Institute of Actuaries of Australia

RISK MARGINS FOR OUTSTANDING CLAIMS LIABILITIES IN HEALTH INSURANCE

Adam Searle and Dimity Wall

Presented to the Institute of Actuaries of Australia
Biennial Convention 23-26 September 2007
Christchurch, New Zealand

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: actuaries@actuaries.asn.au Website: www.actuaries.asn.au

Table of contents

Key words

Health insurance, risk margin, Bayesian, reserving method

Purpose

This paper examines a method for calculating risk margins for health insurance and suggests benchmark risk margins intended to achieve a 75% probability of adequacy.

Synopsis

Many of the current risk margin approaches used by actuaries in other fields rely on assumptions that are not typically valid for health insurance. In particular, assumptions that there are no calendar period effects (those that occur along the diagonal of a typical claims triangulation) or that development factors in subsequent periods for the same cohort are independent, fail spectacularly for health insurance portfolios.

Health insurance is so short-tailed that approximately 35% of claim payments are made within the month they are incurred and 90% of payments are typically made within two months. Changes in the claims processing delay of even a few days at month end can significantly alter the development pattern – with a corresponding catch-up in the next development period. This creates both calendar period effects and significant dependency between subsequent development periods. The volatility in the chain ladder development factors is predominantly caused by changes in claims handling (both provider and insurer) and does not generally suggest that the ultimate liability for individual treatment (or ‘service’) months is particularly volatile.

A typical approach to setting risk margins for health insurers is to look at the historical accuracy of outstanding claim liability projections against actual payments for a number of periods, and to use this as a guide to set risk margins.

This paper formalises a similar approach by exploring a mechanical reserving method and comparing its result against actual payments. We make use of the underlying predictability of health insurance claim costs, to measure and then remove seasonality. We then fit trend lines to the ‘de-seasonalised’ average monthly claim costs per policy. The stability of the claims cost means that the underlying trend is a very accurate indicator of the ultimate liability.

We make use of a Bayesian approach, giving weight to the underlying trend in more recent periods and a chain ladder projection for the older periods. This is similar to a Bornhuetter-Ferguson (“BF”) approach, except our weighting towards the ‘prior’ (the underlying trend) is much greater in recent periods than the BF method would suggest. This is due to the accuracy of the underlying trend and the instability of the chain ladder method for recent periods.

We assess the historic deviations between the actual experience and the ultimate liability projections generated as above for a number of valuation periods. Assessing the reserve generated at each of these valuation dates against the actual outcome (with hindsight) allows an assessment of the accuracy of the model, and calculation of the percentage loading to the central estimate that is required to target various probabilities of sufficiency.

Our analysis used historic payment and exposure data from 19 Australian health insurers. The application of this methodology and standardising of various assumptions across a variety of insurers allows the assessment of benchmark risk margins. This paper presents the benchmarks we have derived at the 75% probability of adequacy.

The paper also provides an indication of the risk margin that might be appropriate in the Liability Adequacy Test (“LAT”) – an assessment of the adequacy of unearned premiums.

Introduction

The requirement for health insurers in Australia to hold risk margins for outstanding claims liabilities on the balance sheet arose with the requirement that reporting for health insurers must be in accordance with AASB1023 ("General Insurance Contracts") from annual reporting periods beginning on or after 1 January 2005.

Although many insurers were implicitly holding risk margins in provisions for outstanding claims liabilities prior to this date, AASB1023 requires that:

'***The financial report shall disclose in relation to the outstanding claims liability:***
(a) the central estimate of the expected present value of future payments for claims incurred;
(b) the component related to the risk margin;
(c) the percentage risk margin adopted in determining the outstanding claims liability […]);
(d) the probability of adequacy intended to be achieved through adoption of the risk margin; and
(e) the process used to determine the risk margin'

There is no prescribed probability of adequacy, however it is noted that AASB1023 states that:

'Risk margins adopted for regulatory purposes may be appropriate risk margins for the purposes of this Standard, or they may be an appropriate starting point in determining such risk margins.'

In practice it appears that many health insurers have chosen to adopt a probability of adequacy of 75% which is the same as the APRA requirement for general insurers.

The potential share market listing of three health insurers may lead to increased scrutiny of health insurers' financial statements, and the assessment of risk margins may become a more significant and topical issue than it has been historically. The impact of Solvency II, IFRS and LAT requirements are also leading to an increase in the attention given to risk margins.

We would like to express our thanks to all the insurers who allowed us to use their data in our analysis. Unfortunately, thanking the insurers by name would be likely to result in the loss of their anonymity….

We are able to express our thanks to various people who provided comments on this paper, including David Torrance, Jefferson Gibbs, Ben Ooi and David Whittle.

The availability and suitability of current methods

The nature of health insurance

Health insurance outstanding claims liabilities are very short-tailed by nature, and many of the risk margin estimation techniques used by actuaries in general insurance are not applicable to health insurance. Case estimates on outstanding claims are rare, and methods used to estimate the outstanding claims provision rely predominantly on analysis of paid development patterns through use of the chain ladder method. Reasonableness adjustments are based on an analysis of projected and historical burning costs (ultimate liability divided by exposure). The exposure measure commonly used is based on the number of insured adults and is typically referred to as the Single Equivalent Unit ("SEU")).

The development pattern is short-tailed for health insurance due to the nature of the benefits insured. For typical general insurance contracts, once an event has occurred (or possibly been identified, or reported, in some instances) that gives rise to a liability, the insurer generally picks up the cost of all future associated payments. For example, an accident at work may involve a series of regular payments in the short term while an employee is unable to work, payment of the cost of any treatment or rehabilitation, and potentially a significant lump sum settlement many years down the line if the employee is never able to return to work.

Health insurance is different in that cover is only provided if an individual is insured with a particular insurer at the date of the treatment, no matter the date of cause of the injury or necessary treatment. The equivalent 'date of occurrence' for health insurance is therefore the date of treatment and the only delay involved is that between the treatment date and settlement of the cost of that treatment with the hospital / service provider. There tend to be fewer disputes over the cost and acceptability of claims than there are in general insurance, and it would be very unusual for a case to be resolved after a prolonged legal battle as occurs in a variety of circumstances in general insurance.

The delay to settlement varies according to the type of benefit – hospital, ancillary and medical. However, these are all very short-tailed when compared to most general insurance classes, and the limitations in standard risk margin methods applies to each of them. The graph below shows a typical development pattern for health insurance claims.

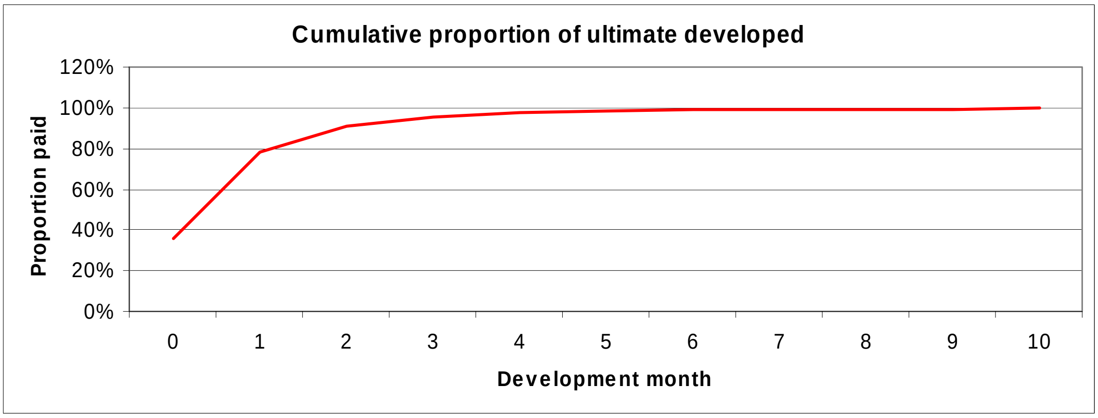

As can be seen from the graph, approximately 35% of claims (by both number and amount) are paid during the month they occur (the treatment month), 80% are paid by the end of the following month and 90% are paid within two months.

Volatility in the claims payment pattern

Processing delays can significantly alter the payment pattern and claims backlogs varying by a week or more are not uncommon. This can be caused by holidays (for example during December / January and other school holidays) or other events. More significant deviations in claims processing can result from changes in claims handling practices such as the implementation of electronic payments or IT system changes.

As the payment pattern is so short, changes in claims backlogs that would be of less significance when setting the reserves for many general insurance classes become more significant for health insurers. The typical outstanding claims reserve for a health insurer is of the order of one month's claim payments. Therefore an additional week's delay in claims processing will alter the required reserve by approximately 25%, and the typical development pattern above will not be appropriate for assessing the reserves. A more significant delay in claims backlogs (such as an IT implementation) may result in the reserves increasing by a factor of 200% or more, and the appropriate development pattern would then be very different to normal.

The chain ladder method frequently produces an unreasonable estimate of ultimate cost per SEU for the most recent service months (accident or occurrence month in general insurance parlance). However the chain ladder method is generally considered to be suitable for use without significant adjustment when estimating incurred amounts for more mature service months.

Given the fluctuations in the development pattern in the early months and the impact this has on the required level of reserves, an assessment of the risk margin by considering the variations in the development pattern would lead to a large risk margin. This is typically the result when the standard methods used to assess risk margins in general insurance are applied to health insurance data. This includes methods such as Mack, Bootstrapping and Stochastic Chain Ladder.

The volatility in the chain ladder development factors is predominantly caused by changes in claims handling (both provider and insurer) and does not generally suggest that the ultimate liability for individual service months is particularly volatile. In fact, the ultimate liability is exceptionally stable due to a relatively high claims frequency and the absence of very large claims.

The definition of a 'claim' can be difficult in health insurance as each individual payment typically relates to a treatment and may be further split according to various components such as theatre fees, surgeon fees, room costs, anaesthetist fees, drugs, etc. Individually each of these is typically of very small cost – even the cost of new drugs is not large compared with large general insurance claims. For claims management purposes however, a claim is often defined as the total cost of a series of treatments relating to the event. This makes sense from a case management point of view, but each treatment within the same event may have a different liability date when determining reserves.

Stability of the ultimate liability

The graph below shows the stability of the burning cost by service month for our 'example insurer'. The 'example insurer' has been derived by adding together the data from a variety of actual insurers, including a scaling of some of the insurers to ensure the data presented here is anonymous.

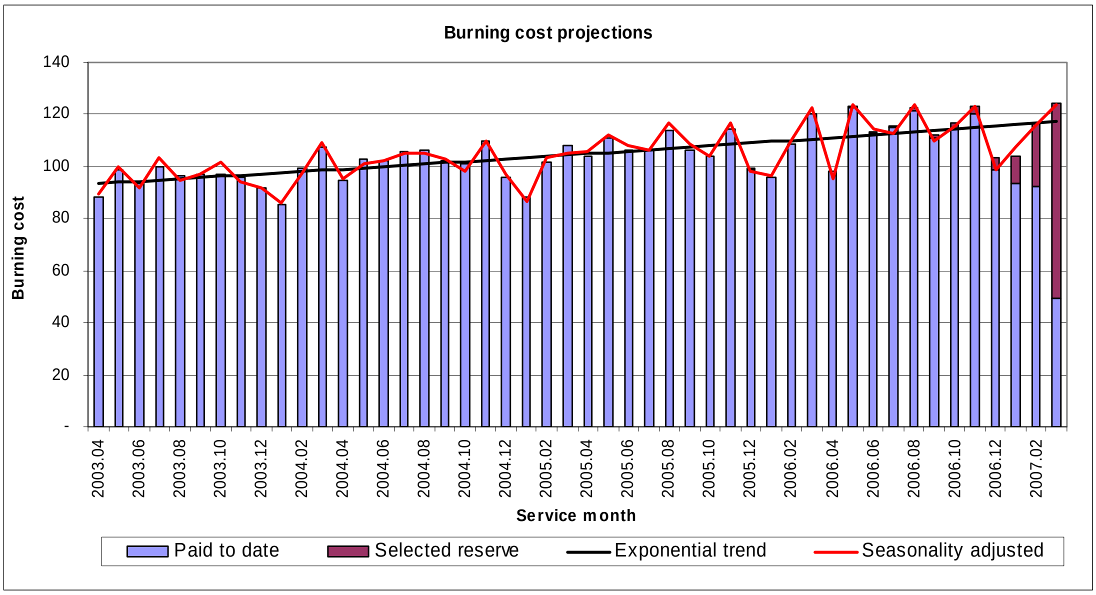

The detail of how the trend lines are determined is discussed in more detail below. However, it is clear that the 'seasonally-adjusted' trend line proves to be an exceptionally good fit to the actual historic data (shown as solid columns).

Given the stability of the ultimate liability, we do not believe that the volatility suggested by applying standard risk margin techniques is appropriate.

Approach

A typical current approach to setting risk margins for health insurance outstanding claims liabilities is to use as a guide the accuracy of outstanding claim liability projections against actual payments for a number of historic valuations.

This paper adopts a similar approach in principle, but formalises the approach. Adoption of a detailed mechanical reserving method allows a less subjective assessment of the required risk margin. Our application of this method across approximately 30 valuations for 19 insurers of varying sizes, results in a set of benchmarks and the suggestion of a methodology that can be used to assess both the central estimate of the liabilities and the risk margin for an individual insurer.

Actuarial judgement

In practice, the chain ladder method is a commonly used method to assess outstanding claims provisions with adjustments for the most recent service months. Actuarial judgement is often used when determining the weight to be applied to the chain ladder method and when adopting a final provision. The methodology adopted for the purposes of this paper is a mechanical approach which attempts to replicate the application of actuarial judgement.

It may be expected that the application of actuarial judgment should result in a more accurate provision than the mechanical method adopted here. This would be due to the actuary taking into account additional factors, such as the claims backlog state, when assessing an appropriate reserve. As a result, the risk margins derived here may be at the slightly conservative end of a range of reasonable risk margins. We are not suggesting that this model can be a complete replacement for an actuary's judgement when assessing reserves. However, subjective judgement is a risky process in itself, and may not necessarily lead to reduced volatility….

It may be appropriate for an individual insurer to have a higher risk margin than those suggested in this paper. This could be the case if their individual circumstances resulted in periods of unusual increased volatility or if claims processing practices varied considerably. This would result in the chain ladder method (which is still relied on reasonably heavily to assess the reserves) being even more unstable for that insurer than is typical in our analysis.

Data

Claims triangulations of 19 health insurers of varying size were used for our analysis. In general, the data included at least four historic years up to March 2007 with payment triangulations split by month of service and delay to payment month.

Although data was provided separately by hospital, medical and ancillary claims types, these triangulations were grouped for our analysis. This reduced the volatility of our reserve projections automatically for any diversification benefits between the various claim types and avoided any subjective assessment of correlation factors. There is currently no requirement to report outstanding claims provisions or risk margins separately by claim type for regulators or accounting standards.

SEU exposure information was obtained for insurers from PHIAC B data.

No case estimate data has been used in our analysis. Case estimate data is not typically used (or available) for health insurers. Given the difficulties in defining a claim and assessing an appropriate case estimate for the outstanding claims liability component of a ‘claim’, it is not obvious how case estimation would be of much benefit to an actuary when assessing reserves.

Our adopted method

The method we have used to determine the outstanding claims provisions reduces reliance on the assumption that the chain ladder development factors are stable. It places heavy reliance on the predictability of the ultimate monthly claim costs for health insurance.

An overview of the steps involved in our analysis are set out below, with more detail provided in the following section.

1. Calculate a monthly by monthly development triangle of paid burning costs (paid to date amounts divided by the SEU for each service month)

2. Project ultimate burning costs using a standard paid chain ladder approach

3. Remove much of the apparent volatility in the ultimate burning costs by standardising the projections for the number of working days in each service month

4. Fit an exponential curve through the standardised ultimate burning costs

5. Calculate any remaining seasonality in the ultimate burning costs for each service month

6. Calculate a ‘seasonally-adjusted’ trend line of ultimate burning costs for each service month by adjusting the exponential fitted line for both working days and seasonality

7. Repeat steps 1-6 for each of the historic valuation months to obtain a set of projections as at various historic dates

8. Given the various historic projections, determine an appropriate weighting between the chain ladder and seasonally-adjusted trend for each development delay

9. Calculate errors between the seasonally-adjusted trend line and the actual ultimate cost and examine these to assess reasonableness

10. Repeat steps 1-9 for each of the 19 different insurers

11. Compare the weightings for each development delay for each of the insurers and assess an appropriate set of general weightings that can be used

12. Calculate the selected reserve at each historic valuation date for each insurer using the general weightings

13. For each insurer, calculate the percentage deviation of the actual reserve required (using hindsight) from the selected reserve at each historic valuation

14. For each insurer, assess the percentage loading required for the reserve to be sufficient in 75% of the historic valuations and compare these across insurers by size of the outstanding claims reserve to determine benchmark risk margins

Each of these steps is considered in detail below.

1 Calculate a monthly by monthly development triangle of paid burning costs (paid to date amounts divided by the SEU for each service month)

For each development month, the cumulative claim payments were divided through by the exposure for that service month. This produced a triangle of 'paid burning costs' on which we based our analysis. The major benefit of using this rather than the raw payments data is that it removes any changes in the ultimate claims projections purely due to changes in membership.

2 Project ultimate burning costs using a standard paid chain ladder approach

The development factors were adopted directly from the underlying data for each insurer. We considered using a standard development pattern across all insurers, but after examination decided that there was sufficient variation between insurers to adopt a different pattern for each.

We recognise that there are some flaws in adopting this development pattern directly. However, the flaws in using the chain ladder method are most significant for the most recent service periods and, as is explained later, the weighting given to the chain ladder method for these periods is quite small.

3 Remove much of the apparent volatility in the ultimate burning costs by standardising the projections for the number of working days in each service month

The number of working days in each month is a good indicator of the ultimate claims cost for a service month. In order to create a more stable underlying trend, we removed the volatility due to working days from the ultimate liability projections for each month.

Essentially, this step involved dividing through the burning cost for each service month by the number of working days in that service month. This resulted in a smoother trend of 'ultimate liability per working day in each service month' as shown below for our example insurer.

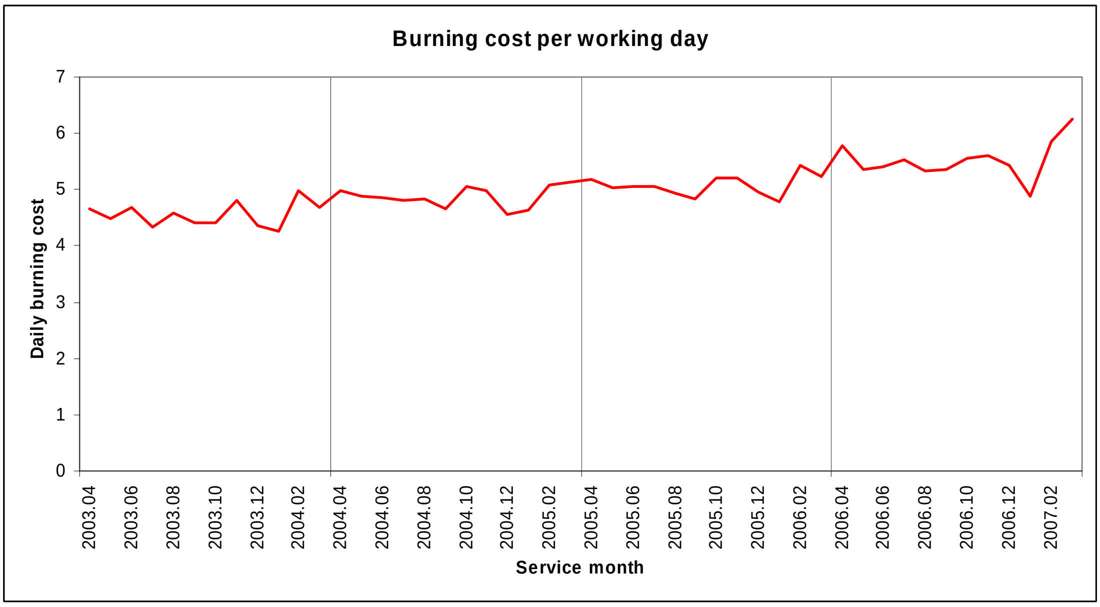

4 Fit an exponential curve through the standardised ultimate burning costs

As can be seen from the graph, some volatility remains in the ultimate liabilities for each month. However, it is clear that there is a fairly stable underlying trend.

We fitted an exponential curve through the standardised ultimate burning costs. An exponential curve was adopted to allow for the effect of compound inflation. When fitting the curve, the data from the most recent five service months were not used as these were believed to be insufficiently developed to be confident that the chain ladder was accurately projecting the ultimate liability. After five months development, the payments are approximately 98% – 99% of the ultimate cost for the month and there is therefore very little remaining volatility in the ultimate projection for the months used in fitting the trend lines.

5 Calculate any remaining seasonality in the ultimate burning costs for each service month

We further reduced the variability of the actual burning costs from the trend line by allowing for the seasonality that exists beyond that explained by the number of working days in the month.

The seasonality could be calculated in a number of ways. Seasonality varies slightly according to the profile of the insurer – for example, due to a different mix of policyholders or differences in the proportion of hospital, medical and ancillary covers. We assessed the seasonality of each insurer separately by considering the percentage deviations for each service month from the fitted trend line. Where there was a consistent deviation in the same service month in each historic year, we interpreted this as seasonality and calculated a seasonality factor by averaging the percentage deviations for all months in the data.

We standardised the seasonality adjustments across the 12 months to ensure the average adjustment we were applying to the underlying trend was nil (to avoid introducing bias in the projections). This standardisation to a nil overall adjustment could be an arithmetic or geometric average depending on how the seasonality factors are applied. We applied a geometric adjustment as our seasonality factors were applied multiplicatively to the underlying trend.

We did compare the seasonality adjustments calculated between insurers to gain reassurance that we were capturing a genuine feature and not merely removing inherent volatility. The seasonality pattern appeared to have the same overall shape and we were therefore comfortable that seasonality was a genuine feature.

The seasonality pattern calculated from the underlying data for our example insurer is set out below. These are multiplicative factors applied to the exponential trend line, after adjusting for the number of working days in each month. Obviously some of these are more statistically significant than others.

Service month	Multiplicative factors
January	92.6%
February	104.4%
March	100.9%
April	106.0%
May	101.3%
June	102.0%
July	99.8%
August	99.5%
September	96.5%
October	101.0%
November	102.4%
December	94.4%

6 Calculate a 'seasonally-adjusted' trend line of ultimate burning costs for each service month by adjusting the exponential fitted line for both working days and seasonality

Our seasonally-adjusted trend line was calculated by applying the seasonality factors to the exponentially fitted trend line and then for each service month multiplying the daily burning cost by the number of working days in that month.

The graph below shows the accuracy of the calculated seasonally-adjusted trend line for our example insurer at the latest valuation date. This does not show residuals for the most recent service months as the payments were too immature for us to be confident that we knew the ultimate claims cost against which to compare our trend line.

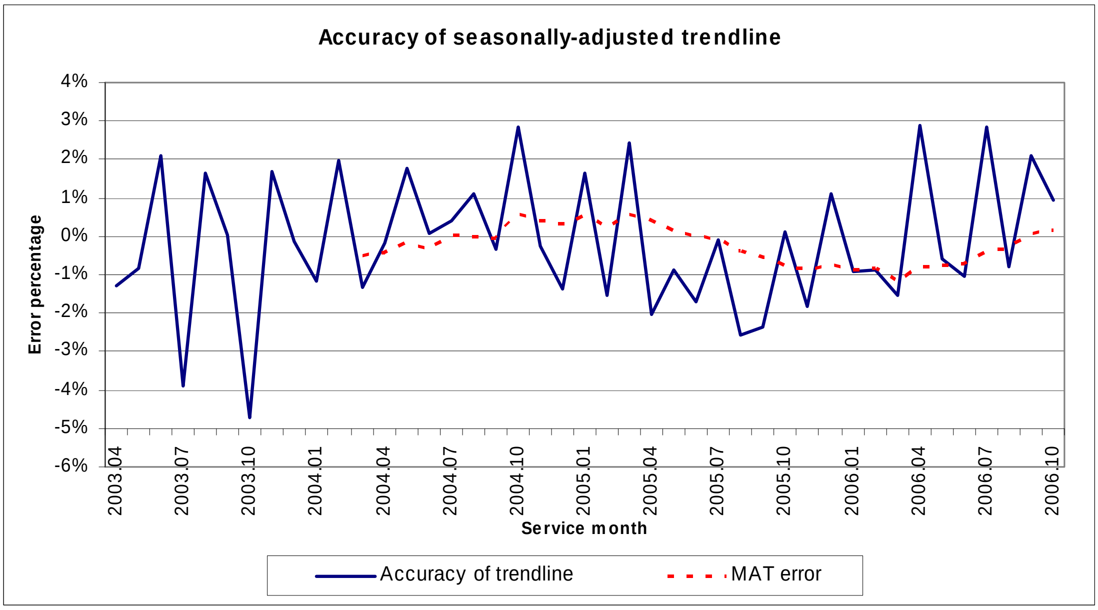

The resulting deviations between the trend and the actual ultimate costs do not appear to show significant bias. The volatility that does exist appears to be fairly random and, in the main, oscillates around a nil deviation. Interestingly, the MAT (moving annual total) trend line shows very little deviation, suggesting that an annual burning cost for our example insurer could be estimated very accurately. This suggests that a risk margin loading for premium liabilities might be very small (expressed as a proportion of the central estimate rather than as an absolute measure). This is examined in more detail later in the paper.

7 Repeat steps 1-6 for each of the historic valuation months to obtain a set of projections as at various historic dates

Repeat all the above steps for each historic valuation month where there is sufficient data. We typically used a minimum of 4 years' worth of data, which enabled us to perform up to 30 valuations (at each of the latest 30 month-ends) for each insurer.

8 Given the various historic projections, determine an appropriate weighting between the chain ladder and seasonally-adjusted trend for each development delay

For each historic valuation month, we have calculated two separate projections of ultimate cost for each service month. One is based on the chain-ladder method and the other is based on an assessment of the underlying trend with adjustments for seasonality and the number of working days in each month.

The next step involved an assessment of what the appropriate weighting is for each development month between these two methods to produce a 'selected' ultimate liability. At each valuation date we have a projection for every development month (i.e. at the 2007.03 valuation month, we have a projection for the 2007.03 service month at development 0 and a projection for the 2007.02 service month at development 1, etc. At the 2007.02 valuation month, we have a projection for the 2007.02 service month at development 0 and a projection for the 2007.01 service month at development 1, etc.). There are therefore up to 30 projections for each of the development months (one from each historic valuation).

Our selected weighting for each development month was assessed by minimising the least squares of the differences between the selected ultimates for each service month at the appropriate development month and the actual ultimates. This produced a series of weightings by development period.

As one might expect the weighting towards the chain ladder method increased as the development period increased and the actual payments became increasingly reliable as a basis for estimating the ultimate liability.

The line on the graph below shows the application of the weightings for our example insurer to the exponential trend line fitted at the most recent valuation date.

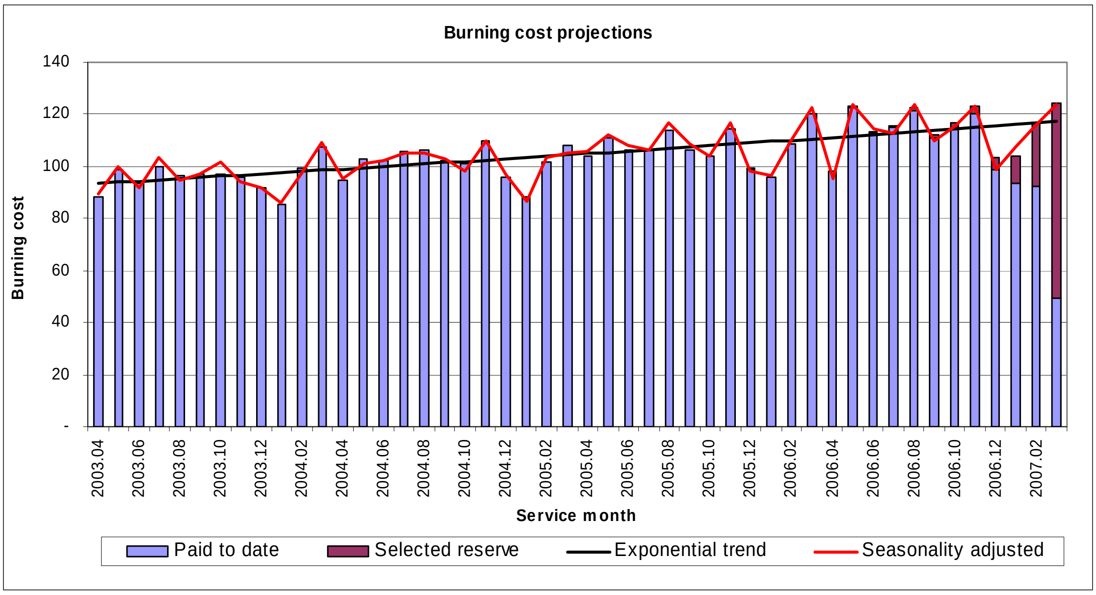

This graph (as shown in the earlier section above) demonstrates how accurately the ultimate claim costs for each service month can be projected.

9 Calculate errors between the seasonally-adjusted trend line and the actual ultimate cost and examine these to assess reasonableness

Our method examined periods which have since reached an ultimate state. This allowed us to examine the errors in our projections (defined as the deviations of our seasonally-adjusted trend from the actual ultimate cost). This included checks to ensure there was no significant bias, examination of the correlation between adjacent months' error terms (serial correlation), and an assessment of the error distribution – including a count of +ve / -ve strings, and total number of +ve versus total –ve deviations.

We performed this in detail during the design phase of our modelling only and on a few insurers' data to ensure that the results appeared reasonable.

The following graph shows the deviations at each delay period (one is obtained for each delay period at each historic valuation point) for our example insurer.

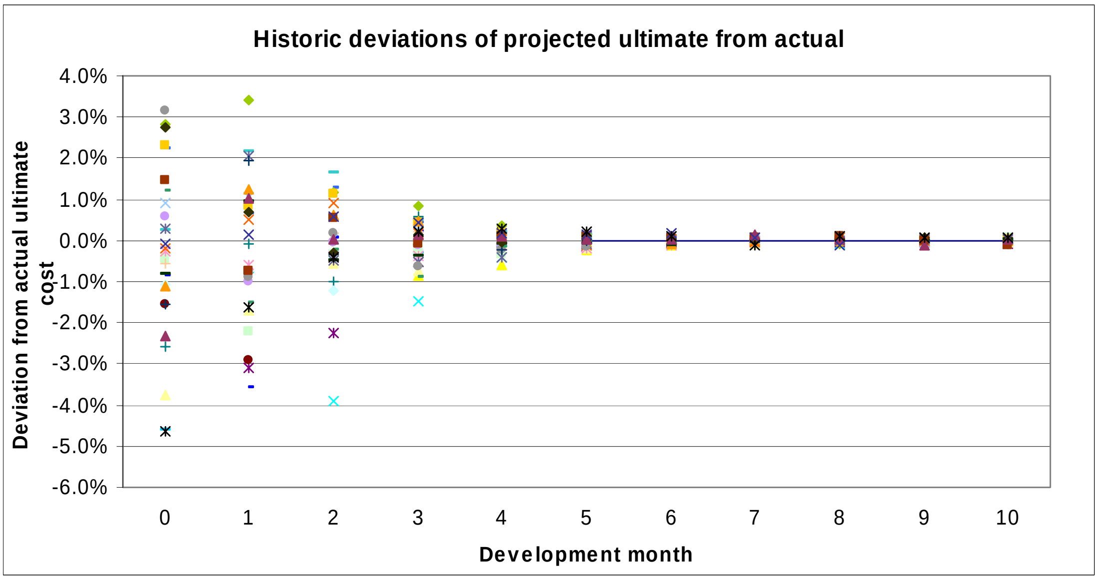

As expected, the selected ultimate is most accurate for more developed periods, with very little error in the projection after 5 months. Even for the early development months, the projection is very accurate relative to most general insurance liability valuations, with all of the projections for the latest service month being within 5% of the actual ultimate cost for that month. As our example insurer was created by combining a number of insurers' data, the volatility shown here is much lower than would be expected for the smaller insurers. However, even for the very smallest insurers we analysed, the most extreme errors in our selected ultimate for delay month 0 were rarely greater than 20%. This is a far more accurate projection than is achieved using purely the chain ladder approach applied to the most recent month.

Step-changes in the claim costs (perhaps due to annual renewals of agreements with hospitals and other service providers or annual reviews of ancillary benefits) did not appear to be very significant for most insurers. However, a potential modification of the methodology would be to allow for these where they are known. This should further improve the accuracy of the seasonally-adjusted trend and may be of most benefit when projecting the future trend line and assessing premium liabilities where recent or future step changes are known to have occurred.

10 Repeat steps 1-9 for each of the 19 different insurers

The above steps were repeated for each of the insurers included in our analysis.

11 Compare the weightings for each development delay for each of the insurers and assess an appropriate set of general weightings that can be used

It would be inappropriate to apply the weightings assessed directly from the insurer itself when assessing risk margins as this would remove some of the inherent volatility and therefore understate the required risk margin. Therefore, we repeated all of the above steps for each insurer and assessed an appropriate weighting for each development period based on a review of the weightings minimising the least squares for each insurer.

It was apparent that the general pattern was for the larger insurers to have greater weightings towards the seasonally-adjusted trend line than the smaller insurers. This makes sense as the larger insurers are likely to have more stable ultimate liabilities by month due to their overall size, so the underlying trend may be more accurate for them in individual months. However, the development patterns are likely to vary for all insurers due to fluctuations in claims processing, which will be an issue irrespective of the insurers' size.

We fitted a logarithmic curve through the weightings by size for each development period. This suggested that appropriate weightings to apply to the trend line over the chain ladder projection were as follows.

	Central estimate of reserve size for outstanding claims ($m)							
Development month	**1**	**2**	**5**	**10**	**20**	**50**	**100**	**200**
4	2.2%	2.9%	3.7%	4.4%	5.0%	5.9%	6.6%	7.2%
3	1.0%	3.7%	7.3%	10.1%	12.8%	16.4%	19.1%	21.8%
2	2.7%	9.5%	18.5%	25.3%	32.1%	41.0%	47.8%	54.6%
1	25.0%	33.4%	44.4%	52.8%	61.2%	72.2%	80.6%	88.9%
0	78.0%	79.6%	81.7%	83.2%	84.8%	86.9%	88.5%	90.1%

Our approach is essentially a Bayesian one, with the trend line providing the prior belief and the chain ladder providing an estimate based on actual data. In principle, this is similar to a BF approach, although our weightings give significantly greater credibility to the trend line than would be the case if a BF method were used.

The BF method gives a weighting to the chain ladder method equal to the expected amount paid to date from the standard development pattern. Given the typical payment pattern for health insurance, this would result in weightings to the trend line of approximately 65% (development 0), 20% (development 1), 10% (development 2), etc.

12 Calculate the selected reserve at each historic valuation date for each insurer using the general weightings

The individual models were then re-worked using the weightings from the general model to calculate the ultimate liabilities at each of the 30 historic valuation periods for each insurer.

13 For each insurer, calculate the percentage deviation of the actual reserve required (using hindsight) from the selected reserve at each historic valuation

A comparison of these 30 projections against the actual outstanding claims liability at those valuation dates (using the benefit of hindsight) enabled us to assess an error distribution for the selected model as percentages of the selected central estimate for each valuation.

It was possible to read off from this distribution what the required percentage loading would be to obtain various probabilities of sufficiency.

15 For each insurer, assess the percentage loading required for the reserve to be sufficient in 75% of the historic valuations and compare these across insurers by size of the outstanding claims reserve to determine benchmark risk margins

We assessed for each insurer the percentage loading that would be required to obtain a probability of adequacy of 75%. A plot of these by the average selected reserve over the period gave a distribution of the risk margin required to obtain a 75% probability of adequacy for insurers of varying size.

Not surprisingly, the percentage risk margin reduces as the insurer size increases. We grouped the insurers according to size and have calculated appropriate benchmark risk margins for four different size categories. These are plotted on a logarithmic chart below.

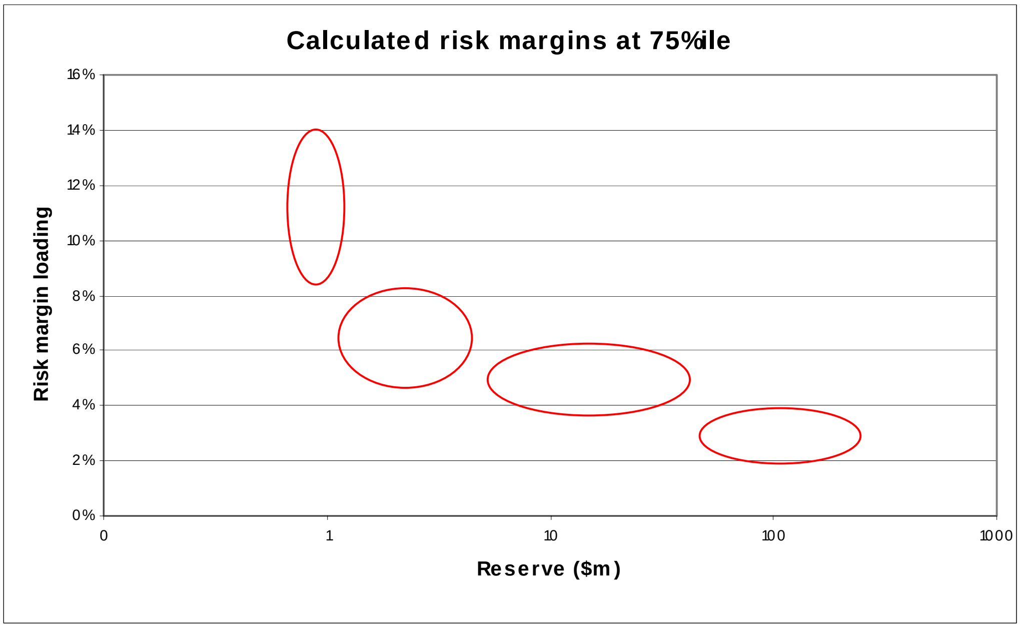

The four circles show the risk margins required at the 75% probability of adequacy for each of our four size groupings. We have not plotted the individual points for insurers as this would remove some of the anonymity of the individual insurers. However, the circles are generally a good fit for the insurers in the relevant size group.

The following table shows our benchmark risk margins by size of reserve to achieve a 75% probability of adequacy.

Outstanding claims central estimate	Benchmark risk margin for health	Tillinghast paper – short-tailed classes
< $1.5m	8% - 14%	20% - 30%
$1.5m - $5m	5% - 8%	20% - 30%
$5m - $50m	4% - 5.5%	8.0% - 18.5%
> $50m	2% - 3.5%	5.5% -13.2%

There is some overlap in the benchmarks between the size groupings and this is to be expected. Generally, we would expect the larger insurers in each group to be toward the lower end of the benchmark risk margin range and the smaller insurers in each group to be toward the higher end. However, there are many other factors which influence the appropriate risk margin as well as size and it is not necessarily inappropriate for an insurer to have a risk margin outside the range we have assessed.

The above table also shows comparisons with the benchmark risk margins proposed in the Tillinghast paper[1] for short-tailed general insurance portfolios. The risk margin benchmarks for health insurance are significantly lower than those for general insurance classes, recognising the greater stability and predictability of health insurance ultimate liabilities.

[1] *Research and Data Analysis Relevant to the Development of Standards and Guidelines on Liability Valuation for General Insurance*

Risk margins for the Liability Adequacy Test ("LAT")

Introduction

AASB1023 requires that:

'The adequacy of the unearned premium liability shall be assessed by considering current estimates of the present value of the expected future cash flows relating to future claims arising from the rights and obligations under current general insurance contracts. If the present value of the expected future cash flows relating to future claims arising from the rights and obligations under current general insurance contracts, plus an additional risk margin to reflect the inherent uncertainty in the central estimate, exceed the unearned premium liability less related intangible assets and related deferred acquisition costs, then the unearned premium liability is deficient. The entire deficiency shall be recognised in the income statement.'

This section sets out our assessment of the risk margin loading (at the 75% probability of adequacy) that would be appropriate to use in the LAT. The data we used was identical to that in the assessment of the risk margins for outstanding claims liability.

Industry interpretation appears to be that health insurers are required to perform the LAT assuming that all current members will remain with the insurer until the next 1 April. At the typical balance date of 30 June, there will therefore be a requirement to calculate the LAT assuming the current membership remains in force for the following nine month period.

This is a simplification of the actual position where some members pay annual premiums and other members will cancel before 1 April. However, for the purpose of our analysis, we have assumed that the LAT is calculated based on a flat membership remaining in force for the following nine month period.

Adopted methodology

For each insurer, we calculated the trend line using data up to each historic valuation point. Strictly, we only used data for service months up to five months preceding the valuation date to ensure that the ultimate projection for those months was accurate enough on which to base a trend.

We then projected the seasonally-adjusted trend line for the nine months following the valuation date. This was simply a case of extending the exponential curve and applying the relevant loadings for working days and seasonality factors.

Having projected this data, it was possible to compare (with hindsight) how accurate this nine month projection was.

The data was more limited for this analysis than for the outstanding claims analysis. This is because we lost the nine most recent valuation dates for each insurer (for the nine most recent valuation dates we did not have actual data against which to compare our future projections).

This reduction in the data meant that we could only assess risk margins for premium liabilities across all insurers and we were unable to provide separate estimates by size. We excluded the data from one of the smaller insurers from the premium liability analysis as the reduced data we had was too sparse and over too short a time frame for us to project a trend line that seemed reasonable.

There was some suggestion in the data that the projections for the larger insurers were more predictable and that a lower risk margin may be appropriate. Similarly, higher risk margins for the smaller insurers may also be appropriate. Intuitively this appears to make sense.

However, on the available data we were unable to assess the relativities with accuracy, although we believe this should be possible given more historical data for the insurers or data for a greater number of insurers.

Results

The table below shows the distribution of the error terms across all insurers.

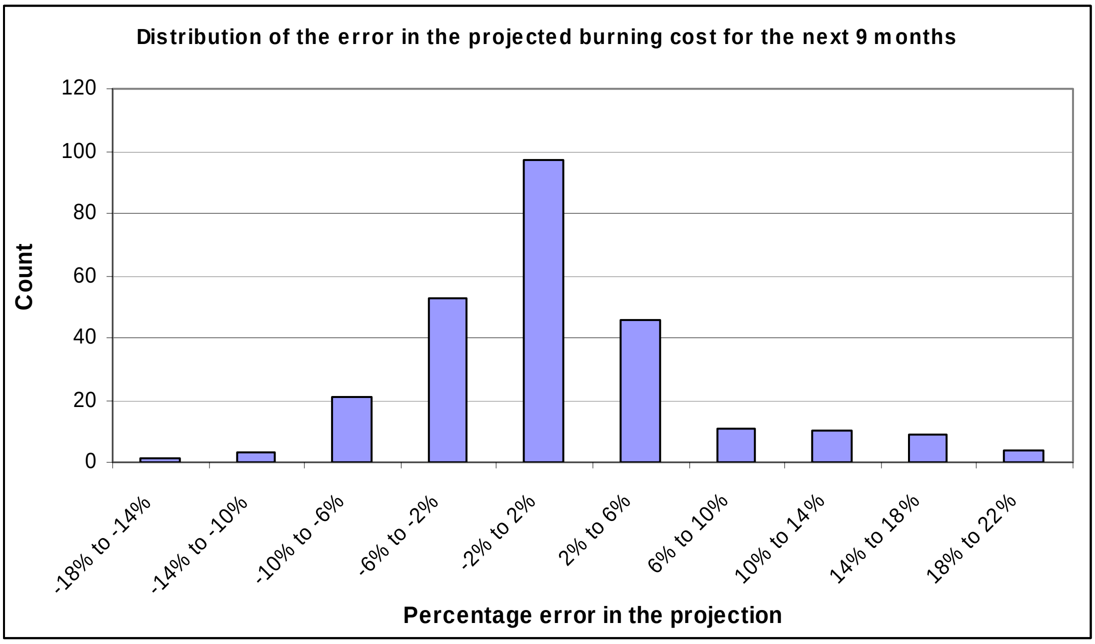

Reading the 75% probability of adequacy from historic projections, suggests that a risk margin loading of 2.8% would be appropriate for the LAT assessed at 30 June. It is worth re-emphasising that this could be significantly higher for smaller insurers (and lower for larger insurers).

Comparison with outstanding claims

It is interesting to note that the risk margin percentage at the 75% probability of adequacy may be lower for many insurers' premium liabilities than it is for the outstanding claims reserve. Initially this may appear strange (and certainly would be strange for many general insurance classes).

This apparent anomaly is due to the risk margin being expressed as a proportion of the outstanding liabilities, rather than the ultimate liability. The ultimate liability for each month, as discussed above, is very stable and this can be projected very accurately. For premium liabilities, the risk margin is expressed as a proportion of the error in the ultimate liability projection, which results in a small percentage loading.

For outstanding claims, the risk margin is expressed as a proportion of the outstanding claims component only (rather than as a proportion of the ultimate liability). As the claims are settled very quickly, the vast majority of the outstanding claims reserve relates to the most recent few months. Even though much of the ultimate liability is settled for recent months, our analysis on outstanding claims suggests that the underlying trend should still be given significant weight when setting the reserves. The uncertainty in the underlying trend therefore translates to a similar uncertainty in the outstanding claims reserve as an absolute amount. Expressed as a proportion of the outstanding claims liability, this results in a larger percentage risk margin.

Limitations

Probability of adequacy

Our method is likely to be suitable for estimating risk margins up to around a 90% probability of adequacy. For risk margins in excess of this level, it is likely that consideration would need to be given to extreme events that are unlikely to be evident in the historic data we have used.

Data limitations

Our analysis was based on the data for only 19 insurers. While we believe we had enough insurers in each of our four groupings to assess benchmark risk margins, an analysis covering all insurers in the Australian market would be likely to give a more complete picture.

Restrictions due to confidentiality

We have decided not to discuss the variability within each group due to potential issues over the confidentiality of individual insurers.

Alternative methods for assessing the central estimate

We believe the risk margin loadings we have presented in this paper are appropriate when using the methodology described above to assess the central estimate of the outstanding claims reserve. A different methodology adopted for the central estimate may result in a higher (or lower) assessment of uncertainty. Our benchmarks are intrinsically based on the reserving model used to assess the central estimate. Where an alternative method is used to establish the central estimate, our benchmarks may be less appropriate. These benchmarks should therefore be used cautiously if a different method is adopted for assessing the central estimate.

Use of data post balance-sheet

We understand that some insurers in the marketplace currently assess their outstanding claims and risk margins using data post the balance sheet date. Given the short tail of the payment pattern, an extra month or two of data would significantly reduce the risk margin required (as an absolute amount – although it would increase as a percentage of the amount still outstanding). However, there is some question over whether this assessment of the risk margin represents a true and fair value of the liabilities as at the balance date, as the uncertainty of the reserve did exist at the balance date and therefore it could be argued that this should be reflected in accounts drawn up at that date. The risk margins calculated in this paper are not intended to be used as benchmarks when using data after the valuation date, although the methodology presented here could be extended to assess these.

Appropriateness of the benchmark for LAT

Our analysis of the risk margins appropriate for the LAT only provided a single benchmark comparison irrespective of size. This figure is unlikely to be suitable as a risk margin loading for small insurers where the variability in the projection is significantly higher.

In addition our analysis considered only the predictability of ultimate benefits (hospital, medical and ancillary combined) and did not consider the impact of risk equalisation and management expenses.

Bibliography

The following papers were reviewed during the preparation of this paper:

A Practical Guide To Measuring Reserve Variability Using: Bootstrapping, Operational Time And A Distribution-Free Approach	*Julian Lowe*	*1994*
Research and Data Analysis Relevant to the Development of Standards and Guidelines on Liability Valuation for General Insurance	*Robyn Bateup and Ian Reed*	*2001*
Stochastic Claims Reserving in General insurance	*PD England, RJ Verrall*	*2002*
Assessing & Monitoring Insurance Liability Uncertainty	*Laurel Kong, Felix Tang, Stephen Robertson-Dunn, Elaine Collins*	*2005*
Report from the Risk Margins Taskforce	*Karl Marshall, Conor O'Dowd, Brett Ward*	*2006*