Institute of Actuaries of Australia

Claim Segmentation, Valuation and Operational Modelling for Workers Compensation

Prepared by Richard Brookes, Anna Dayton
and Kiat Chan

Presented to the Institute of Actuaries of Australia
XIV General Insurance Seminar 2003
9-12 November 2003

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: insact@actuaries.asn.au Website: www.actuaries.asn.au

1 Background

As at 30 June 2002, there were nine licenses for insurers to undertake Workers Compensation business in NSW. These insurers are expected to actively manage the claims on their books and are provided with a range of incentives from NSW WorkCover to do so.

In particular, long tail claims are a significant element of the NSW Workers Compensation liability and as such NSW WorkCover is keen to ensure that the claims in the "tail" are actively managed in such a way as to reduce the Scheme's liability. For the last four financial years (ie starting as at 1 July 1999) WorkCover has been providing financial incentives for effective performance on these claims.

The principle underlying one of the insurer remuneration incentives is the Tail Liability Incentive measure, which aims to measure the change in each insurer's liability for the outstanding "tail" claims between the start and the end of the financial year and provide as remuneration a share of any savings. For this incentive measure, tail claims are defined as being any claims that are greater than one year since injury as at the start of the financial year.

In broad terms, the saving in the tail liability is defined as:

- The opening liability (in dollars as at the end of the incentive year, with the same economic assumptions as the closing liability); less
- Actual payments made in the financial year (in dollars as at the end of the incentive year); less
- The closing liability.

If a particular insurer has a positive result, they have achieved a saving and are entitled to an incentive payment.

The need for an objective allocation of the Scheme liabilities to the insurer level drove the need for a valuation method based more closely on claim characteristics. We describe below a methodology based on claim segments and transitions between these segments. This model structure is not new and can just be considered as a coherent collection of traditional actuarial PPAC and PPCF models. However, for the purpose under discussion we believe that the framework improves on an aggregate valuation methodology because it:

- allows for the transparent use of each insurer's claim profile; and
- is more responsive to this different claim profile and different management practices.

It also explicitly allows for interaction between payment types for the same claim. For instance, a claim which is currently receiving medical payments may or may not be receiving weekly benefits or be expecting a common law payment. As we shall see later this turns out to significantly affect the expected level of future medical payments and the expected time to closure. Finally, the methodology also translates into a clear methodology for monitoring future transitions and payment levels and, in particular, seeing how these react to insurers' claim management initiatives.

2 The methodology

The valuation methodology is based on separating claims into groups of claims with similar characteristics. We will refer to these groups as "segments" or "states". At any point in time each claim belongs to exactly one segment. We consider the historical numbers of claims in each segment and movements between segments from quarter to quarter. We model these movements between segments and call them transitions. We then analyse historical payments by claim segment to project future payments by payment type for each segment separately.

The liability is determined by:

- taking current claims by segments and, adding IBNR claims, using the transition rates to project the future segmentation of claims; and
- multiplying the projected number of claims in each segment in each future period by the expected payments per claim.

The result is a series of cash flows which we then inflate and discount to give the overall valuation result.

2.1 The segments

We used a combination of prior knowledge of the Scheme, and trial and error, to group the claims into segments. The resulting segments were:

- Open deafness claims. These claimants generally only receive a lump sum payment for permanent impairment, and pain and suffering.
- Closed deafness claims.
- Open common law claims. These tend to settle for large lump sum amounts shortly before closure but some weekly, medical and other payments are made in the interim.

- Claims that have been commuted but are not yet closed. These claims are usually waiting for legal costs to be paid before closure.

- Active weekly claims. Those claims where the claimants are receiving income compensation. The claimants are either totally incapacitated, partially incapacitated or undergoing retraining.

- Non-active weekly claims, with active medical payments. These claims have an outstanding weekly case estimate. Medical but no weekly payments were made in the quarter.

- Non-active weekly claims, with weekly the largest case estimate. No weekly payments were made in the quarter but there is an outstanding weekly case estimate which is the largest of the case estimates (by payment type) for that claim. These claims can be claims for which weekly payments are yet to commence but, for tail claims, they are normally claims where there has been an interruption to weekly payments, either due to operational delays (for instance, employer reimbursement schedules) or the claimant has returned to work.

- Other open claims. This is the largest group of open claims but they are less financially significant than the other major groups.

- Closed common law claims.

- Closed commuted claims.

- Other closed claims.

These definitions have a strict hierarchy and the definitions above are given from the top of the hierarchy to the bottom. If a claim falls into the definition of a particular segment then it is automatically excluded from the segments below.

The graph shown here illustrates how the claim segmentation described above breaks up the portfolio of claims in the NSW Workers Compensation Scheme that are open as at 30 June 2003:

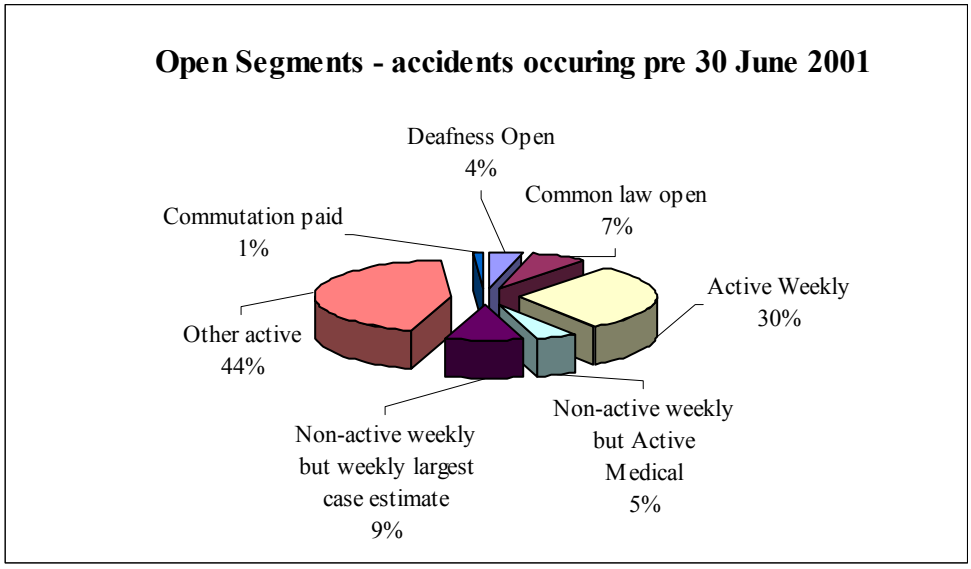

In subsequent work, we derived homogeneous groups of claims statistically using data mining analysis and the resulting groups were a reasonably close match to the segments above.

2.2 Transitions

Having grouped the claims into segments, the next step in the methodology is to determine how claims move between the segments. We ask the questions:

- Where do the claims go to?
- At what rate do they transit?

In order to properly compare numbers of transitions one needs to remove the distortion caused by the differing numbers in each state over time. To deal with this we consider transition probabilities. We define a transition probability as the percentage of claims currently in a state that move to the target state in that development quarter. We consider the transition probabilities to be functions of the current state, the destination state and development period. Considered in this way, the transition rates for this portfolio turn out to be reasonably stable over time and, as with any actuarial model, assumptions can be derived using a combination of historical data, regression techniques and judgement.

Given the above definition of transition probability, the sum of the probabilities of movements into all states, given a particular current state, add to one.

Note also that the transition probability from the current state to the same state is similar to a continuance rate such as may be used in the standard actuarial Payment Per Active Claims (PPAC) model.

As an example, we show modelled transition rates from the active weekly state, by development quarter:

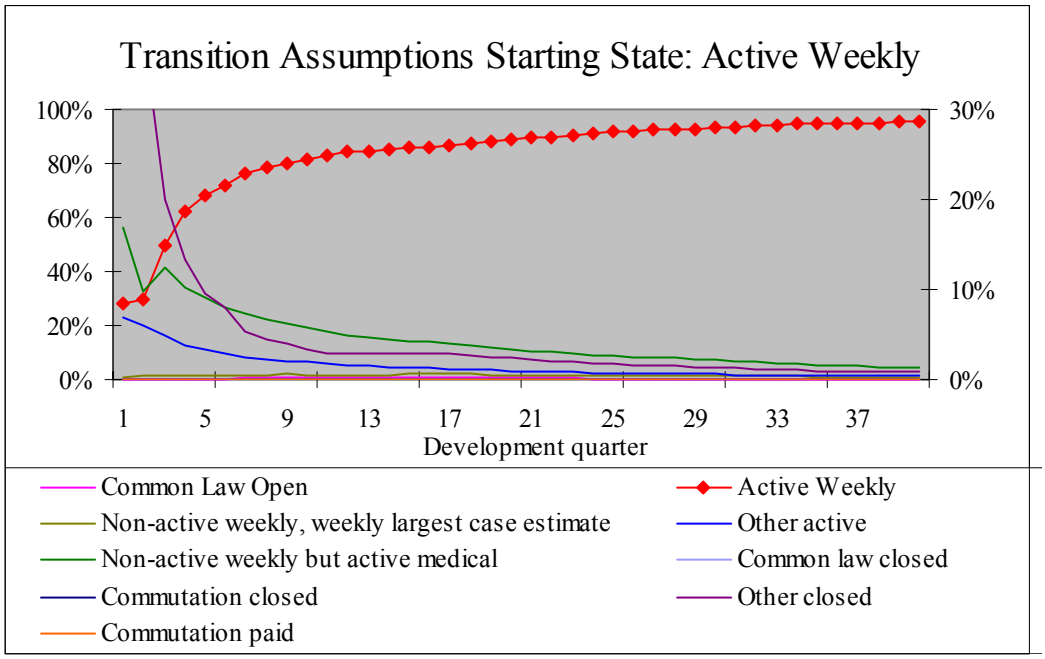

Note that in the above graph, the red line corresponds to an active weekly quarterly continuance rate and this levels out at around 95%. This line on the graph corresponds to the axis on the left hand side of the graph. The transitions out of the active weekly segment correspond to the right hand axis, and it can be seen that the highest transitions early on are into the other closed state, with other significant transitions being into the active medical and other active states.

2.3 Payments

One of the primary aims of splitting the claims into distinct segments is to examine the different payment levels by segment. We have tried to express payment levels so that they relate meaningfully to the claims for which they are being paid. As such, most of the combinations of payment types and segments were analysed as a Payment Per Active Claim (PPAC) type model, that is expressed as payments made per claim per quarter.

However, lump sum settlements (common law and, in the pre-2001 reform environment, commutations) are treated in a way which is close to a PPCF model. Common law payments to open common law claims are analysed as total common law payments on that claim per transition to the common law closed state. This gives an assumption which is gross of partial payments to open common law claims and these need to be deducted from the final liability. Commutation payments are analysed as commutation payments on that claim per transition to the commutation paid state.

The other complication concerns the active weekly claims. In the NSW Scheme there are three distinct types of weekly benefits:

- Total incapacity benefits (which can also be split into the first 26 weeks post injury, and after 26 weeks);
- Section 40 or partial incapacity benefits; and
- Section 38 or partially incapacitated workers not suitably employed while seeking employment or retraining, and receiving total incapacity benefits.

We could have considered claimants receiving each of these payment types as being in a different claim segment and projected transitions between the states. However, due to coding issues, a significant number of claimants appear to be receiving payments under more than one of the weekly types. For this reason, and a reluctance to add to the number of claim segments, we have instead modelled the proportion of the active weekly segment who receive each type of benefit as a function of development period. This assumes that the claims have equal probabilities of transition, irrespective of which of the weekly payment types they are receiving.

The fitted payment levels can be examined in two ways. Firstly, we can consider the levels of the different types of payments made to claims in the same segments, and secondly (and perhaps more interestingly) we can examine levels of payments for the same type of benefit across the different segments. We show examples of both of these below.

Firstly, consider payments per claim in the active weekly state:

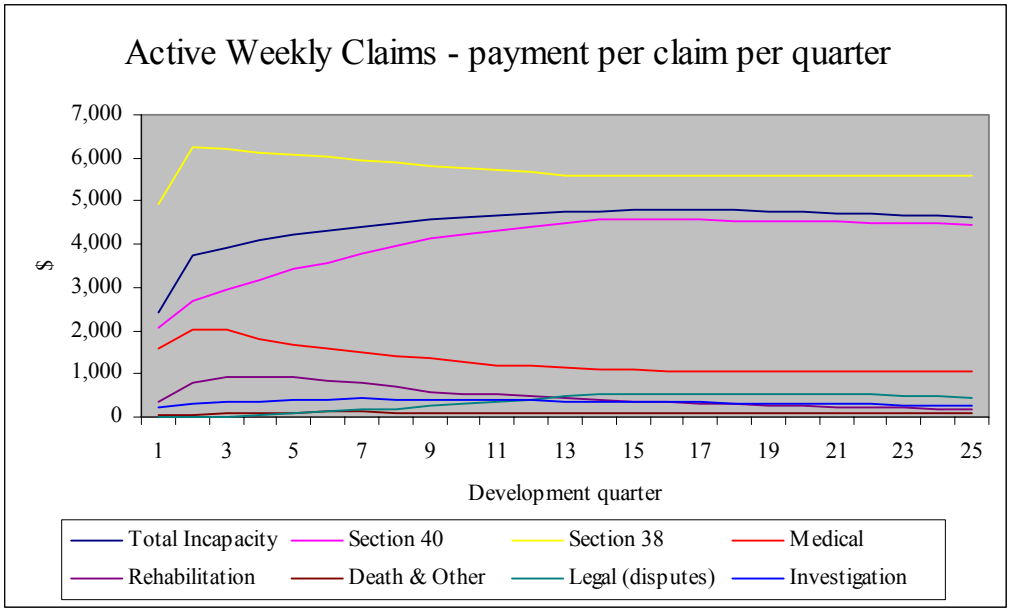

The highest payments are the weekly benefit types, with significant amounts also being paid in medical, rehabilitation, legal (disputes) and investigation.

As noted above, it is also interesting to compare the levels of a particular benefit type across the different claim segments. Consider medical payments which are made to claimants in the deafness open segment, the common law segment, the active weekly segment, the active medical segment and the other active segment. The graph below shows the difference in adopted payment per claim levels for each of these segments:

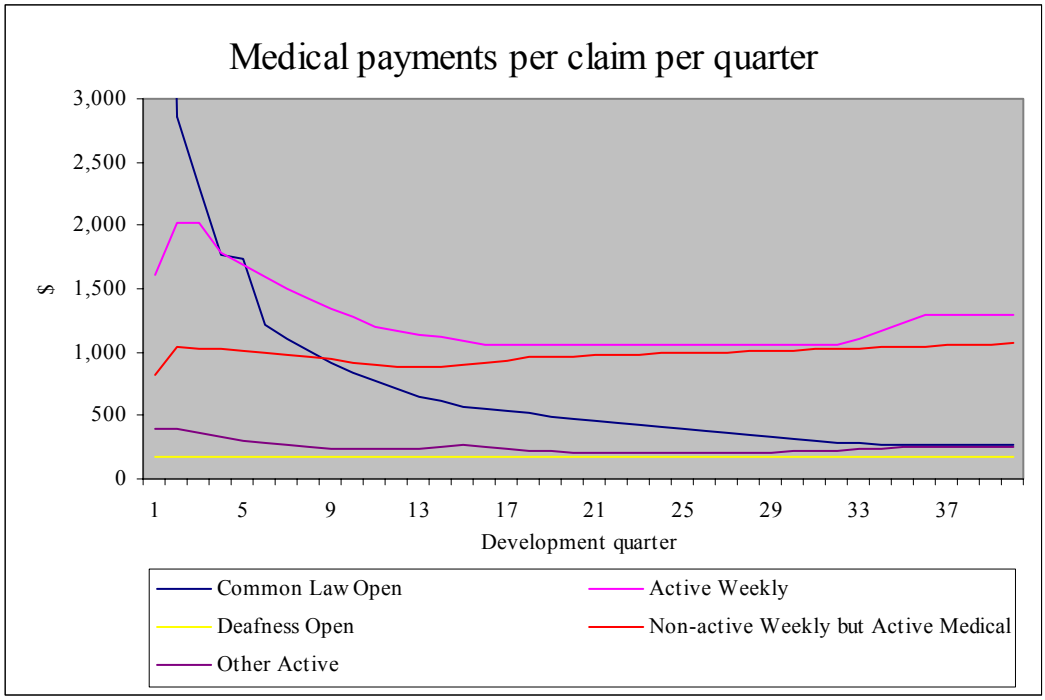

3 Model Limitations

The number of potential assumptions to fit can be daunting. In theory, with 11 states and 10 payment types there are 231 development curves for each of eight insurers! In practice, many of the transitions are insignificant, as are many of the payment types from some of the states. With some judicious grouping of payments types and the relatively mechanistic method for setting the insurer assumptions described in Section 4 below, assumption setting becomes somewhat easier.

There is a perception reported to us that the type of multi-state transition model we have described here is very sensitive to small changes in the transition rates. We cannot comment on this other than to say that we have not found this to be the case for the NSW portfolio. It is true that there are a few key assumptions which must be set carefully but our experience is that the data in these cases is sufficiently stable to support this and that "sensible" assumptions give sensible results.

However, the model is certainly sensitive to the claim profile at the valuation date. This can be seen as either an advantage or a disadvantage, depending on the situation. For very early development periods, where the transition rates are very high, the claim profile can be significantly affected by relatively small process delays and this leads to an instability in the valuation results. The methodology could, no doubt, be modified to overcome this but since we are only dealing development periods of more than one year we have not done this at this stage.

Finally, theoreticians will note that we have essentially assumed a Markov type process here. For instance, if a claim reopens from a closed state into the active weekly state then it is assumed to have the same future experience as an "average" active weekly claim. This can be important if an unusual group of claims are reopened in the quarter before the valuation date for some sort of "one-off" weekly payment but will close shortly thereafter. This valuation model is likely to over-estimate the liability in such a situation.

Notwithstanding these limitations the model apparently gives a good picture of the overall portfolio dynamics. In the case of this particular portfolio, the methodology gives very good agreement with the more conventional actuarial methodology used for the statutory valuations, both in total and across accident periods.

4 Fitting Insurer-Specific Assumptions

In order to remove some of the subjectivity in setting of insurers' assumptions we used the following methodology for translating assumptions set for the Scheme to take account of individual insurers' historical experience.

For a particular assumption, we express the insurer's experience as a proportion of smoothed Scheme experience over a range of development periods. We also assess the variability of this proportion over the range and over time. If the insurer's experience is close to the Scheme assumption and/or highly variable the Scheme assumption is adopted. If the insurer's experience is not too variable and very different from the Scheme assumption individual insurer's data is used to determine the adopted fit. If the insurer's experience is between these two extremes we used a credibility weighted average of Scheme fit and the insurer's experience.

The overall effect is that insurers with stable experience that differs significantly from the Scheme assumptions will most likely be set assumptions based on that experience while insurers with very little experience of their own will have assumptions based more on overall Scheme experience.

To demonstrate this we have created two fictitious insurers, the first is a large insurer with different experience than the Scheme:

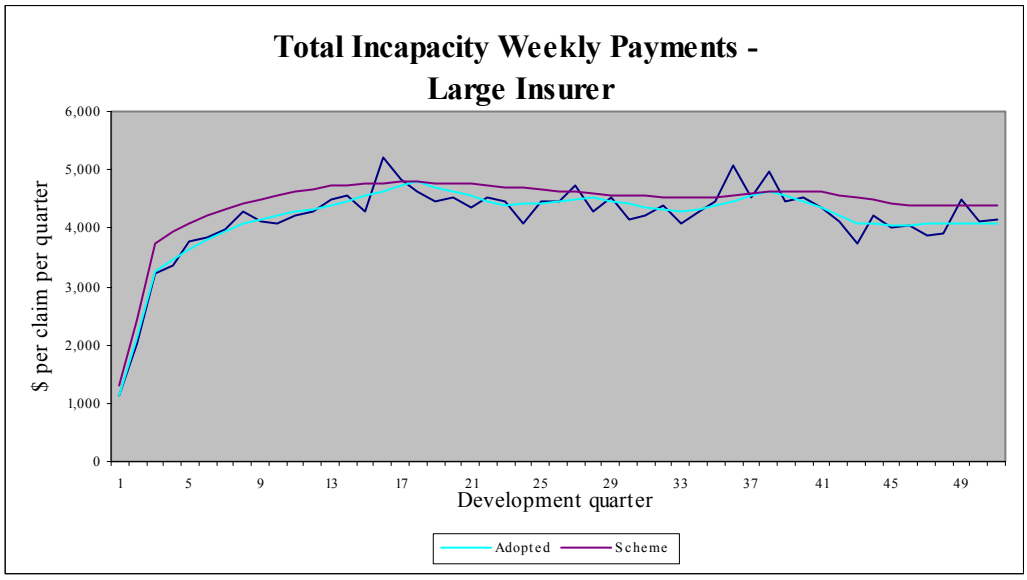

As a comparison we also show here a small insurer, also with different experience from the Scheme. The increased variability however, particularly in the tail, means that we tend to place more reliance on Scheme assumptions:

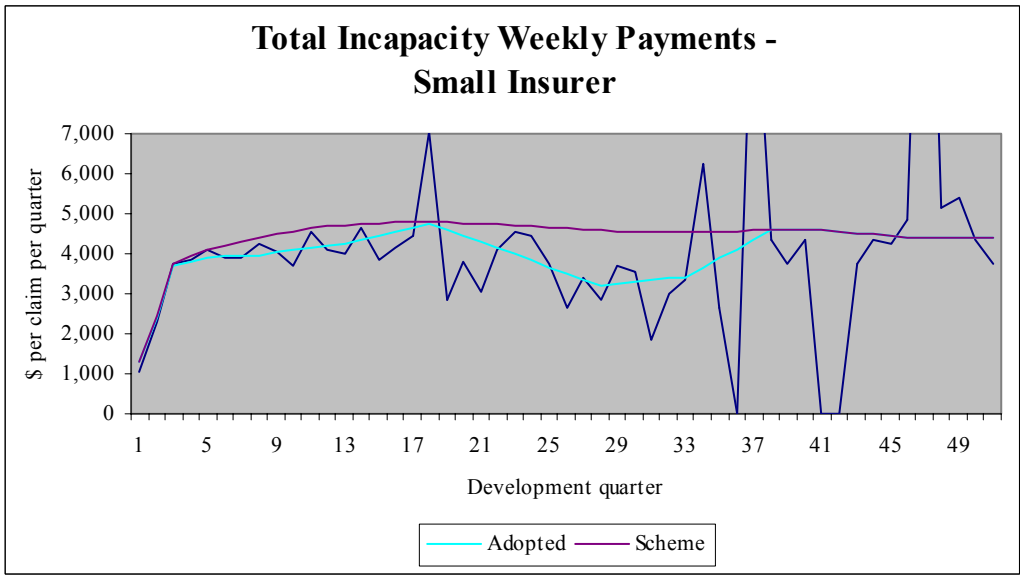

5 Actuarial Liability per Claim

One of the advantages of the framework is that the assumptions can be considered independently from the number of claims. This has a number of implications:

- For any particular segment and any particular quarter of accident we can use the model to determine the average expected liability for a single claim. This allows us to assess the relative values of claims in the different segments.

- Using a pre-calculated actuarial liability per claim by accident period and segment, and knowing the current claims mix, we can quickly determine an outstanding liability estimate. This provides a straightforward way of understanding the value of decisions regarding claim management. There are some dangers in this since we are treating the claim segments for each accident period as homogeneous when this is not really the case. However, the value of a simple link between claim management outcomes and the actuarial liability should not be underestimated.

- The above two points lead naturally to a simple method of monitoring portfolio claim management outcomes and payment levels.

We demonstrate this with an example. At 30 June 2002 (or any time well before 30 June 2003) we can calculate the average inflated and discounted liability per claim at 30 June 2003, per claim at 30 June 2003. Suppose this gives the figures in the following table by claim segment for the June 2000 accident quarter:

Segment	**($)**
Open deafness	6,367
Closed deafness	1,616
Open common law	285,673
Active weekly	101,588
Non-active weekly, with active medical payments	54,446
Non-active weekly, with weekly the largest case estimate	47,635
Other open	22,247
Closed common law	4,122
Closed commuted	0
Other closed	1,796
Commuted but not yet closed	10,427

Note the clear differentiation in terms of value from common law open at one extreme to closed commuted at the other. The two non-active weekly states have similar estimates but the make-up of the liability in terms of payment types is quite different, with (unsurprisingly) the non-active weekly with active medical payments state having a much larger liability for medical payments.

Now suppose that at 30 June 2003 the claim profile for this same accident quarter is:

Segment	Number of claims	Liability ($M)
Open deafness	81	0.52
Closed deafness	439	0.71
Open common law	259	73.99
Active weekly	859	87.26
Non-active weekly, with active medical payments	146	7.95
Non-active weekly, with weekly the largest case estimate	240	11.43
Other open	1,191	26.50
Closed common law	288	1.19
Closed commuted	435	0.00
Other closed	25,220	45.30
Commuted but not yet closed	21	0.22
Total	29,179	255.06

Then as soon as the claim profile is known, the liability based on this profile can be calculated by multiplying the liabilities per claims with the claim numbers, giving a liability of $255M on all claims reported to date, to which we add an allowance for IBNR. Note that almost 20% of the liability is in respect of expected reopenings of closed claims.

6 Links to Monitoring

The context in which this valuation methodology was developed is one of rewarding insurers for actions which reduce the Scheme liability. The methodology is intended to provide a transparent link between positive claim management behaviour and a positive outcome in terms of the actuarial liability. This can be taken one step further and converted into a responsive monitoring tool which can be applied to the portfolio on a monthly basis.

Together with Kris Bruckner at WorkCover we have developed a monitoring tool which is currently being provided to all the insurers so that they can monitor their own portfolios. The tool takes as input the recent claim profile and payment history for the insurer. It provides monthly and year to date outputs in a number of areas:

- Actual versus expected case management outcomes such as exits from and reactivations to the active weekly class, emerging common law claims, numbers of lump sum payments;

- Actual versus expected payments corrected for the actual claim profile and number of lump sum settlements. This enables insurers to separate the issues of profile management and expense control to answer questions such as “I know I have more active weekly claims than expected but am I paying more or less medical payments per active weekly claim and how much has this cost me?”

- A monthly reassessment of the portfolio liability, based on assumptions set at the beginning of the year, and a calculation of whether a profit or loss was made over the month and the year to date.

- An attribution of the profits and losses to the various items of case management and expense control.

Clearly the monitoring tool cannot be relied on to give a definitive estimate of outstanding claim liabilities. The most significant limitation in this respect is that the transition and payment assumptions are not updated over the year and that such an update would change the financial picture. However, provided that the limitations are understood we believe it provides a great deal of useful information.

We will likely present the monitoring tool at a later conference. However, we close this paper with a simple example of some of the information it provides, using the whole Scheme tail portfolio as an example.

The most summarised form of information is a graph comparing opening and closing liabilities and payments:

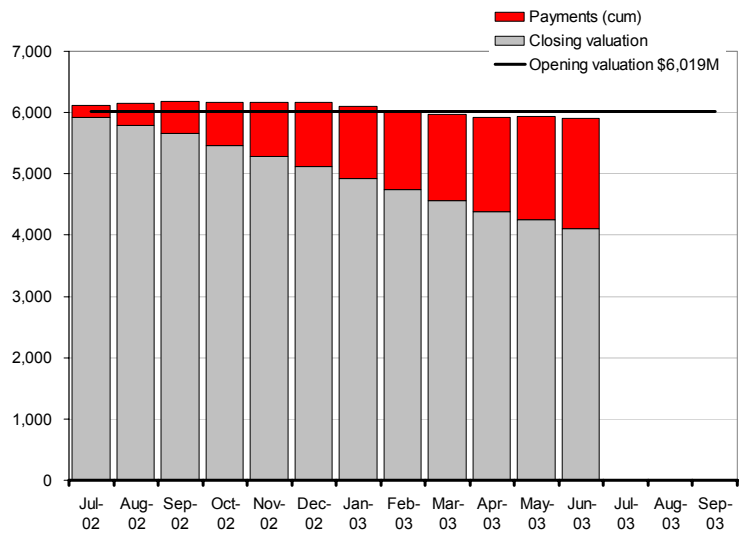

In simple terms, if the columns on the graph fall below the level of the opening valuation then there has been a release of liability over the year to date. The monitoring tool quantifies this more precisely and attributes it by source. In this case we will examine the contribution from the number of active weekly claims. The tool shows the actual and expected number of active weekly claims, together with the number of exits and reactivations:

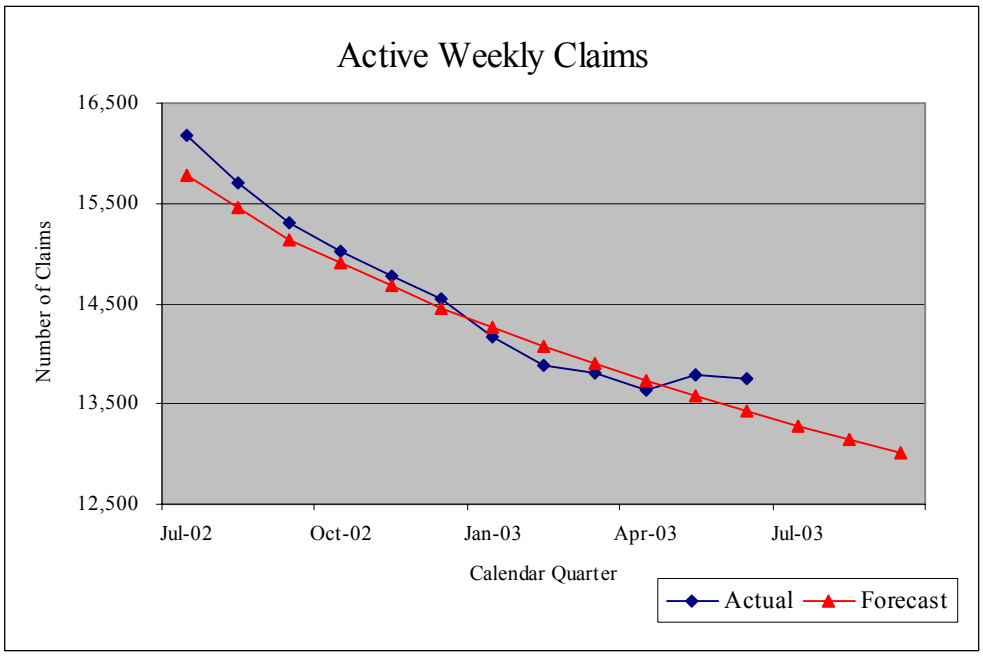

The general pattern is of a greater number of claims than expected in July but a gradual improvement until late in the year. The picture becomes clearer when we look at the exits and reactivations:

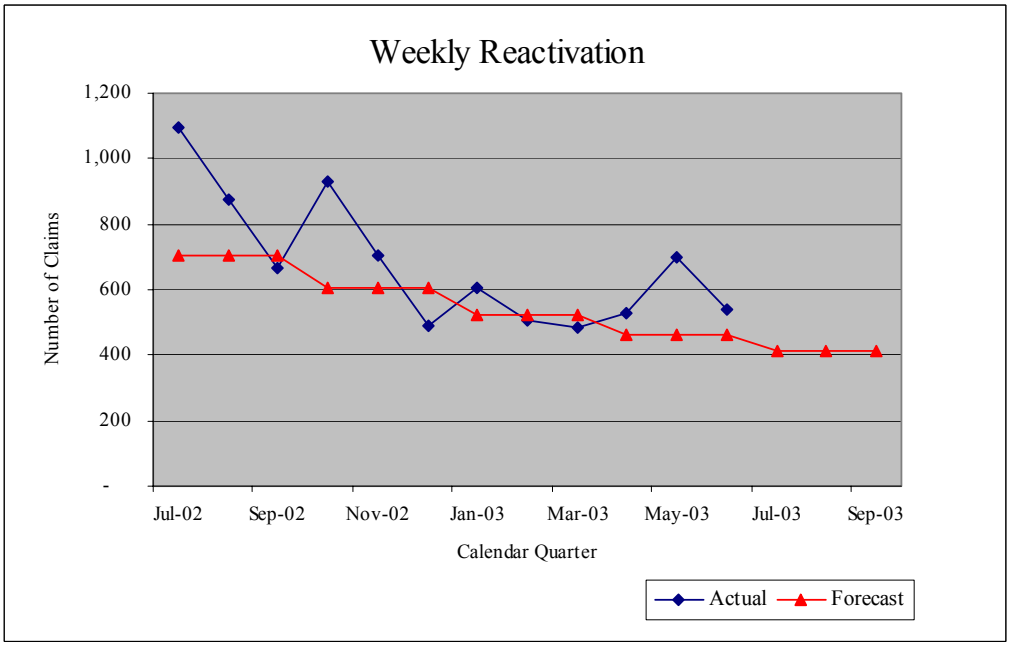

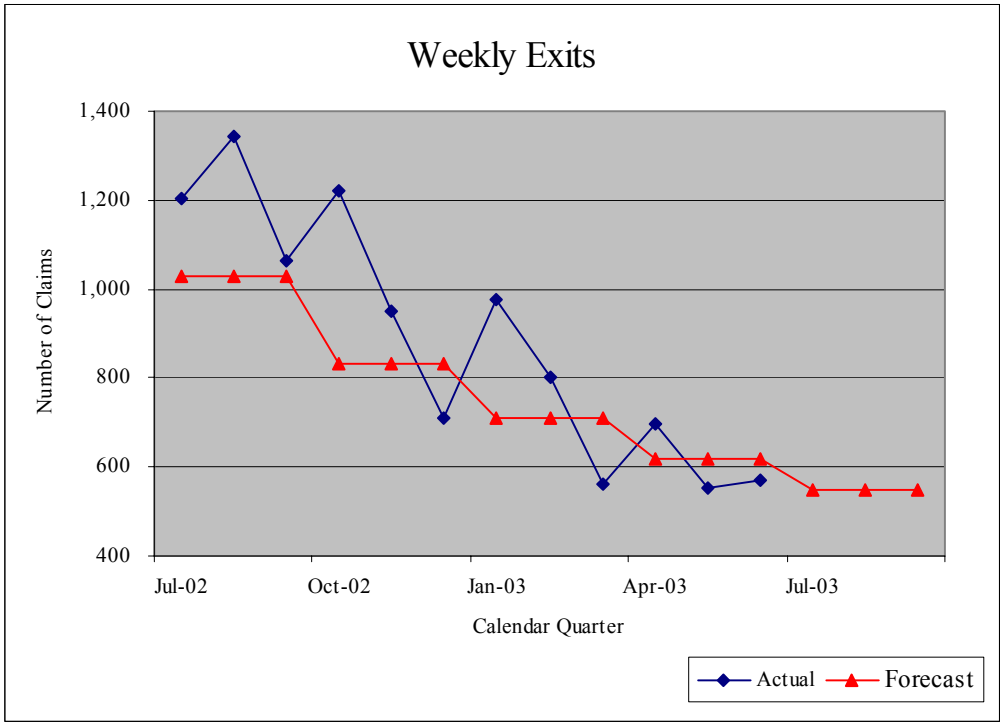

There was a lot of portfolio activity for the first half of the year, with both exits and reactivations being considerably higher than expected. For the net quarter activity reduces somewhat and the final quarter sees more than expected reactivation and less than expected exits. It is difficult to draw conclusions at a Scheme level but a portfolio manager in insurer would be likely to:

- Analyse the reactivations, understand why they have increased, whether they are genuine relapses or "one-off's" likely to exit soon after;
- Understand whether there was any reason underlying the small number of exits in the last quarter.

It is worth noting that, although we cannot show them for confidentiality reasons, the equivalent results for many of the insurers appear are stable and show features which portfolio managers have attributed to various operational issues. The authors had feared that the monthly figures would be too variable to be of much value to management but these fears turn out to be unfounded, and in fact we believe that it is possible to further break down the larger insurers and separately assess teams within these insurers.

The final piece of information we show from the monitoring tool is the impact on the liability of the active weekly experience:

Quarter to:	Savings in Liability ($M)
Sep-02	($12.98)
Dec-02	($0.59)
Mar-03	$5.58
Jun-03	($60.54)