Section B
DATA & FORECASTING

Preamble

This section introduces some of the main building blocks for claims reserving. To begin with, there is the important idea of making a projection of past experience into the future. Since the future never takes the trouble to conform properly with the past, any projection whatsoever will be subject to error. One needs, therefore, to understand the principles which can lessen the likely degree of error, and so bring credibility to the work.

Apart from those principles which make for stability, there is the matter of the data themselves and the actual methods of forecasting. These are not intrinsically difficult matters, but there is a fair amount of detail to be mastered. On the data side, a number of different quantities can be used in the projections, or as supporting evidence — not only claim amounts, but such items also as claim numbers, premium income and loss ratios. They can often be displayed in different ways in the search for pattern and regularity, and the concept of the development table is particularly important here. Then there is the question of data validation, and of how the classification of the risk groupings is to be made.

On the forecasting side, there are some surprisingly simple methods available. It is straightforward, almost intuitively obvious, to look for the average or trend which is present in a sequence of figures. The really vital question to ask is whether the available evidence supports the continuation of such average or trend into future periods. Although far more elaborate types of projection can be devised, it is these simple foundations on which they rest, and which should therefore first be thoroughly understood.

Contents

[B1]
THE PROJECTION OF PAST EXPERIENCE

In claims reserving, the aim is to estimate the future claims experience which is to be expected on the business written to date by the insurer. As a first approach, the values set as case reserves on open claims by the claims handling staff may be used. However, some variation in these before final settlement of a claim is likely to take place, and by definition such values cannot cover the IBNR component of the required reserve. Thus it is usually necessary to go beyond the case reserves. The approach which then emerges, quite naturally, is to look at the insurer's past history of claims experience and to project this forward to the future years.

Taking this approach, the first need will be for suitable historical data. Ideally they will consist of such items as the number of claims reported and the number settled, and the amounts paid out by way of settlement. There will be information on the premiums written or earned, and perhaps other measures of risk exposure such as the number of units covered (e.g. households or motor cars). The data will be classified according to the class or sub-class of business involved, and also by the year of origin (i.e. accident or underwriting year). For each class and year, and for each data element, there should then be a series of figures showing the development with time up to the current date. In addition, figures showing the development of the case reserves themselves may be available, and can also be used as a basis for projection.

The second need will be for a method of projection, and very many of these have by now been devised. They range from the use of simple arithmetic on the familiar triangular arrays of data to the employment of highly sophisticated mathematical and statistical techniques. From the number of different methods available, the problem is to select that which is most appropriate in the given circumstances of each particular case. The Manual's main purpose is to describe the methods, together with their advantages and disadvantages.

Given that the data are available (and there will often be gaps and deficiencies), and that the skill needed for the projection is to hand, a leading question has now to be faced. That is, to what extent is it actually justifiable to project forward the experience of the past on to the future years of development? There can perhaps be no final philosophical answer to such a question. However, the theoretical understanding of statistics and probability, borne out in practical experience shows at least that it is *reasonable* to make such projections.

The projections, however, cannot be done arbitrarily. A systematic approach needs to be adopted. To begin with, for example, the reserver should scrutinise the data with which he or she is presented, or which he or she intends to collect. Apart from the obvious point of its validity and consistency, the stratification of the data into the main business classes and the risk subgroups will be of great importance too. For each subgroup, the larger it is and the more homogeneous the risks it contains, the greater the degree of statistical stability will be. Generally speaking, however, the desiderata of size and homogeneity tend to work against one another, as shown more fully in §B2. The reserver must find, and be prepared to justify, a suitable compromise in the risk classification to be adopted.

A further point of major importance will be to examine the *influences* which have shaped the claims pattern in the past, and how these may currently be changing. Such factors as the volume of business or the rate at which claims are handled, the level of inflation or the legislative climate, can all affect the position. If such influences are properly understood, then significant shifts in the experience may be detected promptly in advance, and taken into account by adjusting the projections. (A fuller discussion of the main influences on the claims pattern follows in §C.)

Above all, in claims reserving it is not sufficient just to take the data and blindly apply the first projection method which comes to hand. At each step, intelligence has to be applied. There are key questions which need to be answered afresh each time a new projection is to be made. A checklist now follows:

a) What historical data are available to the reserver, and how far can confidence be placed in its reliability?

b) To what extent is the homogeneity of the groups in the risk classification satisfactory?

c) What conditions have shaped the past experience, and what significant changes in them can be detected which may affect the future out-turn?

d) What methods of projection are proposed, and are these properly suited to the given circumstances?

To ignore these points is to ignore the whole essence of the work.

<>

[B2]
DATA GROUPINGS: PRINCIPLE OF HOMOGENEITY

The underlying principle of insurance is statistical in nature. In the business of taking over the risks of others, the insurer is best protected by taking on a sufficient number of similar, but independent, risks. The proportion actually becoming claims, and the amounts payable, can then be predicted within manageable margins. Hence an adequate premium can be set with some confidence in advance of the risk period itself. This is a result of what is popularly known as "the law of large numbers", but which appears in statistical theory as the necessary relationship between the variance of a sample and its size.

As with insurance at large, so it is for claims reserving in particular. Stability in projections is to be sought by aiming to work with data groupings each containing a sufficient number of similar but independent risks on the assumption that they determine the characteristics of the resulting claims. The question is, how far should the classification of business be taken in order to produce such individual risk groups? To begin with, there are the main types of business, such as Motor, Property, Liability and so on. These are reflected in the supervisory authority classification, and must therefore be observed for the purpose of statutory returns. Such an initial classification will be desirable also from the point of view of reserving. But the heterogeneity of many of these main classes is such as to make further subdivision essential.

To take the example of Motor, it will certainly be necessary to separate out Private from Commercial business. Private Motor can then be further classified into:

Motor Car Comprehensive
Motor Car Non-comprehensive
Motor Cycle

Again, the division seems necessary, given the different risk combinations covered by comprehensive and non-comprehensive business, and the different characteristics of motor cycle riders as a class from those of car owners.

The subdivision could again go further within each of these three categories. Thus we might use distinguishing features, say, of geographical area, make of car or cycle, age of driver, and so on. The further we take the classification in this way, the greater the homogeneity in each of the resulting risk groups. But the stage can soon be reached where the individual groups lose their statistical credibility. That is, their size (or lack of it) is such as to produce an unacceptably high variance, at least so far as claims reserving is concerned.

In practice, the ideal of homogeneity is not to be pursued with too much rigour. Indeed, often it cannot be so pursued. The data themselves may not permit very much subdivision — e.g. the required fields may not have been inserted into the data-base records in the first place. Again, the time available for the work of reserving will not be unlimited, and to multiply the sub-groups multiplies the work to be done. A sense of proportion has at all times to be kept.

To return to the Private Motor class of business, most insurers would be unlikely to go much further than the 3-way split suggested above. Indeed, some might treat Comprehensive and Non-comprehensive together as a single risk class. Strictly speaking, this is not a desirable combination — the two classes have a quite distinct risk profile and claims run-off. Thus, bodily injury and other third party claims will be the main element on the non-comprehensive side, while the comprehensive will have a more even split between third party and physical damage. Hence, although the overall length of the tail may be the same in both cases, the comprehensive business will show the stronger early development. This will be especially true in the first two years or so, during which time virtually all of the physical damage is likely to have been settled, but comparatively few of the major liability claims.

The justification for taking the two groups together can only be that the proportion of comprehensive to non-comprehensive is reasonably stable, and thought likely to continue so in the future. The patterns may be upset if there is a sharp change in business volume during the course of a particular year of origin.

Generally speaking, with personal lines business, lack of homogeneity would not be expected as a problem. The number of policies will often be large, with the individual amounts at stake relatively small, thus providing good conditions for statistical treatment. But in the commercial lines, where each risk taken on will have its own special characteristics, and where there may be relatively few policies issued, the homogeneity of a class will frequently be in doubt. The problem may be exacerbated by the existence of unique risks for very large sums insured.

The solution adopted may vary with the class of business involved. Thus, in commercial property, the answer may be to rely more on the case estimates than any statistical projection, or at least to use the case estimates as the main source of data for adjustment. In commercial liability, however, such is the length of the tail that case estimates may not help a great deal, except for the older years where some development has already taken place. The reserver will therefore be thrown back on statistical projection, but without the comfort of a firm underpinning. The need for a full and intelligent assessment of the conditions and influences surrounding the business will be all the more important, and no projection should be regarded as sound without such an assessment to back it up.

◇

[B3]
THE CLAIMS DEVELOPMENT TABLE

This section looks at the main claims data that will be needed for the work of reserving, and the format they are likely to take. Suppose that the risk classification is already established, with proper regard to the constraints of size and homogeneity in the subgroups. We then wish to examine the data for a particular group, say as at the 31 December for the current accounting year. What form will the information from the insurer's data-base take?

To begin with, there will be the claim amounts paid out during the course of the accounting year just past. Let us say the total is given as £5,769,000, rounded to the nearest £1,000. In itself, the figure is not very informative, although it can for example be set against comparable amounts for previous years. Even this is scarcely sufficient for the purposes of projection — there is no information on such vital matters as the length of the business run-off, the relative age of the claims being settled, or the true relationship to premium income. What is needed is an analysis of the claims figure by period of origin of the business.

Analysis by Origin Period

The origin period itself needs some attention. It is most commonly taken as a year, but can also be a quarter or even a month for rapidly changing lines. Again, it is common to take *accident year* as the origin for the business. Accident year is the term used to refer to the calendar year in which the occurrence giving rise to the claim took place. It is perhaps the most natural origin, but is not invariably used. Thus, in reinsurance work the origin is more often the *underwriting year,* i.e. the calendar year in which the policy covering the risk was written or renewed. Such a definition enters for the simple reason that it is the normal accounting basis in reinsurance. Finally, the *report year,* the year in which the claim is first registered in the insurer's books, can also be used in some reserving analyses.

Taking accident year as the origin, suppose the overall claims figure breaks down in the following way:

Yr 1	23,000	Yr 5	1,007,000
Yr 2	148,000	Yr 6	1,536,000
Yr 3	422,000	Yr 7	1,889,000
Yr 4	744,000		
			£5,769,000

Year 1 is the earliest accident year for which claims are still being paid out. Then the other years follow in succession until Year 7, which is the year just past. The breakdown is informative, for example in indicating the probable length of the run-off. But more still can be learned by building up the picture with similar claims information from earlier years of account. For example, suppose information is available from the 5 previous years. Then a whole table can be drawn up, which might appear as follows (figures in £000s):

		Year of Payment					
		2	3	4	5	6	7
	<0 †	650	340	110	19	-	-
	0	478	395	272	110	20	-
	1	744	501	442	288	127	23
Yr of	2	1001	854	568	565	347	148
Origin	3		1113	990	671	648	422
	4			1265	1168	800	744
	5				1490	1383	1007
	6					1725	1536
	7						1889
	Total	2873	3203	3647	4311	5050	5769

† Individual data are not known for accident years earlier than Year 0. <0 implies that aggregated data are being given for these years.

The Development of Claims

Patterns for analysis are now beginning to emerge: eg, the volume of claims in the table is increasing steadily as the years progress. But the clearest picture will emerge if we directly compare the development pattern of claims for each successive year of origin. This can be done by examining the rows of the table above. The comparison is made much easier by shifting each row successively one place further to the left. The elements of the lower diagonal, for example, then form the first column of a new table, and so on for the other values. The top axis becomes, instead of payment year, the year of development for the business. The new table is as follows:

		Year of Development						
		0	1	2	3	4	5	6
	<0				650	340	110	19
	0			478	395	272	110	20
	1		744	501	442	288	127	23
Yr of	2	1001	854	568	565	347	148	
Origin	3	1113	990	671	648	422		
	4	1265	1168	800	744			
	5	1490	1383	1007				
	6	1725	1536					
	7	1889						

To read the table, take for example the origin year 2. By the end of the year itself, the claims paid out on the business originating in that year are £1,001,000. Then in the following year, a further £854,000 is paid, and so on, until in the most recent year claims are £148,000. The development years are labelled by the progression 0, 1, 2, 6 along the top of the table. The convention adopted is that development year 0 is just the origin year itself in each case. Then succeeding development years follow in natural sequence. Thus, for origin year 2, there have been five development years following it. The most recent year (the current accounting year) is therefore the development year 5. But for origin year 6, the most recent year is only development year 1, and so on. A useful relationship to note is that:

Year of Origin + Year of Development = Year of Payment

The relationship is quite general, and can be checked by applying it to the cells in the above table.

(The above convention, which is used throughout the Manual, is in common use. However it should be noted that in Lloyd's and the London Market, it is customary to label the development years 1, 2, 3 etc., i.e. starting with "1" instead of "0".)

Rows, Columns & Diagonals

Once the data have been put into this format, it is very suggestive of means for analysing and projecting the claims. Thus, ratios of values along the rows give the development pattern for each individual accident year, and regularities may soon become apparent. Down the columns, the ratios give the trend pattern from one accident year to the next, which again may be revealing. Lastly, the diagonals can be seen to relate to the position in succeeding calendar years, with the lowest diagonal representing the calendar year immediately past. (The sum of the values in this diagonal will take us back to the originally quoted claims figure of £5,769,000.)

The data array of claims of development year against origin year is thus a fruitful one. But there is a disorienting feature that the reserver may come across in practice. It is that the axes of the table can be arranged in different ways. Thus, the two main axes can be interchanged, or the order of the origin years can be reversed, so that the later years come at the top rather than the bottom. Also, some offices prefer to use payment years rather than development years as one of the axes (i.e. they revert to the earlier table above). All possible variations seem to be in use! However, the arrangement shown above is the one most commonly found in the literature, and it will be kept to throughout the Manual.

The Cumulative Claims Table

There is a further variation of the table which is often useful. Rather than looking at the year by year addition to the claims for each year of origin, we may be interested in the *cumulative* development. The cumulative figures are obtained simply by adding the values along each row. In the present table, this cannot be done for years earlier than Year 2, owing to the missing data. But for the years from 2 onward, the process yields the following array:

		Year of Development					
		0	1	2	3	4	5
	2	1001	1855	2423	2988	3335	3483
	3	1113	2103	2774	3422	3844	
	4	1265	2433	3233	3977		
Yr of	5	1490	2873	3880			
Origin	6	1725	3261				
	7	1889					

It will be observed that the data are now in the exact shape of a triangle. Such triangular form is widely used in claims reserving work. Though the form is appealing, it has its deficiencies. For example, the relationship with payment year data is not fully apparent. The diagonals other than the leading one are incomplete, and to improve the connection we need to return towards the parallelogram shape of the previous display.

Apart from the payment year relationship, the parallelogram of data has advantages from the projection point of view as well. Thus the given example contains a fair amount of information for development years 4, 5 and on, whereas in the strict triangle it is scanty indeed. Of course, it may be that data cannot be obtained at all for the earlier years of origin, in which case the triangle will have to suffice. But if the data can be found, the extension to parallelogram form may be well worth the effort.

In the claims reserving literature, a strong convention has arisen involving the use of triangular data arrays. In general, the Manual will follow the convention. But the reader should be aware that it is not an absolute requirement, and can often be dispensed with to good advantage.

<>

[B4]
DATA QUANTITIES

In claims reserving, a number of data items are commonly used in addition to the basic information on claim amounts paid out. These include case reserves, premium income, loss ratio, claim numbers and risk exposures as the main quantities. The information will mainly come from the insurer's data-base, but industry statistics may also be brought in. Availability of data is likely to differ between reinsurance and the direct market. In general reinsurance data will be less full and less up-to-date. In particular, claim numbers as opposed to claim amounts will often neither be known nor obtainable.

The present section describes the main data items, and gives some figures for illustrative purposes. Frequently the data can be set out in the tabular form of the previous section, showing development against year of origin, and this form is used where possible. However, it is well to note that some initial work has to be done to produce such tables or triangles — the data in their raw form are not always so conveniently presented.

Claims Paid

Amounts paid out on claims are by definition the central quantity for reserving purposes. If set out by period of origin, a development table or triangle results, as shown in §B3. The example data are repeated here for convenience, in their year by year form (but omitting years of origin earlier than Year 1):

		Year of Development						
		0	1	2	3	4	5	6
	1	†	744	501	442	288	127	23
	2	1001	854	568	565	347	148	
Yr of	3	1113	990	671	648	422		
Origin	4	1265	1168	800	744			
	5	1490	1383	1007				
	6	1725	1536					
	7	1889						

† Data not available

It is important to be clear as to the definition of the claims payment information. E.g. does it include expense directly attributable to the claims, such as litigation costs and loss adjusters' fees? Does it contain the partial payments on claims not yet fully settled? Are the figures gross or net of reinsurance, salvage and subrogation? Do they need adjustment perhaps because of some reporting or data processing delay? The bare figures given as example in the text do not fully convey the real life complications which the reserver must be ready to handle.

Case Estimates

Such estimates, usually made by personnel from the claims department, are a natural adjunct to the values for the paid claims. At the end of any accounting period, there is bound to be a number of claims still outstanding and the estimates will give a first approximation to their cost. For example, at the end of Year 7 (i.e. the current accounting year), the breakdown of the case estimates by year of origin might be:

≤ Yr 1	31,000	Yr 5	1,796,000
Yr 2	234,000	Yr 6	2,881,000
Yr 3	475,000	Yr 7	3,929,000
Yr 4	969,000		£10,315,000

In evaluating this information, the reserver should again be asking the relevant questions. E.g. is the likelihood of future inflation of claims cost taken into account in the estimates, and if so to what extent? Are the estimates intended to include a degree of conservatism? Do they have an allowance for direct claims expense? And so on. To gain a proper understanding, the reserver should seek contact with claims personnel, and ferret out the definitions of and underlying assumptions in the figures.

If data are available from past accounting years, it will be possible to build up a development table for the case estimates just as it was for the paid claims. Such a table might appear as follows (Year 1 data not available):

		Year of Development					
		0	1	2	3	4	5
	2	1776	1409	1029	606	384	234
	3	2139	1701	1199	809	475	
Yr of	4	2460	1971	1546	969		
Origin	5	3031	2549	1796			
	6	3644	2881				
	7	3929					

The table is similar in form to the paid claims table, but there is a difference in its status. It summarises sets of estimates made at points in time, i.e. the end of each accounting year. The paid claims data, on the other hand, are an accumulation of amounts through the years in question. The case estimates, of course, cannot be accrued in this way. But the values can be combined with those for the cumulative paid claims, to produce a quantity usually known as *incurred claims.* The latter is effectively an estimate of the ultimate loss to be experienced on a given year of origin from known claims at the accounting date, and makes no allowance for the IBNR.

Premium Income

The next data item of importance is the premium income. This provides the essential measure of the volume of business against which the claims are being paid out. It can also be seen as a first measure of the insurer's exposure to risk in the business class under consideration. It is not, however, a pure measure in that it does include a weighting for office expense in addition to the risk premium content. Also, the state of the insurance market will affect the relationship of premium rates to the quantum of pure risk — i.e. in a soft market, competitive pressures may force premiums down to the point where they are scarcely adequate to cover the risks underwritten. But in a hard market the opposite will be true, and the risk element will be well covered.

Premium income will, quite naturally, relate to the year of origin of the business. It will be extracted from the insurer's data-base as either the *earned* premium or the *written* premium for the years in question. The distinction is an important one, and is worth spelling out in detail. Thus, earned premium relates to all policy exposures on which the insurer is liable during a given period (normally a calendar year). E.g. for a policy renewed on 1 April, the earned premium for the current year will be 25% of the previous year's premium plus 75% of the current year's premium, and so on. Written premium, on the other hand, covers all premium income generated in the period in question, whether for new policies or renewals. In the case of the policy renewed on 1 April, the written premium will be 100% of the current year's premium.

The distinction between earned and written premium connects with the choice of either accident year or underwriting year for the claim development analysis. It is essential that the correct combination be used. (Actuarially speaking, it is a matter of correctly defining the exposed-to-risk.) Thus, where the origin for the claims development is accident year, the earned premium definition should preferably be used. On the other hand, for the origin as underwriting year, it is right to use the written premium.

The distinction often accords with the split between direct business and reinsurance. In direct insurance, the combination of earned premium with the accident year is most common. But in reinsurance there is usually little choice in the matter, and the data are often in such a form that only written premium with the underwriting year can be used.

Following our earlier illustration we give some example figures for premium, set out by the year of origin from Year 1 to Year 7:

Yr 1	4,031,000	Yr 5	6,590,000
Yr 2	4,486,000	Yr 6	7,482,000
Yr 3	5,024,000	Yr 7	8,502,000
Yr 4	5,680,000		

For direct business, the amount of the premium, whether earned or written, will usually be known for the end of the accounting year in question apart from an element of "pipeline" premiums relating to late notification of increments and cancellations. Unless these are significant it will not make sense, therefore, to draw up any kind of development table. But in reinsurance, it may take two or three years or even more before the premium is fully reported. Hence a development table can be drawn up, just as for the claim payments, and can be used in a similar way to project the final amount of the premium for any given year. (A further point is that a 3-year accounting system is the norm for reinsurance on the London Market. The system is well described in the London Market references in §O, and will not be further discussed in the Manual.)

Loss ratio

Loss ratio can be defined as the ratio of the ultimate amount of claim to the total premium for a given class of business. Thus it is not a primary data item — indeed it is what the reserver is effectively trying to forecast for the business on the years still open at the accounting date. However, in the past, underwriters may have established norms for the expected loss ratio on given classes of business. They may further be able to estimate how far such norms are likely to be stretched by the conditions more recently prevailing in the market. Such information provides the reserver with an initial set of guidelines against which to test the outcome of his or her projections. It also enables the reserver to extend the range of methods, e.g. in Bornhuetter-Ferguson and related techniques. Finally, where a sequence of values is available the loss ratio itself can be a subject for projection.

When defined as the ratio of ultimate claim to premium, the loss ratio is more precisely said to be the ultimate loss ratio. But there are other forms. Specifically, one may speak of the *paid loss ratio* and the *incurred loss ratio.* Such terms are used to denote the ratio of claim to premium as the business for a given year of origin develops. The paid loss ratio is just the amount paid to date on claims divided by the premiums. It rises from a low value in the early part of a development to reach the ultimate value once all claims for the year in question are settled. The incurred loss ratio is a similar quantity, but in which claim amounts paid to date are supplemented by the current value of claims outstanding.

Claim Numbers

A useful item, giving considerable further knowledge of the development of a year's business, is that of claim numbers. The numbers per se are of value, in giving a measure of the claim frequency. Also, when combined with the data on claim amounts, they enable the average cost per claim to be found. The reserver thus gains a fuller picture of the behaviour of the claims, and a first glimpse of the claim size distribution itself. Unfortunately, however, data on claim numbers are very often not available in the reinsurance field.

During the history of any given claim, there are certain distinct events which can be recognised:

a) Occurrence of the event giving rise to the claim.

b) Reporting of the claim to the insurer, and its recording in the insurer's data-base.

c) Settlement of the claim, either partially or in full.

Claims may thus be counted: i) as they are reported to the insurer and become established as open claims on the books, and ii) as they are finally settled and no longer represent any future liability to the insurer. For a given accounting year, the number of claims reported and the number settled can be defined, together with the number open at the beginning and end of the year. The simple relationship of these quantities is:

No. of claims open at 1 January
plus No. reported in year
less No. settled in year *equals* No. open at 31 December

The relationship can be used either as a check on the data, or to determine one of the quantities, if missing, from the values of the other three. In practice, doubt would most often attach to the number settled, which could be found or verified as:

No. of claims settled in year *equals* No. reported in year
plus No. open at 1 January
less No. open at 31 December

Having obtained the numbers for the current accounting year, the next step will be to divide these according to year of origin, whether this be accident or underwriting year. Then the numbers for preceding accounting years can be set alongside, and development tables produced as described for claim amounts in §B3. These tables will again show the development of business for each successive year of origin, and will be in the familiar triangular (or parallelogram) form. To begin with, three separate development tables can be produced:

a) No. of claims reported in each year, year by year basis.
b) No. of claims settled in each year, year by year basis.
c) No. of claims remaining open at end of each year.

For claims reported and claims settled, the figures can be added along the rows to give the cumulative position for each origin year. Hence two further tables result:

d) No. of claims reported, cumulative basis,
e) No. of claims settled, cumulative basis.

An example would be, for claims reported, year by year basis:

		Year of Development					
		0	1	2	3	4	5
	2	401	84	38	15	7	2
	3	513	90	41	28	12	
Yr of	4	665	88	50	34		
Origin	5	690	93	57			
	6	725	116				
	7	789					

Claims reported, cumulative basis:

		Year of Development					
		0	1	2	3	4	5
	2	401	485	523	538	545	547
	3	513	603	644	672	684	
Yr of	4	665	753	803	837		
Origin	5	690	783	840			
	6	725	841				
	7	789					

Extensive data can thus be developed where claim numbers are available. The interesting question arises as to how claim numbers are to be related to claim amounts when average costs per claim are being calculated. Some natural relationships exist with the quantities of paid and incurred claims, and with case reserves, but their handling requires a little care. The matter is dealt with in main section §H on average cost per claim methods.

With claim number data, as usual, there are some caveats. The figures can be complicated, for example, by claims which had been regarded as fully settled becoming reopened. This may occur, perhaps, because fresh symptoms develop in an injured claimant, or a new statute contains some retrospective effects. Again, some claims may prove to be null and void, and hence be closed with no payment by the insurer. How such circumstances are treated will affect the data and their proper interpretation. It is important, therefore, to know exactly what the claim numbers contain and what they exclude. Additional data to clarify such points as the reopened claims and those settled at nil may well be needed.

Measures of Exposure

As mentioned above, premium income can be regarded as a first measure of the risk exposure. But other measures can be used, and may become appropriate according to circumstance. The principal ones are as follows:

Total of sums insured at risk
Total of EMLs at risk (EML is Estimated Maximum Loss)
Number of policy or insured units earned/written

Different methods might be used to extract the required values from the insurer's data-base. A rough and ready technique for sums insured or EML will be to take the average of the values for the in-force policies at 1 January and 31 December of the year in question. Actuaries and statisticians will recognise this as an application of the census year method — as such, it can of course equally be applied to the policy or insured units. For the latter, however, a more rigorous method will be to extract the earned or written exposure on a policy by policy basis. Such an extraction requires more time and effort, but can be possible where policy files are held on an efficient modern computing system.

Of the above measures, sums insured or EMLs will be more appropriate for the commercial classes of business, particularly commercial fire. But with personal lines, where a large number of similar policies are written, the number of units will generally be better. Eg. in householders' insurance, it would be the number of houses insured for the year, or in motor the number of vehicles, and so on. As with premium income, the earned/written distinction will be important. If a policy terminates on 30 September, say, without renewal, or if a new policy commences on 1 April then in each case the earned exposure will be only .75 of a unit. But for the former policy the written exposure will be nil, while for the latter it will be unity.

Industry Data

Where a new line of business is being marketed, or where a new insurance company is being set up, there is no record of company experience on which to build. The best approach may therefore be to examine sources of data for the insurance industry as a whole.

One source which has grown in value in recent years is that of the insurance company returns to the supervisory authority. Since the 1970s, it has been obligatory for insurers to provide data on the main classes of business, and the main risk groups within class, in development table form. The problem, however, is that companies may use somewhat different definitions for the risk groupings. Also the business classes themselves are too broad to help with the analysis of particular types of policy or types of risk. Finally, it has to be remembered that the purpose of the returns is chiefly to demonstrate the solvency of the insurers. Hence they will not necessarily provide the most satisfactory data for management control of a new business line.

A further source of industry statistics is the Association of British Insurers (ABI). For example, in fire insurance, the ABI runs a market statistics scheme, whose main aim is to produce burning cost values by trade classification. Although the main use of such statistics is in underwriting and ratemaking rather than in reserving, they might assist an insurer able to extract data on its sums insured in forecasting the claims experience. These statistics are available only to those member companies of the ABI who have contributed the relevant data.

<>

[B5]
SIMPLE BREAKDOWNS OF THE CLAIMS PATTERN

The main types of data which are of use in claims reserving have been set out in §B4. It is worth looking at how these data elements connect together logically. A good way of doing so is to take three simple breakdowns which can be applied to the claims figures. The breakdowns are in any case useful to have in mind, both conceptually and practically speaking.

First Breakdown

We use the term *overall loss* to denote the full amount paid out on the group of claims in question, including if need be a component for expense. The first breakdown we wish to apply uses the information on numbers of claims:

$$\text{Overall Loss} = \text{Number of Claims} \times \text{Average Cost per Claim}$$

This formula can be taken to refer to the ultimate position reached on a given class and/or given year of business, or to the development at any point along the way. Normally, it will be used to determine the average cost per claim figures from the available data on losses and numbers of claims. A study of the movements in average cost per claim both by year of origin and year of development can give the reserver a fuller picture of the business being analysed.

Second Breakdown

The second breakdown brings in the exposure information:

$$\text{Overall Loss} = \begin{array}{c}\text{Measure of}\\ \text{Exposure}\end{array} \times \begin{array}{c}\text{Frequency}\\ \text{of Claim}\end{array} \times \begin{array}{c}\text{Average Cost}\\ \text{per Claim}\end{array}$$

The idea here is to go deeper into the claim number information, and replace it with a frequency measure on the earned or written exposure. The frequency will be expressed, e.g. as the number of claims per 100 exposure units. Both the ultimate position and a partial development can be referred to in this way. The information may reveal new characteristics of the business in question, particularly if it is set out in full development table form.

Third Breakdown

There will be those situations, particularly in reinsurance, where only the claims payment data are available and none on the number or frequency of claims. In such cases, a third breakdown comes into its own. This breakdown, which takes premiums as the starting point, is again a very simple one:

Overall Loss = Premium Earned (or Written) × Loss Ratio

For direct business, the premium should have a known value soon after the end of the year of origin. The loss ratio (i.e. in paid loss ratio form) will then develop proportionately as the loss itself progresses towards the ultimate value. It is usually instructive to watch the loss ratio in development, since the comparison with other years of origin can be directly made. For reinsurance, the position is more complicated, since both terms on the right hand side are likely to show a development with time. But if the premium development has any regularity to it, the loss ratio will again be interesting to watch.

These three possible breakdowns of the loss should become familiar to the reserver as part of his or her conceptual basis, particularly when the reliability of a given method of projection is under the microscope. The point about the breakdowns is that the available data may sometimes point up trends or shifts in one or other of the components of the overall loss. In addition, analysis may show that a given reserving factor affects one component in particular, so that its final influence becomes easier to assess. The more that can thus be discovered about the anatomy of a given class of business, the better will be the chance of producing dependable reserving figures.

◇

[B6]
DATA SYSTEMS & VALIDATION

The main source of data for reserving purposes is likely to be the insurer's central computer system. The system will contain the main policy and claim files, together with the company's income and expenditure information. It is likely to be set up as an intricate data-base, with the record files indexed on certain key fields and inter-related with each other in a logical structure. The data-base will be completed by a suite of programs enabling data to be entered, modified and extracted in various forms. A high level query language may also be available, enabling on-line requests for information to be answered with some speed and efficiency.

It is important for the reserver to have a good knowledge of the company's data-base, and to be aware of its limitations. For example, there will be limits to the distinctions which can be made between different types of business — and this will affect the decision on the risk groups which are to be used in the reserving analysis. Other limitations will apply to the type of information which can be extracted. While data on claim payments, and probably also claim numbers, will be readily available, it may be impossible to establish the exact shape of the claim size distribution. If so, it will prevent the reserver from using some of the more complex methods, Reid's method (see Volume 2) being a case in point.

Sometimes certain data desirable for reserving may be made available, but only at a cost. It might, for example, prove necessary to read sequentially all the policy and claim records in the system, perhaps running into many millions of accesses. In such a case, either the expense or the time needed may be prohibitive. Sometimes sampling may be feasible and less costly.

It goes without saying that the reserver should know the exact definition of the data figures produced from the computer system. Thus, are the premiums recorded gross or net of commission? Do the claim number data include those claims which are settled without payment? Do the figures for paid claims include settlement costs, such as loss adjusters' fees and legal expenses? Each item will have its possible variations, and its true particulars must be known.

For reserving, the data situation may well be far from the ideal. Insurance data bases are usually designed in the first place to satisfy accounting and policy renewal purposes — the features desirable for statistical work may come a poor second. The position has been improved by the requirements of the returns to the supervisory authority, in which claims data have to be shown by year of origin, and therefore in a form suitable for reserving. But in general it may be most convenient to initiate the reserving analysis from data which are being produced anyway for the year-end accounts.

Data Reliability

It is essential that a thorough set of checks should be made on the reliability of data used for claims reserving. There is a number of aspects to this, which are dealt with in turn below.

Data Input. The computer software should be such as to incorporate a range of checks on all data that are input to the system. Examples are check digits in policy numbers, to help ensure that the correct record is being updated, and validation tests on the dates, currency codes, monetary amounts, etc. being entered.

Data Processing. Given some familiarity with the system, the reserver will be able to check on procedures used to extract the reserving data and arrange them in amenable form. He or she should ensure that all relevant records and business groupings have been included in the data, and check for deficiencies caused, say, by the late processing of reinsurance accounts.

Reconciliation of Data. Wherever possible, data should be reconciled with revenue accounts and details of policy and claim movements. An example would be to take for each year of origin the cumulative claims paid to the end of the current year, deduct the respective amount for the previous year, and check the result against the claims paid figure in the current year's accounts.

Other Checks. Further evidence can be gained in a number of ways. It can be useful to examine a sample of the claims files themselves, to throw more light on the anatomy of the business and the completeness of the data. Again, discussions with both claims staff and data processing staff may help to expand the picture and give advance warning of any new difficulties in the pipeline.

To sum up, the reserver's aim should be to examine critically each stage in the data production cycle. Nothing should be taken for granted, and efforts should be made to prove the degree of reliability of the data, to understand their content and test their reasonableness.

<>

[B7]
FORECASTING: SIMPLE AVERAGES & TRENDS

Methods of projection in claims reserving form the main subject of the Manual. But before discussing the particular methods, it is useful to do some basic groundwork. In projections of past experience into the future, the essential problem can be expressed as that of extending a time series. Thus, suppose a chronological sequence of claim amounts, claim numbers or development ratios is given. Taking the last, the data might read:

1.057 1.053 1.059 1.062 1.059 1.066 1.064

We now want to extend this sequence, say over the next 3 periods. How should this be done? There can be no foregone conclusion as to what is right, but two simple methods immediately suggest themselves:

i) To take an average
ii) To further a trend

Taking the first of these, the simple average of the 7 figures in the example works out at 1.060. Hence the extension would be:

1.060 1.060 1.060

Taking the second case, the trend, plotting the figures does suggest there could be a slight upward movement. A line can be fitted graphically:

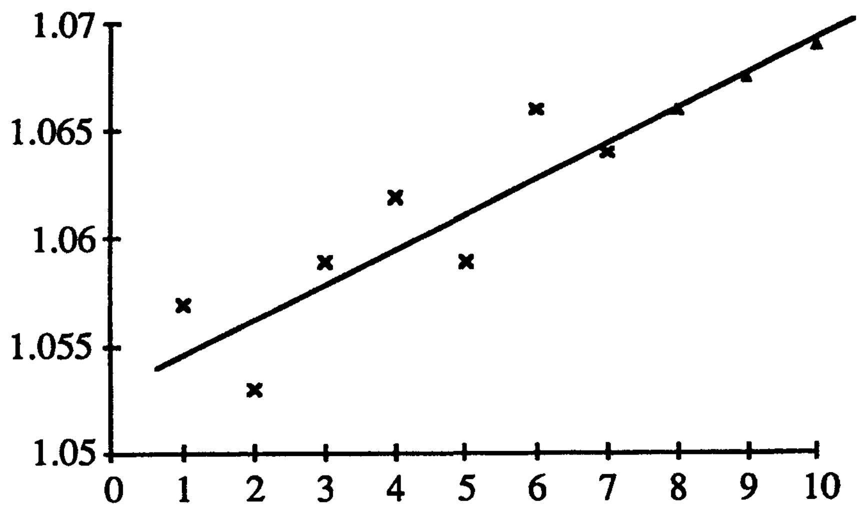

This gives as the extension:

1.066   1.0675   1.069

which is an appreciably different result.

How should a choice be made? The assumption in both cases would be that there is a strong underlying pattern to the data, which is being disturbed by random variations about a mean. The difference is that in the first case, the mean is taken to be static, while in the second it is slowly increasing.

Given these distinct assumptions, the answer as to which one to use can only be found in the light of other knowledge. But if there is a proper appreciation of the business situation, then evidence for or against the real existence of a trend may be readily apparent. Eg. evidence from the claims department may support the hypothesis that settlement patterns are slowing down from year to year. Hence a trend in the loss development ratios is certainly to be expected, as against a static value.

Variations on Averaging

The choice between taking an average or a trend is perhaps the major one to be made. It is not the whole story, however. To take the example above, a simple average over the 7 years' figures was used. But many other variations would be possible within the theme of averaging. Some of the main ones are given below.

a) Curtailed Averaging Period

The use of *all* past data in compiling the average may be inapt. If the business and the influences on it are rapidly changing, then figures from as long ago as 7 years may be quite irrelevant. Hence a shorter averaging period should be chosen, say 3 years. The example figures would then yield the projection:

1.063   1.063   1.063

b) Exceptional Values Excluded

It might be more reasonable to exclude any aberrant figures from the average, particularly if these can be explained by known, exceptional influences. Again, the highest and lowest figures in any given sequence could be excluded as a matter of course, with the aim of producing a moderated value for any average.

The given example is fairly well behaved, and excluding the highest and lowest values leaves the same average of 1.060 as in the first trial above.

c) Weighted Averages

The past years' data can be given different weights, normally with higher weights for the more recent years. The rationale here is that the more recent the data, the greater the weight should be placed on their relevance for the future. In the given example, relative weights of:

1 2 3 4 5 6 7

could reasonably be used. The weighted average then works out at 1.0615, with a corresponding projection forward. (The weights have been chosen here in a simple arithmetic progression — but weights in geometric progression, or some other intrinsic relationship, could also be used.)

d) Claim or Exposure Weighted Averages

If the ratios given are, say, development factors on paid claim data, then they can be weighted according to the claim values from which they were originally derived. Thus, if the last three years' data only are used, and if the claim amounts concerned are:

800 1150 1300

then the weighted average comes to the value: 1.0635. This type of weighting is in fact commonly used in the chain ladder projection. As well as paid claim values, exposure measures for the years of origin can also be used as weights. Suitable measures might be the number of policy units exposed, or the earned or written premium for the year.

<>

[B8]
MATHEMATICAL TRENDLINES

This section deals briefly with the mathematical aspect of fitting a trendline to a given set of data points. In the example of §B7 the trendline was fitted by eye, using a simple graph. A more satisfactory method, however, from the theoretical point of view, is to fit a mathematical line or curve. The most common standard adopted is to find that line which minimises the sum of the squares of the deviations of the observed points

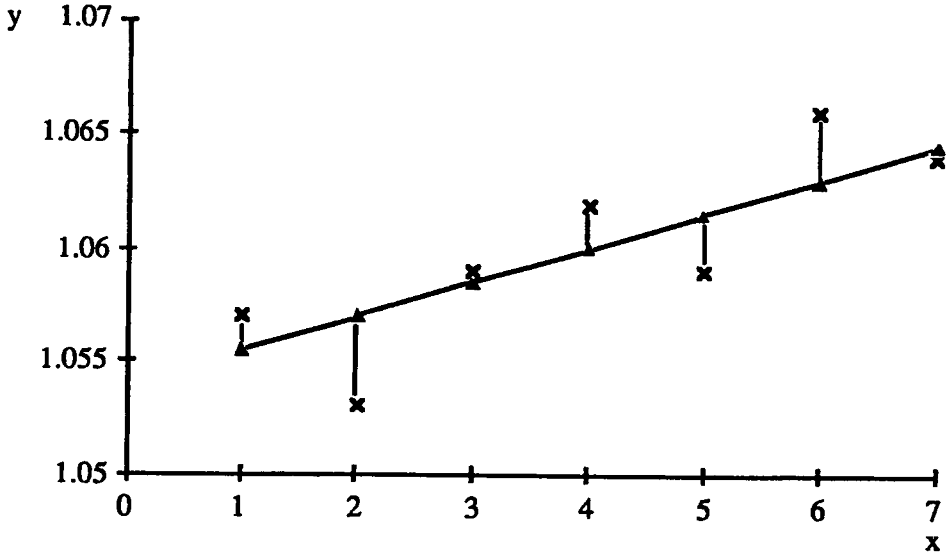

Here, the line to be fitted is:

$$y = ax + b$$

where *a, b* are constants to be determined. It is taken that there are *n* points, with co-ordinates (x_i, y_i). The quantity to be minimised is thus:

$$\sum_i (ax_i + b - y_i)^2$$

Partial differentiation with respect to a, *b* respectively gives the equations:

$$\sum_i (ax_i + b - y_i) \, . \, x_i = 0$$

$$\sum_i (ax_i + b - y_i) = 0$$

Let $\bar{x}$, $\bar{y}$ be the mean values of the x_i and the y_i. The second equation immediately transforms to:

$$n \,.\, (a\bar{x} + b - \bar{y}) = 0$$

Hence:

$$b = \bar{y} - a\bar{x}$$

Substituting this value back into the first equation gives:

$$\textstyle\sum_i \{a \,.\, (x_i - \bar{x}) - (y_i - \bar{y})\} \,.\, x_i = 0$$

and hence also:

$$\textstyle\sum_i \{a \,.\, (x_i - \bar{x}) - (y_i - \bar{y})\} \,.\, (x_i - \bar{x}) = 0$$

Thus a is found as:

$$a = \textstyle\sum_i \{(y_i - \bar{y}) \,.\, (x_i - \bar{x})\} \,/\, \sum_i (x_i - \bar{x})^2$$

These formulae evaluated for the main example give a, b as:

$$a = .00168 \qquad b = 1.0533$$

Hence the trendline is:

$$y = .00168\,x + 1.0533$$

Evaluating y for $x = 8, 9, 10$ gives the required projection:

1.0667 1.0684 1.0701

This compares with the values earlier fitted by eye of:

1.066 1.0675 1.069

The difference between the two sets of estimated values is not very great in this case, but can sometimes be quite marked. The advantage of the mathematical trendline is that it provides a fully reliable procedure for making the fit, i.e. one not subject to individual bias.

Fitting an Exponential

The assumption so far has been that any fitted trend should be a straight line. That is, the trend will show increasing values in arithmetic progression as the years pass by. But sometimes to assume a trend which progresses by *geometric* ratio may be more appropriate. The mathematical procedure is then to fit an exponential curve rather than a straight line to the data.

The simplest means for this is to convert the y-values on to a log scale, and then carry out the linear fit as before. In short period projections, the switch to the exponential may often not affect the results greatly. However, over a longer period, the influence of the geometric factor will very much become apparent.

This again highlights the importance of the choice of forecasting method. Even with relatively well-behaved data, such as those in the given example, appreciable differences in the results soon become apparent. The only way to make an informed choice of method is to be cognisant of the business conditions and influences which are currently making themselves felt. The assessment of such influences is taken up in the next main section, §C, of the Manual.

Deeper Waters

Beyond the simple functions dealt with above, there are many and more complex mathematical functions that can be used for trendline purposes. Indeed, the whole theory of Curve Fitting and Time Series Analysis can be brought into play if so desired. These subjects are touched upon in some of the methods described in later parts of the Manual, but are beyond the scope of the present section.

<>