Institute of Actuaries of Australia

Uncertainty-based framework for setting industry premium rates in workers' compensation

Prepared by Ivan Lebedev and Harry Rao

Presented to the Institute of Actuaries of Australia
12th Accident Compensation Seminar
22-24 November 2009
Melbourne

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

Uncertainty-based framework for setting industry premium rates in workers' compensation

Ivan Lebedev and Harry Rao

Abstract

The paper systematically examines key aspects of a typical industry premium rates review, from the selection of a suitable relativity measure to the approaches for dealing with sparse data. Using an empirical claim size distribution and Poisson distribution for the number of claims, we develop a statistical model of industry cost-ratios (claims costs/(remuneration x AWE)). This model is used to assess the experience period length requirements for estimating premium relativities and to validate the estimates against actual observations. We also discuss a simple practical approach to dealing with industries with limited claims experience which may provide an alternative to the hierarchical credibility modelling.

Key words: industry premium rates, workers' compensation, claims cost distribution, uncertainty, adaptive period-length

1 Introduction

Charging different premium rates from different industries aims to achieve a fair distribution of the total cost of workers compensation, whereby employers in more hazardous industries with higher cost of claims pay more than their counterparts in less hazardous industries. While premium rating generally is a rather non-trivial exercise, two features more or less specific to workers' compensation create additional challenges: long-tail nature of liabilities and shortage of claims experience as many industries generate only a small number of claims per year.

In order to calculate premium rates one has to decide what period of historical experience to include in the analysis, how to define 'claim cost' (eg use burning cost, fixed period paid-up cost or full incurred cost), at what level to cap large claims and how to deal with industries with insufficient amount of claims experience. In practice, such decisions are usually made either because they simply appear reasonable or even more commonly because 'it has been always done this way'.

This paper seeks to clarify this situation and develop a better understanding of all fundamental elements of premium rate calculation – claim size, aggregate claims costs, underlying and observed relativities, confidence intervals for estimation errors, etc. To our surprise, we found the literature on this topic in the workers compensation context virtually non-existent.

Reviewing industry relativities every year is generally considered absolutely essential and is justified by the need to respond to the emerging experience. It appears that there exists a commonly shared view that industry relativities vary over time, ostensibly in response to improvements in OH&S standards. However, the results of this paper suggest that at an industry level year-to-year changes in claims costs are quite random and for most industries systematic trends, even if they exit, are impossible to measure due to the lack of data.

A standard approach to premium rates review is to analyse all industries over the same fixed historical period (e.g. last 3-5 years) and then use hierarchical credibility approach for industries with insufficient experience. In this set-up, the rates for small industries are almost totally dominated by the rates of industry groups higher in the hierarchy. The method developed in this paper suggests first using industry's own experience to the largest extent (going back as far as possible if needed) and only then utilising collateral information. This minimises the reliance on collateral data that may be unreliable because industries comprising a particular industry group can have quite different risk characteristics.

The paper is organised as follows. Section 2 presents a quick tour of claims experience at the industry level in order to set the scene for the subject of this paper. Section 3 formally sets out industry premium rate calculation as a mathematical problem and discusses the choice of claims costs measure and paid-up period length. In Section 4 we model the aggregate claims cost as a compound Poisson process with an empirical claim size distribution derived from observations. The key result of this section is the approximate relationship between the observed number of claims and the coefficient of variation of the underlying relativity measure. Section 5 presents the adaptive experience period-length approach and validates it using quantile-quantile diagram test. In Section 6 we discuss a method of 'minimal disturbance' that can be used as a simple alternative to credibility weighting and validate it using quantile-quantile plots. Finally, Section 7 summarises our results.

2 Claims cost experience at industry level

To set the scene for developing a mathematical framework for industry premium rate setting, it is worth having a quick tour of actual industry-level experience to better understand the object of this study. In what follows the word 'industry' will refer to a group of employers with the same South Australian WorkCover Industrial Classification (SAWIC) code. At present SA economy has approximately 400 active SAWIC codes, some of which are large while others may only have a handful of employers. Given a large number of industries, it is not practically possible to present the historical experience of them all. The approach we have taken here was to draw a random sample of three industries form each of top-level industry divisions. By being random the sample gives an unbiased view of the typical industry-level experience[1].

---

[1] The industries randomly selected for Appendix A excluded those that had so few claims that there was nothing to show on the graph. For some industry divisions (e.g. 'Low-risk trade' the proportion of such industries is quite large).

The graphs given in Appendix A show industry cost-ratios (cost of claims/(remuneration/(52xAWE))) for accident years from 1995 to 2006. The cost ratios are based on two-year paid-up costs capped at $72,000 and adjusted for Scheme-wide improvement in claims costs as explained in Section 5.2. Also shown on the graphs are the 80%-probability confidence intervals for each individual year and the 12-year total experience. The confidence intervals are based on the statistical approach presented later in the text. Cost-ratios for accident years with less than 10 claims are not shown in the graphs.

In our view, the key feature seen on the graphs in Appendix A is that at an industry level systematic trends in cost ratios are normally not observed. In fact, in most cases the cost-ratios for individual accident years are visually indistinguishable from the ones that can be randomly generated by stationary stochastic process whose parameters have not changed over the years (the output of such a process is shown for comparison at the end of Appendix A).

The industry-level cost experience suggests that viewing industry premium rating exercise as a means of pro-actively responding to emerging experience or even pre-empting future trends is rather idealistic. Given the inter-annual variability of cost ratios, a more appropriate objective for the annual premium rate review appears to be assimilating new data into the existing experience in order to more accurately determine the constant underlying relativities.

3 Industry premium rate setting as a mathematical problem

3.1 Overview of industry premium rate calculation procedure

Calculating industry premium rates typically involves apportioning the target average premium rate (APR) across industries to find so-called base premium rates and converting the set of base rates into final, or table, rates that satisfy various constraints such as the cap on rate itself, the cap on the change of rate from previous year, etc. This paper will only focus on finding the base rates as the transition from base to table rates is a mechanical exercise that is quite straightforward.

Introduce the following notations:
L = APR,
W_i = projected remuneration for industry *i*,
ρ_i = rate relativity for industry *i*,
r_i = base premium rate for industry *i*.

In industry premium rating, the APR and projected remuneration are considered given and the main task is to determine the relativities. Once the relativities are found, the base premium rates are calculated from the simple relationship

$$r_i = L\frac{\rho_i \sum W_k}{\sum \rho_k W_k} \qquad (3.1)$$

It is easy to see that the premium rates given by (3.1) are proportional to relativities and in total provide the target aggregate premium collection given by $L\sum W_k$.

3.2 Underlying cost ratio

So far we have treated relativities ρ_i merely as a set of numbers whose ratio represents the relative riskiness of any two industries. In practice, relativities must be associated with some sort of an outcome characteristic that measures industry riskiness with respect to claims costs. Consider first an idealised case where all industries are large enough to give statistically significant experience and the claim payments for a given accident year have fully run off. If we had known the ultimate cost of claims from the start, then the most equitable way of sharing this cost between the industries would be in proportion to cost-ratios

$$C_i = \frac{s_i}{U_i} \qquad (3.2)$$

where s_i is the total observed ultimate cost of claims and $U_i = W_i / (52 \times AWE)$ is the inflation-adjusted remuneration, or approximate number of Full Time Equivalent (FTE) units.

Assume now that the ultimate costs are still known, but not all industries have statistically significant experience. In this case it can be argued that the aggregate cost should be shared in proportion to the expected value of cost ratios,

$$\kappa_i = \frac{E[S_i]}{U_i}, \qquad (3.3)$$

where S_i is the random variable that represents the total ultimate cost of claims.

To illustrate the difference between (3.2) and (3.3) consider a situation where in a particular accident year a given industry incurred a very high cost of claims. If the industry is large, then with the benefit of hindsight one can say that for this accident year the industry should have been charged a high premium rate. However, if the industry is small and its theoretical mean cost ratio is low, then even in hindsight one would not charge a higher rate because the bad experience was random and not characteristic of the industry's true riskiness.

In what follows, the theoretical value κ_i will be referred to as the underlying cost ratio, as opposed to the observed cost ratio C_i given by (3.2).

3.3 Choice of historical data for relativity estimation

Calculating industry premium rates for the forthcoming year involves two distinct tasks:

- Estimating past industry relativities from the historical data, and
- Forecasting industry relativities in the target year from their past values

Whilst forecasting is typically done in a very primitive way (i.e. one simply assumes that future relativities will be the same as most recently observed ones), the need to forecast puts important restrictions on the choice of historical data. Long paid-up periods more accurately estimate ultimate costs, but introduce a lag between the historical period over which the relativities are being estimated and the target year over which they will be applied. Short paid-up periods have the opposite effect.

A common method to "overcome" this problem is to calculate the costs over the whole lower part of the run-off triangle, the approach sometimes called the 'burning cost' method. The burning cost calculation utilises all available data and appears to be a good compromise. However, in our view this approach suffers from major problems as discussed below.

Consider a burning cost calculation over N years. In this method, the exposure is the total remuneration over this period and the outcome is the total paid-up cost of claims incurred over these N years. It is clear that in this set-up each accident year makes an equal contribution to the aggregate exposure (assuming year-to-year remuneration changes are minor), but unequal contribution to the aggregate cost, as illustrated in the following table (based on total Scheme costs).

Net incremental claim payments ($ 'million)

	DevYr						
AccYr	**1**	**2**	**3**	**4**	**5**	**Paid-up cost**	**Percent**
200406	50	93	95	55	46	339	32%
200506	52	101	78	69		300	28%
200606	52	96	86			234	22%
200706	51	94				145	14%
200806	53					53	5%
Total						1,071	100%

Imagine now a situation in which a particular industry undergoes a recent expansion, so that its remuneration increases in the year 2008, whilst the ultimate cost of claims per unit exposure does not change. In these circumstances we will see an artificial reduction in the burning cost ratio of this industry, because the denominator (exposure) will increase substantially more than the numerator (paid-up cost of claims). Conversely, an industry that undergoes contraction (due to real economic reasons or movement of employers into Comcare or self-insurance) will have an artificially increased burning cost ratio. These examples demonstrate that whilst attempting to include all available experience, the burning cost method introduces a bias and may respond to changing experience in the wrong way.

The other major shortcoming of the burning cost is that it is virtually intractable from the statistical view point, because this method mixes the costs from accident years at different stages of development that have very different statistical properties.

In view of the above shortcomings, it is proposed that using a fixed paid-up period length for all accident periods is preferable to the burning cost calculation. The disadvantage of this approach is that it introduces a time gap between the most recent accident period that can be used in the analysis and the target accident year for which the relativity is being set. However, we may expect that changes in the underlying industry relativities occur on much longer time scales than changes in industry sizes. Therefore, a method that does not react to changes in industry size but has a time gap is preferable to a method that has no time gap but over-reacts to industry size changes. Thus, the analysis in this paper will be based on fixed paid-up period length.

3.4 Choice of paid-up period length

The optimal paid-up period length should satisfy two opposing requirements: it should be long enough so that the corresponding resulting relativities would approximate the relativities based on the ultimate cost and short enough to allow using relatively recent experience. One way to quantify the required paid-up period length is to examine the convergence of cost ratio relativities as a function of paid-up period length.

To perform this analysis we considered the cost of claims from the 1997 accident year and calculated claims relativities using paid-up periods varying from 1 to 40 quarters. At each choice of paid-up period, we calculated base premium rates based on the corresponding relativities and the APR=3%, using formula (3.1). The convergence of premium rates was measured by the difference between base premium rates at each choice of paid-up length and at the paid-up length of 40 quarters. The difference, in turn, was expressed as the total cross-subsidy between industries, relative to the total premium collection. In mathematical terms, this can be expressed as follows:

$$\text{Aggregate cross-subsidy} = \frac{\sum Xs_i(d)}{\sum W_i}, \tag{3.4}$$

where $Xs_i(d)$ is the cross-subsidy provided by industry i at paid-up period length of d quarters, given by

$$Xs_i(d) = \begin{cases} W_i(r_i(d) - r_i(40)), & \text{if } r_i(d) > r_i(40) \\ 0, & \text{otherwise} \end{cases} \tag{3.5}$$

Here, $r_i(d)$ is the premium rate that corresponds to the relativity measured with paid-up period of d quarters.

The results of this calculation are shown in Figure 3.1.

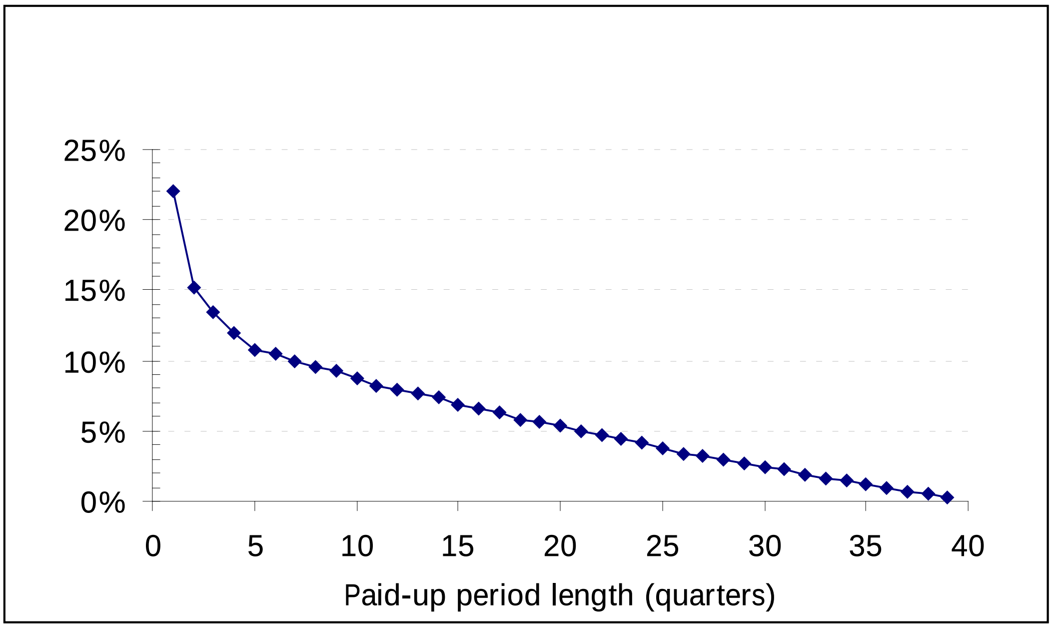

Figure 3.1. Aggregate cross-subsidy as a function of paid-up period length for accident year 1997.

Figure 3.1 shows that if we had initially set the premium rates based on 1 quarter paid-up costs, rechecked the relativities 10 years later using 40 quarters paid-up costs and revised the premium rates, then relative to the revised premium rates the initial set of rates would have a cross-subsidy of 22% of the total premium collection.

It is seen that the convergence of relativities is quite slow and even 10 years after the accident year cost-ratio relativities still evolve. Clearly, the relativities will continue changing even beyond development quarter 40 which is the end point for our analysis. Nonetheless, it is clear that a paid-up length of less than about 6 quarters is not acceptable as the resulting relativities are too unstable.

On the balance, the paid-up period of 9 quarters, or 2 years, appears to be an overall reasonable choice. Suppose that we are calculating industry premium rates for 2009/10 accident year using the data as at 31 December 2008. In this case, the latest developed accident quarter ends on 31 December 2006, which is 3 years prior to the middle of the target accident year on 31 December 2009. Intuitively, a lag longer than 3 years will not be acceptable from the practical considerations. At the same time, the 9-quarter paid-up period is long enough to lie beyond the initial zone of rapid accuracy improvement seen in Figure 3.1.

There is one more consideration to suggest that even without the practical issues related to the time lag, a very long paid-up period may not provide the most accurate estimate of relativities. This is related to the fact that the real aim of the industry premium rates calculation is not to estimate the incurred costs per se, but rather to estimate true underlying relativities from the observed paid-up cost ratios. Longer paid-up periods increase the variance of claim costs, which in turn increases the statistical error of estimating true mean total cost of claims from the sample mean cost. This matter will be explored in more details in Section 4.

On the basis of this reasoning, for the analysis presented in this paper we have chosen the fixed paid-up period length of 9 quarters.

3.5 Target accuracy

Suppose that we aim to set the base rates r_i with the relative accuracy of *x*%. What accuracy do we need to achieve for the relativities ρ_i?

To answer this question we shall use the first-order log-differential expansion

$$\frac{dr_i}{r_i} = \sum_j \left(\frac{\partial \ln r_i}{\partial \rho_j}\right) d\rho_j \tag{3.6}$$

Taking the partial derivative of (3.1) yields

$$\frac{d \ln r_i}{d\rho_j} = \frac{\delta_{ij}}{\rho_j} - \frac{W_j}{\sum \rho_k W_k}, \tag{3.7}$$

where δ_{ij} =1 if $i=j$ and zero otherwise.

Substituting (3.7) into (3.6) and using (3.1) we get

$$\frac{dr_i}{r_i} = \sum_j \left(\delta_{ij} - \frac{W_j}{\sum W_k}\frac{r_j}{L}\right)\frac{d\rho_j}{\rho_j} \tag{3.8}$$

Equation (3.8) shows that the change in base rates for a given industry *i* is due to two effects: the change of industry's own relativity (represented by the first term) and the collective effect of changes of relativities of all other industries in the economy, represented by the second term).

It can be seen that the second effect is relatively insignificant. Indeed, the term $W_j / \sum W_k$ is the relative size of industry *j* relative to the whole economy and is quite small (in 2007/08, the maximum of 2.51% was for SAWIC 633601 (Technical services nec), followed by 2.49% for SAWIC 511101 (Road freight transport)). The ratio r_j / L of the industry base rate to the APR can exceed 1, but typically will be contained within a factor of 4, which means that the whole expression $(W_j / \sum W_j)(r_j / L)$ will not exceed 10%. Although the number of terms in the sum in (3.8) is large (about 500 in South Australia), their total can be expected to be small because they will have different signs and will cancel each other. This leaves the change of industry's own relativity as the dominant effect, so we can approximately write

$$\frac{dr_i}{r_i} \approx \frac{d\rho_i}{\rho_i}, \tag{3.9}$$

which shows that the relative error in estimation of relativities should be the same as the accuracy targeted for specifying the premium rates.

The ballpark figure for the desired accuracy of premium rates can be estimated from the published table of rates. The table has the step of 0.1%, which represents the relative accuracy of 0.033 relative to the APR of 3%. As we shall see later from the statistical analysis, such a high accuracy for estimating relativities can hardly be achieved even for the largest industries.

3.6 Mathematical formulation of industry premium rate-setting problem

It will be assumed that:

- the APR and projected remuneration are known;
- the paid-up period is fixed at 9 development quarters since injury quarter

Under these assumptions, the industry premium rates-setting reduces to the following problem:

Estimate underlying cost ratios $\kappa_i = \frac{E[S_i]}{U_i}$ *for the target accident year with the target accuracy, where* S_i *is the 9-quarter paid-up cost of claims and* U_i *is proxy number of FTE units.*

4 Statistics of workers' compensation claims process

The aggregate claims cost for a given industry over a particular period is given by

$$S = \sum_{n=1}^{N} X_n \tag{4.1}$$

where N is the random number of claims for a given industry and X is the random paid-up cost of an individual claim. Thus, the aggregate cost S follows a compound distribution which depends on both claim number and claims size distributions. The observed value of aggregate cost represents a sample from this distribution, on the basis of which we seek to estimate $E[S]$. In order to understand the properties of the distribution for S, we shall first develop the empirical claim size distribution and then use Panjer's recursion (Hart et all, 2007, p 143) to calculate the compound distribution.

The ultimate aim of this section is to investigate the convergence of the compound distribution to the limit given by the Central Limit Theorem and to establish confidence intervals for cases with small sample size when the Central Limit theorem does not apply.

4.1 Probability distribution of claim size

For the start, let us consider claim costs for all industries, without specifically restricting claims to a particular SAWIC. For this analysis, we selected claims over a 7-year period from accident years 2000 - 2006. Claim payments were first expressed in real terms using the ABS Average Weekly Earnings index and further adjusted for superimposed inflation using the average claim cost index based on Scheme-wide experience, as discussed in more detail in Section 5.2. For calculating the distribution, the minimum claim size was set at $100 and the maximum was capped at $72,000[2]. Claims below the threshold were excluded. The histogram of the distribution is shown in Figure 4.1. The key statistics are summarised in Table 4.1.

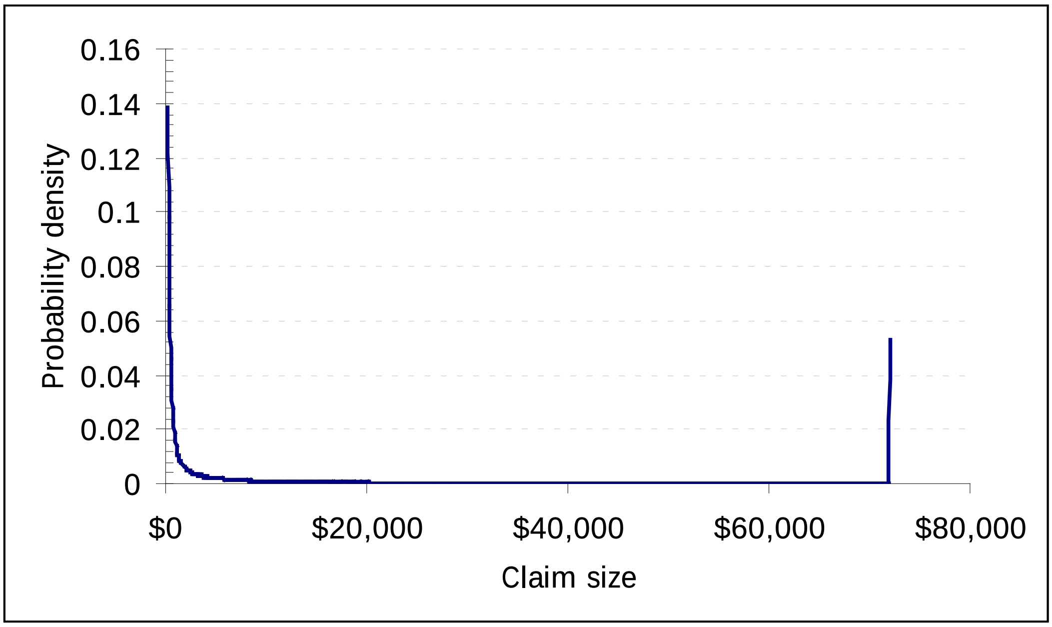

Figure 4.1. Observed claim size distribution – all industries combined. The claim cost was capped at $72,000.

Table 4.1 Key statistical characteristics of claim size distribution.

Statistic	Value
Mean	$8,160
Std deviation	$18,340
Coeff of variation	2.23
Median	$600
10% percentile	$100
90% percentile	$26,500

One can see that the distribution of claim size is very skew. Claims below $1,000 account for approximately 60% of all claims, yet the mean value of the distribution is $8,160 and 5% claims lie at or above the $72,000 cap.

Suppose we want to calculate the average claim size by taking the sample mean. Given the form of the distribution, it is not at all clear how many claims need to be averaged in order for the Central Limit theorem to work and bring the sample mean distribution to the bell-shape. This question can be answered by numerically computing a series of convolutions of claim size distribution with itself, using a classical formula for the distribution of a sum of *n* identically distributed (i.i.d.) variables. The results are given in Figure 4.2

[2] The role and choice of the cap are discussed in Section 4.4.

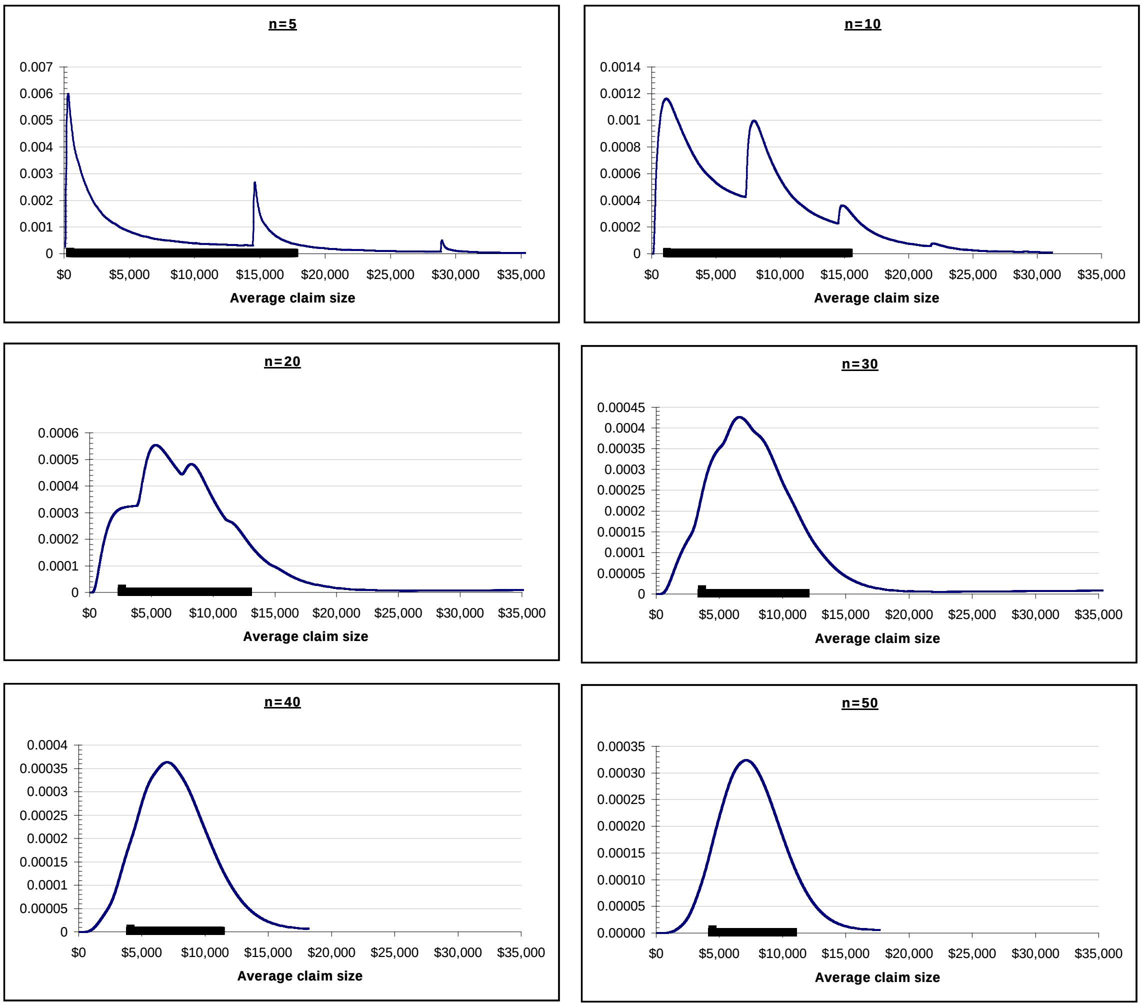

Figure 4.2 Empirical distribution density for the sample average of *n* i.i.d. variables subject to claim size distribution shown in Figure 4.1. Thick line indicates the interval between 10th and 90th percentiles.

It is seen that the distribution takes the familiar Gaussian shape at approximately n=40. For n<10 the distribution is so highly skew that any meaningful statistical estimation is impossible.

The distributions considered above were global in the sense that they included claims from all industries. Consider now claim size distributions for individual industries. Table 4.2 gives the estimated sample mean, standard deviation and coefficient of variation for all industries that had 500 claims or more over the 7-year period covering accident years 2000 – 2006. Figures 4.3 – 4.4 show the scatter-plots of estimated standard deviation and coefficient of variation (CV) versus sample mean for different industries.

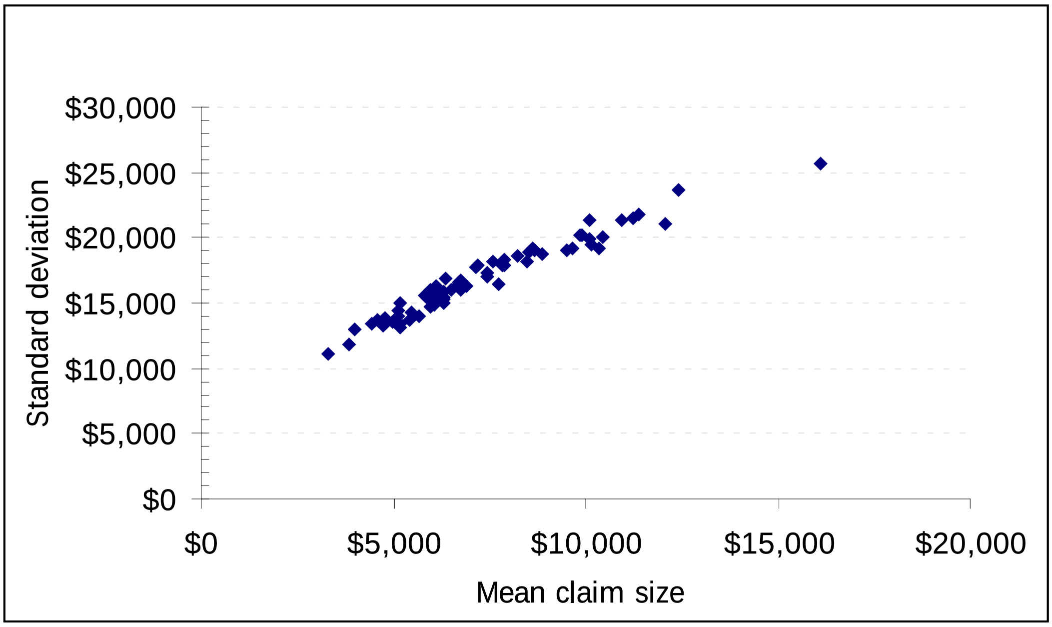

Figure 4.3. Scatter-plot of standard deviation against mean claim size for industries given in Table 4.2

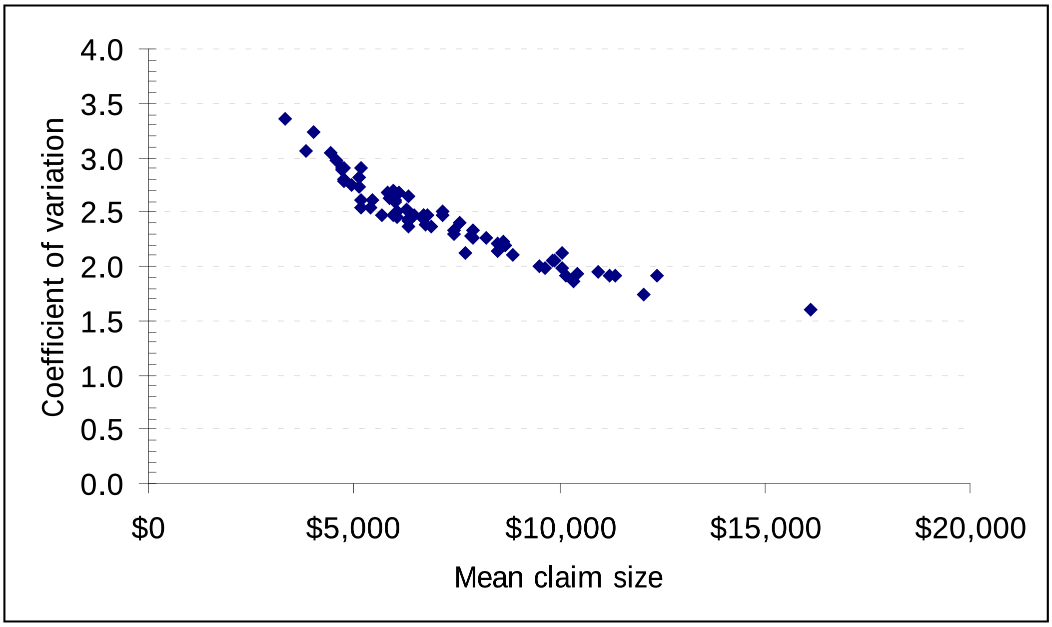

Figure 4.4 Scatter-plot of coefficient of variation against mean claim size for industries given in Table 4.2

Table 4.2 Key claim size statistics for industries with 500 claims or more.

		No claim size cap			Claim size caped at $72,000		
SAWIC	Claim count	Avg claim size	Std dev of claim size	Coeff of variation	Avg claim size	Std dev of claim size	Coeff of variation
323201	1,191	$3,983	$15,909	4.0	$3,316	$11,130	3.4
218801	1,276	$4,466	$16,266	3.6	$3,842	$11,776	3.1
336101	584	$4,534	$16,350	3.6	$4,009	$12,979	3.2
486101	1,629	$5,997	$24,445	4.1	$4,417	$13,408	3.0
314101	1,895	$5,842	$21,230	3.6	$4,580	$13,652	3.0
486106	1,352	$6,251	$24,845	4.0	$4,726	$13,708	2.9
485301	1,029	$5,534	$18,615	3.4	$4,727	$13,609	2.9
488601	1,585	$5,286	$16,540	3.1	$4,753	$13,255	2.8
315301	1,252	$5,457	$18,184	3.3	$4,766	$13,372	2.8
473601	1,202	$6,609	$31,521	4.8	$4,773	$13,899	2.9
294201	505	$5,621	$17,174	3.1	$4,954	$13,581	2.7
424401	925	$6,460	$22,889	3.5	$5,128	$14,426	2.8
336901	1,151	$6,432	$21,776	3.4	$5,138	$14,043	2.7
264401	810	$6,044	$19,140	3.2	$5,153	$13,442	2.6
335201	560	$5,844	$17,622	3.0	$5,161	$13,100	2.5
486501	902	$6,569	$22,147	3.4	$5,196	$15,055	2.9
914405	506	$5,987	$17,018	2.8	$5,403	$13,708	2.5
211505	3,910	$6,446	$19,977	3.1	$5,471	$14,265	2.6
540601	561	$6,634	$23,340	3.5	$5,685	$14,031	2.5
424301	1,633	$8,062	$27,569	3.4	$5,822	$15,642	2.7
263401	846	$7,491	$23,794	3.2	$5,865	$15,447	2.6
488101	1,082	$7,011	$22,482	3.2	$5,943	$14,652	2.5
474201	515	$7,995	$25,701	3.2	$5,959	$16,029	2.7
314201	655	$7,431	$22,712	3.1	$5,991	$15,558	2.6
347401	2,031	$7,272	$22,193	3.1	$6,001	$15,700	2.6
511401	605	$7,909	$28,522	3.6	$6,064	$15,181	2.5
323401	1,420	$6,811	$18,882	2.8	$6,066	$14,917	2.5
316601	511	$7,744	$24,544	3.2	$6,095	$16,283	2.7
830505	601	$7,262	$20,648	2.8	$6,271	$15,833	2.5
335701	1,322	$7,144	$20,144	2.8	$6,310	$15,398	2.4
823301	634	$7,465	$22,069	3.0	$6,311	$15,264	2.4
914401	599	$7,125	$19,425	2.7	$6,325	$14,993	2.4
254101	1,341	$7,217	$20,556	2.8	$6,330	$15,887	2.5
316801	871	$8,032	$25,380	3.2	$6,340	$16,811	2.7
216101	1,049	$7,685	$22,146	2.9	$6,487	$16,014	2.5
253501	1,318	$8,188	$23,597	2.9	$6,688	$16,483	2.5
472801	1,063	$8,489	$25,457	3.0	$6,702	$16,542	2.5
476901	854	$8,158	$23,356	2.9	$6,705	$16,376	2.4
923301	1,265	$7,736	$21,184	2.7	$6,738	$16,052	2.4
213101	879	$7,794	$21,338	2.7	$6,772	$16,778	2.5
472805	517	$7,918	$21,655	2.7	$6,873	$16,292	2.4
849101	8,137	$9,448	$28,347	3.0	$7,160	$17,722	2.5
849102	1,833	$9,392	$27,681	2.9	$7,169	$17,927	2.5
211501	688	$8,999	$24,354	2.7	$7,417	$16,969	2.3
923101	2,290	$9,121	$26,128	2.9	$7,435	$17,311	2.3
287401	817	$9,871	$29,970	3.0	$7,585	$18,171	2.4
814101	1,811	$8,691	$21,209	2.4	$7,726	$16,382	2.1
512201	976	$9,809	$27,141	2.8	$7,849	$17,871	2.3
411301	1,176	$11,748	$35,484	3.0	$7,885	$18,334	2.3
923201	2,467	$9,560	$25,430	2.7	$7,896	$17,903	2.3
424201	883	$11,099	$30,747	2.8	$8,227	$18,577	2.3
476401	1,111	$9,577	$22,955	2.4	$8,496	$18,203	2.1
296301	1,187	$11,031	$29,417	2.7	$8,521	$18,840	2.2
934001	593	$10,759	$28,400	2.6	$8,620	$19,132	2.2
412201	1,886	$12,376	$35,240	2.8	$8,684	$19,092	2.2
13401	874	$10,761	$27,590	2.6	$8,846	$18,689	2.1
830501	834	$11,387	$26,352	2.3	$9,532	$18,991	2.0
13601	594	$11,951	$30,559	2.6	$9,635	$19,126	2.0
20601	573	$13,605	$39,602	2.9	$9,858	$20,160	2.0
424801	1,193	$13,794	$38,750	2.8	$9,880	$20,213	2.0
411101	886	$12,603	$29,774	2.4	$10,082	$19,908	2.0
849501	579	$15,050	$39,376	2.6	$10,086	$21,380	2.1
814301	5,373	$11,672	$25,620	2.2	$10,146	$19,473	1.9
18201	510	$14,209	$36,187	2.5	$10,362	$19,242	1.9
486401	546	$12,111	$26,326	2.2	$10,442	$20,104	1.9
412101	576	$16,578	$42,857	2.6	$10,942	$21,352	2.0
638907	979	$14,045	$31,227	2.2	$11,221	$21,552	1.9
511101	5,287	$16,750	$41,928	2.5	$11,381	$21,731	1.9
638701	1,725	$14,319	$31,271	2.2	$12,052	$20,993	1.7
424901	678	$16,987	$37,032	2.2	$12,394	$23,643	1.9
423101	516	$21,955	$41,303	1.9	$16,105	$25,680	1.6

It is interesting to note that the shape of the relationship CV(mean) shown in Figure 4.4 is consistent with a bimodal distribution of claim size. Consider a random claim size that takes only two values, a small value (say, $500) with probability P and a large value (say, $20,000) with probability (1-P). In the workers' compensation context, we can associate the small value with medical-only claims and the large value with claims that require income support.

Figure 4.5 shows the CV(mean) relationship for this distribution where P varies in the range from 0.8 to 0.92. It is seen that the CV(mean) relationship for this simple distribution is very similar to the observed one given in Figure 4.4.

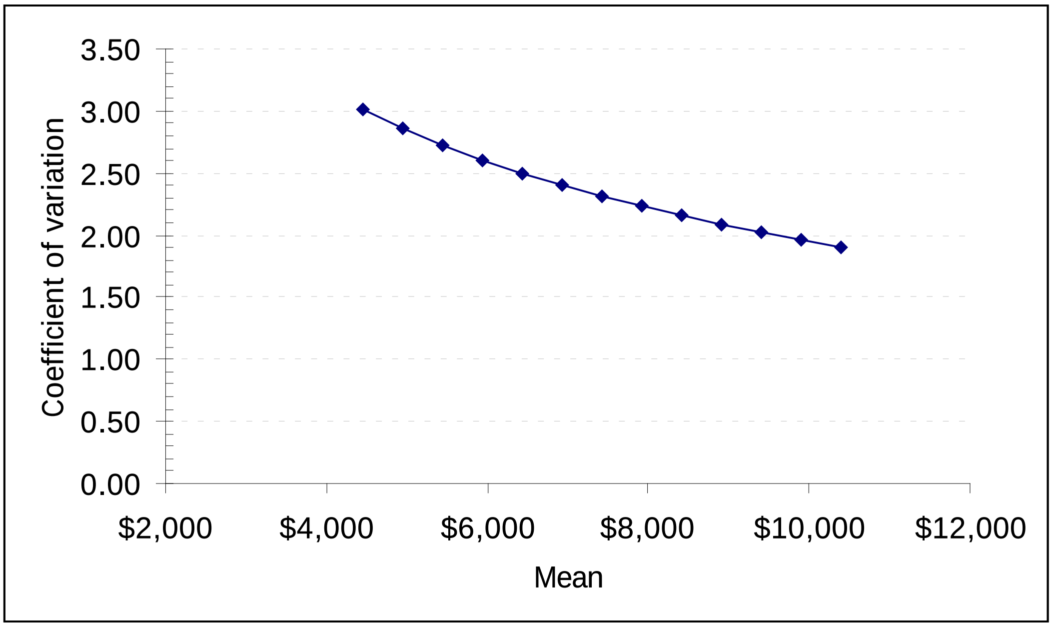

Figure 4.5 CV(mean) function for a simple bimodal claim size distribution.

This similarity between figures 4.4 and 4.5 suggests that the key difference between the industry-specific claim size distributions may be the proportion of claims that do and do not require income support, rather than the difference in the costs of these types of claims.

While the analysis of inter-industry claim size distributions is an interesting topic in itself, the important conclusion for the purposes of this paper is that notwithstanding the fact that industries with larger average claim size have smaller CVs, the CVs for individual industries are of the same order of magnitude as the CV for the global distribution. This indicates that one can use the global distribution for ball-park estimates of confidence intervals and related matters.

4.2 Probability distribution of aggregate claim cost

The aggregate claim cost, given by expression (4.1), depends on both the number of claims and the size of individual claims. It is reasonable to assume that each worker is exposed to a certain claim risk, so that the total risk exposure for the industry is given by the total number of employees and the number of claims is a Poisson variable,

$$N \sim Poisson(fU)\,, \qquad (4.2)$$

where $U = W/(52\ AWE)$ (AWE – average weekly earnings for SA) is the proxy number of full-time workers and *f* is the claim occurrence rate. Under this model, each industry *i* has its own claim rate f_i and will generate a random number of claims N_i .given the exposure U_i. It should be noted that in reality claim occurrence rates may vary between individual workers and employers within the industry, which would lead to over-dispersion relative to the Poisson distribution. The implications of this are discussed in Section 7.

Using the Poisson distribution assumption together with the empirical claim size distribution and assuming that claim frequency and claim size are independent, one can calculate the distribution density for the aggregate cost in the two regimes: pre-Normal (small number of claims) and Normal (large number of claims). Consider each of them separately, starting from the Normal regime.

4.2.1 Aggregate cost under the normal approximation

If the number of claims is large, then the aggregate cost S will be normally distributed. The expected value and variance of the compound distribution can be found from the well-known formulae:

$$E[S] = E[N]E[X], \quad Var[S] = E[N]Var[X] + Var[N](E[X])^2 \qquad (4.3)$$

For the Poisson distribution,

$$E[N] = fU, \quad Var[N] = fU \qquad (4.4)$$

Using (3.3) and (3.4), we can find CV[S] as follows:

$$CV[S] = \frac{\sqrt{Var[S]}}{E[S]} = \frac{\sqrt{fU(Var[X] + (E[X])^2)}}{fU\,E[X]} = \frac{1}{\sqrt{fU}}\sqrt{1 + (CV[X])^2} \qquad (4.5)$$

Note that the term fU is the expected number of claims. In practice, our best estimate of the expected number of claims is the actual observed number of claims, *n*. Hence, given the observed aggregate claims count and cost for a particular industry over a certain exposure period, the estimate for *CV[S]* can be found as

$$CV[S]_{est} = \frac{1}{\sqrt{n}}\sqrt{1 + (CV[X])^2} \qquad (4.6)$$

If we were looking at the CV of the total cost of n claims, where n is a fixed number, then under the normal approximation the result would be a well-known expression $CV[X]/\sqrt{n}$. It is seen that the additional variability brought about by the randomness of claims count increases *CV[S]* to the value given by (4.6).

The results of the previous section show that CV[X] for individual industries ranges from 1.6 to 3.4. Taking the global value CV[X]=2.3 as a representative number gives an approximate relationship that can be applied to all industries

$$CV[S]_{est} \approx \frac{\sqrt{1 + 2.3^2}}{\sqrt{n}} = \frac{2.5}{\sqrt{n}} \qquad (4.7)$$

Relationship (3.7), together with the Central Limit Theorem that states that for sufficiently large *n* the distribution of *S* is normal, allows one to calculate confidence intervals for the estimates of *E[S]* based on the observed values of *S*. In terms of the relative accuracy,

$$\varepsilon = \frac{s - E[S]}{E[S]},$$

an estimate based on *n* claims is accurate to within $\pm z_\alpha CV[S]$, where z_α is a factor that depends on the confidence level (eg, z_α =1.96 for 95% confidence level, z_α =1.28 for 80% confidence level).

Suppose that a particular industry had 1,000 claims in a given accident period and has the total cost of claims *s*. Then the theory suggests that its relativity estimate is accurate to within $\pm\frac{2.5\times1.28}{\sqrt{1000}}=10\%$ at the 80% confidence level. Note the difference with the 3% target accuracy implied by minimum step in the current table of industry rates.

4.2.2 Aggregate cost for small number of claims

The analysis of convergence to the normal distribution in the previous section shows that the Central Limit Theorem applies at *n* of about 40. For smaller number of claims we need to explicitly calculate the distribution for S and estimate confidence intervals from that distribution.

As in the previous section, assume that $N \sim Poisson(fU)$ and the X is distributed as shown in Figure 4.1. Under these assumptions it is possible to calculate the resulting compound distribution for the aggregate cost S using Panjer's recursive formulae (Hart et all, 2007, p143). The results of this calculation for various expected number of claims are shown in Figure 4.6. Table 4.3 gives the values of 10^{th} and 90^{th} percentile points for these distributions relative to the mean.

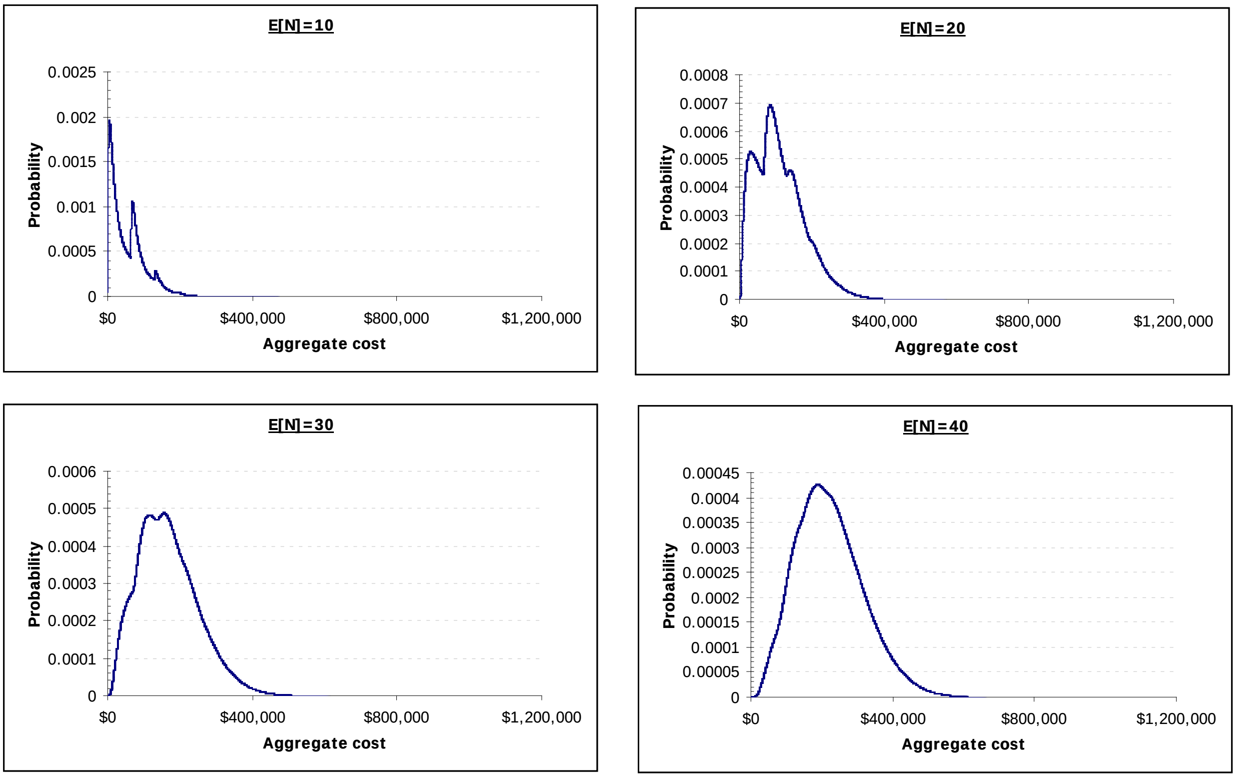

Figure 4.6. Distribution density for aggregate cost S for different values of mean claims count. At E[N]=40 the distribution approaches the Normal distribution.

Table 4.3 The 10th and 90th percentile points relative to the mean

E[N]	10th percentile/Mean	90th percentile/Mean
10	11%	223%
20	27%	183%
30	40%	167%
40	48%	158%

The results in Table 4.3 show that if cost ratio for a given industry is based on a sample of 40 claims, then the true underlying value of the cost ratio may vary anywhere between 50% and 160% of the observed value. The accuracy is even worse for smaller values of claims count.

4.2.3 Combined estimates of confidence intervals

Combining the 10th and 90th percentile points for aggregate cost distribution in both the small claim count case and the large claim count case gives a complete picture of the estimation accuracy, shown in Figure 4.7.

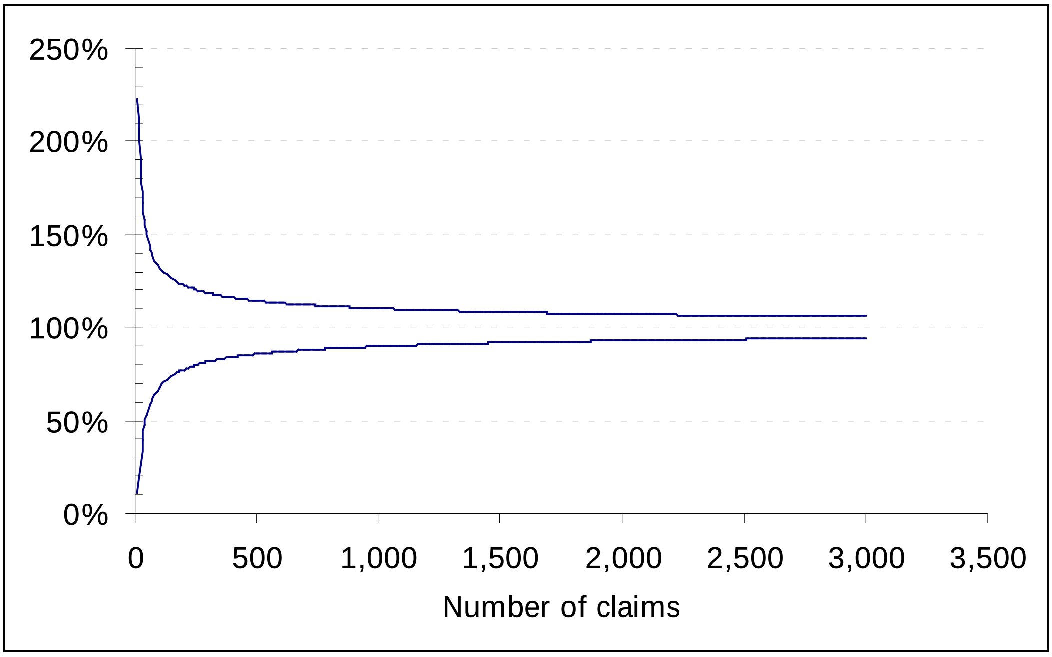

Figure 4.7. Combined confidence intervals. Lower curve – 10th percentile/mean, upper curve – 90th percentile/mean

Note that even an observation based on 3,000 claims is only accurate to within $\pm 6\%$!

4.3 Estimation accuracy and minimum experience period

The analysis of claims size and aggregate cost distributions has shown that achieving a meaningful accuracy in estimation of the underlying true relativity (as measured by the cost ratio) requires aggregating claims costs over a large number of claims. However, in reality, the number of claims generated by a single industry is rather limited. Indeed, the South Australian Scheme as a whole incurs about 20,000 claims per year and has about 400 industries, which gives on average about 50 claims per industry per accident year. The actual distributions of the number of claims, for a single accident year (2006) and for a five-year period (2002 to 2006) are given in Tables 4.4.

Table 4.4 (a). The distribution of the number of claims per industry for accident year 2006

Number of claims	Number of industries
<10	160
10 to 19	75
20 to 29	33
30 to 39	33
40 to 49	22
50 to 99	48
100 to 199	29
200 and more	15
Total	415

Table 4.4 (b). The distribution of the number of claims per industry for accident years 2002 to 2006

Number of claims	Number of industries
<10	81
10 to 19	37
20 to 29	33
30 to 39	23
40 to 49	29
50 to 99	69
100 to 199	66
200 to 299	38
300 to 399	26
400 to 499	9
500 to 599	7
600 to 699	7
700 to 799	9
800 to 899	4
900 to 999	6
1000 and more	15
Total	459

These results, together with the error analysis of the preceding subsection, put industry relativity estimation at an interesting angle. We can see that even for the largest industries, a single accident year not enough to accurately estimate the underlying true cost ratios. Moreover, even the aggregate experience of five accident years is not enough to achieve a good accuracy, since for all but 15 largest industries the estimation accuracy will be worse than $\pm 10\%$ and for 81 industries there is no data to make any possible estimation at all.

Suppose we need to measure the industry's underlying cost ratio with a given accuracy (say, 10%). Our error analysis shows that this accuracy requires approximately 1,000 claims. In general, given the target accuracy level each industry has a certain minimum experience period that is required in order to achieve the accuracy. For some industries, the minimum experience period may span many years, while for others it may be much shorter.

In a typical forecasting situation, one aims to base the prediction on the most recent observations. However, the minimum experience period puts a limit on the time scale at which we can observe the underlying cost ratios and on our ability to resolve changes and respond to them.

4.4 Claim size cap and 'uncertainty principle'

Claim size cap is a standard mechanism used to minimise the volatility of observed cost ratios. The justification for the cap is that it removes very large claims whose incidence is random and does not represent the true riskiness of those industries that happen to have produced such claims in a particular study period. The key trade-off here is that a large cap increases volatility of results but more correctly attributes the costs to those industries that systematically produce large claims, while a small cap has the opposing effect.

The framework based on the target estimation error allows one to view cap selection from a different angle. A large cap increases the variance of claim size distribution, which in turn increases the variance of aggregate claim cost. For a given target estimation error, this increases the minimum experience period. Hence, we are dealing with a kind of an 'uncertainty principle': for a given level of estimation accuracy, there is a limit to our ability to localise the measurement in time and to account for all claims costs. If so desired, one can set the cap at a level that covers, say, 95% of all claim costs; however, under this choice one would have to include in the analysis very distant experience which may not be representative of the current risks and costs.

In the current analysis, we used the cost cap of $72,000 which accounts for approximately 80% of the uncapped cost. (ie, the sum of all capped claim costs = 80% of the sum of all uncapped claim costs). This leaves 20% of costs unallocated to individual industries, yet even under this relatively low cap the minimum experience period requirements are very demanding.

While it is natural to ask how one finds the optimal claim size cap, it is unlikely that a formal mathematical criterion can be developed. A practical approach may be to plot the graphs of the average (across all industries) minimum experience period and the number of industries with insufficient experience against the cap size and make an executive decision that weighs the time lag against the accuracy of cost allocation.

5 Adaptive period-length approach to calculation of industry cost ratios

5.1 Basic idea

In a typical industry premium rating study all industries are analysed over the same experience period. However, the uncertainty-based framework shows that each industry has a certain minimum experience period required to estimate its underlying cost ratio with the target accuracy. This logically suggests an adaptive approach, in which the experience period for each industry is of the minimum length for this particular industry. Under this method, for each industry we will rely on the recent data to the maximum possible extent given the target accuracy. This approach is easy to implement in practice – starting from the most recent sufficiently developed accident year, for each industry one adds prior accident periods until the total number of claims exceeds the value required by the accuracy target.

The following table shows the distribution of the number of industries by the minimum experience period length for the target accuracy of $\pm 10\%$ at the 80% significance level, which requires approximately 1,000 claims.

Table 5.1 The minimum experience period required to achieve target estimation accuracy for the underlying cost ratio. The last accident period end date is 31 December 2004.

Minimum experience period (y)	Number of industries	Proportion
1	0	0%
2	3	1%
3	1	0%
4	4	1%
5	4	1%
6	7	1%
7	11	2%
8	6	1%
9	11	2%
10	2	0%
11	4	1%
12	13	3%
13	6	1%
14	5	1%
15	16	3%
16	10	2%
All available experience	383	79%
Total	486	100%

These results show that even at the modest accuracy target of $\pm 10\%$ at the 80% significance level, only a handful of industries can be assessed on the basis of recent experience and more than three quarters of the industries require using all available historical data.

5.2 Scheme-wide superimposed inflation and claim frequency trends

Using adaptive period-length requires adjusting the data for Scheme-wide changes in claim frequency and claim size over time. Without such adjustments, trends in cost ratios can lead to biases in cost-ratio estimates for individual industries.

Figure 5.1 shows the graphs of Scheme-wide indices for claim frequency, average claim size and cost ratio.

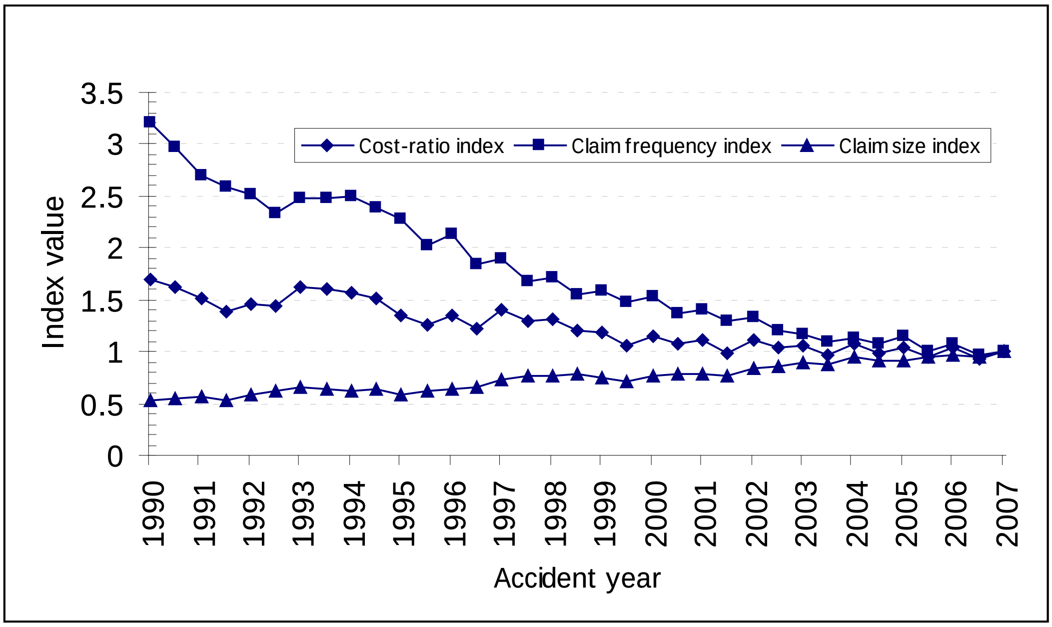

Figure 5.1 Scheme-wide indices. Claim frequency and cost ratio are expressed per number of proxy FTE units. Claim size refers to the two-year paid-up cost capped at $72,000. Claims below $100 were excluded.

It is seen that over the 17-year period South Australia have seen a dramatic reduction in claim frequency, accompanied by a growth in the average cost of claim. The net impact on the cost ratio was a reduction.

In order to eliminate these trends from the data used for relativity analysis, past observations were adjusted as follows:

- Industry exposure from a given accident period was divided by the frequency index
- Claim amounts from a given accident period (already expressed in current dollar values) were divided by the claim size index.

These adjustments ensure that small industries that must rely on their historical experience are not disadvantaged against large industries whose cost ratios are calculated from more recent data and, therefore, are lower due to the downward trend of claims costs.
Note that the graphs in Appendix A and claim size distributions discussed in the previous section were based on adjusted data.

5.3 Testing of proposed methodology

Using experience from almost 20 years back for estimating cost ratios for a period 3 years away from the most recent observation is a radical proposition that requires testing. The following test seeks to check whether the cost ratios calculated using adaptive experience period lengths provide estimators that are consistent with the probability distributions of claim numbers and aggregate costs developed in the previous section.

The following diagram shows the experience period and target period in the accident quarter – development quarter space.

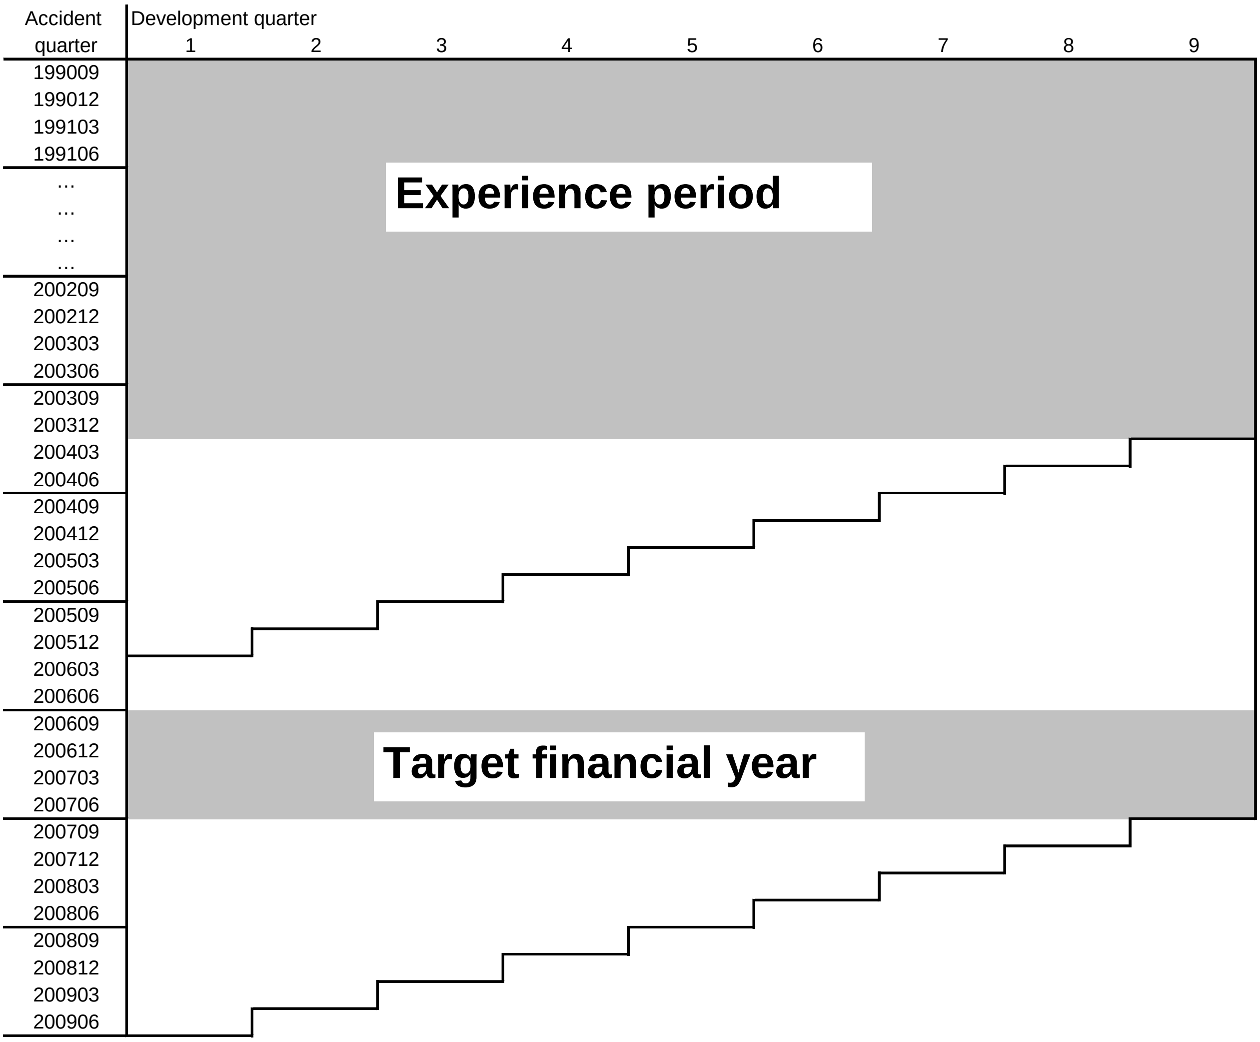

Figure 5.2. The experience period used for estimating industry cost ratios for the target financial year 2006/07. In order to set premiums for 2006/07 we would use the data cut as at 31 December 2005, hence the latest accident quarter developed to development quarter 9 is December 2003. Note that 2006/07 is the latest accident year that can be analysed using the data available at the time of writing this paper (June 2009 data extract).

In practice, cost ratio estimation involves calculating cost ratios for each industry from its own experience period length and then assuming that the same cost ratio applies over the target accident year. In probabilistic terms, we can think of the calculated cost ratio as being the mean of the distribution that, as we believe, will apply over the target period. While the actual cost ratio for the target period for each industry will be different from the mean, the realised distribution of observed cost ratios should be the same as the theoretical one with the estimated value of the mean. Whether this is true of false can be tested using the quantile-quantile plot.

In mathematical notation, the test operates as follows. For each industry i denote the observed claim frequency (number of claims per number of proxy FTE units) as f_i and the average claim cost as $\overline{X}_i$. For the same industry denote the actual exposure over the target period by U_i and the actual aggregate cost of claims by S_i. Finally, denote the probability density of scheme-wide claim size distribution (shown in Figure 4.1) by $g_0(x)$ and its expected value, equal to \$8,160 (refer to Table 4.1) by μ_0.

Assume that the observed frequency and average claim cost are the true parameters of the compound Poisson distribution for aggregate claims cost in the target year. For a given industry *i*, this distribution has the following attributes:

Poisson parameter: $\lambda_i = f_i U_i$ (5.1)

Claim size distribution: $g(x, \mu_i) = g_0(x \frac{\overline{X}_i}{\mu_0})$, (5.2)

where μ_i is the underlying mean claim size for industry *i*.

Note that (5.2) implies that the claim size distribution for a given industry *i* is a scaled version of the scheme-wide distribution. This form is consistent with the assumption of constant coefficient of variation that has been made earlier. On this basis, (5.2) appears to be a reasonable parameterisation.

Given (5.1) and (5.2), for each observed value of S_i , we can find the corresponding probability distribution by first scaling S_i by the factor $\overline{X}_i / \mu_0$ and then choosing one of standard compound Poisson distributions with parameter $\lambda = \lambda_i$ (the examples for λ=10, 20, 30, 40 are shown in Figure 4.6).

For values of Poisson parameter greater than 40 the compound distribution can be approximated by the normal distribution $N(\eta_i, \sigma_i)$ with the parameters given by

Mean=$\eta_i = \lambda_i \overline{X}_i$ (5.3)

Standard deviation=$\sigma_i = \eta_i \frac{2.5}{\sqrt{\lambda_i}} = 2.5 \overline{X}_i \sqrt{\lambda_i}$ (5.4)

Note that formula (5.4) is based on (4.7) and assumes that claim size distributions for all industries have the same coefficient of variation equal to 2.3.

Using empirical compound distributions for industries with the expected number of claims in the target financial year less than 40 and $N(\eta_i, \sigma_i)$ for industries with the expected number of claims greater than 40 one can calculate the quantile for each observed aggregate claims cost S_i. For λ_i<40 this can be done numerically as described above, while for λ_i>=40 the quantile has the analytical expression:

$$P(S_i) = \text{Prob}(S \le S_i) = \Phi(\frac{S_i - \eta_i}{\sigma_i}), \quad (5.5)$$

where $\Phi(z)$ is the distribution function for the standard normal distribution $N(0,1)$.

The figure below shows the quantile-quantile plot of observed claims cost S_i, *i*=1, 2, … for all industries in the Scheme. The value on the y-axis shows the proportion of all industries for which the theoretical value $P(S_i)$ was less than the hurdle value shown on the x-axis.

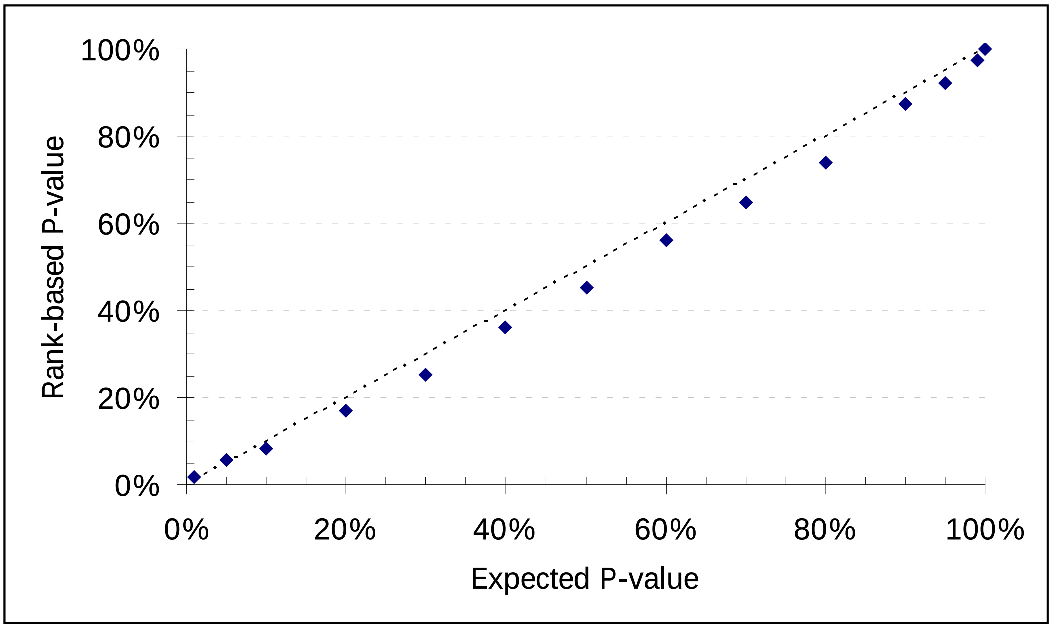

Figure 5.3. Quantile-quantile plot of observed claims costs in accident year 2007. There were 359 industries included in the rank calculation. Industries that had less than 40 claims over the entire experience period or $\lambda < 1$ were excluded.

For a perfectly consistent prediction model, the predicted and actual quantiles should fall on a straight line. While the graph in Figure 5.3 shows a minor bias, the agreement is nonetheless remarkably good. Note that the rank calculations above included industries with the expected number of claims λ as small as 1, for which the scaling assumption (5.2) is potentially problematic. If the quantile calculations are restricted to industries with λ>=20, then the agreement is nearly perfect, as it is shown in Figure 5.4

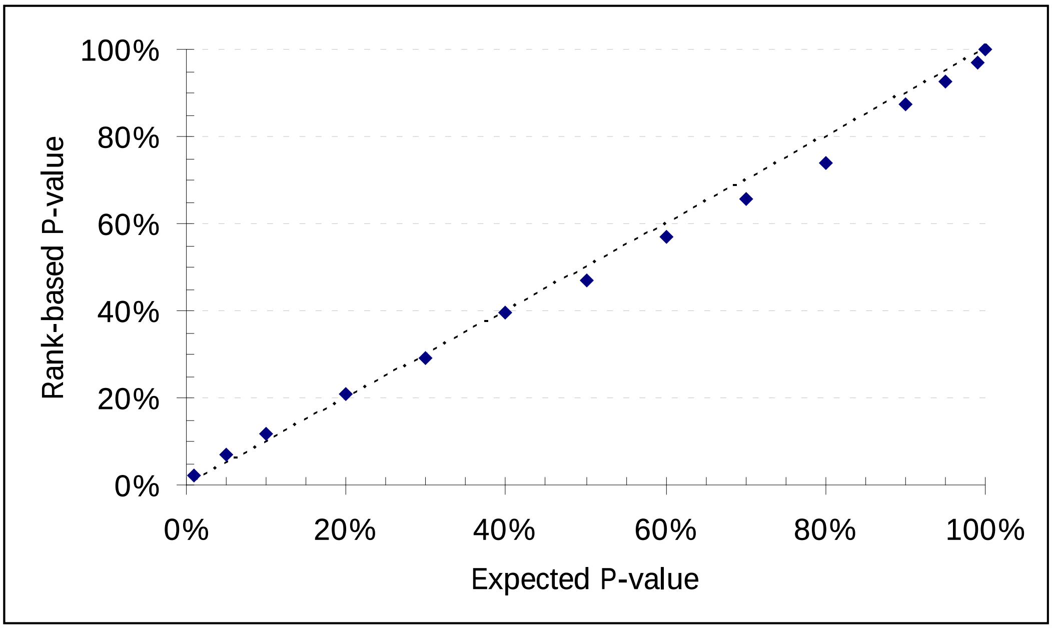

Figure 5.4. Quantile-quantile plot of observed claims costs in accident year 2007. There were 172 industries included in the rank calculation. Industries that had less than 40 claims over the entire experience period or $\lambda < 20$ were not included.

One of key steps in applying the adaptive period-length method is adjusting past exposure and costs using Scheme-based indices, as described in Section 5.2. The quantile-quantile plots in Figures 5.3 and 5.4 indicate that this adjustment works in the sense that estimates based on the adjusted historical data are consistent with the more recent data. However, it is instructive to see what happens if these adjustments are not done. The quantile-quantile plot for this case is shown in Figure 5.5. As should be the case, that the straight line has been transformed into a convex curve which clearly shows the presence of the estimation bias.

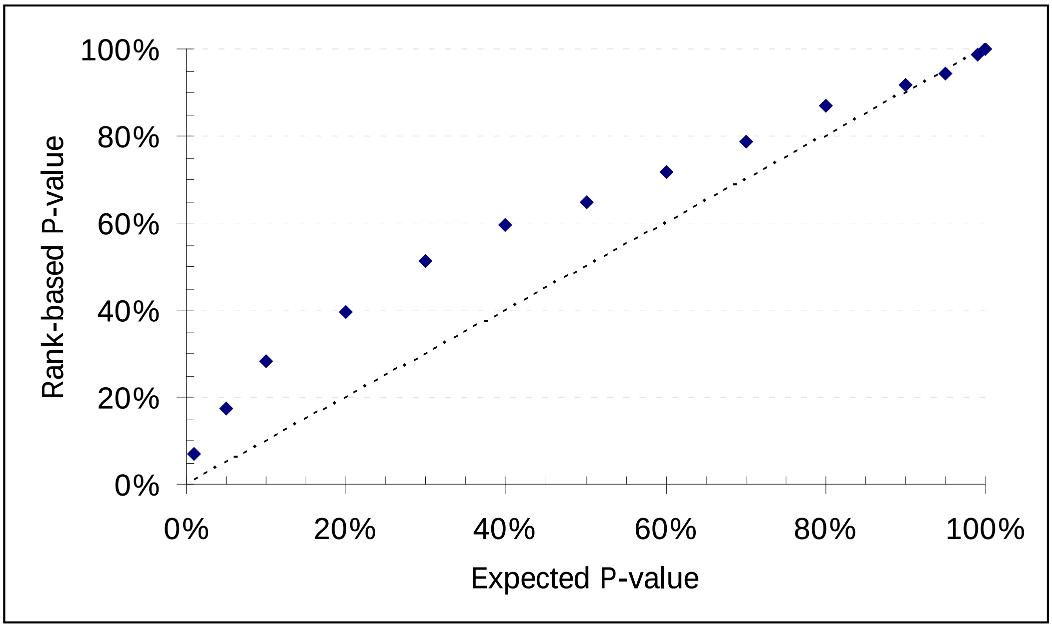

Figure 5.5 The same test as shown in Figure 5.3, but applied to data unadjusted for Scheme-wide trend in claim frequency and average size.

To summarise, it appears that the statistical test described above supports the proposed methodology. While the test results for industries with small expected number of claims are not perfect, this is relatively unimportant as in practice the relativities for such industries are likely to be heavily based on collateral information through the use of credibility weighting or the approach proposed in the next section.

6 Minimal disturbance approach to dealing with sparse data

The lack of data in some of rating cells is a well-known issue in premium rating work. It is commonly addressed by using the experience of higher-level 'parent' cells in the hierarchical credibility framework. A comprehensive reference to contemporary credibility theory can be found in the book of Buhlmann and Gisler (2005).

The hierarchical credibility method forms the estimate of the sought quantity X (risk premium, cost ratio, etc) for a given rating cell as a weighted sum of observations in the cell itself and category-average values higher in the hierarchy. While this set-up is intuitively appealing and straightforward, ascribing particular weights to the elements of the credibility sum is not at all simple. The theory that addresses this question derives the credibility weights from the estimates of variation of X within different levels of the hierarchy. Undertaking such estimation at minimum requires having the time series of X for each rating cell in the hierarchy and making a number of strong assumptions about various risk parameters being i.i.d and conditionally independent.

The reality of worker's compensation is very different. As it is shown in Table 4.1, three quarters of the industries require 20 years of data to produce a single estimate, so having a time series is out of question. There is not enough data to verify the key assumptions of the underlying theory.

In these circumstances the hierarchical credibility modelling of industry premium rates becomes an ad-hoc exercise, where the credibility weights are found in such a way as to provide the right amount of year-to year movement in premium rates. Too much weight given to self-experience makes premium rates too volatile, whereas too much weight given to industry groups makes premium rates too stagnant. While this may be a reasonable pragmatic approach, from the conceptual view point it is hardly satisfactory as the resulting premium rates do not have a fully objective justification.

The consideration of estimation errors developed earlier in this paper enables an alternative way of dealing with insufficient data. The proposed approach is to set the rate for a given industry equal to that of the broader industry division unless the industry's own experience is statistically significantly different. In the latter case, the industry rate is set at the edge of the confidence interval nearest to the industry division value. This process can be thought of as minimising the disturbance of premium rates within the portfolio by keeping each industry rate as close as possible to the corresponding division without violating the statistical significance test. Several examples of applying this procedure are discussed below.

The graph in Figure 5.5 shows the cost ratio range (at 80% confidence level) for SAWIC 489401 "Newsagents, stationers, etc" together with the average cost ratio for its industry division "Low-risk trade". The measurement for each target financial year ending 30 June YYYY is based on the accident period ending on 31 December YYYY-4. Since the number of claims to-date is less than 1,000 required for the target accuracy, the estimates for each target financial year use all available experience. It is seen that as the number of claims in the experience pool increases, the range of true underlying cost ratio values becomes narrower. The SAWIC 'breaks away' from its industry division in 2005 when its own underlying cost ratio becomes different from the division-average cost-ratio in a statistically significant way.

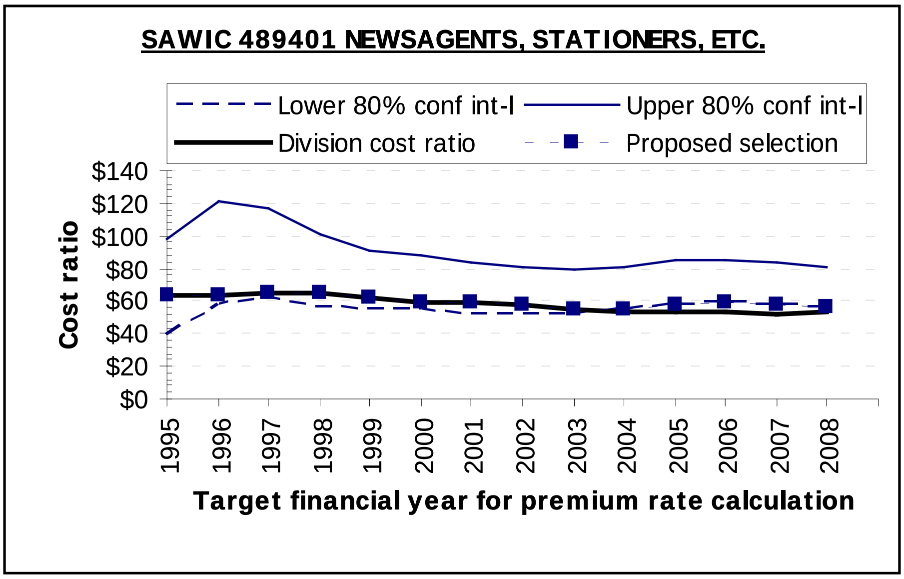

Target FY	Claims count
1995	55
1996	83
1997	112
1998	130
1999	164
2000	188
2001	200
2002	217
2003	236
2004	257
2005	283
2006	298
2007	322
2008	336

Figure 5.5. An example of applying the proposed minimum disturbance approach to calculating SAWIC cost ratio for setting premium relativity. The table to the right shows the number of claims involved in the estimate for each target financial year.

The graph in Figure 5.6 gives an example of a case where a small industry has the cost ratio that is significantly different from that of its industry division. The industry “Beer, ale and stout manufacturing” has the annual remuneration of approximately $50m. In relative terms this is a small industry (the total remuneration for the “Manufacturing” division is in excess of $2b) and under the credibility framework is unlikely to have any substantial weight given to self-experience. However, the uncertainty-based approach shows that even taking into account the limited amount of claims experience available for this industry, its true underlying cost-ratio is below the industry-average somewhere by a factor of 4 to 5. In these circumstances it appears fair and equitable that this industry's premium should be scaled accordingly.

Finally, the graph in Figure 5.7 gives an example of the situation where the minimum disturbance approach leads to a very large movement in the cost-ratio. It is seen that the industry “Oil and fat manufacturing” generates between 0 and 2 claims per accident year. The dramatic movement between 2006 and 2007 was due to a single claim that reached the maximum amount of $72,000. In the credibility framework such an event would be quite unnoticeable due to the small weight given to self-experience of this small industry. However, under the proposed approach we see the full size of the cost-ratio movement once the division -average value moves outside of the confidence interval range. Whilst this is a shortcoming of our method, it can be partially addressed by requiring that industries must accumulate a certain minimum number of claims before being allowed to 'break away' from their division.

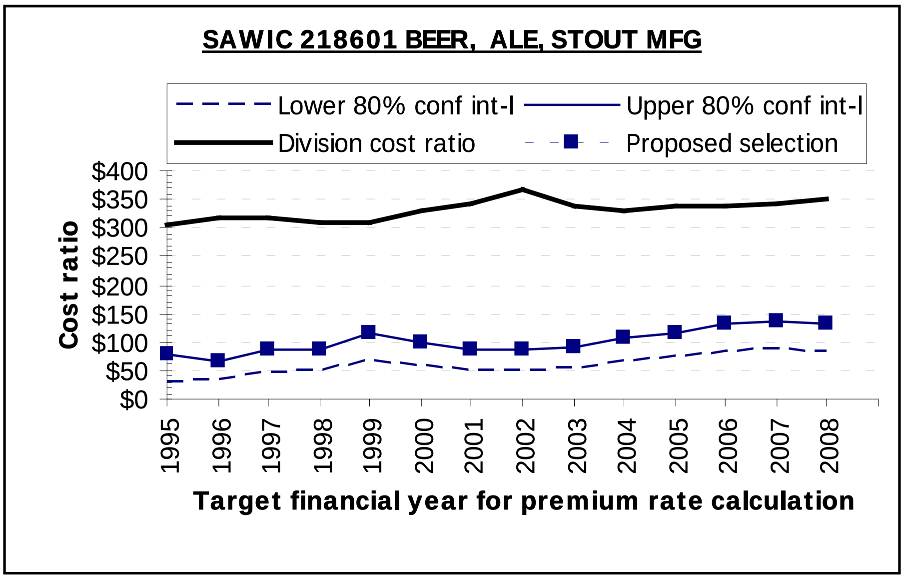

Target FY	Claims count
1995	54
1996	78
1997	97
1998	126
1999	139
2000	141
2001	143
2002	150
2003	159
2004	173
2005	189
2006	205
2007	209
2008	216

Figure 5.6. An example of applying the proposed minimum disturbance approach to calculating SAWIC cost ratio for setting premium relativity. The table to the right shows the number of claims involved in the estimate for each target financial year.

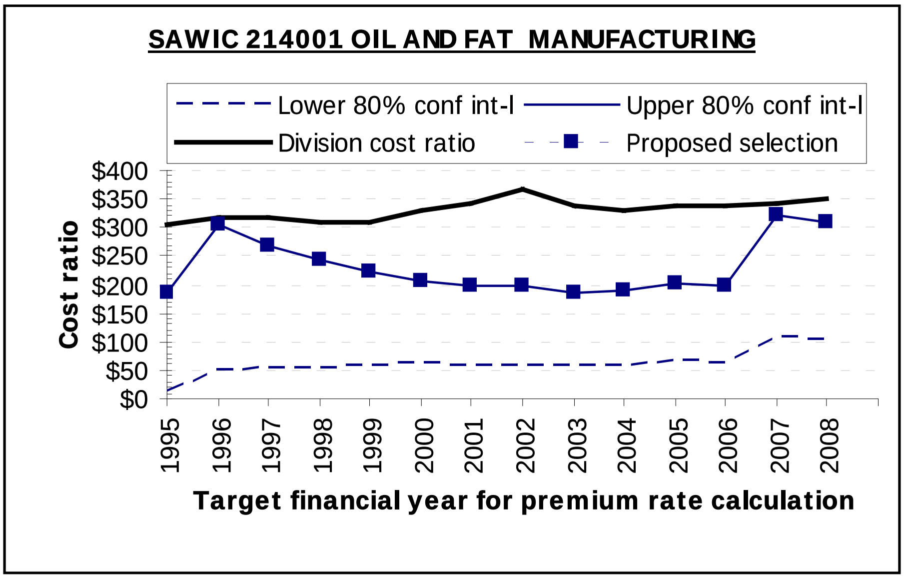

Target FY	Claims count
1995	13
1996	22
1997	26
1998	29
1999	35
2000	39
2001	39
2002	40
2003	41
2004	43
2005	44
2006	44
2007	45
2008	45

Figure 5.7. An example of applying the proposed minimum disturbance approach to calculating SAWIC cost ratio for setting premium relativity.

A key shortcoming of the proposed procedure is that for each industry the final estimate taken from either the corresponding industry group average or the edge of the confidence interval range is sub-optimal in the maximum likelihood sense. This may bias the estimators or make them less accurate than the simple observed means that they replace. To check whether this is the case we used the same quantile-quantile graph test described in the previous section. In order to apply this test, one has to independently specify the estimated Poisson parameter and average claim cost for the compound claims cost distribution. In the case when the industry cost ratio was taken from the edge of the confidence interval, we found these values by independently interpolating between industry parameters and industry division parameters.

The test result given in Figure 5.8 shows that the quantile-quantile curve is very close to the straight line. This indicates that the estimates obtained from the minimum disturbance procedure are, on average, unbiased and provide the accuracy consistent with the assumed statistical distributions.

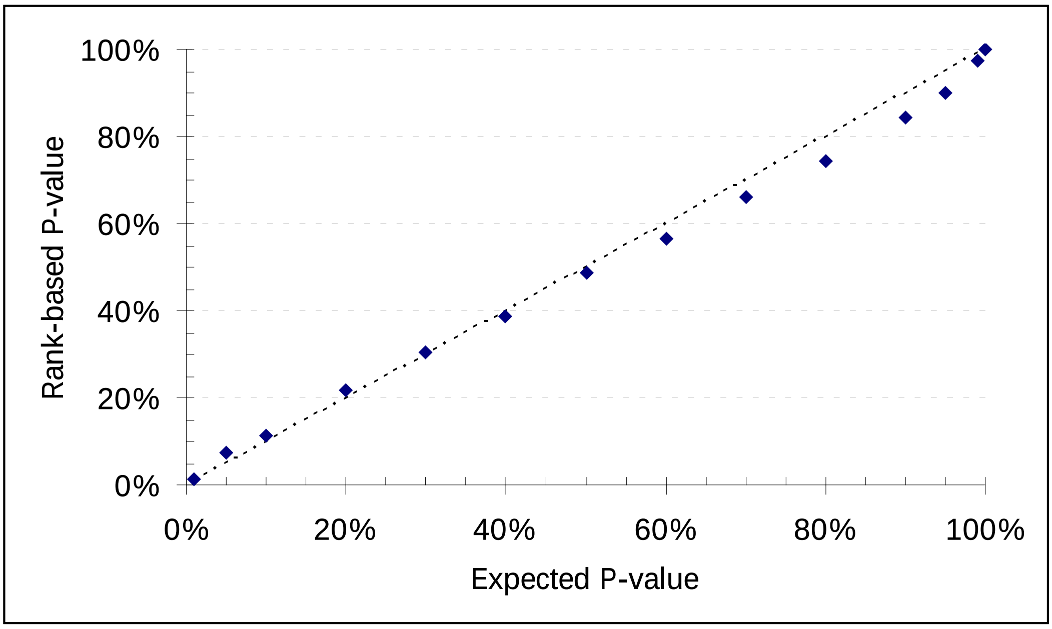

Figure 5.8 Quantile-quantile plot for claim distribution parameters estimated from the minimum disturbance method. The minimum number of claims allowed for self-experience was set at 100.

For comparison, Figure 5.9 shows the same quantile-quantile test in the situation when the distribution parameters for all industries were set on the basis of their industry division.

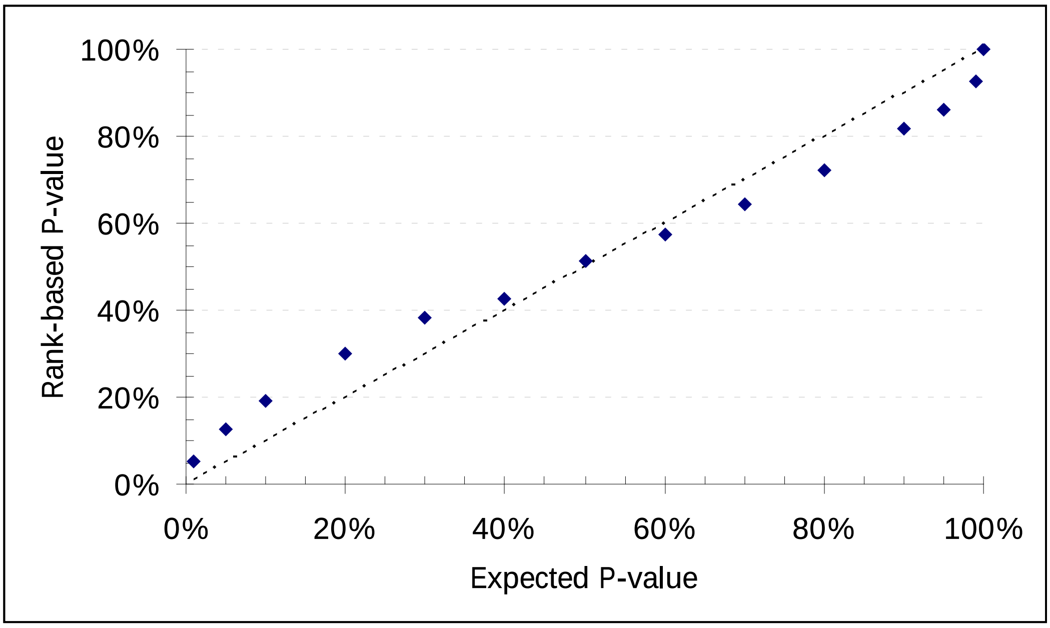

Figure 5.9 Quantile-quantile test when all industries are given industry group-average characteristics. The deviation from the straight line shows that in this case the variance of observed values is above the expected random variation because industry group-averages are, in general, poor estimators for individual industries.

7 Discussion

In this section we review the key results of this study and discuss to what extent they depend on the assumptions made and whether they can be transferred to other jurisdictions.

Stochastic model of aggregate claims costs and quantification of estimation accuracy

In our view, the key result is the realisation how really inaccurate are the estimates for underlying cost ratios for most industries. The assumptions made along the way were:

- Claim numbers are Poisson-distributed. In reality, negative binomial distribution would be more appropriate because claim occurrence rates can vary between individual employers and workers within the industry and the annual variability of claim numbers for a given industry appears greater than what is consistent with the Poisson process. However, all things being equal this should lead to an even greater volatility of aggregate claims costs. Hence, this assumption has not been restrictive.
- Independence of claim frequency and claim size. In South Australia, we know that over the years the number of all claims has decreased whilst the proportion of claims that involve income support have increased. Therefore, at the long time scale this assumption is definitely false. On shorter time scales there is also a possibility for frequency/claim size interaction, because employers may seek to under-report minor claims to avoid premium penalty under the bonus/penalty scheme which has changed over the years. All things equal, an interaction between claims frequency and size would lower the variability of aggregate costs. It is impossible to tell to what extent this effect is significant. We expect that in combination with our under-estimation of claim number volatility the net impact is small.
- Assumption of that claim size distributions for all industries have the same coefficient of variation. This is not an essential assumption and if required can be dropped, so that the error estimates for each industry rely on its own CV. As we have seen, the range of CV values for individual industries is rather small, so all key conclusions would remain the same.

Finally, while the results presented in the paper are based on the claim size distribution specific to SA, they can be readily reproduced for other jurisdictions by using their claims data to build the claim size distribution.

Adaptive experience period-length approach and its validation

In our view the key results were both the adaptive approach idea and the proposed validation methodology using quantile-quantile plots.

One of the challenges specific to setting of workers' compensation premiums is that it is virtually impossible to compare the rating basis with subsequent experience – not only one needs to wait until the bulk of claims costs for a given accident year run off, but also the realised relativities are so volatile that the random error makes it hard to assess whether the premium relativities implemented in that year were accurate. The quantile-quantile test overcomes this last difficulty by looking at the whole ensemble of industries and checking whether in aggregate the assumed relativities were appropriate. We are not aware of other techniques that achieve the same objective.

Minimum disturbance methodology

This idea has no sound theoretical basis and as such is probably the weakest part of the paper. It is, therefore, natural to ask why bother with an ad-hoc method when there exists an established hierarchical credibility framework that does a similar job.

In our view, apart from disconnect between the theory and practice discussed earlier, the main shortcoming of the credibility method is its heavy reliance on the assumed hierarchical structure of industries. The key premise of this method is that industries within a given higher-order group share similar risk characteristics. This is something that is quite hard to verify, particularly if the industries in a given group are small and have limited claims experience. Moreover, in practice it is common to use the existing industry classification based on the type of activity rather than the underlying claims risk. As a result, credibility weighting applied to small industries potentially may lead to situations when they are consistently charged unfair premiums just because of being incorrectly aggregated into a higher-order group.

In contrast, the minimum disturbance approach has a relatively little reliance on the assumed structure. Once an industry demonstrates that its experience is significantly different from that of its group, it is no longer bound to it. Unfortunately, the shortcoming here is the need to manage increased year-to-year volatility of premium rates for those industries that have broken free from their group.

It remains to be seen whether or not the minimum disturbance idea can be applied in practice. At present we plan to implement it for the forthcoming premium rates review in South Australia, however this will require more hindsight testing and developing a practical approach for managing the transition from credibility-based rates to the new ones.

8 References

Buhlmann H & Gisler A, 2005, *A course in credibility theory and its applications*, Springer

Hart D G, Buchanan R A & Howe B A, 2007, *Actuarial practice of general insurance, 7th edition*, Institute of Actuaries of Australia.

Appendix A

Agriculture, Forestry and Fishing

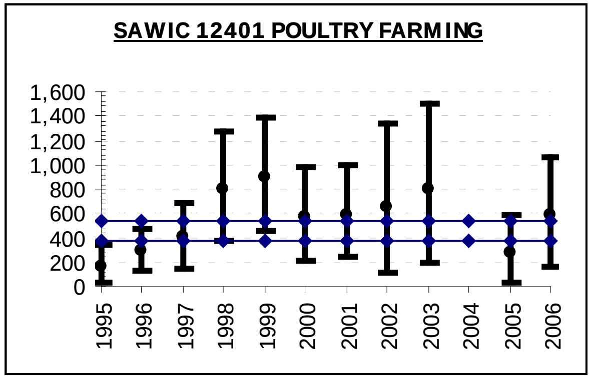

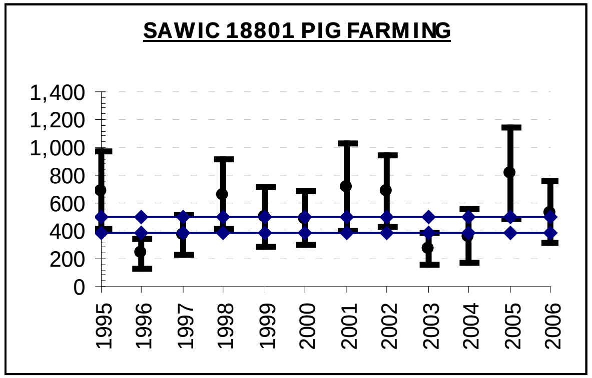

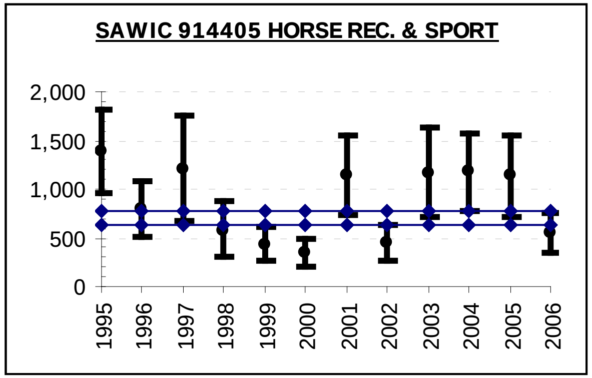

Mining

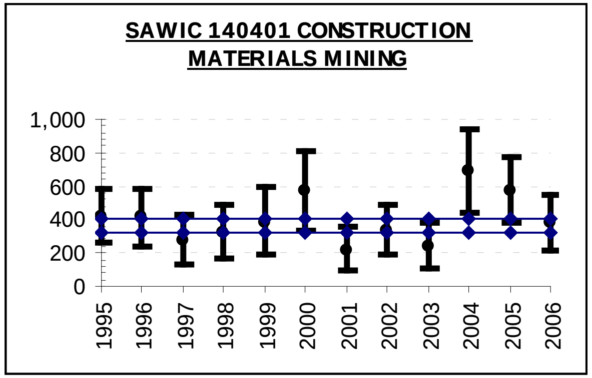

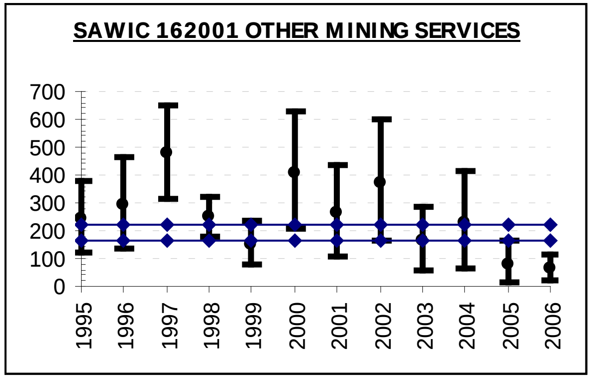

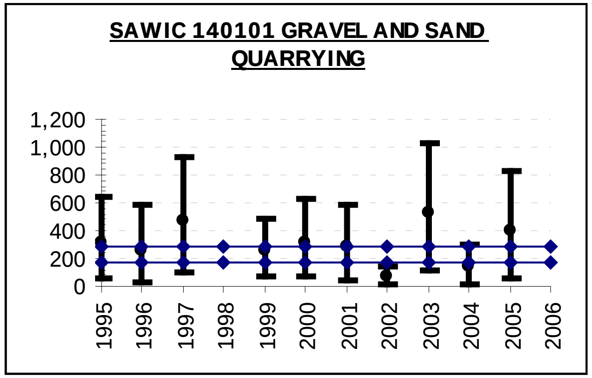

Manufacturing

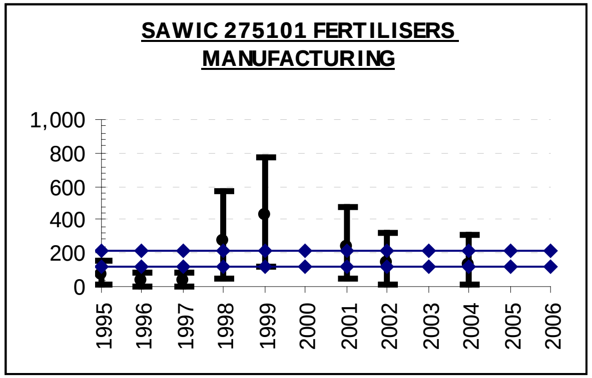

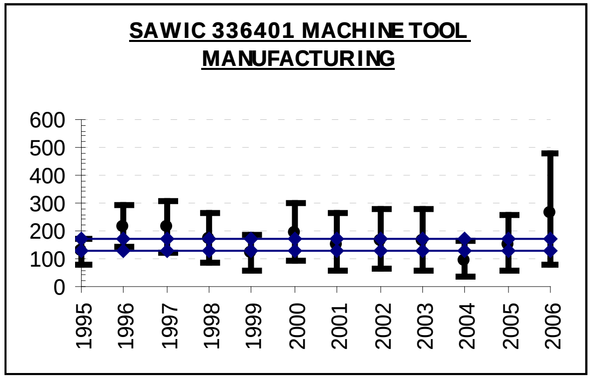

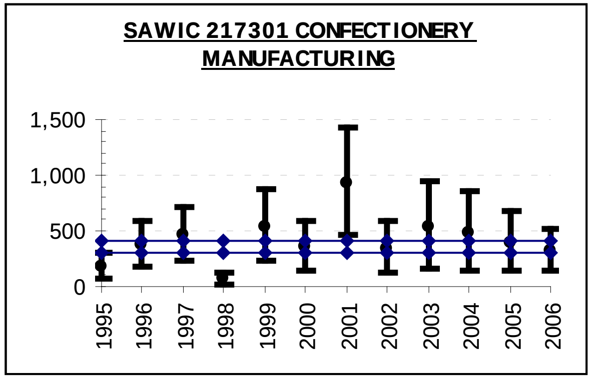

Low risk Manufacturing

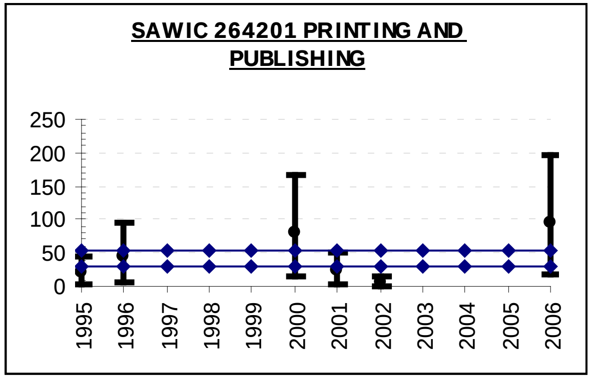

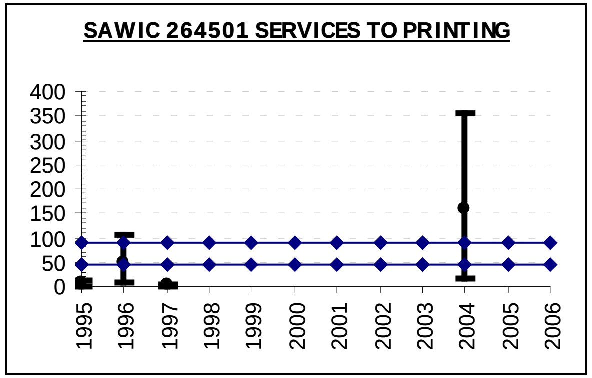

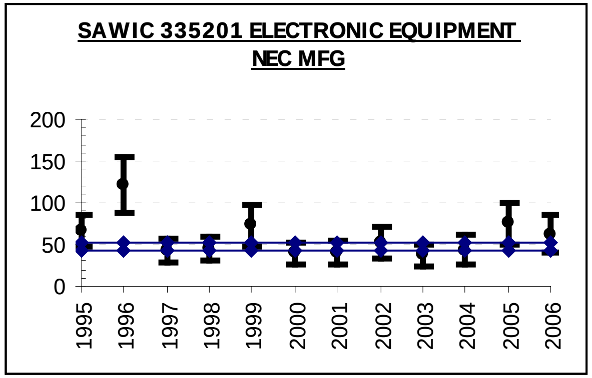

Construction

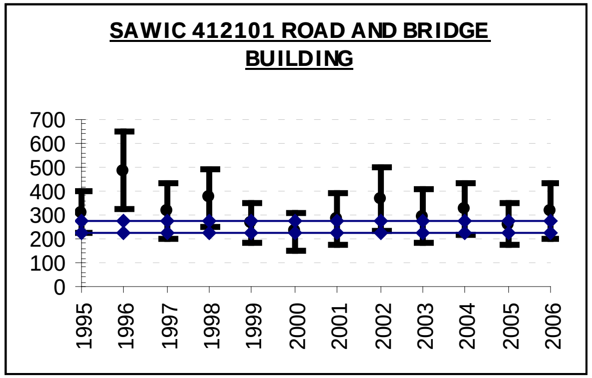

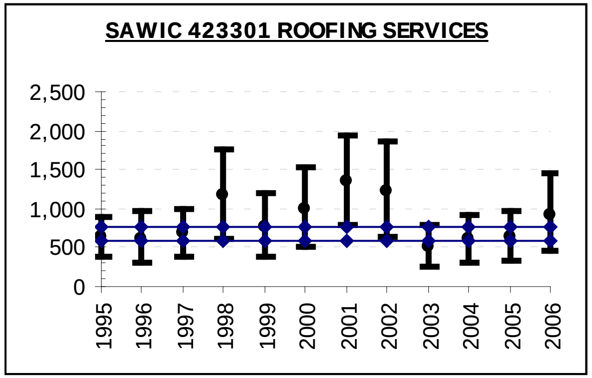

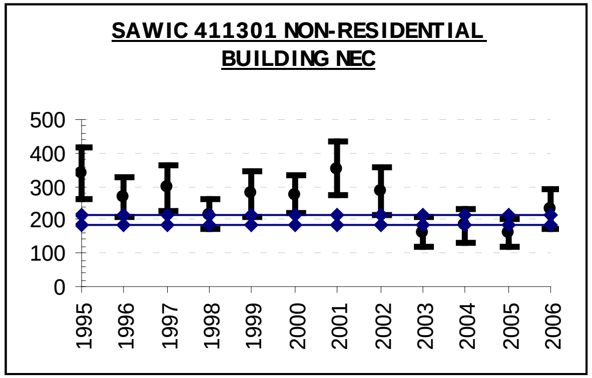

Wholesale and Retail Trade

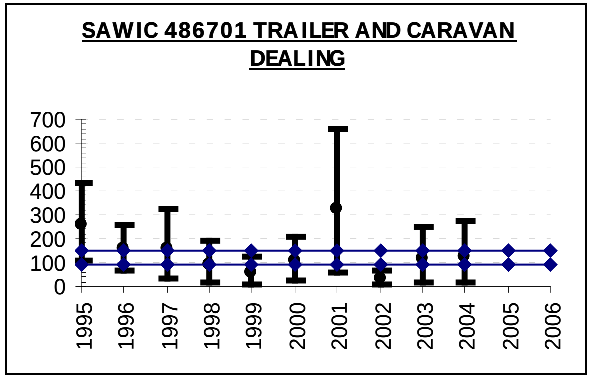

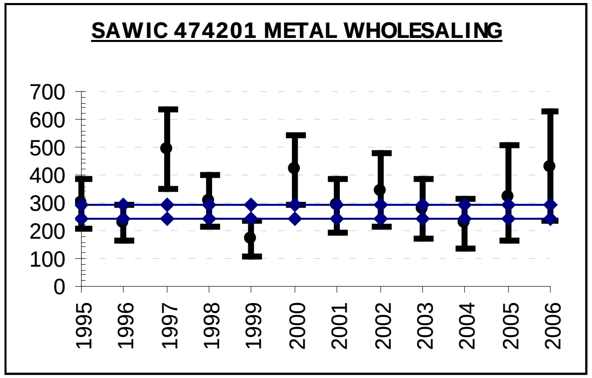

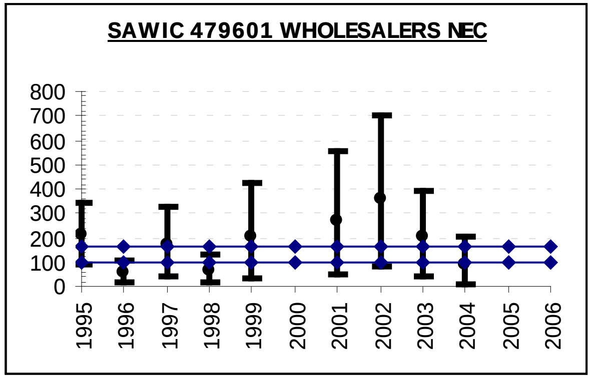

Low-risk Trade

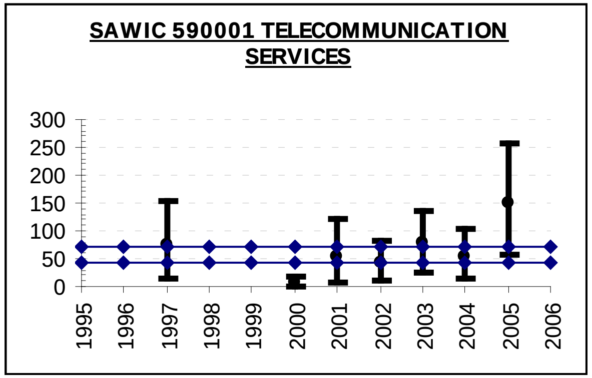

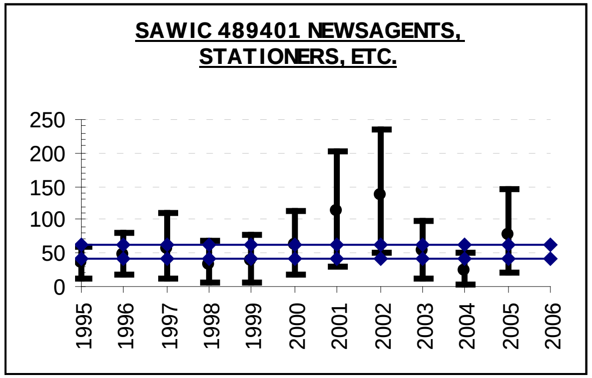

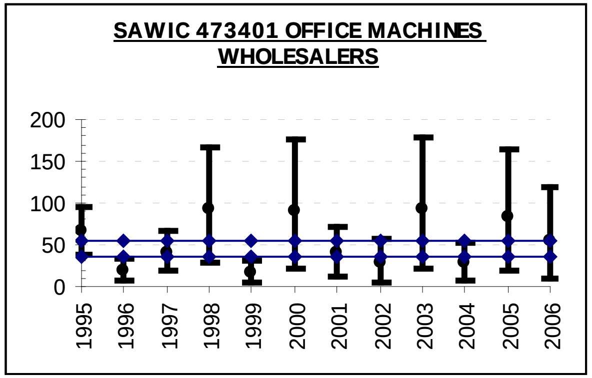

Transport & Storage

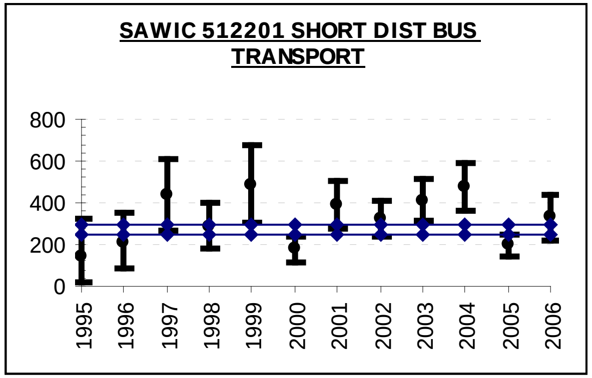

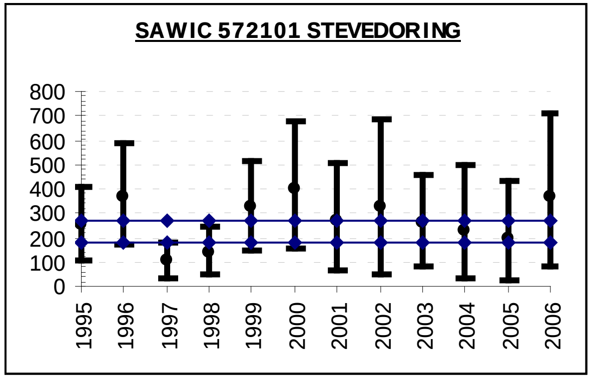

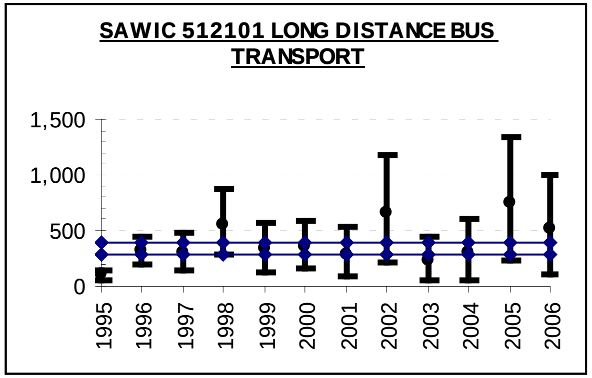

Finance Property & Business Services

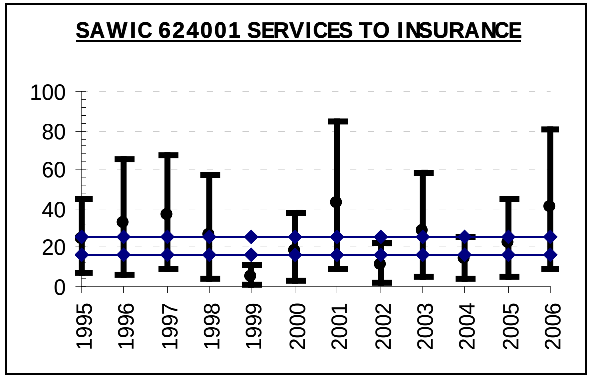

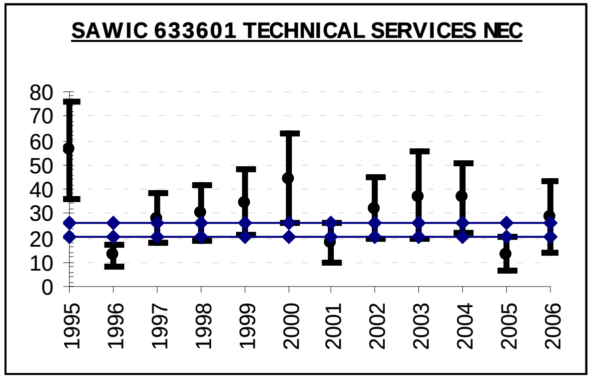

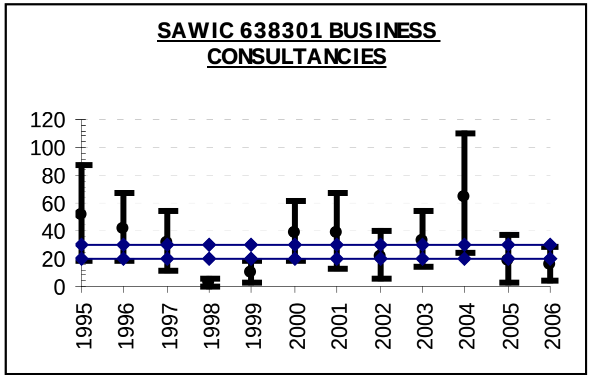

High-risk Services

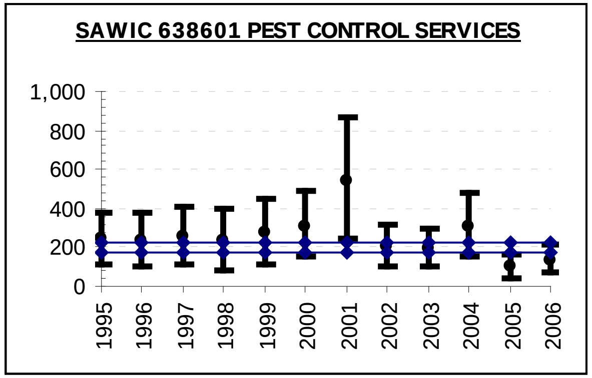

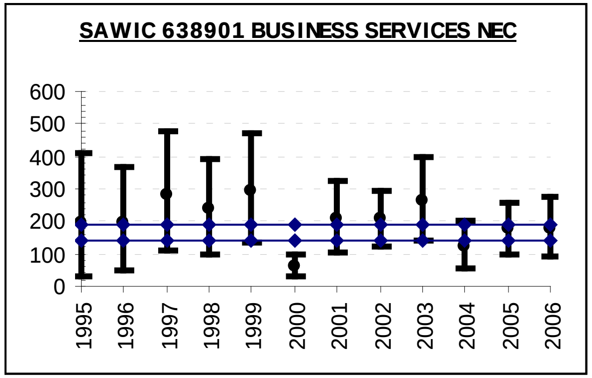

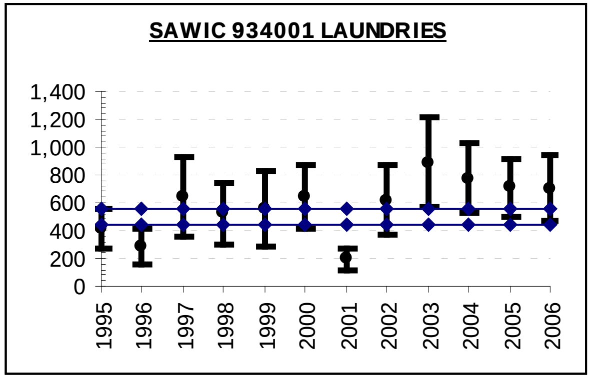

Community Services

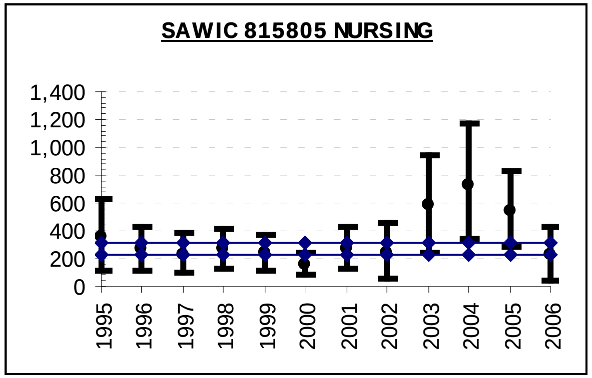

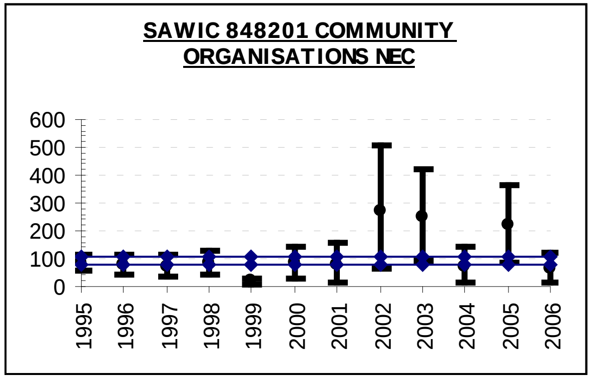

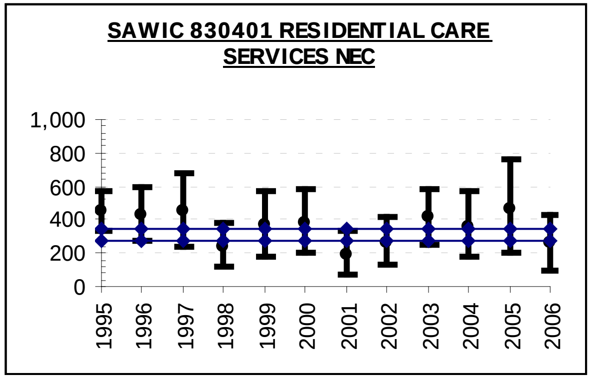

Skilled Professionals

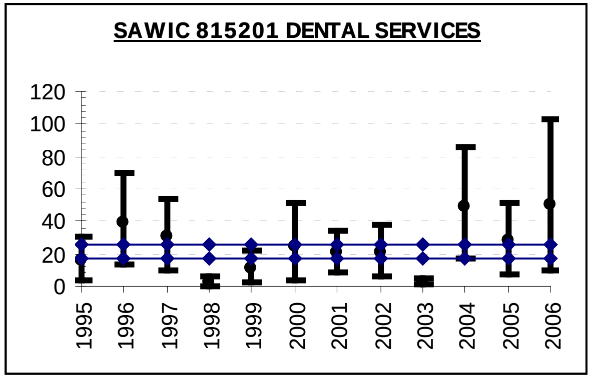

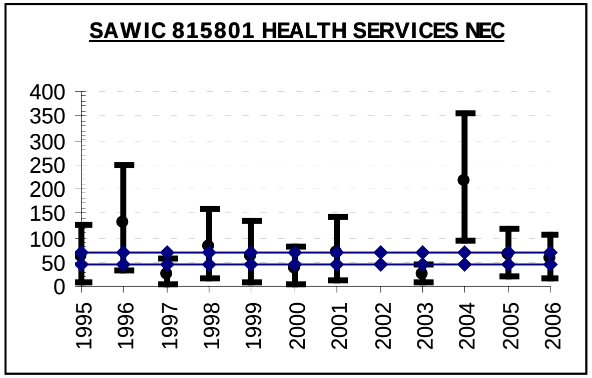

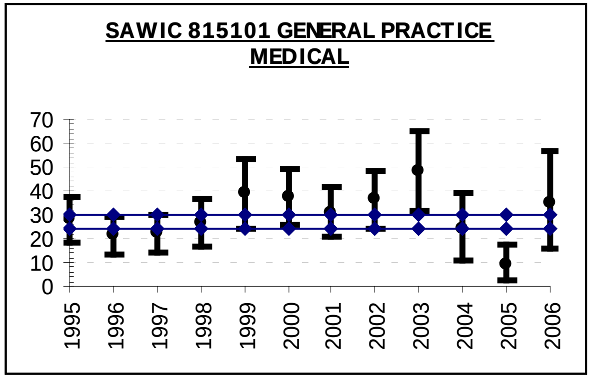

Recreation Personal Other

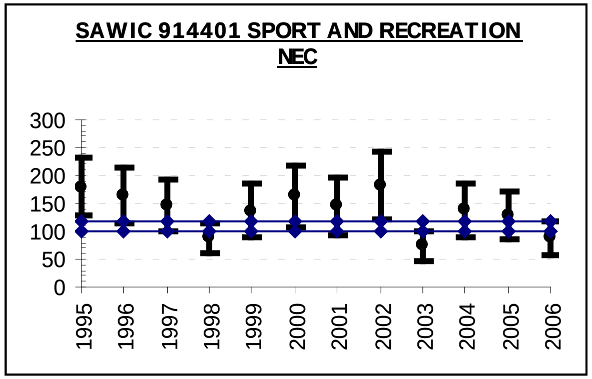

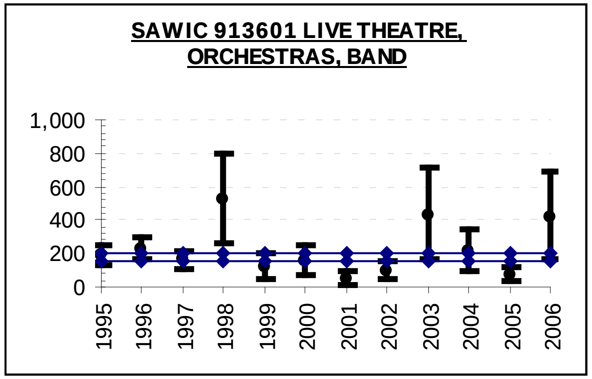

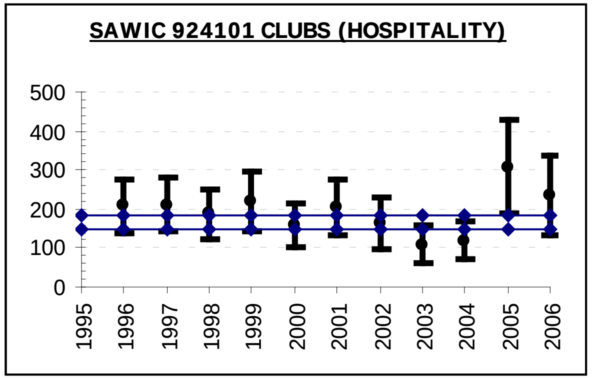

Generated using Poisson-distributed ($\lambda = 50$) number of claims and claims costs from the distribution in Figure 4.1