Individual claim modelling of CTP data

Prepared by Gráinne McGuire

Presented to the Institute of Actuaries of Australia
XIth Accident Compensation Seminar 1-4 April 2007
Grand Hyatt Melbourne, Australia

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

Abstract

The modelling process for an Australian CTP data set is discussed here. In the last number of years, the Scheme in question has been subject to falling claim frequency, which has the most impact on lower severity claims. With the mix of claims by severity changing through time, an overall model, ignoring claim severity will, in the case of frequency reductions in lower severity claims, underestimate the claim size. Thus, a model incorporating severity is required and the building of a GLM for finalised claim size incorporating severity is discussed here.

The Scheme has also been subject to major legislative change in the last few years. Issues surrounding the modelling of these changes are discussed.

Keywords: CTP data, Generalised linear models, injury severity, legislative changes

1. Introduction

The use of Generalised Linear Models ("GLMs") to model Compulsory Third Party ("CTP") Insurance data was discussed in Taylor and McGuire (2004). This paper found that a model of individual finalised claim size built using a GLM had many advantages over a chain ladder model constructed in the traditional actuarial way for the data in question, leading to a more parsimonious, interpretable model

The model discussed in that paper was based on data to September 2003. In the intervening years to June 2006, new experience has accumulated. Two issues have arisen that suggest that significant reworking of the model is required. Firstly, the Scheme has been subject to falling frequency rates. Such a trend was present before 2003, but claim frequency was observed to fall sharply between 2002 and 2004. A reducing frequency rate can be associated with the removal of smaller claims from the experience, leaving the larger, more serious claims. If this is the case here, then the significant changes in claim frequency would lead to changes in the mix of claims and thus, to claim size. The model presented in the earlier paper does not include any terms recognising the mix of claims; therefore the change in mix may be recognised as superimposed inflation. A simple extrapolation of past superimposed inflation is unlikely to give the correct result.

Secondly, legislation came into effect at the end of 2002. This legislation has a wide-ranging impact, affecting the types of permitted claims as well introducing statutory limits on some of the damages that can be claimed. Such legislation has a significant effect on the claims experience. However, modelling it is not a simple matter since it would be expected to affect claims of differing seriousness to different extents. Further, as at June 2006 (the date at which the data are current), there are only 3 ½ years of experience under the new legislation. Given the long-tailed nature of CTP, this experience is still relatively immature. This must be addressed in any projections.

The aim of this paper is to extend the finalised claim size GLM in Taylor and McGuire (2004) to incorporate claim severity to enable the model to deal appropriately with the changing mix of claims and the 2002 legislation. It presents an almost realistic case study of the modelling of a CTP data set; some details have been suppressed but these do not affect the arguments in this paper.

Out of scope of this paper is the consideration of other possible predictors of claim size (e.g. litigation, employment status, injury types etc). Models including such predictors are discussed in stochastic case estimation models (e.g. Brookes and Prevett, 2004) or in the type of models discussed in Taylor, McGuire and Sullivan (2006).

The structure of the paper is as follows: in Section 2, the data used are described. Section 3 gives more details on the motivation driving the extension of the model to incorporate claim severity. The modelling and projection process is discussed in Section 4. A closing discussion is given in Section 5.

2. Data

The data consist of Compulsory Third Party ("CTP") insurance claims in one state in Australia. The insurance is underwritten by private sector insurers, who are required to submit their claims data to a centralised data base.

The data used in this paper are extracted from that data base. The data base comprises a unit record claim file containing the following pieces of information:

- Date of injury;
- Date of notification;
- Injury codes and claim severity;
- Various other claim characteristics (e.g. legal representation, litigation indicator etc);
- Histories of
  - Finalised/unfinalised status (some claims reopen after having been designated finalised) including dates of change of status;
  - Paid losses by payment type;
  - Case estimates.

The scheme of insurance commenced in its present form in September 1994 and the data base contains claims with dates of injury from then. It is current at June 2006.

Since the purpose of this paper is to demonstrate claim size modelling using GLMs (e.g. for loss reserving or pricing purposes), analysis will be limited to finalised claims. Therefore, with some exceptions, data are only required in respect of finalised claims. The main exception is that, for the purposes of estimating operational time, the ultimate numbers of claims to be notified in each accident quarter have been estimated outside this paper and have been taken as given (but note the commentary in Section 5).

Wherever paid loss amounts are used, they have been converted to 30 June 2006 dollar values in accordance with past wage inflation experienced in the state concerned. This is done to eliminate past "normal" inflationary effects on the assumption that wage inflation is the "normal" inflation for this type of claim. Henceforth, any reference to paid losses will carry the implicit assumption that they are expressed in these constant dollar values.

Of course, claims inflation differs from wage inflation from time to time; the excess of claims inflation over wage inflation is referred to as superimposed inflation ("SI"). The modelling of this is discussed in Section 4.2.4.

As indicated in Section 1, the purpose of this paper is to extend the individual claim modelling introduced by Taylor and McGuire (2004) to incorporate claim severity into the model. Thus the dependent variables of interest are the individual sizes of finalised claims.

Let

$i =$ accident quarter ($i = 1, 2, ..., 48$ where 1 = Sep94, 48=Jun06);

$j =$ development quarter ($j = 0, 1, ..., 47$);

$k =$ calendar quarter of finalisation ($k = 1, 2, ..., 48$ where 1 = Sep94, 48=Jun06);

$s =$ injury severity

$N_i^s =$ ultimate number of claims in accident quarter i for severity s;

$G_{ij}^s =$ number of claims finalised to date from accident quarter i for severity s up to and including the current claim;

$t = \frac{G_{ij}^s}{N_i^s}$ = operational time of claim.

Then the $\{F_{ikt}^s\}$, representing the total size of a claim, of injury severity s, from accident period i, finalising in calendar quarter k with operational time t, are the quantities of interest in this paper.

Injuries are recorded using the Abbreviated Injury Scale ("AIS") codes. AIS codes consist of a five digit code giving the injury type (for example 70101 corresponds to a cervical whiplash injury). A sixth digit gives the severity of that injury. This may take values from 1 (least severe injury type) to 5 (catastrophically injured) and 6 (death). There are also some administrative types of injury code, with severity level 9.

Up to five injury codes may be applied to any one claim. The measure of injury severity that has been used is the "Maximum Injury Severity" ("MAIS"). This is calculated as the maximum of all the injury severities associated with that claim. Note, however, that a severity of 9 is ignored if there are other severities associated with that claim that fall between 1 and 6. Therefore, a claim only is assigned a MAIS of 9 if it only has injury codes with severity 9.

A further level of injury severity is "Blank". As the name suggests, this exists when there is no injury information associated with a finalised claims. Such cases are rare.

In this paper the term "severity" will generally be used in preference to MAIS. Thus, a statement that a claim is "severity 2" means that its MAIS is 2.

3. Motivation

3.1 Severity

Not surprisingly, claims of differing injury severities can have very different claim sizes. As an example, histograms of the distribution of claims sizes (corrected for wage inflation) are given in Figure 3.1 for claim severities 1 and 5.

Figure 3.1 Claim sizes for severities 1 (top) and 5 (bottom)

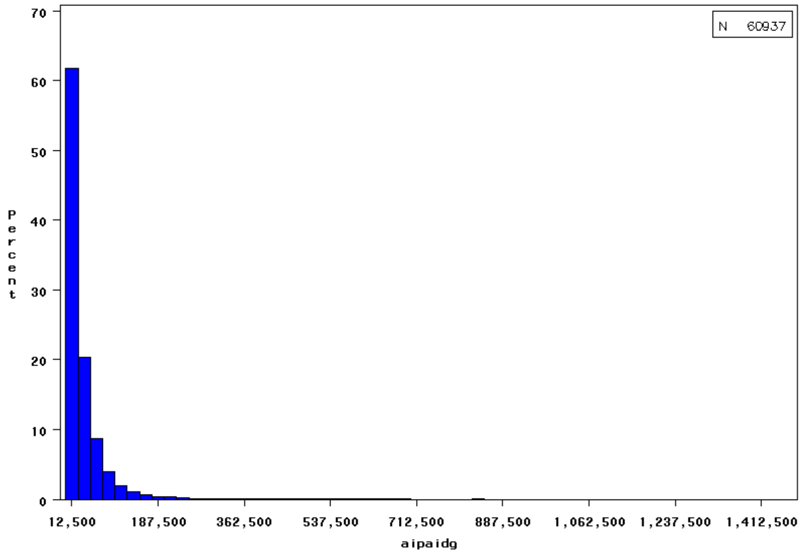

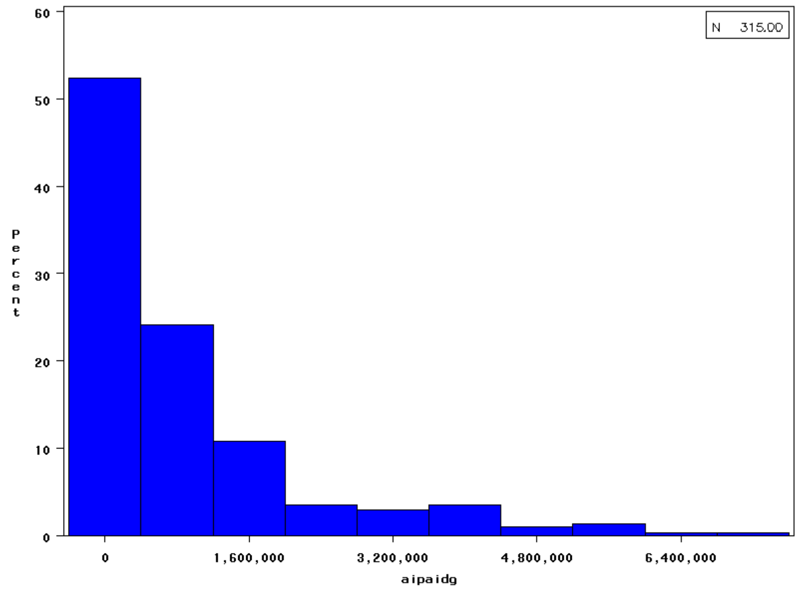

Taylor and McGuire's (2004) model for finalised claim sizes ignored the effect of claim severity, instead taking the approach of modelling all claims together. This approach is satisfactory if the mix of claims remains constant throughout the modelling period and any projection period (and if a suitable error distribution can be found). If, however, the mix of claims is changing, then the overall average claim size will increase and decrease as the relative proportion of lower severity claims decreases and increases.

To illustrate this point further, Figure 3.2 plots the claim frequency by injury severity both on a linear and log scale. It is apparent that the severity 1 frequency saw a significant increase in 1998, followed by a reduction from 1999 to 2001 and an even more significant decrease from 2002 to 2005. Some reduction is also seen in other severities (e.g. severity 2), though not to the same extent.

Figure 3.2

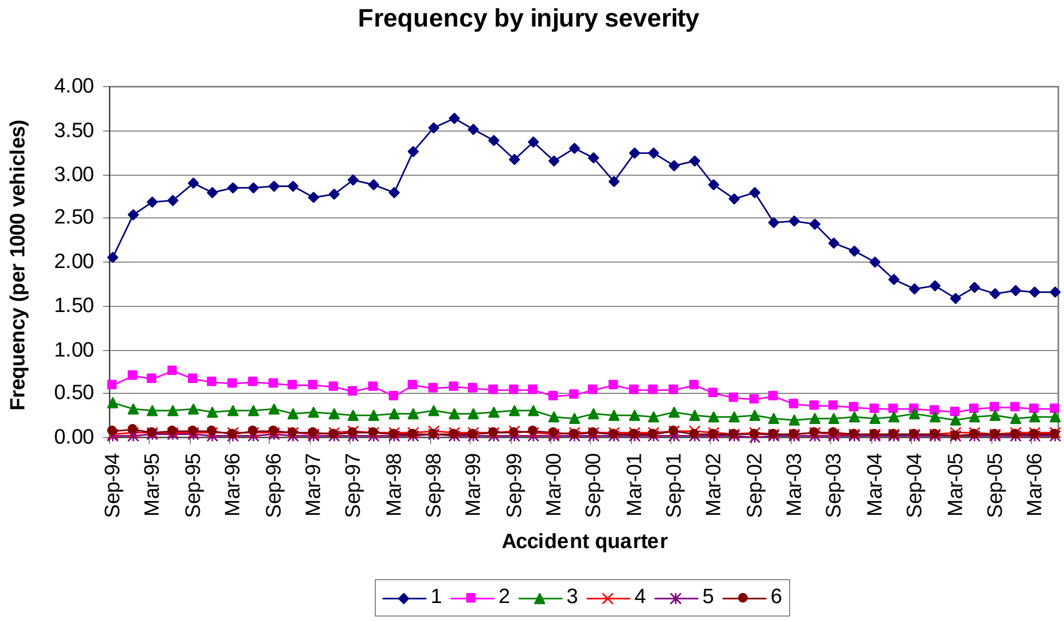

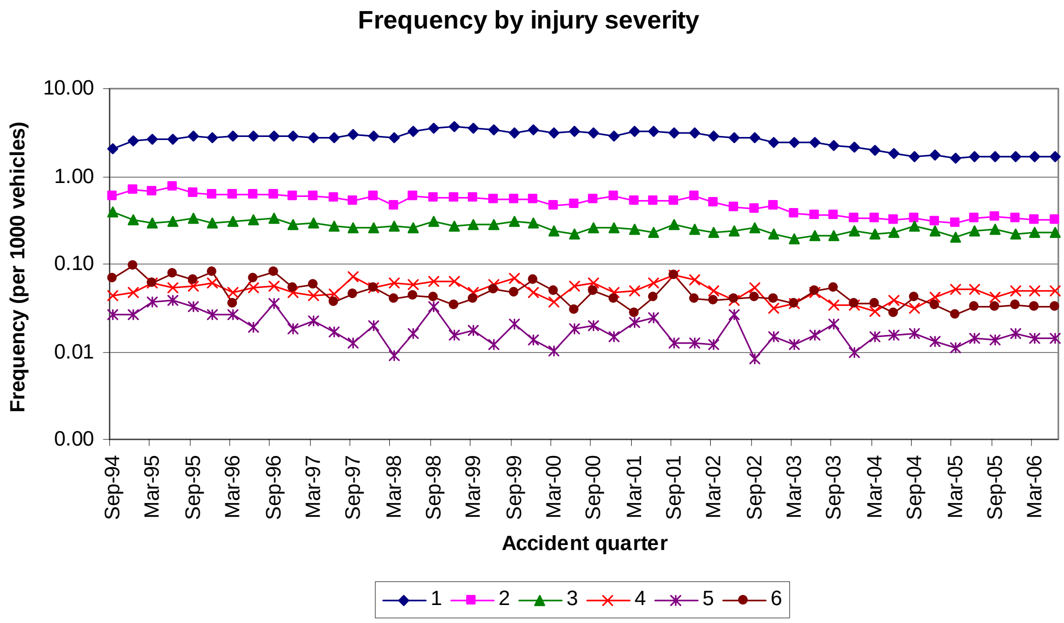

Table 3.1 below presents an illustration of the effects of changing mixes in claim frequency on overall claim size. The claim sizes presented in the tables are the simple averages of all claims finalised to date. The frequency by accident year is the actual estimated frequency in that year. Rather than displaying the actual frequencies, the proportion of claims in each severity is displayed to highlight the changes in the mix of claims from year to year. The final column contains the overall average claim size based on the severity specific claim sizes (the same for each year) and frequency (varies by year).

Table 3.1 Illustrative claim size by accident year

Accident year ending 30 June	Proportion of claims in each severity 1	2	3	4	5	6	Other	Overall Average size
1997	72%	15%	7%	1.2%	0.6%	1.5%	2.6%	57,110
1998	73%	13%	6%	1.5%	0.4%	1.1%	4.0%	53,594
1999	75%	12%	6%	1.2%	0.4%	0.9%	4.3%	52,140
2000	74%	12%	6%	1.2%	0.4%	1.1%	5.1%	51,501
2001	75%	13%	6%	1.3%	0.5%	0.9%	2.6%	53,494
2002	75%	13%	6%	1.5%	0.4%	1.2%	2.0%	54,049
2003	77%	12%	7%	1.3%	0.4%	1.3%	1.7%	53,261
2004	75%	12%	8%	1.2%	0.6%	1.4%	1.2%	57,093
2005	71%	13%	10%	1.9%	0.6%	1.4%	1.3%	61,939
2006	71%	14%	10%	2.0%	0.6%	1.4%	1.3%	62,698
Average size of all finalised claims	32,187	69,167	161,647	339,372	906,672	110,734	19,966	

The effect on claim size of varying mix of claims is apparent from Table 3.1. Recall from Figure 3.2 that the severity 1 frequency increased in the calendar year 1998. This corresponds to increased proportion of Severity 1 claims and the fall in claim sizes in the table above in the financial years ending June 1998 and 1999. In 2003 to 2006, the relative proportion of Severity 1 claims decreased while that for severities 2-6 generally increased. This is reflected in an increasing claim size, particularly from 2003 to 2004 and from 2004 to 2005.

The conclusion is that, for this particular data set with its varying mixes of claim severity, it is necessary to take claim severity into account in any model to firstly, understand the claims process and secondly, to make accurate projections.

3.2 Legislative effects

Two pieces of legislation have impacted the scheme since its commencement in September 1994:

- Legislation introduced in the 2000Q4, among other things, required mandatory conferences and restricted access of plaintiffs to legal costs for claims below a certain amount;

- Legislation placing restrictions on the types of claim as well as introducing statutory limits (reducing some of the damages that could be claimed) came into effect in late 2002.

In both cases, effects are mainly seen on the smaller (less severe) claims whereas the larger claims remain unchanged (or are expected to in the case of the second piece of legislation, but to date experience of large claim settlement is too sparse to make any definitive statements).

The 2000Q4 legislation is quite specific in its application to plaintiff legal costs. Thus, the modelling of these effects is quite straightforward. Section 4.2.3.1 discusses suitable model terms for recognising this legislation.

The latter piece of legislation is more wide-reaching and uncertain in its effects than the former and any modelling approach is necessarily more complicated. Table 3.2 and Table 3.3 present some data summaries showing the impact of this legislation.

Table 3.2

Accident	Average claim size values (in 30/6/06 values) in development quarter								
quarter	0	1	2	3	4	5	6	7	8
	$	$	$	$	$	$	$	$	$
Sep-02	4,214	3,924	7,255	17,142	23,051	27,608	29,819	34,735	38,020
Dec-02	9,799	3,713	7,670	13,351	17,125	18,590	27,397	28,418	39,862
Mar-03	1,430	2,972	3,210	4,968	6,600	9,476	10,249	18,004	14,679
Jun-03	2,851	1,854	2,767	3,163	5,514	7,195	10,772	20,490	23,987
Sep-03	128	1,679	3,051	4,887	7,192	10,227	13,946	16,479	24,030
Dec-03	1,099	1,604	3,280	3,692	5,992	6,098	11,952	19,127	32,322
Mar-04	-	2,354	4,415	2,621	3,786	8,802	13,667	22,697	25,314
Jun-04	495	2,846	2,746	3,923	4,563	12,713	14,161	25,069	36,179
Sep-04	408	1,296	2,186	4,267	6,125	10,688	19,161	32,930	
Dec-04	815	1,190	3,882	5,058	5,845	12,976	18,057		
Mar-05	1,783	1,966	3,818	4,653	6,900	10,265			
Jun-05	896	1,899	3,287	3,927	7,244				
Sep-05	1,732	2,347	3,743	6,130					
Dec-05	2,367	2,567	3,900						
Mar-06	1,266	2,274							
Jun-06	2,367								

Table 3.2 gives a summary of average claim size in the first 9 development quarters. The new legislation came into effect during December 2002, so March 2003 is the first quarter with the effect in full. A noticeable reduction in claim size is apparent from the table. However, not all of this reduction is a genuine fall in claim sizes. Table 3.3 presents the operational time in the same time period and it is apparent that claim finalisation rates have slowed down, with operational time post-legislation being approximately one quarter behind where it was pre-legislation. (Note that this is not a problem for either an overall GLM or a severity specific GLM so long as these are based on operational time rather than development quarter).

Table 3.3

Accident quarter	Operational time in middle of development quarter								
	0	1	2	3	4	5	6	7	8
	%	%	%	%	%	%	%	%	%
Sep-02	0	1	5	10	16	23	32	40	46
Dec-02	0	1	4	8	13	21	29	37	43
Mar-03	0	1	3	6	11	17	24	30	35
Jun-03	0	1	3	7	12	18	24	30	36
Sep-03	0	1	4	8	13	18	23	29	36
Dec-03	0	2	5	10	15	20	25	32	38
Mar-04	0	1	4	8	13	18	25	31	38
Jun-04	0	1	5	9	14	20	27	34	41
Sep-04	0	1	5	9	15	20	26	33	
Dec-04	0	2	5	9	15	20	26		
Mar-05	0	2	5	9	14	20			
Jun-05	0	2	5	10	15				
Sep-05	0	2	5	10					
Dec-05	0	2	6						
Mar-06	0	2							
Jun-06	0								

A further complication in the analysis of the results in Table 3.2 is that the post legislation time period coincides with the period of rapidly reducing claim frequency (refer back to Figure 3.2). Thus the changing mix of claims is also likely to be playing a role in the average claim sizes in Table 3.2.

The first legislative effect was captured in the Taylor and McGuire (2004) GLM, which was based on data as at September 2003. A reduction was applied to all claims with operational times between 0% and 35%. This reduction decreased with operational time; thus it was at its maximum at t = 0% and decreased to 0% at t = 35%. This represents the wearing off of the restriction to access to plaintiff costs with increasing claim size. However, with the introduction of severities into the model, it would be expected that this effect could be refined. For example, it would not be expected to impact claim sizes for Severity 5 (catastrophically injured people) at all, but might affect most severity 1 claims.

Although the second piece of legislation was in force at September 2003, its effect was not included in the previous model as there was insufficient claims experience at the time. Since this paper uses data as at June 2006, it is possible to incorporate its effect, at least for the lower severities. It should also be clear from the above discussion that the only possible way to correctly model this part of the experience is through the use of a severity specific model. In this way, changes in claim size due to the legislation may be separated from the change in overall claim size induced by the changing mix of claims.

4. Modelling

In this section the severity based GLM is developed. The section opens with an overview of an overall model (i.e. ignoring severity) of claim size, together with its implications for projected claim size. Next the incorporation of severity in the model is considered. Finally, the legislative effects discussed in Section 3.2 are modelled.

4.1 Overall claim size model

Taylor and McGuire (2004) describe the process for fitting an overall GLM to these data (at an earlier point of time). Those interested in details of model fitting are referred to that paper. Here, the overall model at June 2006 is simply presented. This is very similar to the model given in the earlier paper, but with modifications in light of recent experience as well as the incorporation of the 2002 legislative effect. The model takes the following form:

$$
\begin{aligned}
E\left[Y_r\right] = \exp\ \ & \alpha + \beta_1 t_r + \beta_2 \max(0,10-t_r) \\
& +\beta_3(\max(0,\min(t_r-10,10)) + && \text{[operational time effect]} \\
& +\beta_4 \max(0,30-t_r) + \beta_5 \max(0,t_r-90) \\
& +\beta_6 I(t_r > 98) \\
& +\beta_7 I(k = \text{March quarter}) && \text{[seasonal effect]} \\
& +\beta_8 I(k < \text{1996Q4}) + \beta_9 \min(k,25) && \text{[finalisation effect]} \\
& +\beta_{10}\mathrm{I}(39 \le \mathrm{k} \le 44) + \beta_{11}\ \mathrm{I}(\mathrm{k} \ge 45) \\
& +\beta_{12} \min(9,\max(0,k-25) && \text{[finalisation x} \\
& +\beta_{13}(\max(0,\min(t_r-10,10))\min(k,25) && \text{operational time} \\
& +\beta_{14} t_r \min(k,25) && \text{effect]} \\
& +\beta_{15}(\max(0,\min(t_r-10,10))\min(9,\max(0,k-25) \\
& +\beta_{16} t_r \min(9,\max(0,k-25) \\
& +\beta_{17} I(i > \text{2000Q2})\max(0,30-t_r) && \text{[first legislative effect]} \\
& +I(i > \text{2002Q4})\begin{bmatrix} \beta_{18} + \beta_{19}\max(0,27-t_r) + \\ \beta_{20} I(k > 37) \end{bmatrix} && \text{[second legislative effect]}
\end{aligned}
\qquad (4.1)
$$

The claim sizes that result from this model (assuming no future superimposed inflation) are plotted below in Figure 4.1.

Figure 4.1

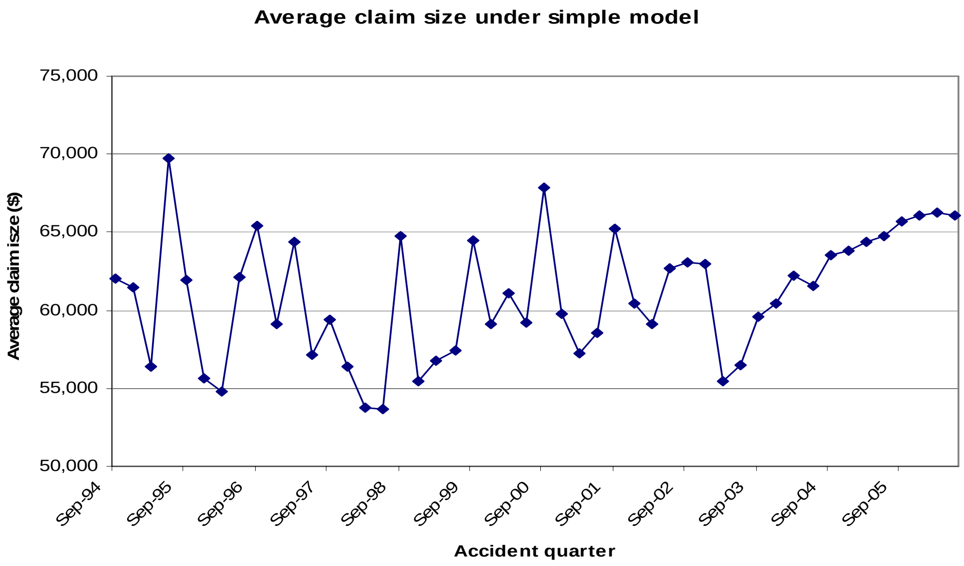

4.2 Extending the model to incorporating severity and legislative effects

Extending the model to severity involves, in principle, considering up to 7 distinct models (the severities 1 – 6, and the "other" category [this consists of claims missing severity codes or coded as severity 9, but is only a small proportion of the data, typically less than 2%]). Two possibilities exist for this work:

- Develop seven separate models, one for each severity; or
- Build one GLM for all severities but with model terms differentiated by severity.

For the latter case, the same error distribution must be applicable to all severity classes. Assuming this is the case (it is important to note that it is not necessarily so; therefore this assumption should be carefully checked before modelling), building one model does result in a more difficult modelling process since interactions between claim severity and other effects must be considered.

However, in the hands of a competent modeller, it should take considerably less time than the fitting of several separate models. Further, an all-inclusive model permits the use of information over two or more severity classes where that experience is relevant. For example, it might be observed that severity classes 2 and 3 have similar superimposed inflation trends. The full model can use the information in both of these classes to yield a more accurate estimate of this trend than could separate models for classes 2 and 3.

In this particular case a further complication arises in that the 2002Q4 legislation has had a significant impact on claim sizes from 2003 onwards. Thus, the most recent 3 ½ years of the scheme may be expected to look systematically different from the early years. Attempting to incorporate severity and the legislation at the same time is a large task for any modeller, no matter how experienced. Therefore, the following process is proposed:

- First, model the experience to 2002Q4 only. Add in severity into this model;
- Compare the actual experience post 2002 with the expected experience from this model. Thus, the expected experience is what would be expected if no legislative changes had occurred. This yields guidelines for how the pre-2003 model should be adjusted to take account of the changes;
- Incorporate the 2003 – 2006 data, adding model terms to include the legislative changes.

A final point to note is that the operational time associated with each claim should be specific to that severity. In other words, severity 1 claims should have operational times running from 0% to 100% depending on their order of finalisation rather than being allocated an operational time based on the finalisation of claims of all severities. If the latter case were used, then typically lower severity claims would occupy the lower operational time points and the higher severity claims the higher points. This would lead to some confusion between operational time and severity effects.

4.2.1 Selecting the error structure

The starting point for any claim size model is to choose the error structure (i.e. the distribution) used to model the data. Log normal and gamma are common choices. Other possibilities include quasi-likelihoods from the Exponential Distribution Family ("EDF"). Indeed, the overall model described in Section 4.1 uses an EDF(2.3) error structure. Readers are referred to Taylor and McGuire, 2004 for definition and details of this, but to put this into context, a gamma distribution corresponds to an EDF(2), inverse Gaussian to EDF(3) so an EDF(2.3) is longer-tailed than gamma but not as long-tailed as an inverse Gaussian distribution. The EDF family has the advantage over the log-normal model that data transformations are not required. Therefore, the problems associated with bias corrections may be avoided.

Details of how to select an error distribution using residual plots are given in the earlier paper. In short, the process involves fitting an initial model with a particular error structure and examining the residual plots for any evidence of problems (such as heterogeneity [e.g. "fanning" out or in] of residuals). A satisfactory set of residual plots is evidence that an appropriate error distribution has been selected.

Figure 4.2 presents the residual plots from using a gamma error structure for these data, incorporating severity into the structure. As well as residual scatterplots and histograms, a Quantile-Quantile (QQ) plot may be used to assess model fit. QQ plots are powerful tools for assessing distributional fits – a straight line indicates a good fit. In Figure 4.2, it is seen that the QQ plot does deviate from a straight line, indicating some problems with the distribution choice.

Some may argue that the problems highlighted in the QQ plot should not be ignored and that it may be evidence that the data are from a mix of error distributions and therefore require separate modelling. However, segmentation of these data into two data sets (low severity claims [Severities 1 and other] and high severity claims [Severities 2 - 6]) did not produce residual plots that improved on those presented in Figure 4.2.

Figure 4.2 Selecting an error distribution

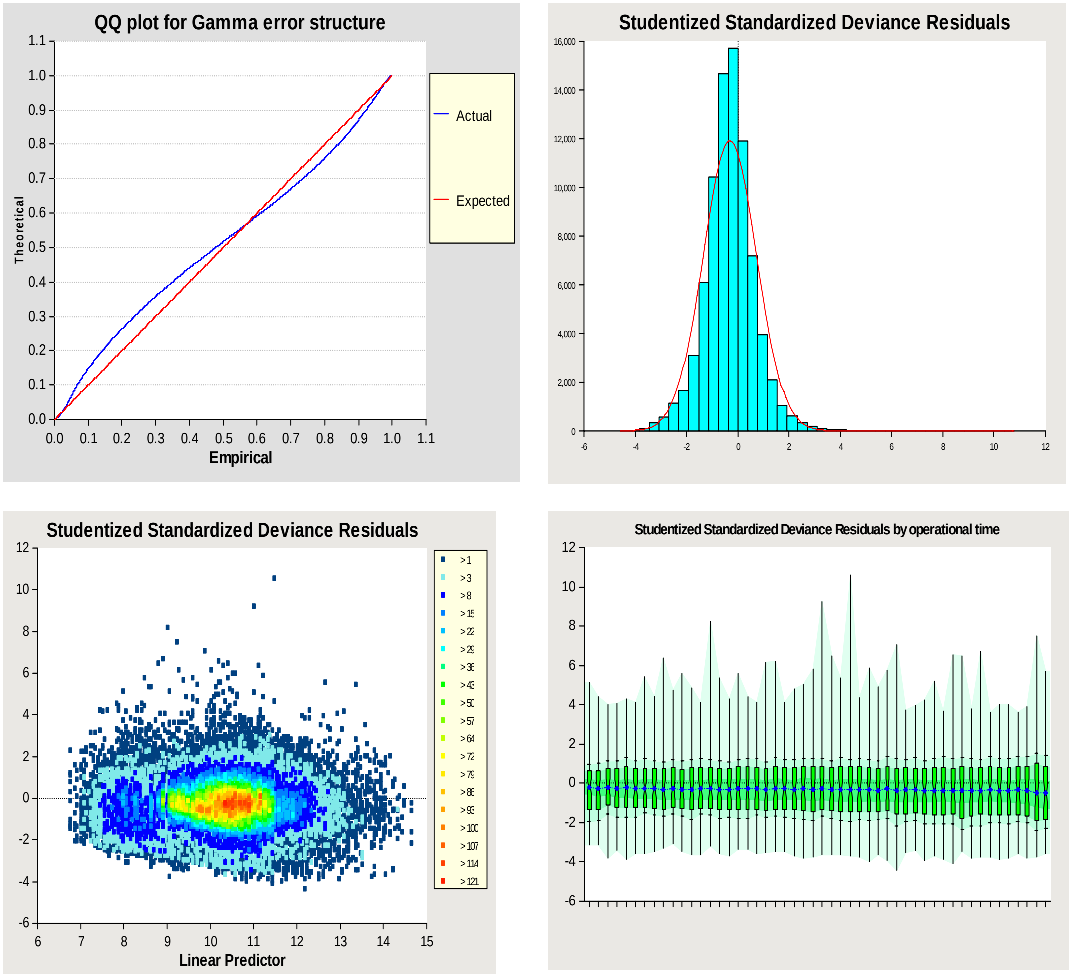

Since the other residual plots look satisfactory and segmentation of the data does not improve the results, modelling all finalised claims together appears warranted. As well as examining a gamma distribution, EDF(p) for p=close to 2 were also examined (e.g. p=1.8, 2.2). None of these showed any improvement over the gamma distribution; thus a gamma error was chosen.

The selection of an EDF(p) distribution with $p<2.3$ is unsurprising; by factoring in severity the model now has distinct groups with very different averages and thus the distribution does not require as long a tail to capture the high claim sizes as the overall model did.

4.2.2 Modelling data up to the legislative change

The process used here to fit the model involves firstly incorporating the main effects (severity, operational time, finalisation quarter) as categorical factors.

Thus, operational time might have 50 levels (one for each 2% band), finalisation quarter might have a different level fitted to each individual quarter etc. The continuous effects (operational time, finalisation quarter [which represents any superimposed inflation effect]) may then be simplified by representing them by continuous terms.

Figure 4.3 Operational time

Effect fitted = L_optime (0-100%) + L_optime (0-6%) + L_optime(6-20%) + L_optime (80-100%) + I(optime > 98%)

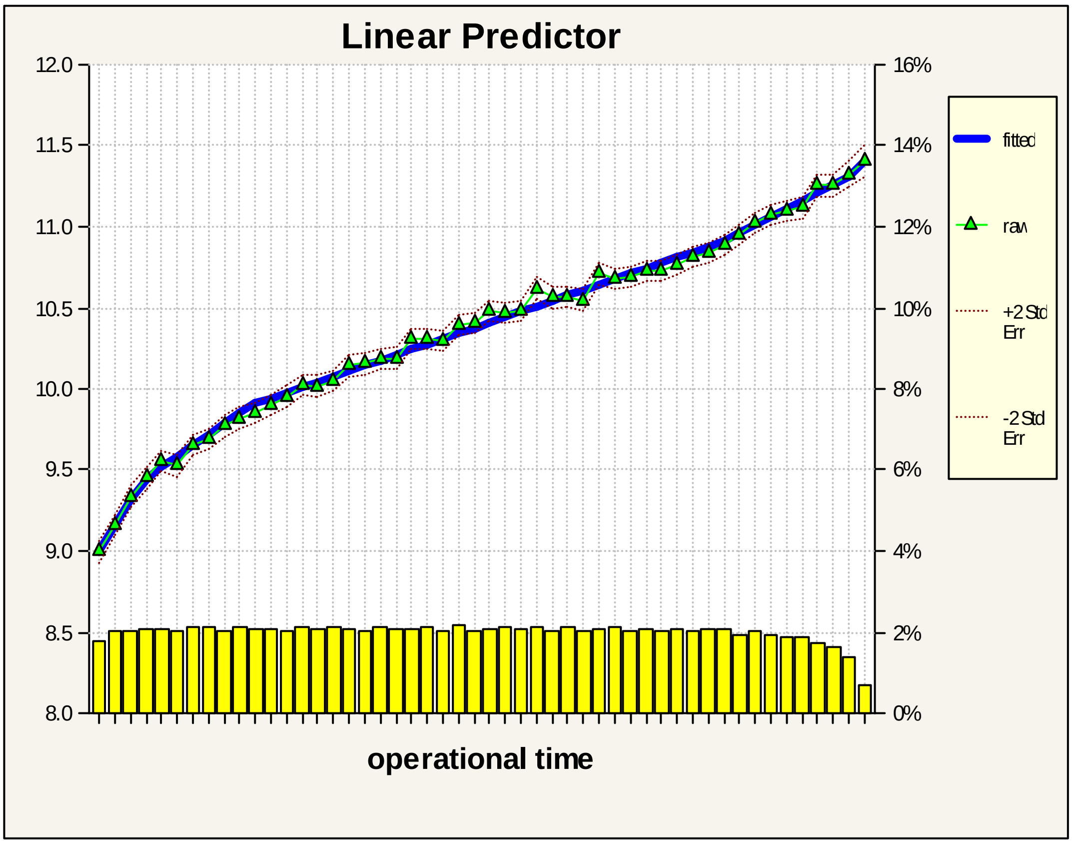

Figure 4.3 above and Figure 4.4 below display the results of this process for operational time and finalisation quarter respectively. The fitted effects are shown in both graphs, with the convention that "L_" before a factor represents a linear (i.e. continuous) effect involving that factor while "I(.)" represents a jump factor applying to the identified levels of that factor. For example, "L_optime(6-20%)" represents a linear effect applying to operational times from 6% to 20%, with the line flat up to 6% and flat from 20% onwards, while I(optime>98%) means that a jump effect is included for operational times greater than 98%. Thus the fitted effects listed in Figure 4.3 fit a series of linear terms to operational time with knots (turning points) at 6%, 20%, 80% and a jump at 98%.

In both graphs, the yellow bars are a measure of exposure, the green line represents the fit at categorical levels while the blue line is the smoothed fit. The purple dotted lines show the error bounds (two times standard error) on the initial categorical fits.

Figure 4.4 Finalisation quarter

Effect fitted = I(finqtr < Mar97) + I(finqtr = Mar04) + I(finqtr = Jun04-Jun05) + I(finqtr > Jun05) + I(finqtr = March) + L_finqtr(Sep94-Sep00) + L_finqtr(Sep00-Dec02)

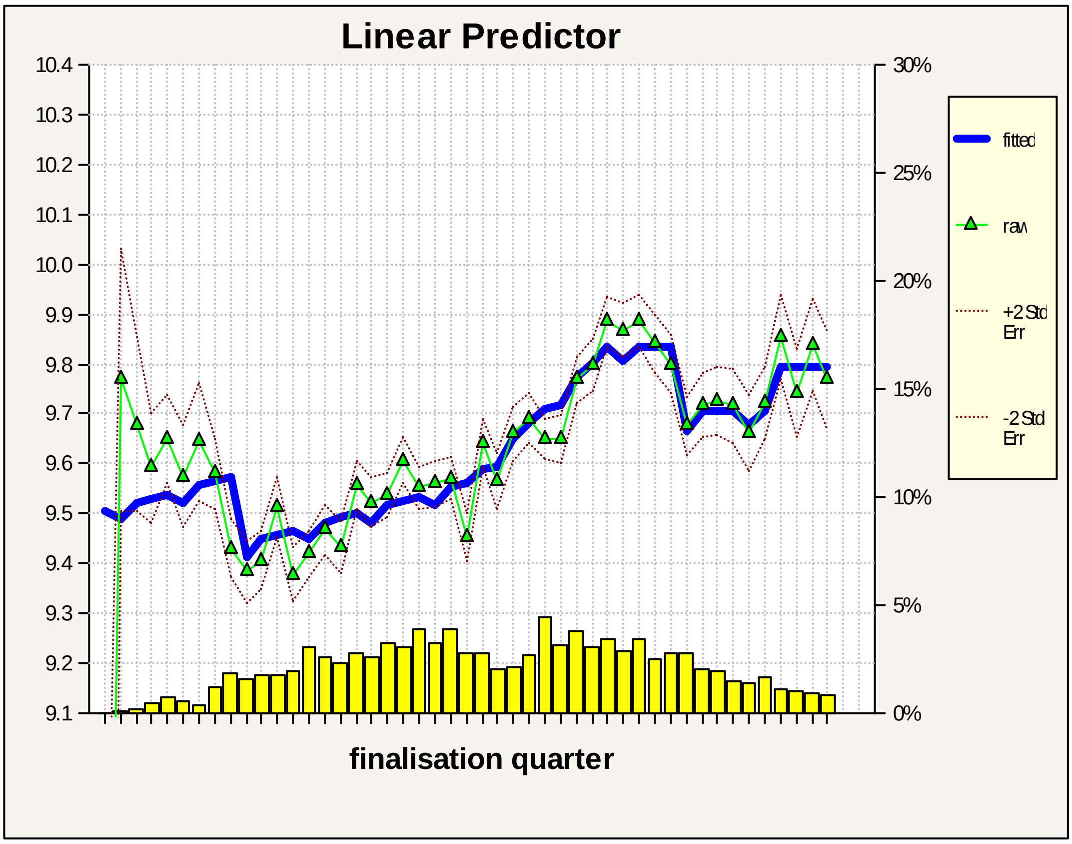

It is seen that the fitted operational time curve is a good fit to the raw (categorical) values, the fitted line generally staying within the error bounds of the raw fit. The finalisation quarter line (which includes a seasonal effect, as well as some one-off reductions or increases, e.g. at Sep05, as well as linear effects to September 2000 and from September 2000 to December 2002) looks more variable, but again is generally contained within the error bounds. It does deviate from the apparent trend in early finalisation quarters but no attempt has been made to fix this (this would have little impact on the model results due to its early occurrence and the low amounts of data in this part of the experience).

The finalisation quarter trend above (the estimate of past superimposed inflation) would be difficult to extrapolate into the future given its shape. Judgement would be necessary to ascertain appropriate superimposed inflation assumptions for these data. Such questions are beyond the scope of this paper; the results presented herein assume no future superimposed inflation.

The next step in the fitting process is to consider how the model might differ by severity. Initially, one might be interested in how the operational time effect (see Figure 4.3) might vary for each of the severities. A first step to investigating this is to include interactions between the raw operational time effects and each severity. This is shown in Figure 4.5 below. Note that the severity 1 effect is smooth since severity 1 is the "base" level for the severity class and therefore, in the presence of interactions the fitted effect for severity 1 is simply the main operational time effect, fitted above in Figure 4.3. However, with the inclusion of interactions for the other severity classes, this effect is modified to reflect severity 1 experience only.

Figure 4.5 Operational time effect split by severity

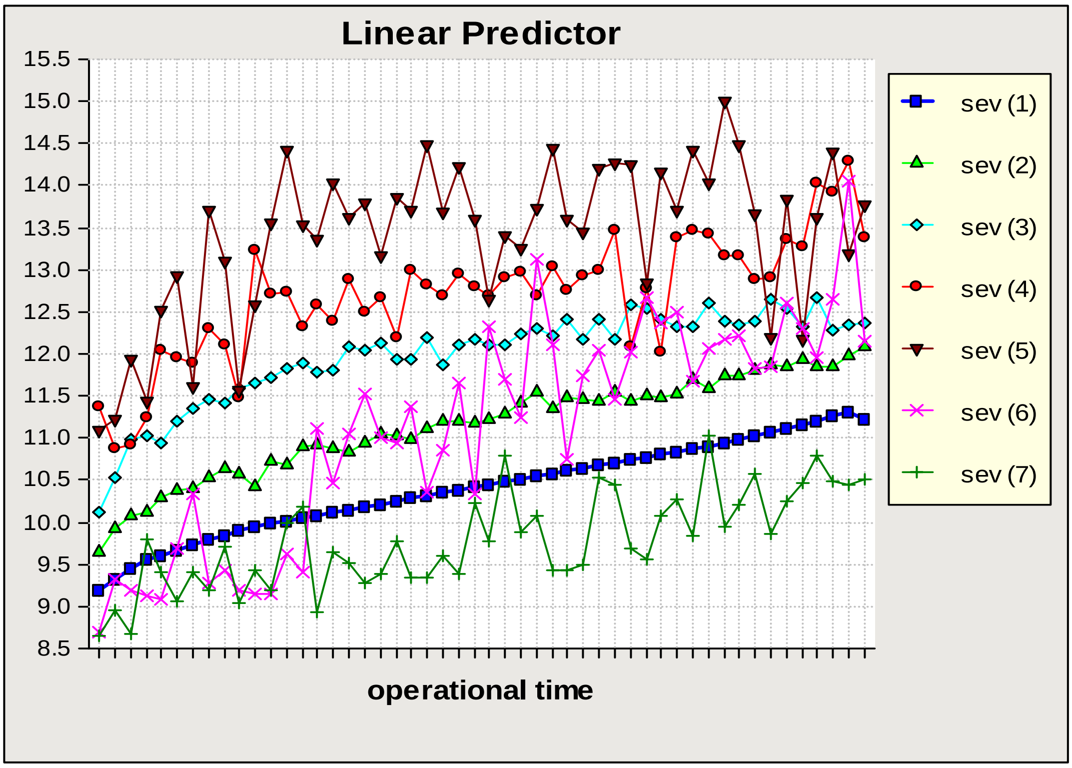

Since the model has already incorporated simplified (i.e. continuous rather than categorical) effects for operational time, a first step to appropriately modelling interactions might be to interact severity with some of those continuous operational time effects. Here, severity is interacted with the linear operational time effect that runs across all operational times (i.e. from 0% – 100%). This interaction is shown in Table 4.1.

Table 4.1 Significance testing of interactions

Interaction	parameter	standard error	approx t-test
sev (2).L_optime(0-100%)	0.0010	0.0004	**2.51**
sev (3).L_optime(0-100%)	-0.0012	0.0006	**-1.98**
sev (4).L_optime(0-100%)	0.0004	0.0015	0.27
sev (5).L_optime(0-100%)	-0.0029	0.0024	-1.22
sev (6).L_optime(0-100%)	0.0215	0.0013	**16.30**
sev (9,BL).L_optime(0-100%)	-0.0038	0.0009	**-4.18**

Approximate t-tests (i.e. parameter mean / parameter standard error) may be used to identify significant interactions. T-test value of 2 or above (equivalently -2 or below) indicate significant parameters. From Table 4.1, it is seen that the interaction with severities 2, 6 and 9+blank are significant while that with severity 3 is borderline. This suggests dropping the interactions with severities 4 and 5. For the time being, the interaction with severity 3 is retained; this will be reviewed at the end of the model fitting process.

Figure 4.6

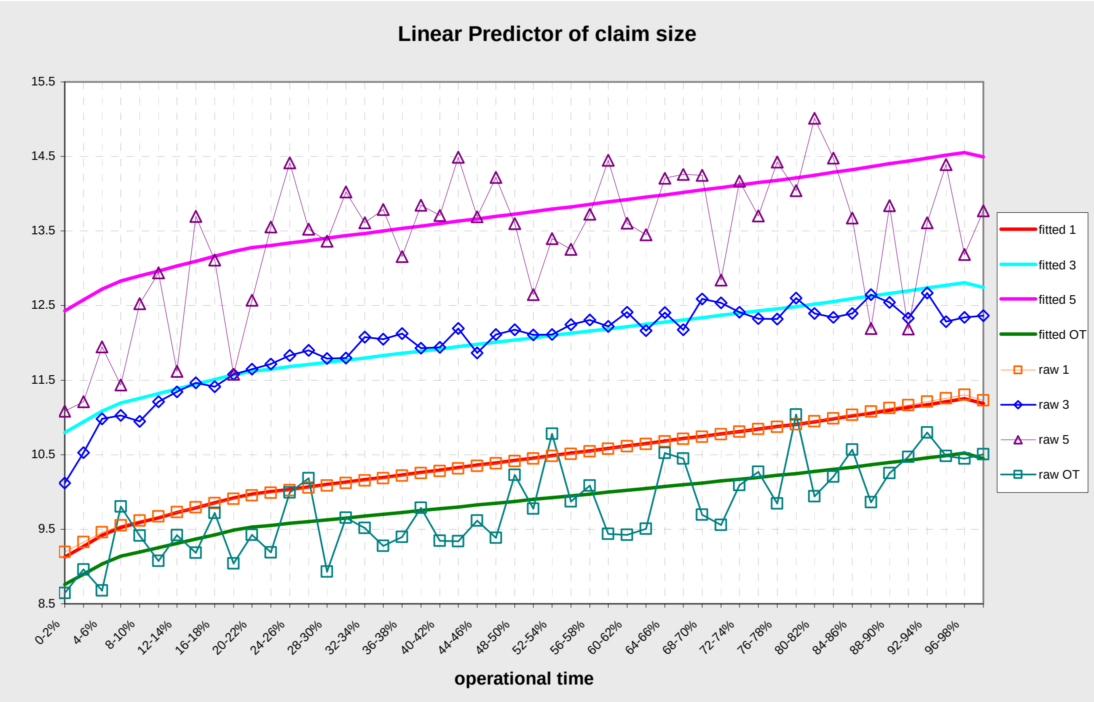

Figure 4.7

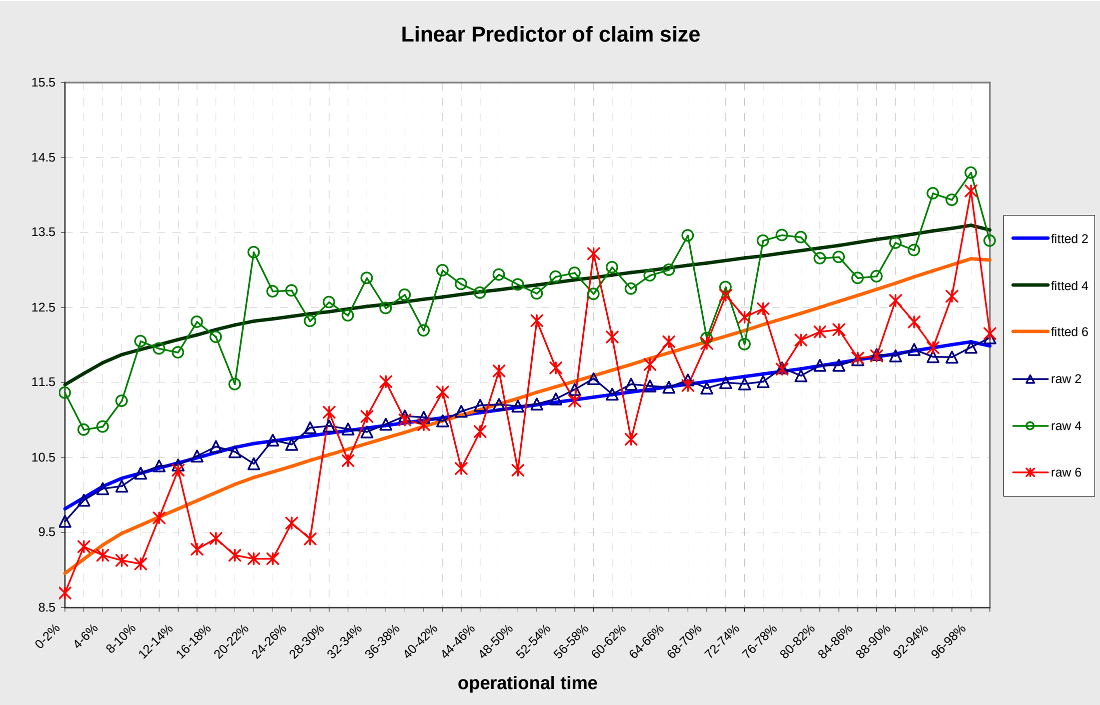

Figure 4.6 and Figure 4.7 above display the results of this modelling exercise. Generally the correspondence between the raw and fitted effects is good. There are perhaps some areas of divergence but these are not serious. Thus, on the grounds of both statistical significance and usefulness of the models (in other words, a model which fits every last bump and twist is at best unhelpful and at worst, fitting noise rather than genuine effects), these simplified interaction effects are accepted.

Other statistical tests that may be used to determine the acceptability of simplified models include Chi-square tests (for nested models) and the Akaike Information Criterion (AIC). Any text book on GLMs will describe these tests (for example, McCullough and Nelder, 1989). By the Chi-square test, a model is not significantly worse if it has a high p-value (traditionally 5% or higher). Using AIC (which offsets the increased deviance in a simpler model against the increased parsimony, i.e. fewer parameters), a model with a lower AIC is preferred.

Table 4.2 Further statistical tests

	AIC	Model 1 optime . severity	Model 2 L_optime(0-100%) .severity	Model 3 L_optime(0-100%) .sev (2,3,6,9+Bl)
Model 1	1379463.727		pv=0.0% (Chi-Sq) AIC diff = 567	pv=0.0% (Chi-Sq) AIC diff = 563
Model 2	1380030.992			pv=48.4% (Chi-Sq) AIC diff = -4
Model 3	1380026.954			

Based on the statistics, both Model 2 and Model 3 (the simplified models) are worse than Model 1 (the model with the full set of operational time and severity interactions (339 parameters). However, both the number of parameters and the graphs in Figure 4.6 and Figure 4.7 suggest that this model is not useful. Model 3 is preferred since it is more parsimonious than Model 2 (note the lower AIC value), and is not significantly worse (by the Chi-Square test).

The important point to take from the above discussion is that statistical tests, particularly the Chi-square and AIC, should be used to guide the model fitting process rather than an absolutes. Clearly, Model 1 is over-fitting and is not useful in practice. These tests are most useful to aid the modeller in deciding between a number of reasonable parsimonious models.

Other testing processes may be used to discriminate between a number of competing models. One example is the learn and test process. Here the data may be split into a learning data set (say containing about 70 - 75% of the data) and a test data set containing the remainder of the data. The split should be random, but should ensure representation of all factor levels in each data set to the extent that this is possible. The model is constructed on the learning data set, then applied to the testing data set to check its validity.

Once severity and operational time have been modelled, the process continues. The next interaction that might be examined is that between finalisation quarter and operational time. The logic behind this interaction is that superimposed inflation (captured by finalisation quarter) might be expected to vary with operational time. Typically smaller claims may be subject to higher rates of superimposed inflation ("SI") than larger claims. In fitting this interaction, reference may be made to the overall model (see Section 4.1) from which it is observed that SI has a triangular shape – increasing to a peak at 10%, and thereafter decreasing. A similar shape is likely to be suitable here.

The severity model has the additional complexity of possibly different rates of superimposed inflation so three-way interactions would be needed here to test and model any significant differences between severities. This particular data set has the additional feature of one-off jumps in the claim size by finalisation quarter (refer back to Figure 4.4). Again it would be expected that these jumps might vary by severity.

Following this process, the model now consists of factors involving severity, operational time, finalisation quarter, with appropriate interactions.

4.2.3 Legislative effects modelling

4.2.3.1 2000Q4 legislation

The current model has factors involving finalisation quarter and operational time (which is of course a mapping of the development quarter). Legislative effects tend to be accident period effects. As any modeller knows, it is necessary to be extremely careful when including effects from correlated factors like accident, development and finalisation period. Thus, modelling a legislative effect is not simply a case of putting in the accident period factor and seeing what results, but is rather a case of applying judgement to fit suitable effects.

Figure 4.8   2000Q4 legislative effect

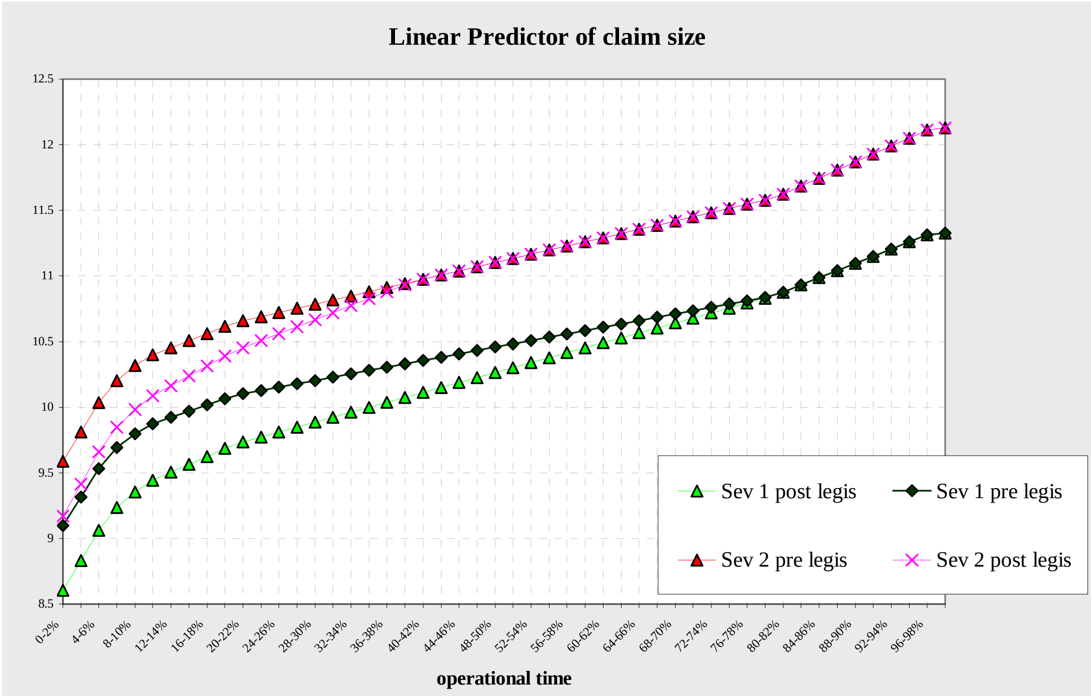

This is best illustrated with reference to the first of the pieces of legislation discussed in Section 3.2. This legislation, introduced in 2000Q4, has the effect of limiting legal (plaintiff) costs on small claims. Costs are not recoverable for claims which settle for under $30,000 and are partially recoverable for claims between $30,000 and $50,000 (on an increasing scale). Therefore, it would be expected that claim sizes, previously under $50,000 would be reduced by a decreasing amount (decreasing with operational time) rejoining the main operational time curve at the point where claim size is approximately $50,000.

The first point to note is that this will not apply to all severities. For example Severity 5 (catastrophically injured) is unlikely to have claim sizes that low. Secondly, the operational time joining point will depend on the severity. Severity 1 claim sizes are, on average, lower than Severity 2 sizes. Thus, the point of operational time at which this effect wears off should be lower for Severity 2. In fact, inspection of claim sizes leads to a joining point at approximately 80% for Severity 1 and 40% for Severity 2. Figure 4.8 above displays the fitted results (as a function of log(estimated claim size)) for Severities 1 and 2.

"Actual vs expected" triangles may be used to validate these legislative model terms. See the following section for a definition and some examples. The Chi-square test and AIC may also be used.

4.2.3.2 2002Q4 legislation

This legislation came into force in December 2002. Thus, the first full accident quarter of the Scheme under its influence is March 2003. As indicated in Section 3.2, this has a wide-ranging effect on claims. It would be expected that its effects would vary for claims of different severities. Therefore, the legislation should be considered separately for each claim severity. The largest class (Severity 1) is the logical beginning point.

Perhaps the best way to visualise the legislative effect for this severity class is also using a triangle, by accident and development quarter of actual/expected ("A/E") values where the actual values are the actual finalised claim sizes while the expected values are from the model fitted omitting accident quarters from 2002Q4 (but applied to accident quarters from March 2003 onwards).

Figure 4.9 Severity 1 Actual/Expected [pre-2002Q4 legislation model]

Mar-03		85%	56%	93%	52%	43%	38%	48%	44%	51%	73%	83%	81%	82%
Jun-03			73%	41%	46%	35%	34%	49%	50%	76%	80%	78%	59%	
Sep-03		51%	76%	45%	59%	41%	54%	50%	69%	64%	92%	83%		
Dec-03		36%	75%	41%	45%	42%	59%	60%	92%	89%	90%			
Mar-04		109%	43%	39%	36%	40%	59%	92%	85%	123%				
Jun-04	26%		47%	35%	30%	50%	49%	78%	87%					
Sep-04		7%	46%	41%	34%	51%	66%	84%						
Dec-04		32%	117%	42%	46%	67%	77%							
Mar-05		51%	42%	46%	47%	61%								
Jun-05		19%	64%	48%	74%									
Sep-05		54%	44%	48%										
Dec-05		38%	94%											
Mar-06		43%												
Jun-06														

Figure 4.9 displays this triangle of A/E values. Note that the A/E values are colour-coded (blue where actual is less than expected and pink where the actual is greater). Since the 2002Q4 legislation was intended to reduce claim size, the dominance of blue is not surprising. More important are the values of A/E since these will provide information on the effect of the legislation. Looking at these it is seen that the reductions relative to their pre-CLA estimates wear off with increasing development quarter (as shown by generally increasing A/E values). This is not surprising, for it indicates that larger claims are not affected as much by the legislation.

A more subtle, but important point may be ascertained by looking at the pattern of A/E values in any development quarter. Ignoring variations from statistical noise, the A/E values generally increase looking down each development quarter column.

A possible explanation is that the legislative effects appear to wear off over time (note this wearing-off is distinct from that affecting larger claims discussed above). This could be a general erosion of benefits as the relevant parties get to grips with the legislation. Alternatively it may be due to the so-called frequency effect – whereby falling frequency means that the smaller claims are being removed.

Figure 4.10 Frequency of Severity 1

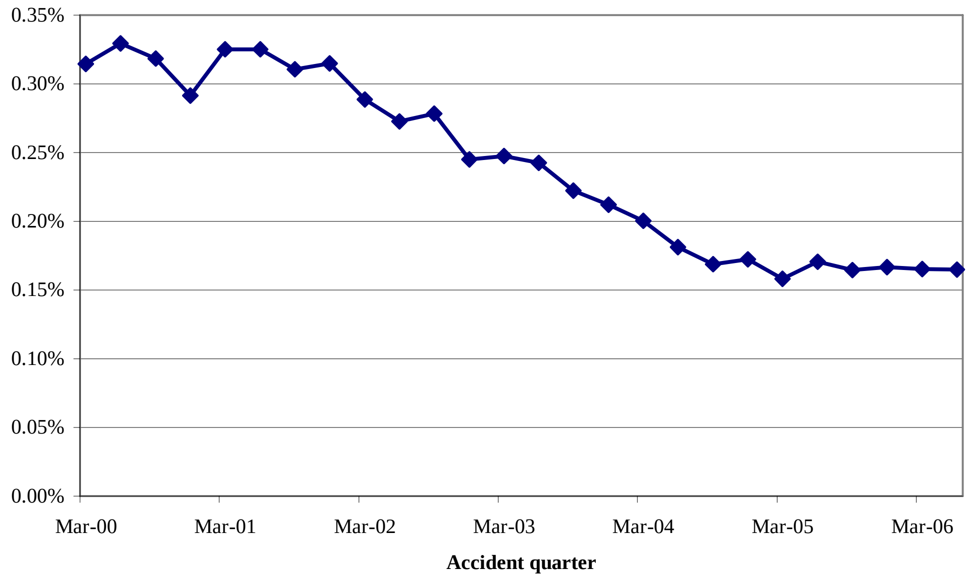

Figure 4.10 displays the frequency of Severity 1 since March 2000. The frequency is generally decreasing, most strongly from December 2001 to March 2005. Taking this and the results in Figure 4.9, the following argument may be made:

- Frequency of Severity 1 has been falling for some time. There may be an additional decrease due to the 2002Q4 legislation;
- It is possible that the additional decrease is due to the legislation having removed some of the less severe claims from the experience. Therefore, the claims that are left are the more severe of Severity 1 claims;
- Thus, it would be expected that, all other things being equal, the average claim size of Severity 1 would increase. This can happen even in the case where the legislative effect still does reduce claim sizes – the claims that are left are larger, so even reduced by the legislation as they may be, they may still be typically larger than claims settling at similar times before the legislation and before the frequency reduction;
- This would be reflected by apparent wearing-off of the legislative savings, with more wearing-off as the frequency reduces – the pattern that is seen in Figure 4.9.

The discussion to date in this section suggests that the modeller is faced with finding appropriate model terms for severity 1 to:

- Model a reduction starting in accident quarter March 2003 due to the 2002Q4 legislation. This reduction in claim size affects small claims more than large claims so it must wear off across development time as claims get larger;
- Further, the size of the reduction must be allowed to reduce with claim frequency.

The question arises of how much of the frequency effect is attributable to the legislation and how much might be a general reduction in claim frequency. The latter would be expected to affect all claims equally and not alter claim size; the former removes the small claims and therefore increases claim size. The model in this paper assumes that the entire frequency effect is attributable to the legislation. Given the presence of long term trends in frequency, this may be an unduly pessimistic stand; some further discussion on this point is given in Section 4.2.5.

Figure 4.11 Fitted Severity 1 2002Q4 legislative and frequency effect for different accident quarters

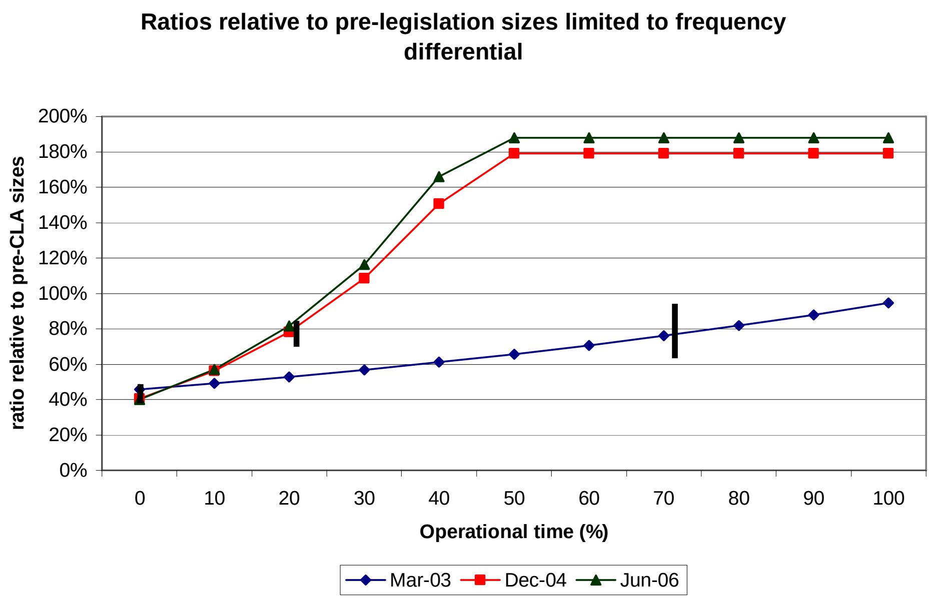

I(accqtr > Dec02){1 + log(pre-legis Sev 1 freq/Sev 1 freq in accqtr) [1 + 100-optime] }

Figure 4.11 displays the model terms used to model these effects as well as the shape of these effects. Taking these terms together, they lead to a reduction for each accident quarter which is at its maximum at 0% operational time and wears-off with operational time. Combined with this is a frequency effect captured by the relative frequency of the pre-legislation Severity 1 frequency to the frequency in the post-legislation accident quarter; the log of this frequency is used since a log link is used in the claim size model. The frequency effect leads to a reduction that is smaller for lower frequencies of Severity 1.

The resulting legislative effects (expressed as a ratio of post-legislation claim size to that pre-legislation) are given in Figure 4.11. The flattening out of the extrapolated December 2004 and June 2006 curves is a manual adjustment, not incorporated in the model terms. The reasons for this are discussed later.

The black bars on the graph in Figure 4.11 show the limits of the actual data. For example, for accident quarter March 2003, there are finalised claims with operational time of approximately 70%. However, for accident quarter June 2006, there are actually no finalised claims so observed operational time to date 0%. Therefore, due to the sparseness of experience to date at accident quarters with low frequencies, it is necessary to apply judgement to decide what the legislative/frequency effect should be at operational times beyond that observed in the data. To illustrate the importance of this, Figure 4.12 displays the effects extrapolating them directly from the model. It is unreasonable to believe that the frequency effect could lead to claim sizes more than ten times higher than their pre-2003 counterparts, no matter how large the frequency effect.

Figure 4.12 Post-legislation and frequency effects for severity 1 with no limits for different accident quarters

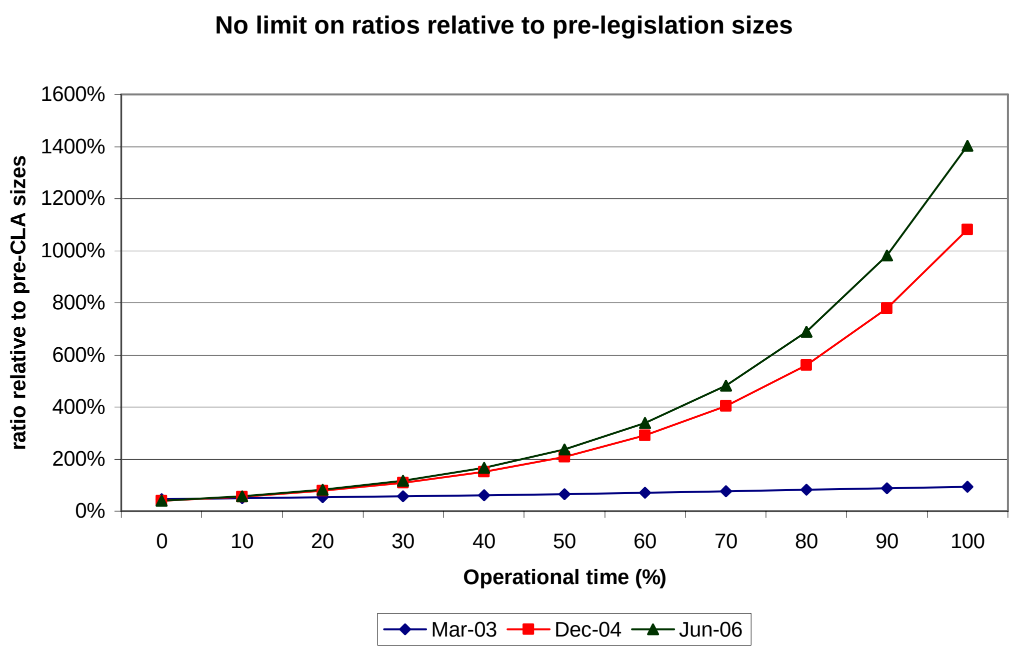

To assess what should be done to extrapolate the frequency and legislative effects, an illustration of the frequency effect may be helpful. The situation depicted in Figure 4.13 is a simple version of the frequency effect. The blue line represents the original claim size curve by operational time. Suppose the frequency drops by 30% and this involves the removal of the smallest 30% of predicted claims only (correspondingly the first 30% of operational time). In this case, this 30% of claims disappears; the remaining claims are those that fall between 30% and 100% operational time. Thus, the new claim size curve will be a distorted version of the original curve from 30% to 100% where the original 30% operational time point becomes the new claim size at 0% operational time. The old claim size at 100% operational time is still the new claim size at 100%. The old curve is then stretched to cover the entire length of operational time, with both curves eventually meeting at 100%.

Figure 4.13 A simple illustration of the frequency effect

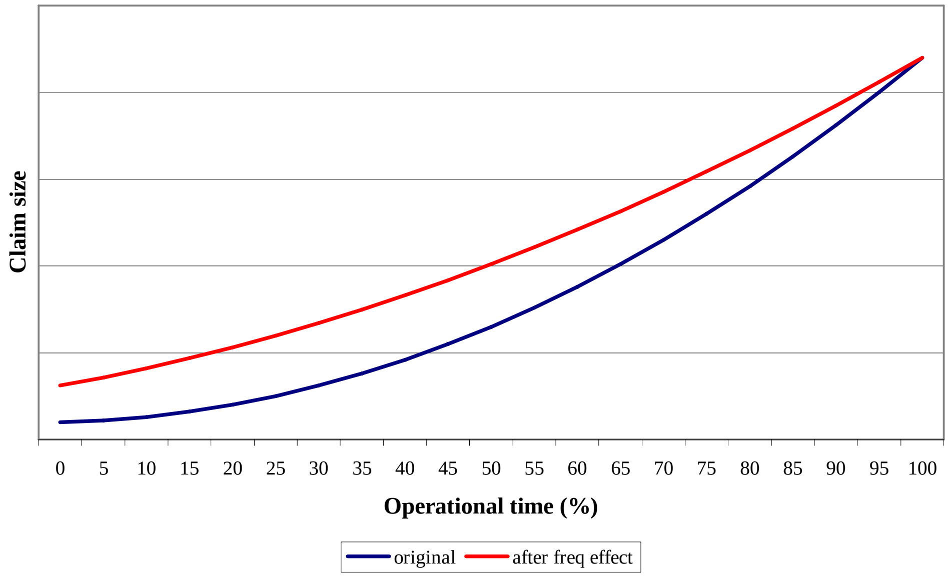

Therefore, this means for the model above that any ratios of post and pre frequency effect claim sizes should not get large indefinitely, but instead, increase for a time (recall that as well as the frequency effect, there is a reducing legislative effect on claim sizes), then start to reduce again, before the two curves meet again at 100%.

Referring back to Figure 4.11, it is clear that this is not something that cannot be accommodated within the model using the post legislation operational time – the experience is too young for the quarters most affected (i.e. the most recent quarters with the largest drop in frequency relative to the pre-legislation time period). Instead, the following procedure is suggested:

1. Accept the model results so long at the relativities between pre- and post-legislation claim sizes are sensible. This may be assessed by comparing the ratio of the Severity 1 frequency in a typical pre-legislation time (given the long-term trends in frequency, the average frequency in 2002 is a reasonable choice) with the frequency in that accident quarter. If this ratio is 180% (say), then a reasonable choice is to assume that the maximum relativity between pre- and post-legislation claim sizes is 180% (representing the maximum increase in average claim size in the extreme case where only the smallest claims have been removed);

2. Apply these relativities to produce post legislation claim sizes. However, these sizes cannot be correct for the entirety of operational time since the frequency effect dictates that eventually the pre- and post-legislation curves must meet (note this assumes that the legislation effect on its own does not reduce the largest claims, if it did then the two curves would never meet, all other things being equal, the post legislation curve would always be end up below the pre-legislation claim size curve);

3. To deal with this, use the pre-legislation claim size curve. Make the assumption that if the frequency is x% lower than before the legislation,

then the post-legislation curve is the stretched pre-legislation curve from x% to 100%, modified by the legislative savings effects. The mapping between the post-legislation operational times (running from 0% to 100%) and the pre-legislation times (from x% to 100%) can easily be determined. Thus, a stretched curve may be calculated for all new operational time points.

4. The estimated claim size effect is then a combination of these two lines; initially claim size is drawn from the model-based post legislation curve. This incorporates both the frequency effect and the legislative reductions. For most accident quarters, this curve crosses the (stretched) pre-legislation curve. At this point, any legislative reductions are assumed to no longer exists, and from there the selected post-legislation curve follows the pre-legislation curve.

Figure 4.14 Example of the projection process for Severity 1

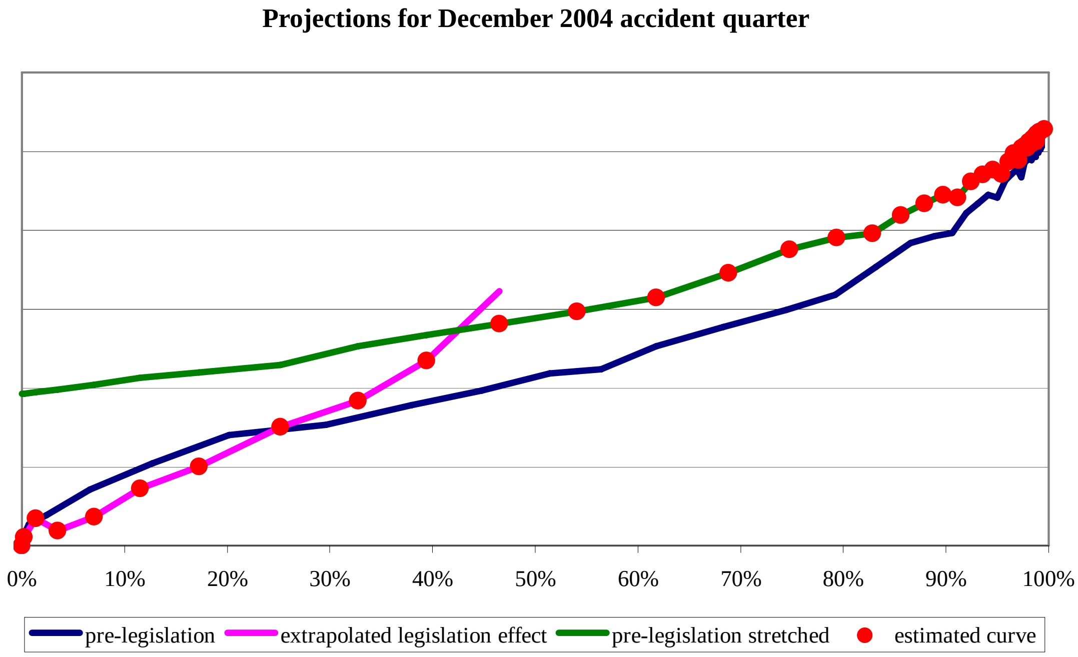

Figure 4.14 is a graphical depiction of this process for the December 2004 accident quarter. The estimated curve for December 2004 is a combination of actual experience to June 2006, the modelled post-legislative December 2004 curve, joining with the stretched pre-legislation accident quarter curve. A similar process is applied to all accident quarters. Note that it is possible for the modelled curve and stretched pre-legislation curve not to cross (this happens for the June 2003 accident quarter, for example). In this case the modelled curve is used across all operational times.

Modelling the legislation effect for the other severities is simpler. There is sufficient data only to examine an effect for Severity 2. Figure 4.15 shows the actual/expected figures for Severity 2, where the expected results are based on the pre-2002Q4 legislation model. Like Severity 1, this shows a reduction in claim size which wears off over development time (equivalently operational time). Unlike Severity 1 however, there does not appear to be a frequency effect, i.e., the reductions appear stable looking down development quarter columns. This ties in with the Severity 2 frequency (refer back to Figure 3.2); this has reduced post-legislation but there are no strong trends in the data. Thus, the modelling and projection of Severity 2 claim sizes is considerably simpler.

Figure 4.15 Actual/Expected for Severity 2

Mar-03	7%	49%	54%	56%	56%	30%	70%	92%	71%	60%	104%	102%	66%	78%
Jun-03	78%	12%		34%	37%	30%	42%	48%	122%	90%	63%	102%	69%	
Sep-03		14%	22%	89%	57%	45%	61%	59%	64%	65%	70%	64%		
Dec-03			68%	40%	49%	13%	49%	61%	61%	50%	79%			
Mar-04		49%	20%	38%	28%	64%	57%	57%	52%	73%				
Jun-04		48%	49%	56%	26%	23%	72%	91%	105%					
Sep-04		25%	21%	34%	24%	44%	67%	68%						
Dec-04	13%	17%	64%	73%	28%	48%	49%							
Mar-05		22%	41%	47%	44%	41%								
Jun-05	4%	13%	33%	40%	48%									
Sep-05		47%	33%	41%										
Dec-05	41%	39%	33%											
Mar-06		75%												
Jun-06														

Figure 4.16 A/E for Severity 2 after modelling the legislative effect

3-Mar	16%	116%	128%	130%	127%	63%	135%	165%	117%	88%	138%	125%	74%	80%
3-Jun	182%	28%		79%	86%	67%	84%	87%	200%	136%	86%	130%	84%	
3-Sep		33%	51%	204%	123%	89%	110%	97%	94%	86%	85%	72%		
3-Dec			160%	93%	107%	26%	89%	101%	90%	66%	97%			
4-Mar		115%	48%	88%	64%	139%	114%	100%	83%	104%				
4-Jun		113%	115%	131%	56%	46%	126%	143%	148%					
4-Sep		59%	49%	78%	51%	85%	116%	108%						
4-Dec	30%	40%	150%	170%	60%	96%	88%							
5-Mar		51%	96%	104%	94%	81%								
5-Jun	10%	31%	79%	93%	105%									
5-Sep		110%	77%	94%										
5-Dec	96%	92%	78%											
6-Mar		174%												
6-Jun														

Figure 4.16 displays the A/E triangle that results from modelling the severity 2 legislative effect. With its scattering of blue (A<E) and pink (A>E) and reasonable values of A/E, it is seen that the model fits well.

Figure 4.17 Severity 2 relativities with pre-legislation claim sizes

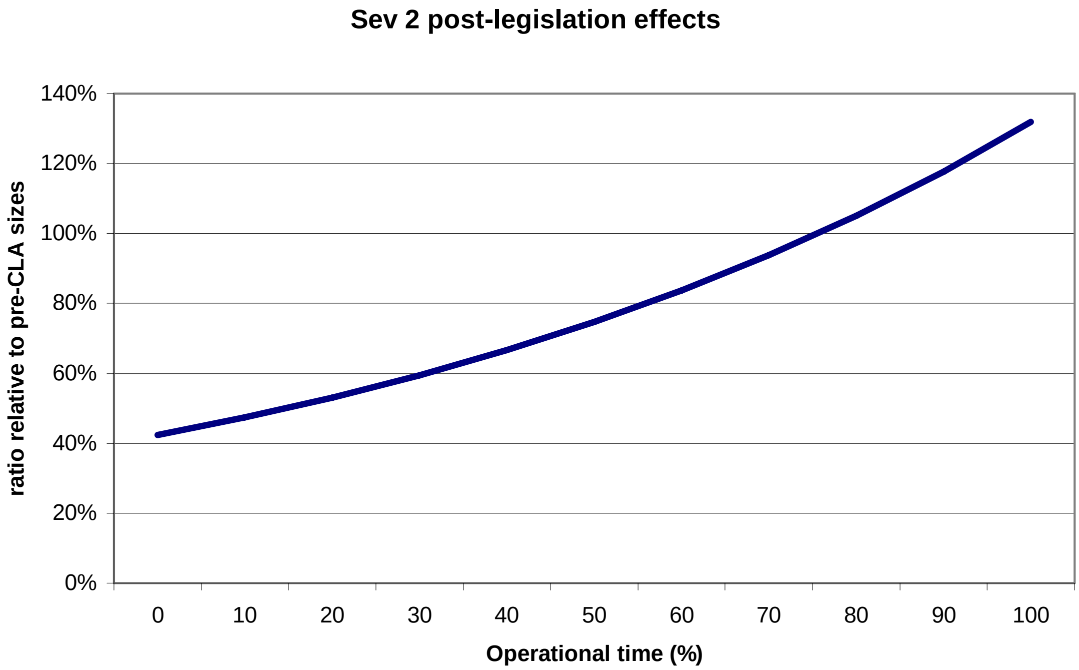

I(accqtr > Dec02) {1+ [100-optime]}

Figure 4.17 above displays the modelled relativities. Note that claims sizes at post-legislation operational times up to approximately 65% have been observed in the data. Therefore, any legislative effects above 65% are extrapolations. When selecting a projection scenario for Severity 2, consideration must be given to whether there is a frequency effect (as for Severity 1, but in this case the frequency effect may involve a one-off increase in claim size at the start of the legislation) or not. In the absence of a frequency effect, the maximum relativity of post-legislation to pre-legislation claim sizes would be 100%. If a frequency effect is believed to exist for Severity 2, the calculations akin to those for Severity 1 should be considered (i.e. imposing a maximum relativity based on the relative frequencies and merging the result with pre-legislation claim sizes). Taking all the post-legislation accident quarters together, the frequency ratio is approximately 140%, suggesting that a similar approach to that for Severity 1 may be required.

As stated above, there is insufficient data to attempt to discern the legislative effects for the higher severities. Therefore, for these severities an individual term I(accqtr > Dec02) is fitted to each. This helps to prevent any post-legislation claim sizes from contaminating the estimation of pre-legislation claim sizes. Negative parameter values resulting from this process may suggest that there have been reductions due to the legislation. However, in the absence of sufficient amounts of data, a prudent course may be to ignore this when projecting claim sizes.

4.2.4 Finalising the model of all data from September 1994 to June 2006

To date, the following model components have been examined:

- A comprehensive model of data from September 1994 to December 2002 (i.e. before the recent late 2002 legislation), including the claim size operational time curve and superimposed inflation effects, all tailored to each severity. Note this model also includes the 2000 legislation.
- A series of modelled relativities linking the post 2002Q4 legislation claim sizes with those from before the legislation. Note that these relativities are derived from a model of all the data.

Figure 4.18 Finalisation quarter effects by severity

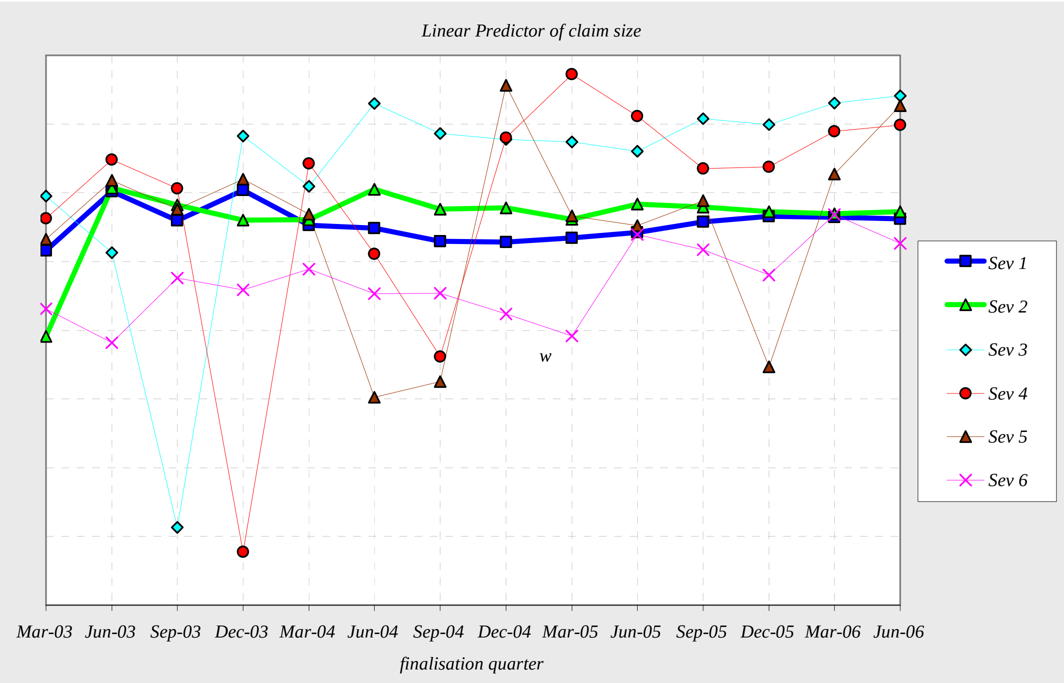

All that remains now is to examine whether there are any superimposed inflation effects peculiar to the post-legislation era. Figure 4.18 displays the results from fitting categorical finalisation quarter effects to each severity from accident quarter March 2003 onwards.

In early finalisation quarters (i.e. calendar years 2003 and 2004), there is very little data. For the higher severities, even from 2005 there is still a very small amount of data. Thus, any real effects would most likely be seen in severities 1 and 2 (highlighted in Figure 4.18). Severity 2 appears reasonably flat in 2005 and 2006 suggesting no superimposed inflation trends there. There appears to be a dip in Severity 1 between June 2004 and June 2005 but no trend. The dip may be included in the model to better estimate past claim sizes, but its inclusion or exclusion is unlikely to greatly affect the projections.

The final model contains 52 parameters. Between them these provide separate models for each severity, taking into account past superimposed inflation and two legislative changes. It is not particularly helpful to quote the model in full here; various terms in it fitted have illustrated in Sections 4.2.1 and 4.2.3 and the final version is a severity specific model of a form similar to the overall model from Section 4.1. What is perhaps more helpful is to contrast this model with others in terms of its parsimony. A traditional chain ladder model for each severity, if such models could capture the intricacies of this data (which is doubtful, see Taylor and McGuire, 2004 for some discussion on this topic) would have at least 96 parameters per severity (this assumes that different sets of chain ladder factors are not needed for accident quarters before and after December 2002). Another option might have been to build separate GLMs for each severity. The mechanics of modelling are easier admittedly in this case, but it is considerably less parsimonious in that it makes no allowance for similar trends between severities. Further the estimation of any such shared trends is less efficient as each model only uses that data specific to its severity.

4.2.5 Model results

The actual/expected triangle for the final model is displayed in Figure 4.19. For display purposes the numbers have been suppressed, while the 2000Q3 and 2002Q4 legislation have been marked with lines. It is seen that this triangle looks satisfactory.

Figure 4.19 A/E for final model, all severities

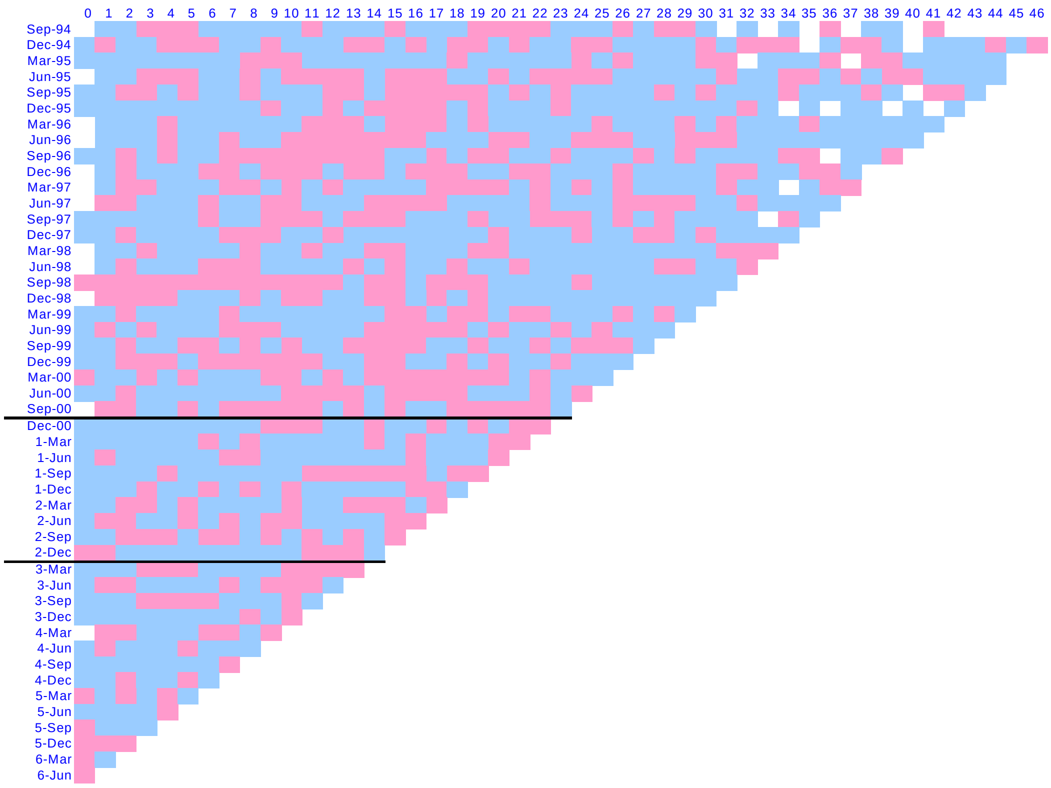

A summary of actual/expected by accident, development and experience (calendar) quarters is given in Figure 4.20. This generally looks satisfactory.

Figure 4.20 Actual/expected summarised by accident, development and experience quarters

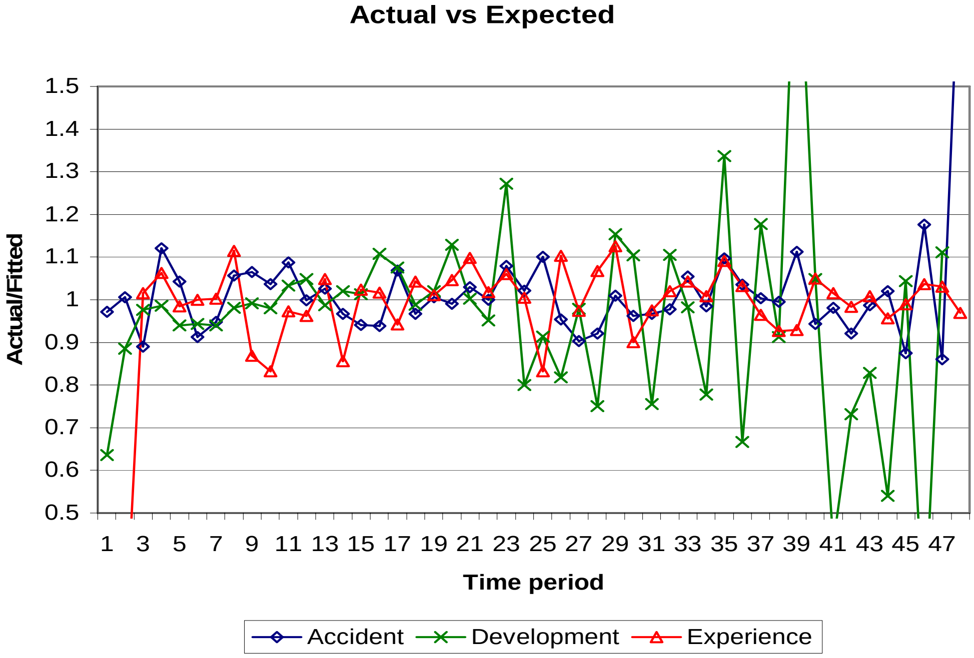

Figure 4.21 Comparison of claim sizes from overall model and severity specific model

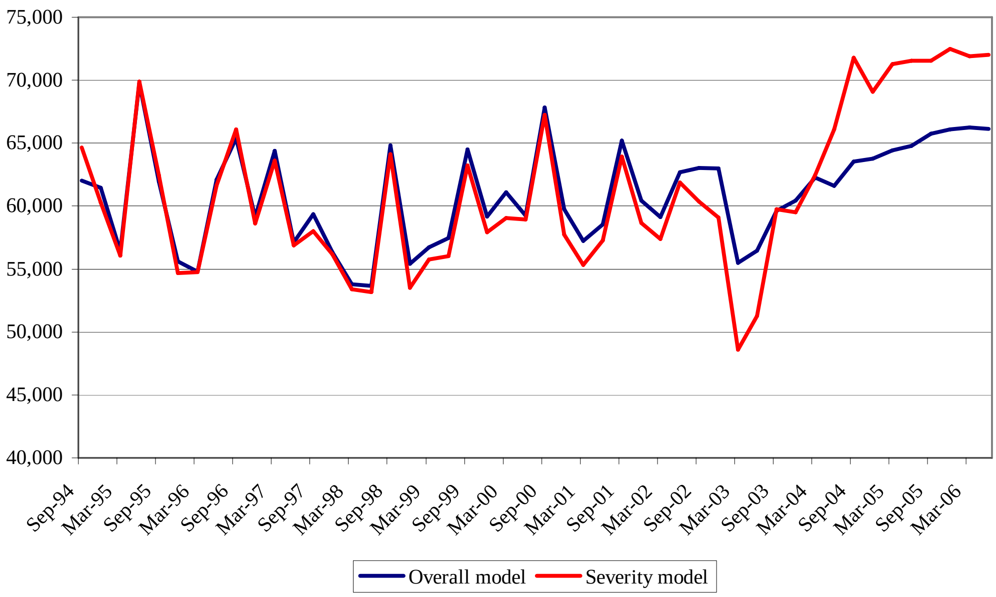

Figure 4.21 displays the average claim sizes, as projected by the overall model (Section 4.1) and the severity specific model (Section 4.2). For older accident periods, the two models naturally give very similar results since there is a lot of experience for these accident quarters. However, in recent accident quarters, the results diverge, with the severity model significantly higher than the overall model. This result was foreshadowed in Section 3.1 where the effect of

claim frequency changes on the average claim size was discussed. Note that both sets of results exclude any future superimposed inflation.

In Section 4.2.3.2, it was assumed that all the frequency reductions since 2003 in Severity 1 (and 2) were attributable to the 2002 legislation. In fact, referring back to Figure 3.2, there appear to be long term trends in the frequency in these severities and others that pre-date the legislation. As indicated in Section 4.2.3.2, assuming that all the frequency reductions apply to small claims only may be a pessimistic stance in the presence of long term frequency reductions (which would generally have no effect on claim size).

This being the case, judgement may be used to lower the claim sizes for Severities 1 and 2. This would take into consideration the long terms frequency trends, considering the attribution of reductions since 2003 between the long term trend (no effect on average claim size) and frequency reductions due to the legislation (assumed to increase claim size).

5. Discussion

This paper has discussed the fitting of a severity based claim size model for CTP data. It has covered such issues as how to use interactions to produce a model that takes into account different features of different severity classes.

Further, the data used in this paper was subject to major legislative change in late 2002. Ways of modelling this experience, including extrapolating the immature experience to produce claim size projections have been discussed.

Outside the scope of this paper has been any consideration of the subsidiary modelling that is required to use a severity specific claim size model. A number of non-trivial issues must be faced before severity specific claim sizes may be used. Perhaps the most significant is to do with future changes in injury coding. Until a claim is finalised (and in theory, not even then since it may reopen), the severity assigned to a claim is subject to change as more information is received about a claim. Thus, deriving ultimate claim numbers (and hence frequencies) in each severity is not simply a case of forecasting IBNRs for each severity (no easy job in itself, but a severity differentiated GLM of claim numbers may be built following similar principles to those outlined above). Rather it is also necessary to estimate out how many reported claims and how many IBNRs make transitions from one severity to another. Thus, models of transitions between injury severity levels are required.

The ultimate numbers of claims in each severity are required as inputs for the severity specific PPCF model since each severity must be assigned its own operational time. The severity specific model, being a PPCF model, also requires projections of future claim finalisations, to correctly allow for any future superimposed inflation effects or to estimate the payment pattern.

Splitting the data up into different severities has greatly increased the homogeneity of each group within the model, and therefore leads to modelling improvements (as suggested by the comparison of projected claim sizes in Figure 4.21). However, with approximately 70% of all claims, Severity 1, in particular, may still be subject to significant heterogeneity. Indeed the graph in Figure 3.1 suggests that there are a wide range of claim sizes in this group. It may be desirable to split Severity 1 into less severe claims and more severe claims. There are various possible ways for doing this. Some examples include the use of injury code data or legal status (e.g. legal representation, litigation). Whiplash claims, for instance, are a large grouping within Severity 1. If case estimates by head of damage were available, then a split by the level of general damages estimates associated with each claim may be another alternative. Any further splits would lead to more groupings (thereby increasing the model complexity) but the same modelling principles would still apply.

The current model is unsatisfactory in that incorporation of the frequency effect in Severities 1 and 2 requires two sub-models model (refer to the discussion of Figure 4.14), while a further judgemental adjustment is required to take account of any long term reducing trends in the claim frequency. It would be desirable to have an integrated process, producing claim sizes directly from one model; this is the subject of current work.

Although the model discussed here is an individual claim model, it is not a Stochastic Case Estimation ("SCE") model (e.g. Brookes and Prevett, 2004) since projections are only available on an aggregate basis, not at an individual level. Some discussion of the relationship between models of this type and SCEs is given in Taylor, McGuire and Sullivan (2006).

Acknowledgments

I would like to thank Greg Taylor, who played a leading role in a lot of this work. I would also like to thank Richard Brookes for his comments on earlier drafts of this paper.

References

Brookes, R and Prevett, B. (2004). Statistical case estimation modelling – an overview of the NSW Workcover model. Institute of Actuaries of Australia, Xth Accident Compensation seminar, 2004.

McCullough, P and Nelder, JA (1989). Generalised Linear Models, Second edition. Chapman and Hall, London.

Taylor, G and McGuire, G (2004). Loss reserving with GLMs: a case study. Institute of Actuaries of Australia, Xth Accident Compensation seminar, 2004.

Taylor, G, McGuire, G and Sullivan, J (2006). Individual claim loss reserving conditioned by case estimates. Published on the website of Institute of Actuaries, UK, at

http://www.actuaries.org.uk/files/pdf/library/taylor_reserving.pdf