriate rates. Also derives any other calendar year effects within the data. Relatively complicated. Unstable.

Methods of estimating IBNR values directly

Method and References	Description of the Model	General Comments
12. Simple ratios (general) I3–4	IBNR value = previous period IBNR value × current value of a quantity related to IBNR / previous period value of the same quantity [various quantities can be used]	Very simple. Requires little information. May be the only method if claims data are scanty or unreliable. Insensitive to changes in the relationship between IBNR and the chosen quantity. Extremely limited applicability.
13. Simple ratios (Tarbell) I5	As 12. with the quantity defined as the product of the number of claims reported in the last 3 months and the average size of claims reported in the last 3 months.	Very simple. Too crude for medium or long tailed business.
14. Average cost per IBNR claim I6	IBNR value = number of IBNR claims × average cost per IBNR claim	Makes use of additional information. Difficult to apply if claims are relatively few in number.
15. Loss ratio (step-by-step) I7–8	IBNR value = IBNR emerged at delay *d* + premium × (IBNR emergence/premium at delay *d*+1 + IBNR emergence/premium at delay *d*+2 + ... + IBNR emergence/premium at last delay) [various ways in which IBNR emergence can be defined]	More suited to long tailed business.

Other methods

Volume 1 also includes methods

for tail fitting	E13
for estimating reserves for reopened claims	K2–3
for estimating reserves for claims expenses	K4–5
for estimating reserves allowing for discounting	L2–4

◇