THE PHILOSOPHY OF RESERVING

by

Robert Buchanan

Abstract

This paper is about how liability provisions should be set when there is material uncertainty. Conventional accounting practice is deterministic. A stochastic approach is needed. Since a single figure is needed in general purpose financial statements, the provision should include the value of uncertainty. There is a conflict between the measurement of solvency and profitability. The best estimate for provisions is a utility estimate, which is intended to provide a typical general user with a figure which will lead to the same decisions as a full understanding of the amount and uncertainty of the liability. Estimates should incorporate proper allowance for all known material influences, including discounting. Estimates made in isolation from previous estimates can be very erratic. A time series approach to estimation is needed, in which the estimates successively adopted only vary to the extent justified by the reliability of the corresponding isolated estimates.

Bob Buchanan is a Fellow of the Institute of Actuaries of Australia and of The Institute of Actuaries in the U.K. He is a Subscriber in the Casualty Actuarial Society. In 1980, he became the second actuary in Australia employed full-time by a general insurance company. He now operates a consulting practice in Australia, specialising in accident compensation business. He is co-author of the Australian textbook, *Actuarial Practice of General Insurance* and has served as chief examiner in general insurance and as chairman of examiners for the Australian Institute.

The author wishes to thank the primary reviewer, Mr Roger Bovard, for his insightful questions during the review process. He also wishes to apologise for the fact that the paper is written in Australian, rather than American, English.

Introduction

This paper is about the philosophy and practice of setting provisions, in respect of uncertain future liabilities, for inclusion in an entity's published accounts and statutory returns. While it is written in the context of liability insurance, for an audience of casualty actuaries, the principles discussed are quite general and are equally applicable to other forms of insurance and, indeed, wherever liabilities are subject to material uncertainty.

Clearly, published and statutory provisions must comply with the relevant accounting standards and legislation. An earlier version of this paper was written in the context of Australian accounting standards. While this version does contain reference to accounting standards, it is intended to be more general and is more concerned with what those standards should be than with how to comply with specific standards.

The paper draws two main conclusions. The first is that the utility of the published provision for uncertain liabilities should equal the expected utility of those liabilities. Because utility varies between individuals and is notoriously ill-defined in practice, this does not lead to a clear prescription for the calculation. If, however, we assume that the relevant individuals are generally risk averse, it follows that the "best estimate" of the appropriate provision is greater than the expected value by an amount which is proportional to the uncertainty and which increases with increasing risk aversion.

The second is that the estimation process should be approached as the estimation of a time series, rather than as a series of isolated estimates. What this means in practice is that, starting from a proper estimate of the probable profitability of a particular cohort of business, at the time it is written, the estimated profitability of that business should vary over time only to the extent justified by subsequent experience, as it emerges. This implies a form of Bayesian estimation in which successive isolated estimates are given high weight if they are highly reliable and low weight if they are highly uncertain.

Valuation and Reserving

Valuation is the process of financial analysis used to gain an understanding of value. It may be applied to the entity as a whole or to selected parts, such as assets or liabilities. Since the aim is understanding, context is vital for valuation. In particular, if asset and liability valuations are to be compared, they must be compatible. It is also necessary to a proper understanding that the valuation basis be realistic. In particular, the valuation should include all reasonably foreseeable components and should (whether explicitly or implicitly) make appropriate allowance for all factors which might reasonably be expected to have a material impact on the value.

Reserving is the process of selecting a figure to be shown as the reserve or provision for liabilities. This figure is normally based on a valuation but, being a single figure, must necessarily convey a more limited understanding. There may also be criteria other than understanding. It is usual for management to seek to create a good impression, while accounting and legislative constraints can also introduce distortions. Ideally, though, the selected provision should convey, as nearly as possible, the same understanding as the valuation to which it relates.

Accounts

Accounts can be broadly divided into general purpose accounts and special purpose accounts. Special purpose accounts are prepared for and to the specifications of a particular user or group of users, while general purpose accounts are published for the information of anyone who may be interested. Standards are needed for general purpose accounts, partly to ensure that the accounts for each entity give a true and fair view of that entity and partly to ensure comparability between entities. This paper is primarily concerned with general purpose accounts. Where special purpose accounts are published or otherwise generally available, however, these should also comply with the standards applicable to general purpose accounts.

The concept of a true and fair view of a complex entity is a difficult one, particularly in the face of material uncertainty. Accounts are prepared as a basis for users to make financial decisions which can usually be expressed in simple terms such as "yes", "no", "how much" or "which one". Vanishingly few users are equipped to make such decisions on the basis of fully detailed financial information. Most can only cope with a few, key, summarised figures, particularly when there is a choice between a number of entities.

While there are many questions which may be asked, most users of general purpose accounts have a particular interest in one or both of two key questions.

How solvent is the entity?

How profitable is it?

Conventional accounting practice has developed in the context of a deterministic world-view. "Bean-counting" requires beans which can be counted. This is true for entities which have well defined tangible assets and liabilities, but sits uncomfortably where there is material uncertainty. Where the uncertainty is high, as is the case with liability insurance, a stochastic approach is needed, in which uncertainty is explicitly recognised. Because relatively few accountants are trained beyond simple descriptive statistics, this can easily lead to misunderstanding.

One particular problem relates to the inclusion of a margin in liability provisions. ain a deterministic world, the amount required can be known and any margin distorts the accounts, to the detriment of a true and fair view. This is anathema to accounting purists. As discussed below, however, a provision for uncertain liabilities should have a better than even chance of adequacy. This has usually been presented either as a central estimate plus margin, which is seen as a distortion, or an undiscounted provision, which is a distortion.

The resolution of this problem is to recognise that the provision must allow for the value of uncertainty. The provision can then be seen as comprising the value of expected payments, plus the value of the uncertainty of those payments. This semantic change gets rid of a term which is likely to cause misunderstanding. It also provides a conceptual basis for quantifying the amount by which the provision should exceed the value of expected payments.

Solvency

Broadly speaking, solvency is a comparison between the value of an entity's assets and the value of its liabilities. This comparison can be expressed in a number of forms, including the difference between the two values, their ratio and, as I shall discuss further, the probability that the assets will prove adequate. While this comparison is simple in concept, there are a number of complications in practice.

One such complication is the scope of the assets and liabilities to be compared. In particular, how should future transactions be treated. One aspect of this is goodwill, which is based on the ability to derive future profits from existing relationships. In insurance, there is the question of funding. At one extreme, it is possible for government bodies to run compulsory insurance schemes on a pay-as-you-go basis while, at the other, it is possible to pre-fund catastrophes.

It is also not always clear how various forms of finance should be treated. Subordinated debt is a borderline case, as are a variety of forms of financial reinsurance.

A related matter is the context of the comparison. Unless the assets and liabilities are perfectly matched, their relative values depend on circumstance. In particular, there are significant differences between:

forced liquidation;

the actual run-off of an insurer which has ceased business;

the actual run-off of a closed portfolio by a continuing insurer;

the notional run-off of existing business by a continuing insurer.

Relative values are also affected by a multitude of external factors, such as interest rates, economic conditions, legislation, court decisions, etc. Some of these may be known at the valuation date, while others contribute to uncertainty. For risk enterprises, such as insurance, there may also be significant inherent uncertainty.

Uncertainty is important in two ways. First, it means that values are not known, but can only be estimated. This raises the question of how the estimation is to be done and what sort of estimate should be used. The three best known estimates are the mean, the median and the mode. For financial valuation purposes, the mean is usually the most appropriate but there are often subjective elements which cannot be properly classified but which most probably favour the mode.

Uncertainty also has its own value. An uncertain income is generally regarded as less valuable than a certain amount, equal to the expected value. This may be seen in the fact that riskier investments generally offer higher yields, even after allowing for the probability of default. The relationship also extends to expenditure, regarded as negative income. People on the whole are prepared to pay to reduce uncertainty. This, and the fact that uncertainty is less than additive, provides the economic basis for insurance.

Timing is another important factor. A given amount, paid now, is worth more than the same nominal amount payable at some future date. The best way to allow for this is to assess future payments in terms of their present values, using a suitable interest rate or series of interest rates. For a meaningful assessment of solvency, it is necessary to use the same discount factors for assets and liabilities.

More generally, for a meaningful comparison of asset and liability valuations, it is necessary that they should be compatible. Ideally, both should be conducted by the actuary, using a single basis for both. More usually, the asset values are determined independently of the liability valuation, in a way which does not relate easily to the liability valuation basis. Nevertheless, to the extent to which there are equivalent assumptions underlying the two bases,they should be the same.

This is easiest if assets are taken at market value. In this case, the discount rates used for the liabilities should be market rates. These rates, in turn, relate to market expectations of future economic and social conditions. These expectations can be seen most clearly in the pronouncements of leading economists and market analysts around the balance date. If the income streams for assets and liabilities are highly correlated, then the other economic assumptions for the liability valuation should be closely based on market expectations. This would be the case if index-linked bonds or equities are held as a hedge against inflation. Where there is little or no correlation, as is the case for a matched portfolio of government bonds, there is more scope for actuarial judgement, but any significant departure from market expectation needs good reasons.

Profitability

In the simplest case, with no uncertainty and all assets and liabilities due instantly, the profit arising in a period is equal to:

> assets minus liabilities at the end of the period;
> minus
> assets minus liabilities at the start of the period;
> plus
> profit paid out during the period.

When we add payment delays, it becomes necessary to consider the discount rates at the start and end of the period. The profit arising can be divided into the inherent profitability of the business and profit arising from the change in discount rates. If assets and liabilities are matched or nearly matched by term, this second component will be small.

Uncertainty adds a third component, arising out of the uncertainty in the opening and closing estimates. Under some circumstances, this can swamp the inherent profitability. For example, it would not be unreasonable for a liability insurer to expect a profit of 5% of premium and to require provisions equal to 5 times premium, with a standard error of those provisions of 20%. A change of 5% (quarter of the standard error) in the liability valuation basis is equal to 25% of premium or 5 times the expected profit.

To obtain the best possible estimate of profitability, the opening estimates should be reworked, with the full benefit of the extra period's data which is available for the closing estimates. This removes the (sometimes large) variations arising from changes in the actuary's interpretation of the experience. If the inherent variability is large, it may also be desirable to extend this reworking back for several periods in order to view profitability over a longer period.

This implies that the profit emerging in a period should be presented in two parts - that which comes from the inherent nature of the business and that which arises because successive estimates are not fully comparable. Such treatment, however, may well cause more confusion than it resolves.

Presentation of Uncertainty

Where a quantity is known, it can be characterised by a single number. Where it is uncertain, however, two or more numbers are needed. It is the nature of balance sheets and revenue accounts, however, that only a single number may be used. If a single number must be used to characterise an uncertain quantity, there is a whole range of possibilities. In my opinion, the value chosen should be the one which will lead to the same decisions about the allocation of scarce resources as would a full consideration of the uncertainties involved. This, of course, is not really possible, since those decisions will be made by a variety of persons, each of whom has a different perspective. What is possible, however, is a compromise value which takes a middle position.

In mathematical terms, the number which should be used is the number which has a utility equal to expected utility of the uncertain quantity. Computationally, this may be of little help, since the probability distribution may not be well known and real-life utility functions are notoriously ill-defined. Nevertheless, it is possible to come to some broad conclusions about the relationship between the uncertain quantity and the number to be used when it is recognised.

Another problem is that, for some classes of general insurance, the uncertainty of the value of the liabilities can be several times greater than the expected profit for a year. Where this is the case, if no steps are taken to limit variability, any comparison of results over time is essentially meaningless.

This leads to a potential conflict between the measurement of solvency and the measurement of profit. In accounting terms, the question is whether it is more important that the balance sheet or the revenue account should present a true and fair view.

The main conflict is between presenting a true and fair view of solvency, and a true and fair view of profitability. Solvency measurement calls for consistent treatment of assets and liabilities, both of which should reflect, as well as possible, the situation at the time of measurement. Consistency from time to time is not important and can, indeed, make it harder for the measure of solvency to respond to changing circumstances.

Profitability over a period is the difference between two successive solvency measurements. If this is to have any meaning, these measurements must be consistent over time. Because profitability is a difference over time, consistency between assets and liabilities is less important, but does no harm.

In brief, solvency measurement must be responsive to changing circumstances, while profitability calls for stable measurement. In changing circumstances, therefore, successive good measures of solvency will give an unreliable picture of profitability, and vice versa. This conflict is essentially irreconcilable. It is the direct result of attempting to measure (a deterministic concept), quantities which, because they are stochastic in nature, can only be estimated. The best we can hope for is a compromise.

It is my impression that accounting theorists favour the solvency or balance sheet view maintaining, incorrectly in my opinion, that the profitability or revenue account view will look after itself. It is abundantly clear, however, that investors and potential investors, who are arguably the principal users of listed insurers' accounts, concentrate, arguably to an unhealthy extent, on the revenue account and on profitability. The balance sheet is of lesser concern to investors unless it is weak or there is the prospect of a takeover. It is therefore vital for a listed insurer to optimise its measurement of profitability. The other main group of users is insurance policyholders and potential policyholders. They are most likely to be concerned with solvency. Insurance regulators are also largely concerned with solvency.

Measurement of Solvency

The starting point for the measurement of solvency should be the actuary's central estimate of the outstanding liability. This should anticipate future trends to the extent that the actuary considers them credible. As a result, estimates for solvency purposes should be more volatile than those made for profit measurement.

Solvency provisions should exceed the corresponding central estimates. The reason for this is that the central estimate makes allowance for what is expected, but not for uncertainty. Uncertainty, nevertheless, has a value which must be provided for. A reasonably widely accepted standard in Australia is that there should be an 80% probability that the provisions adopted should prove adequate.

Measurement of Profit

Insurance profit is mainly earned as reward for carrying risk. (A smaller component relates to the provision of other services, particularly claim management.) Because of the nature of risk, the amount of profit in respect of a particular policy cannot be finally known until that policy has expired and any resulting claims have been finally resolved. Since this period can extend for many years after the policy is issued, the question of how profit should be recognised over time is an important one.

In general terms, insurance profit can be divided into two parts. The first part can be called expected profit and is the profit which would emerge if everything goes as expected. This should be recognised in proportion to the delivery of the services for which it is earned. In general insurance, the most important of these is risk. The corresponding part of expected profit should be recognised as the insurer is released from risk. This implies that the outstanding claim provisions used to measure profit should include a component in respect of the associated risk. This is also true of the profit in respect of future claim management.

The second part of insurance profit can be called experience profit. This arises out of the random or other variation between what was expected when the business was written and what actually happens. In principle, the profit arising out of these differences should be recognised as the events (or non-events) giving rise to them occur. In practice, the best we can hope for is that they will be recognised as the insurer becomes aware of them. (In extreme cases, backdating is also possible.)

Actual payments during a period are relatively straightforward. When the amount paid is different from what was expected, the difference flows straight to profit or loss. Timing differences do not affect profit or loss, except for interest adjustments. Estimate changes relating to new information also flow to profit or loss. Other differences require more thought.

If, for example, an apparent pattern in the data leads the actuary to change the expectation of future events, should this be recognised in profit in the current period, because the change in expectation is a current event, or deferred until the expected change is (or is not) realised? My feeling is that the argument in favour of immediate recognition is strengthened if there is an apparent causal basis for the change and as evidence of the change accumulates. If on the other hand, there is nothing but a recent apparent trend, it is probably preferable not to extrapolate this for profit measurement purposes.

Actuarial Valuation

Much general insurance actuarial practice to date is based on a deterministic view of reality. That is, it is implicitly assumed that future transactions where the amount is unknown, can be adequately described by a single number, the estimate, and that sums, products and other functions of these amounts can be described by the corresponding functions of the individual component estimates. This is a good approximation if the relative variability of the amounts involved is low, but becomes increasingly doubtful as variability increases.

When variability is high, the deterministic approach can give highly misleading results, particularly when non-linear functions are involved. For example, if the average claim size is $50,000, a pure deterministic approach implies that every claim costs $50,000 and that there would never be any claims against a reinsurance treaty excess of $50,000. In practice, however, it is not unrealistic to find reinsurers charging about 15% of the premium for reinsurance excess of $1,000,000 for such an average claim size.

Where variability is high, a stochastic approach is needed, which takes into account the range of possible values of each amount, together with their respective probabilities.

One consequence of the stochastic approach is that we see that variability has a value. In a deterministic world, it would be sufficient to set up provisions equal to the expected value of outstanding claims. In the real, stochastic world there is roughly an even chance that such provisions will prove inadequate. If an insurer were to consistently set provisions on such a basis and release all the "profit" disclosed thereby, so that no surplus is retained, it would be only a matter of time before it failed. Only about one in a thousand such insurers would survive ten years.

It is for this reason that an insurer must charge a profit margin in its premiums and include a margin over the actuary's central estimate in its provisions. Even this statement, however, is a holdover from the deterministic approach. It implies that the best estimate of the value is the value of the central estimate. As we shall see, this is not necessarily the case.

Another consequence is that it is seldom possible to accurately measure anything in general insurance. The best we can do is estimate it. These estimates are of varying quality, depending on the variability and uncertainty of the quantity to be estimated.

Estimation

In the context of this paper, estimation is the actuarial process of forming a view as to the value and uncertainty of (a specified part of) the policy liabilities of a general insurer. An understanding of uncertainty is an essential outcome of the estimation process. Even where formal statistical/time series analysis techniques are not used, this is essentially an exercise in statistical/time series analysis and projection.

Central Estimate

This is a statistical concept, albeit a rather fuzzy one. It is based on the fact that most probability distributions show a degree of clustering, often around a single value. A central estimate is a measure of the location of that clustering. For any particular distribution, there are many possible central estimates. The three best known are the mean, the median and the mode.

The mean of a probability distribution is the weighted average of all possible values, weighted by their respective probabilities. It is also called the expected value. Every observable probability distribution has a mean.

The median of a probability distribution is the value with equal probability that a random observation will be higher or lower. If this coincides with a gap in a discrete probability distribution, convention places the median at the middle of the gap. Every observable probability distribution also has a median.

The mode is the most likely value. This is well defined only for discrete distributions. For continuous distributions, it is necessary to define *most likely* more closely. The conventional definition looks at the probabilities in small, evenly spaced ranges of values. For monetary values, it is usually preferable to work with ranges based on ratios. This can give a very different result. The mode can also vary, depending on the width and nature of the ranges used. For workers' compensation for example, the claim size distribution is, strictly speaking, a discrete distribution, with possible values at one cent intervals. On this basis, the mode is almost certainly the cost of a standard medical consultation. There are a few other popular values, but most other possible values have very low probability. More usually, it is thought of as a continuous distribution. With one cent intervals, this is the same as the discrete distribution. With wider intervals, the picture is much the same, with a few popular ranges, of which the most popular in Australia is about $50. With very narrow percentage ranges, the picture is not changed greatly, but, as these are widened, the distribution starts to show some regularity, with a peak at a much higher value, possibly around $5,000, which is more in accord with a general perception that there are a lot of moderate sized claims.

I have laboured this point because there can be substantial differences between the mean, the median and the mode and because I have a perception that, where subjective judgement is used, this is most likely to produce modal values. For most distributions found in general insurance practice, the mode is the lowest of the three, with the mean the largest. For financial purposes, however, I believe that the mean is the most appropriate central estimate. Using modal values, particularly if two or more modal values are multiplied together, can result in underestimation.

Another feature of central estimates is that, in general, the central estimate of a function is not the same thing as that function of the central estimate(s). The expected value of the product of two variables, for example, is only equal to the product of their expected values in special cases. In general insurance work, it is usually higher.

With highly non-linear functions, such as excess-of-loss reinsurance, the difference can be overwhelming. The excess point is almost always set well above the expected value, so that the amount of the excess at the expected value is zero. The probability distribution, however, always extends above the excess point and the expected value of the excess can be a substantial fraction of the expected gross cost, while the median and the mean are both zero.

Best Estimate

This phrase is commonly used in discussions of estimation and setting provisions. But what does it mean? What should it mean?

As with central estimate, there is no single, clear mathematical definition of the term. Even if there were, this would not necessarily be what was intended in non-mathematical discussion.

What seems to be meant, in the context of setting provisions, is that the best estimate is the best that the actuary can do at that time, on the basis of the available information. In many cases, different methods of analysis have suggested different values and the actuary has had to choose a particular value or combination of values. It is my impression that this choice of values is usually made with little or no regard to the purpose for which the estimate is required.

I would contend, firstly, that the purpose is a vital element in this process and that the estimate should be the best estimate for this purpose. This, of course, raises the question of what is best and for what purpose. What are the desirable qualities of estimates which are to be used as the basis for published provisions?

The ideal is, of course, that the estimates should always, in retrospect, prove correct. If we can develop estimates with very low uncertainty, this takes precedence over all other considerations. It also satisfies them. There is, however, a lower limit to uncertainty, which is set by the nature of the business being valued. Liability business tends to be highly uncertain, and this uncertainty can be greatly increased when the rules are changed.

If certainty is unattainable, we look for objectivity, stability, responsiveness and lack of bias.

By objectivity, I mean that different actuaries should come to the same answer, if given the same information.

Stability means that the answer should not vary unduly with changes in the information on which it is based. This is a particular problem when the answer is highly uncertain. Answers which jump up and down from one valuation to the next are an indication of a lack of stability. For an isolated estimate, this property is closely related to minimum variance. This is the requirement that, if many samples are drawn from a given probability distribution, the estimation procedure should be such as to minimise the variance of the resulting estimates around the true value.

Responsiveness means that the answer should fully reflect real changes in circumstances. Answers which never change or which move in the same direction over several valuations are usually indicative of a lack of responsiveness. Responsiveness and stability are directly opposed and it is necessary to find an appropriate compromise when the answer is highly uncertain.

Bias can be approached in two ways. The first is personal bias and should not be allowed. The second is statistical bias. In this second sense, a lack of bias means that the answer should not be systematically higher or lower than it purports to be. That is, if many estimates of a particular type are made, they should behave in the way expected of that type of estimate. Estimated means should, on average, equal the subsequent observed values, estimated medians should split 50/50 above and below the observed values, etc. For skew distributions, such as are often found in casualty insurance, there is a conflict between minimum variance and lack of bias.

These criteria are easily met if there is little uncertainty about the underlying claim process. Even if there is substantial uncertainty, it is not difficult to arrive at a suitable compromise if the uncertainty is well understood. The real problems arise when, as is often the case for long-tail business or the rules are changed, the true extent of the uncertainty is not known.

There is also a problem which arises in the choice of economic parameters. These are subject to strong external influences and the choice of values for future projections is essentially subjective. There is also a high degree of subjectivity when the rules are changed and the new rules are yet to be tested in practice.

The purpose of the estimate is an important factor in how the inevitable conflicts should be resolved. The best estimate will always be a central estimate. But which one? Even more importantly, of what should it be a central estimate?

Utility Estimation

The amount shown in an insurer's accounts as the provision for outstanding claims is intended to represent the value of those claims. If this is to be a fair representation, the single figure shown should lead users of the accounts to the same conclusions as would a full understanding of the nature, and particularly the uncertainty, of those claims. This leads directly to the concept of utility estimation.

The mathematics of utility estimation is discussed briefly in the Appendix. The principal conclusion is that, if we assume that users of an insurer's accounts are generally risk averse, the utility estimate of the value of outstanding claims is greater than the monetary expected value. To a first approximation, if we assume exponential utility, this difference is proportional to the coefficient of risk aversion and also to the variance of the value of outstanding claims.

In order to apply the theory, it is necessary to know both the probability density function of the liability and the utility function. Both present problems.

While certain components of the probability density function can be estimated using standard statistical techniques, other parts are less tractable. The subjective judgements which must be made in estimating the liability are not properly susceptible to statistical analysis. Nevertheless, I believe that, with the aid of more subjective judgement, it is possible to arrive at a probability distribution which is a not unreasonable representation of the liability and its uncertainty.

The choice of utility function is also difficult. Even for a single individual, utility cannot be measured directly. It must be inferred from the choices made by that individual. Those choices are often influenced by non-financial considerations and can appear to be inconsistent. The usual response to this problem is to take a broad approach, usually exponential utility as described in the Appendix, and to select values which broadly reflect the observed choices.

Utility functions also vary from individual to individual. As a basis for entries in general purpose financial statements it is clearly necessary to adopt a compromise utility function which broadly reflects the general users of those statements. One way of approaching this problem is to look at the values actually recorded in the accounts. It should be possible, for example, to infer a compromise coefficient of risk aversion from provisions actually held by insurers which are generally thought to follow good accounting practice.

Preferred Approach for Provisions

Rather than approach the valuation of outstanding claims as a succession of isolated estimation problems, my preference is to use an approach which draws on the concept of Bayesian estimation and is related to the Bornhuetter-Ferguson approach to valuation. How this may best be done is most easily explained by considering how the actuary’s perception of the cost of a block of claims varies over time.

Before the business is written, there is an expectation, based on market conditions and previous experience, of the total cost per unit exposure. This can be decomposed into expectations as to expenses, claim frequency, reporting pattern, claim cost, run-off patterns and economic and social conditions. While these may or may not be explicitly formulated as part of the rate setting process, it is generally possible to develop them.

Over the term of the policy, the actual claim frequency will emerge (but will not necessarily immediately become apparent). The other components will emerge over a rather longer period, until the last claim is finally settled.

At any particular stage, we will have:

> payments to date, which are known (payments);
>
> case estimates of the outstanding liability on reported claims, which can be quite reliable for short-tail claims, but which vary substantially in quality for long-tail claims, as better information comes to hand; (case estimates);

future development on reported claims, which is unknown, but can be estimated on the basis of past experience (IBNER);

claims which have occurred but have not yet been reported (IBNR); and

future claims within the term of unexpired insurances (unexpired risk).

Over time, claims move up this ladder until, finally, the last payment on the last claim is made and the true cost is known. During this period our initial expectations are:

replaced by fact (payments and, to some extent, case estimates); and

modified in the light of emerging experience.

Volatility in the experience can, likewise, be divided into three parts:

the emerging differences between actual and expected payments;

the emerging differences between actual and expected case estimates; and

modifications in our expectations about the unknown elements of IBNER, IBNR and unexpired risks.

Clearly, a difference between actual and expected payments should flow into profit or loss in the period in which it occurs, to the extent that it represents a real difference in cost, rather than timing. Similarly, differences between actual and expected case estimates should, with a further qualification, be recognised as they occur. The qualification is that case estimation standards can change. Such changes ought not be allowed to distort profitability.

Modifications to expectations are more complicated. To the extent that expectations as to future costs change to offset timing differences and changes in estimation standards, they should be recognised, so that the actual payments and case estimates can be used without adjustment.

Other modifications in our expectations arise because of variations in the experience and in how we interpret that experience in the light of external conditions. It is these modifications which are most likely to cause large variations in our expectations. Changes in interpretation can cause particularly large swings.

If there is no subjective content in the valuation (or if any subjective content remains unchanged), then the Kalman Filter provides exactly what is needed. For a long-established, stable class of business, it provides a theoretically sound basis for reacting to the random shocks which may form part of the stable pattern. In essence, successive one-off estimates are weighted in inverse proportion to their variances.

A number of major classes, however, require substantial subjective input. One subjective element which is unavoidable is the choice of valuation approach. A major change in the environment usually also requires a subjective response, if the future implications of that change are not known. Even if the decision is to make no special adjustment in response to a significant change, that is a subjective decision.

As a result, successive central estimates are an uncomfortable mixture, in which the measurable aspects of the basis should progress fairly smoothly in response to the accretion of new data and the erosion of old, while the subjective aspects can change in a far more erratic fashion. Excessive volatility can also arise when, because of rapid environmental changes, older data is of limited relevance and the actuarial analysis is based on a relatively short period. I refer to such estimates as "isolated" or "one-off" central estimates.

These subjective decisions can have a major impact on the results of a valuation. For example, payments in two closed portfolios of which I have experience have, over the last eight years, both risen way above the levels expected under previous valuation models. In both cases, the behaviour of the case estimates gives limited guidance, in one portfolio, because there was no history of case estimates, and in the other, because the estimation procedures have changed substantially. The valuation results vary widely, depending on whether the higher payments are seen as selective settlement of the larger cases or as a more generous approach by the judges. To resolve this, it is necessary to rely heavily on the views of management.

Under such circumstances, it is not possible to apply the theoretical approach, except on the basis of subjective assessment of some vital parameters. A more pragmatic approach is indicated. This is to break the present strong connection between the isolated central estimate and the provisions to be shown in the accounts.

This link appears to be based on the unspoken assumption that the isolated central estimate is a reliable estimate of the value of the outstanding liability. Where this is true, the strong linkage is appropriate. The relevant accounting concept is that when the value of an asset or liability can be reliably measured, it should be brought to account and that, if it can not be reliably measured, it should either be ignored or shown in the notes as a contingent liability.

For long-tail classes, however, it would seldom be fair to describe an isolated central estimate as a reliable measurement, merely the best that can be done on the basis of the available information. This unreliability can be partly seen if successive estimates of the incurred cost of a cohort of business are compared. Unreliability is, however, often understated for cohorts with a large outstanding liability, because estimators are reluctant to make large changes unless the evidence is overwhelming.

Accounting doctrine states that liabilities must be recognised if they can be measured reliably. It is not clear on what to do if the best available estimates are unreliable. On the one hand, one might say: no reliable measurement, therefore no provision. This approach is commonly adopted outside insurance for contingent liabilities which are unlikely to eventuate. In insurance, this would be silly, because it would lead to outstanding claim provisions way below the probable cost.

A second approach, which is now normal in insurance, is to use the unreliable estimates as if they were reliable measurements. This is the source of considerable spurious volatility in insurance profits.

A better alternative would be to defer judgement on the emergence of profit or loss where this would depend on unreliable estimates. That is, where the outstanding liability at a point of time cannot be measured reliably, use the estimate, as at that point of time, which was (or would have been) made when the business was written (the a priori estimate), modified as appropriate for actual intervening experience to give a "time series estimate". This is essentially the same approach as is used toward unearned premiums. The a priori estimate is no more reliable than an isolated estimate as an estimate of actual cost. Reliability, however, is not the point here. If a reliable estimate is not possible, we need a series of estimates which do not express their unreliability in unjustified volatility and which converge to a reliable estimate as this becomes possible. How the a priori estimate might be determined is discussed in the following section.

In short, the incurred cost of each cohort of business should be estimated when the cohort is written. This estimate of the incurred cost should only be changed over time to the extent that is justified by the reliability of subsequent isolated central estimates of the liability.

A Priori Estimation

The starting point for a priori estimation is the basis on which premiums have (or should have) been calculated. From this, it is possible to project the expected liability for outstanding claims, per unit of exposure, at each duration. While the premium basis does not usually go into this detail, it is possible to dissect this into unreported claims, future development on reported claims and reported incurred cost, less payments to date. These can all be expressed in terms of current values as at the date the business is written, undiscounted, together with interest and inflation assumptions.

The then current value projections for unreported claims and future development should be brought up to date for actual inflation, but not otherwise disturbed. Actual reported incurred costs and payments, suitably adjusted for inflation, should be substituted for the corresponding projected values. The effect of these adjustments is to use the actual experience for things that have happened to date and the premium basis (in current values) for what is still in the future.

If assets are valued at (now) current market values, the current value projection should be adjusted for inflation according to current expectations and discounted at a current market rate. For an amortised bond valuation, it may be appropriate to inflate and discount at the original rates. Other asset valuation bases are probably best allowed for by re-scaling the values corresponding to market values.

Assessment of Uncertainty

The approach which I adopt to the assessment of the uncertainty of actuarial estimates in general insurance is based on a three way division of sources of uncertainty.

The first of these is stochastic variation, the random variation from policy to policy and from claim to claim which is the basis for the need for insurance. It also includes external factors, such as the weather, which operate randomly on groups of policies, or even on whole portfolios. Given a reasonably large body of data, this is amenable to normal statistical analysis. I do not explore these techniques in this paper. This component of uncertainty generally becomes relatively less important as the volume of business increases. If the only source of variation lies at the policy level, an inverse square root rule applies.

The second, economic uncertainty, relates to the choice of inflation and discount rate assumptions. It is possible to limit exposure to interest rate changes by matching investments to the projected cash flows, but there is only limited scope for hedging inflation in those classes where claim payments are formally or otherwise linked to inflation. This form of uncertainty operates on all policies or claims in the same way. The relative uncertainty does not vary with the volume of business.

The third relates to unanticipated changes to the claim process and to changes which may be anticipated in a general sense, but which cannot be reasonably quantified in advance. These include legislative changes, "landmark" court decisions, the emergence of new insured perils and changes in community attitudes and expectations. Again, the relative uncertainty does not vary with the volume of business.

Economic uncertainty is only partially amenable to statistical analysis. It is possible to analyse the variability of past interest and inflation rates, but this analysis shows that there are systematic variations in both the rates and their volatility, which appear to be based on political, social and economic factors operating over periods of several years, with occasional sharp changes. Since we are dealing with run-off patterns which also have a mean term of several years, it is necessary to combine subjective judgement with the evidence from the historical record.

The impact of a given change in economic conditions on liability estimates increases with the mean term. The likelihood of a substantial unexpected drift also increases with duration. I am presently using a rule of thumb, for Australian conditions, which attributes a "standard error" of about 2.5% multiplied by the mean term.

Unanticipated changes are even less subject to statistical analysis. In Australia, there is a long history of changes which I refer to collectively as "superimposed inflation". While some of these are gradual, there are also (retrospectively) identifiable shocks and short bursts of rapid change, interspersed with periods when nothing much happens and even periods of falling costs. A particular concern arises when there is a significant change to the legislative basis of one of the liability classes. In an extreme case, this can lead to a completely new set of payment patterns and levels of cost. Until the legislation has been tested, these can be extremely uncertain. I have also observed that the early experience is usually favourable, while claimants learn how to manipulate the new rules to their advantage, and that the payment pattern is initially very slow, because claimants and their advisers are reluctant to settle claims on the basis of untested rules.

I have found no better basis for assessing uncertainty due to unanticipated changes than my subjective assessment of the confidence which I feel in my valuation assumptions. Unlike some of my colleagues, I feel reasonably comfortable in quantifying such uncertainties in terms of a percentage "standard error". The usual alternative is to illustrate these uncertainties by means of sensitivity analysis, leaving the user to assess the plausibility of the various scenarios shown.

Most of this analysis is usually carried out on a class by class basis and often for portfolios within a class. To assess the uncertainty for the entity as a whole, it is usually necessary to combine the uncertainties for different classes. A similar calculation is needed when the actuarial assignment is restricted to just one class or portfolio.

Random variation is often largely uncorrelated and can be combined simply by summing variances. Where there are material correlations, these can be analysed using standard statistical techniques.

Economic uncertainty is highly correlated across all classes, but varies in its relative impact from class to class. A weighted average of the various coefficients of variation is a good starting point. It may be appropriate to adjust this if a particular class is subject to unique influences, such as the impact of currency rate changes.

Other changes need to be considered on their merits. There may, for example, be a high correlation between all classes based on tort, but little relationship between these and property classes. In some cases, different legislation may mean that costs change independently in different states. Given that this component of uncertainty is largely subjective to start with, probably the best that can be done is to subjectively choose a position between zero and full correlation.

The final step is to combine the three components. As they are largely independent of each other, this can be done by summing variances.

Discounting

If the aim is an understanding of the solvency and profitability of the enterprise, then it is necessary to make realistic allowance for all known features of the experience. Where the valuation is performed in parts, it is essential that the separate valuations be comparable.

An important feature of the experience of any insurer of long-tail risks is that it receives premium income well in advance of the claim payments to which that premium income relates. This timing difference is resolved by investing in assets which are intended to cover the claim payments when they fall due. Thus we have assets generating a stream of future income and liabilities generating a stream of future expenditure. While these two payment streams, with their associated uncertainties, comprise a full answer to the valuation problem, it is not a particularly helpful one. Normal practice is to summarise into two figures, the value of assets and the value of liabilities.

For a meaningful comparison between these two values it is necessary, at the very least, to allow for the time value of money in respect of any differences between the two streams. This is most conveniently done by discounting each stream to a present value, using the same set of realistic discount factors, and making due allowance for the value of the respective uncertainties.

In practice, the asset income stream is seldom explicitly projected. Rather, asset values are determined in accordance with a variety of conventions. The most appropriate of these is market value. If the Efficient Market Hypothesis is accepted, this is the present value of the expected future income stream, discounted in accordance with the market consensus as to the time value of money and with allowance for the market consensus assessment of the value of uncertainty.

For a comparable liability valuation, it is relatively easy to deduce a set of discount factors which give a reasonable approximation to the market consensus as to the time value of money. One approach is to derive a "risk-free rate" (or series of rates by duration) from the observed yield curve for federal government bonds.

Under current U.S. rules, good quality bonds, which comprise the largest part of the assets of most property and casualty insurers, are valued on the basis of purchase price, amortised to maturity. If these bonds are matched by term to the liabilities and new purchases are made as each cohort of premiums is received, then the corresponding liabilities should, in principle, be discounted on the basis of the interest rates in force when the premium was paid.

For other approaches to asset valuation, it is more difficult to infer a coherent discount basis consistent with the asset values adopted, and often conceptually impossible. The market value, however, is usually known. Provided this is not grossly different in total from the value adopted, a reasonable approximation is given by valuing the liabilities on a basis appropriate to the market value of assets and then rescaling in proportion to the asset value adopted.

An alternative view is that liabilities should be discounted, not for the time value of money, but for the actual yields expected on the investments held. This is usually evaluated in terms of the "risk-free rate", adjusted on the basis of historical performance relative to the "risk-free rate" and in the light of any changes expected in future investment strategy. My own view on this is that:

> the difference between historical yields and the corresponding "risk-free rates" is partly random and partly a reflection of the market consensus assessment of the value of uncertainty and, therefore, improperly offsets the risk discount in the market value of assets; or, equivalently
>
> this procedure has the effect of capitalising investment profits which relate to future investment performance.

I am also uncomfortable that the value of liabilities should be a function of the assets backing those liabilities.

While liability estimates should always be determined on a discounted basis, consistent with the asset valuation, wherever the effect of discounting is material, there is a strong tradition of adopting undiscounted provisions. This practice is usually justified either on the grounds that it is conservative or that the lack of discounting is a proxy for something else which should be, but is not explicitly, allowed for in the provisions.

True conservatism, in the sense of adopting a provision in excess of what is a reasonable value, serves only to make understanding more difficult. This cannot be justified. Conservatism is also confused with the value of uncertainty. This is the proxy justification and fails because the discount is unlikely to be an appropriate value for uncertainty.

The problem with the proxy argument is that, if something should be allowed for, it should be allowed for explicitly. The main candidates for this dubious honour have been expenses, inflation, claim development and IBNR claims. While there is a general correlation between interest and inflation rates, the relationship is not close enough to lead to any confidence that the discount and inflation rates will even approximately offset. About all that can be said of the others is that they are usually of the opposite sign to the discount which should be shown. Using an unsatisfactory proxy leaves open the likelihood that the value adopted will be quite inappropriate.

It is also argued that the subjective case estimates, which are sometimes all that is possible, particularly for unique and/or amorphous large claims, have no associated payment pattern and therefore cannot be discounted. This lack of discounting is sometimes used, with a form of the proxy argument, as a reason for ignoring the sometimes alarming escalation which is possible for such claims. While explicit inflation and discounting are not applicable to such claims, the subjective assessment process should recognise both the capacity for escalation and the time value of money.

In some jurisdictions, there may be an explicit ban on discounting. In the face of such a ban, if the required calculation leaves out other elements of the liabilities, the proper response is to make a separate calculation, discounted appropriately and allowing for everything, as a check.

Mass Torts

Mass torts, along with certain other large, unique and/or amorphous claims, present a special problem in relation to discounting. Essentially, this problem is that there is no adequate experience on which to quantify the probable cost or to develop a payment pattern.

In principle, it is clear that estimates for such claims, which are likely to involve long delays, should be discounted. This has led to concern that discounted estimates for such claims are likely to prove inadequate.

This concern is well founded, but discounting is not the problem. The problem is to make appropriate allowance for claim inflation. Historically, courts have shown an alarming capacity to award larger sums than could reasonably have been expected. Any estimate for tort claims needs to take this into account. The arithmetic is rather simple, even if the parameters may be highly uncertain.

The first step is to consider the likely outcome of the claim. There are three main elements.

When is it likely to be resolved (including appeals)?

How likely is it that the claim will succeed?

What will it cost if it does?

For the sort of claim we are considering, the answers to these questions will be highly subjective. It may be appropriate to consider a variety of possible outcomes and attach a probability to each. It is usually easiest to assess the cost in terms of current standards of court awards.

Next, we construct an "expected" (in the statistical sense) payment pattern and select inflation and discount assumptions, to give the present value of the projection.

In Australia, there is a well-established track record of escalation of tort awards at perhaps 8% to 10% per annum more than wage inflation when averaged over many years. Allowing for a gap of between 2% and 3.5% between interest and wage inflation, it is clear that, if long settlement delays are expected, there needs to be a substantial uplift to any case estimate which is expressed in terms of current court standards.

Conclusion

The aim of general purpose accounts should be to fairly represent the affairs of the entity reported on. This is also true of a range of special purpose accounts, including statutory accounts, which are available to the general public. Historically, accounting practice has developed in the context of a deterministic world-view. This is not inappropriate for enterprises where the main elements of the accounts are not subject to great uncertainty. For insurance, however, uncertainty is the very essence of the business and accounting rules which work well in a deterministic world break down. It is necessary to account for the value of uncertainty and to recognise that there is a tension between solvency and profitability measurements.

It is the contention of this paper that the best estimate to be shown as a provision for uncertain liabilities is a utility estimate, reflecting the generally risk averse attitude of the market. Such an estimate properly reflects both the value of expected future payments and the value of the uncertainty surrounding those payments. While there is a theoretical basis for deriving utility estimates, some of the parameters for that derivation are not well defined and, as a practical matter, approximate and subjective methods may need to be used.

A suitable compromise between the measurement of solvency and profitability is provided by a time series approach to estimation, in which successive valuation bases vary to the extent justified by the additional experience arising between valuations.

It is also vital, for a good understanding, that no material factor in the experience be ignored. In particular, there should be full allowance for inflation and other sources of claim escalation, and projected payments should be discounted on a basis consistent with the asset valuation adopted. The most natural basis for the valuation of assets is market value.

GLOSSARY

Bornhuetter-Ferguson

This is a valuation technique which is widely used for non-proportional reinsurance and other classes where the experience is too variable to support any reasonable estimate of the IBNR liability (including development of reported claims). It is described by Bornhuetter and Ferguson in *The Actuary and IBNR, PCAS LIX*, 1972.

The first step is to estimate how the reported incurred cost develops over time. This pattern is expressed as the fraction of the total claim cost which remains IBNR at each duration.

The second step is to estimate the expected ultimate loss ratio for each cohort of business, at the time it was written. These two ratios are multiplied to give the expected IBNR cost at a particular duration as a fraction of the premium. This is added to the case estimates for reported claims to give the total outstanding liability.

This result progressively modifies the claim cost, as it should have been expected when the business was written, as better information comes to hand in the form of claim payments and case estimates.

Kalman Filter

This is a technique for updating previous estimates to take new observations into account, so as to minimise the uncertainty of the new estimate. In the simplest case, where the observations show random variation about a fixed, but unknown, expected value, the updated estimate is a weighted average of the previous estimate and the latest observation, weighted in inverse proportion to their respective variances. This can be written

$$E_t = (1 - Z_{t-1}) \times E_{t-1} + Z_{t-1} \times A_{t-1}$$

where

$$\frac{Z}{(1 - Z)} = \frac{\sigma_A^2}{\sigma_E^2}$$

More generally, the same formula applies, but the arguments are matrices.

Bayesian Estimation

This is essentially the same concept as the Kalman Filter but is commonly thought of where the previous estimate may be an *a priori* estimate rather than one based on direct observation of the variable to be estimated. In this case the new estimate is a weighted average of the *a priori* estimate and the observed value(s), weighted in inverse proportion to their uncertainty. This latter term is used, rather than variance, because of the subjective nature of the weights used.

Utility

Utility theory is a branch of economics which attempts to provide a mathematical explanation of the choices which economic entities make in the allocation of scarce resources. The idea is that the entity has a sliding scale of utility which reflects the attractiveness of different levels of monetary wealth.

If, as is thought to be usual, the entity is risk averse, the utility of an extra dollar decreases as its total wealth increases. Conversely, a risk seeking entity ascribes greater utility to the last dollar than the first.

Utility is not particularly helpful in a deterministic context. If the amounts are known, more is always more than less. Where it comes into its own is when uncertain amounts must be compared. If the probability distribution of the outcome and the utility function are both known, we can calculate the expected utility. The course to follow is the one which leads to the highest expected utility.

We can also invert the utility calculation to find the fixed value which has a utility equal to the expected utility of an uncertain outcome. This is the *utility value* of that uncertain outcome.

Variability and Uncertainty

To a large extent, these terms are used interchangeably. There is, however, a distinction which may be drawn.

A quantity is variable if it may randomly take any one of a range of possible values. in principle, a variable quantity has a probability distribution. Variability essentially refers to future events.

A quantity is uncertain if its value is not known. This includes variability, but also extends to quantities which are knowable. In principle, this second aspect of uncertainty is not associated with a probability distribution, but it is often helpful to describe uncertainty in terms of probabilities.

For example, between the survey period and the date it is released, the Consumer Price Index is a fixed quantity which is uncertain. It is not variable, because the information on which it is based exists. In this case, the value is uncertain because that information is not generally available.

One-off and Time Series Estimation

By the term one-off or isolated estimate, I mean an estimate which is formed on the basis of the information available at the time the estimate is made, but with no reference to previous estimates.

A time series estimate, on the other hand, also has regard to the previous estimates including, in the context of this paper, *a priori* estimates based on no direct observation of the cohort of business being estimated.

When isolated estimates are highly uncertain, it is likely that they will also be highly erratic, jumping up and down from time to time for no particularly compelling reason. Time series estimates are much more stable, not because they are any less certain, but because each new estimate only departs from the previous estimate to the extent justified by the reliability of the corresponding isolated estimate.

APPENDIX

Utility Estimation

If a particular individual has the utility function $u(w)$, where w is wealth, and has initial wealth W_0, then that individual will react in the same way to a fixed provision P and to an uncertain liability with probability distribution function $L(x)$, of the utility to that individual of W_0 - P is equal to the expected value of the utility of W_0 minus the liability or, in mathematical terms

$$u(W_0 - P) = \int_{-\infty}^{+\infty} u(W_0 - x)\,\mathrm{d}L(x) \tag{A.1}$$

Perhaps the most common function used by economists for this purpose is

$$u(w) = -\mathrm{e}^{-rw} \tag{A.2}$$

where r is referred to as the coefficient of risk aversion. If we substitute (2) in (1), the left hand side becomes

$$\begin{aligned} & -\mathrm{e}^{-r(W_0 - P)} \\ = & -\mathrm{e}^{-rW_0}\,\mathrm{e}^{rP} \end{aligned} \tag{A.3}$$

and the right hand side

$$\begin{aligned} & \int_{-\infty}^{+\infty} -\mathrm{e}^{-r(W_0 - x)}\,\mathrm{d}L(x) \\ = & -\mathrm{e}^{-rW_0} \int_{-\infty}^{+\infty} \mathrm{e}^{rx}\,\mathrm{d}L(x) \end{aligned} \tag{A.4}$$

so that

$$\mathrm{e}^{rP} = \int_{-\infty}^{+\infty} \mathrm{e}^{rx}\,\mathrm{d}L(x) \tag{A.5}$$

or

$$P = \frac{1}{r}\ln\left(\int_{-\infty}^{+\infty} \mathrm{e}^{rx}\,\mathrm{d}L(x)\right) \tag{A.6}$$

Provided r is positive, this is greater than the expected value of the liability, which is

$$\int_{-\infty}^{+\infty} x\,\mathrm{d}L(x) \tag{A.7}$$

Normal Distribution

If we now assume a normal distribution with mean μ and standard deviation σ, we have

$$L(x) = \frac{1}{\sqrt{2\pi}\,\sigma} \int_{-\infty}^{+\infty} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \, dx \qquad \text{(A.8)}$$

Substituting this in (5), we get

$$\begin{aligned}
e^{rP} &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} e^{rx} \, e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \, dx \\
&= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} e^{-\frac{1}{2\sigma^2}(x^2 - 2x(\mu + r\sigma^2) + \mu^2)^2} \, dx \\
&= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} e^{\mu r + \frac{1}{2}r^2\sigma^2} \, e^{-\frac{1}{2}\left(\frac{x - (\mu + r\sigma^2)}{\sigma}\right)^2} \, dx \\
&= e^{\mu r + \frac{1}{2}r^2\sigma^2} \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} e^{-\frac{1}{2}\left(\frac{x - (\mu + r\sigma^2)}{\sigma}\right)^2} \, dx \\
&= e^{\mu r + \frac{1}{2}r^2\sigma^2}
\end{aligned} \qquad \text{(A.9)}$$

so that

$$P = \mu + \frac{1}{2} r\sigma^2 \qquad \text{(A.10)}$$

Other Distributions

More generally, we note that

$$\ln\left(\int_{-\infty}^{+\infty} e^{rx} \, dL(x) \right) = K(r) \qquad \text{(A.11)}$$

where $K(r)$ is the cumulant generating function

$$K(r) = \kappa_1 r + \frac{1}{2}\kappa_2 r^2 + \frac{1}{3}\kappa_3 r^3 + \ldots \qquad \text{(A.12)}$$

we can substitute (11) and (12) in (5) to obtain

$$P = \kappa_1 + \frac{1}{2}\kappa_2 r + \frac{1}{3}\kappa_3 r^2 + \ldots \qquad \text{(A.13)}$$

and, since the first three cumulants are equal to the first three central moments,

$$P = \mu + \frac{1}{2} r\sigma^2 + \frac{1}{3} r^2 \gamma \sigma^3 + \ldots \qquad \text{(A.14)}$$

where μ, σ and γ are the mean, standard deviation and skewness respectively.