

Introduction

by Gary G. Venter

Selected Papers from Variability in Reserves Prize Program

This volume contains a selection of the papers submitted for the Committee on the Theory of Risk prize on how to measure the variability of loss reserves. Due to various constraints, not all of the submitted papers are included. Several of the excluded papers contained good analyses of reserving, but did not specifically address measuring variability. Others had some promising ideas not fully worked out into calculations. Hopefully these will be refined and submitted for publication in other venues.

Introduction

Several types of variation need to be accounted for to get a representative distribution of loss liabilities. Random fluctuation of the data around the expected value is generally called "process risk." Possible errors arising from estimating the mean, process variance, or parameters of any fitted model can be called "parameter risk." The standard error combines these two elements into a variance measure, and this is calculated in several of the papers. "Model risk" (sometimes called "specification risk") is an additional element of uncertainty arising from the possibility that the model assumptions themselves may be incorrect. A few papers attempt to quantify this as well.

The papers included here fell into three categories: Methods based on variance of link ratios; methods based on the collective risk model; and methods based on parametric models of development.

Methods Based on Variance of Link Ratios

Each age-to-age factor is a mean of several observed factors, so a variance can be calculated as well. Adding an assumption that the observed factors are samples from a lognormal, and that the ages are independent of each other, make the age-to-ultimate factors also lognormal, with readily computable variances. Both assumptions are possible to check, and adjustments can be made if they are too far off. The result is a distribution for the estimated liability for each accident year. Independence is important in that the product of the expected values is the expected value of the product for independent factors.

To get the distribution for the entire reserve, the distributions for the different accident years can be added by simulation or by matching moments of the sum. Normal, lognormal, and shifted gamma distributions are possible candidates for the summed distribution. Another one, not discussed in these papers but used by at least one committee member, is the shifted loglogistic $G(x - x_0) = F(x) = x^2/[x^2 + b^2]$. The moments for F are given by $E(X^n) = b^n/(n/a)!(-n/a)!$. The variance and coefficient of skewness from these are the same for G, being unaffected by the shift x_0 , and so they can be used to match the parameters a and b. Then x_0 is computed as needed to get the right mean for G.

Measuring the Variability of Chain Ladder Reserve Estimates, Thomas Mach

This paper tied for second prize in the competition. It contains a detailed discussion of what assumptions underlie the development factor (often called “chain ladder”) method; i.e., the assumptions that make this method optimal, and how to test if they hold. This includes a test for correlation of age-to-age factors as well as for other assumptions of the chain ladder method. Standard errors are measured without assuming age-to-age factors are lognormal, but age-to-ultimate factors are assumed to be lognormal in any case. The version of the chain ladder here uses all observed factors to compute mean age-to-age factors, but the formulas can be converted to apply to using only the last n diagonals by just using the last n terms of those sums indexed from 1 to I-k.

Unbiased Loss Development Factors, Daniel M. Murphy

Variations of link ratios are derived from loss development triangle data using regression statistics measuring both process and parameter risk. Regression is presented as a generalized procedure which can be used to model age-to-age factors from loss development triangles. Many techniques currently in use can be viewed as types of regression models.

Murphy describes some of the main regression assumptions and illustrates how these assumptions can be tested and used to select an appropriate model. He then describes a recursive calculation of variances of ultimate losses based on the regression statistics. Although the introduction discusses three models frequently used to estimate loss development factors (weighted average development, simple average development and geometric average development), the calculation of variances is

presented in detail for only two models: The least squares linear and the least squares multiplicative models. In actual practice, actuaries generally use the weighted average development or simple average development to estimate age-to-age factors. Using the paper's approach, the variances for the models more commonly used can be derived also, but the reader may need to derive the formulas from basic principles or refer to formulas (i.e. for weighted average factors) contained in an appendix.

Murphy presents the calculation of variances for multiple as well as single accident year ultimates. His formulas assume independence between development ages. Finally, he invokes the t-distribution to derive confidence intervals for the ultimate and the outstanding losses. In order to support the use of the t-distribution, he requires a further assumption that the variances are constant across development ages, which would need to be checked against actual data.

In addition to providing variance and confidence interval formulas, Murphy also uses a simulation procedure introduced by Stanard (PCAS 1986) to evaluate the bias and variance of eight development factor estimates. It would also have been informative if the simulations had been used to test the accuracy of the confidence interval estimates.

Correlation and the Measurement of Loss Reserve Variability, Randall D. Holmberg

An important issue for the development factor approach is potential correlation of link ratios. If they are correlated, the product of the age-to-age factors is not an unbiased estimate of the age-to-ultimate development, and the variance of the age-to-ultimate factor is understated. This paper provides a method to measure and adjust for correlation. The author suggests a simplified model in which the correlation ρ between a given age-to-age factor and the subsequent age-to-ultimate factor is constant for all ages. He then shows how this correlation can be estimated, and how it affects the reserve mean and variance. For the latter, an assumption on the distribution of the factors is made to simplify the computation, and here the uniform distribution is assumed. However, it would not be difficult to change to another distribution, just by plugging its density function and domain of definition into two integrals. The significance of the single ρ assumption is difficult to evaluate, and this area needs further support. The sensitivity to the distributional assumption would also be useful to

know. It may in fact be possible to avoid the distributional assumption by using moment formulas for correlated variables, e.g.,

$$E(AB) = E(A)E(B) + \text{Cov}(A,B), \quad V(AB) = V(A)V(B) + E(A)^2V(B) + E(B)^2V(A) + \text{Cov}(A^2,B^2) - \text{Cov}(A,B)^2 - 2E(A)E(B)\text{Cov}(A,B).$$

The paper at least touches upon several other important issues in the variance calculation. For the variance of tail factors it raises the possibility of expressing the standard deviation of a development factor as a linear function of $\ln(\text{factor}-1)$, and applying this to the last actual factor and the tail factor. Correlation among accident years and among lines of insurance is also estimated.

Variability of Loss Reserves, Robert L. Brown

The effects of parameter risk and correlation among companies are illustrated in this paper, which looks at historical variability in loss reserves for a large sample of companies. Reserve adequacy for the entire sample showed a cyclical variation over time that would not be observed for a like sample of independently fluctuating companies. Reserve adequacy was found to vary by size of company as well, but the largest identifiable influence was consistent variation among companies: Some tended to be more adequately reserved than others over considerable periods of time, even taking into account all other impacts.

Methods Based on the Collective Risk Model

The basic idea of using the collective risk model to measure variability in loss reserves, as outlined in Roger Hayne's paper in the 1989 *PCAS*, is to estimate frequency and severity distributions for outstanding claims, and combine these to get an aggregate loss distribution for those claims. Hayne originally did this separately for reported and unreported claims.

A Method to Estimate Probability Levels for Loss Reserves, Roger M. Hayne

The earlier work by Hayne is expanded to include parameter uncertainty. This is broadly defined to include not just uncertainty about the parameters of a given model, but also the variability that can arise from using different modelling approaches. Significantly greater uncertainty in the reserves is found when this is taken into account. Hayne presents a detailed illustration of his procedure using a professional liability data set from the Berquist-Sherman 1975 *PCAS* paper and for an auto liability

data set from the Advanced Case Study of the 1992 CLRS. These present straightforward techniques for estimating parameters of the claim count and severity distributions and require only a modest amount of data. The severity distribution parameters here are determined somewhat judgmentally.

It should be noted that the use of individual claim information would produce superior parameter estimates, although such information often is not available. It should also be noted that the selected lognormal coefficients of variation appeared to be low for this line of business. (However, Hayne notes that his example illustrates only one of many ways of selecting parameters and he provides some reasoning for his parameter selections). The most innovative contribution of this paper is the use of the results of different methods of estimating reserves to derive the mixing parameter for the severity distribution. This allows the actuary to incorporate specification error into the estimation of loss reserve variability. Once the claim count, severity, contagion and mixing parameters are estimated, the Heckman-Meyers procedure is used to compute the aggregate probability distribution for the loss reserves. Simulation could also be used to implement this approach.

A Note on Simulation of Claims Activity for Use in Aggregate Loss Distributions, Daniel K. Lyons

This paper suggests using severity distributions for both paid and case incurred losses at different valuations and annual probabilities of claims moving from one severity class to another (a transition matrix) to project claim movement over time. The severity distributions are incorporated into a simulation which 1) simulates the number of claims for a year, 2) simulates the report lag for each claim, 3) simulates the movement in each claim's value over time until an ultimate value is reached, and 4) works backward from the ultimate value of the claims to simulate their paid value. By simulating many years of data distributions of paid, incurred and outstanding losses can be produced. The procedure described in the paper could be used to approximate the process which underlies loss development when the losses are aggregated.

The author illustrates his method using severity distributions and transition matrices which have been judgmentally selected; i.e., not based on real data. To actually apply this technique, one would have to construct actual severity distributions and transition probabilities using techniques not described in the paper. The author's example applies to outstanding losses at the beginning of an acci-

dent or policy period, before any losses have been reported. While he mentions that in real life reserving situations, the actuary would need to determine outstanding loss severity distributions on a conditional basis, he does not describe how to do this.

Methods Based on Parametric Models of Development

The chain ladder method is characterized by having a separate level for each accident year and a separate degree of development for each lag. Thus models that have a parameter for each accident year and for each lag are regarded as statistical versions of the chain ladder. This can now be seen to be somewhat of a misnomer, as the assumptions of the usual chain ladder, as outlined in the Mack paper, are significantly different from these models. The logarithms of the losses in an incremental claims triangle (paid in year, for example) may meet the assumptions of regression analysis, which then can be used to estimate model parameters and provide variances.

Statistical Methods for the Chain Ladder Technique, Richard J. Verrall

This paper, which took first prize in the competition, gives a comprehensive presentation of the use of regression models to estimate loss development. It also lays out an interesting approach to adjusting lognormal maximum likelihood estimators for bias, and shows how to construct some Bayesian estimators relevant to the model. The paper does not note, however, that adjusting the MLE of the lognormal mean for bias involves some controversy, with different authors advising upward or downward or no adjustment. The Bayesian estimates discussed include estimation of runoff, estimates for the analysis of variance model, and relation to credibility theory. Relations to the chain ladder method are also discussed, and an excellent list of references is provided.

Probabilistic Development Factor Models with Applications to Loss Reserve Variability, Prediction Intervals and Risk Based Capital, Ben Zehnwrith

Loglinear versions of chain ladder, Cape Cod, and separation method are all outlined. The paper, which tied for second prize, also addresses models that allow for changing parameters over time or smoothing of parameters to avoid multicollinearity. It contains a general discussion of statistical forecasting methods, and sufficient detail is given that many of the examples presented can be reproduced by the reader.

While many of the assumptions made are explicitly listed, there are a number of assumptions that are either unstated, or appear to be statements of fact. For example, in Section 5.0 there is a statement that the logarithm of paid losses at the earliest stage of development has a normal distribution with a mean α and variance σ^2 . This is an assumption implicit in the main model here, but is not necessarily true in general. The assumptions about inflation also need to be carefully evaluated. Inflation is assumed to affect all payments in a given calendar year equally, but in fact losses at different stages of settlement might be affected differently.

Even though the model assumptions may not apply for every data set, this paper gives a comprehensive discussion of methods for fitting a regression model to development data and the testing of such a model for goodness of fit.

IBNR Reserve under a Loglinear Location-Scale Regression Model, Louis Doray

Most authors who use regression to model loss development assume that the initial data (incremental paid losses for example) is lognormally distributed. They take the logarithm of the initial data, and fit linear regression models to the logged data. The logged data is then normally distributed, and the error term (the difference between the fitted values and the logged data) is also normally distributed—hopefully with a reduced variance and zero mean.

The main thrust of this paper is to explore four possible distributions of the error term other than the normal distribution. In each case it presents the mathematics needed and tests the model against a common data set. Maximum likelihood estimation is used to estimate the regression coefficients. Because of the complexity of the various distributions used, and the need for various second derivatives of the log-likelihood function, the use of a computer algebra system would help implementation. Various issues regarding goodness of fit and bias of estimates are discussed. Possible bias of the maximum likelihood estimates is not discussed. The paper does suggest incorporating interest rate risk in presenting interval estimates for discounted reserves.

A comparison of estimates made by a regression method to estimates made using the chain ladder method shows that if the regression model is correct, the chain ladder method underestimates re-

serve needs. The correctness of the regression model is not verified, however. The issue of moving back to estimates of the unlogged data is also not addressed.

A Generalized Framework for the Stochastic Loss Reserving, Changsoeb Joe Kim

This approach measures variability by using goodness of fit from time series (ARIMA) models. It may, however, require a great deal of stability across accident years. The author uses a two dimensional auto-regressive procedure to estimate future incremental payments on a loss development triangle. The application of the procedure assumes that a constant auto-regressive parameter(s) applies to all accident years (which appears to be reasonable) and a constant parameter(s) applies to all development ages (which may not be reasonable). Thus, the payment at age 10 is assumed to be the same proportion of the age nine payment as the payment at age two is of the age one payment. This technique does not address the “tail” problem, or the estimation of payments at development ages greater than that in the historical data. Because the number of observations in most triangles is relatively small and time series techniques generally require a large number of points, the author uses standard loss development factors to convert the triangle into a matrix and derive initial values for the fitting process. Formulas presented for the n-year-ahead variance of the two dimensional auto-regressive process can be used to compute confidence intervals, presumably by using the standard normal distribution, but this is not explicitly stated. (It should be noted that the formula given for the one year ahead variance appears to actually be the formula for the two year ahead forecast variance).

Outstanding Issues

Several issues are still not addressed and could benefit from further research:

1. What techniques are appropriate for which situations and what kinds of data? For instance the regression techniques seem to require relatively stable, homogenous data. The development factor methods require enough observations in each column for a reasonable estimate of the variance of factors in the column. The collective risk model methods require estimates of claim count and severity parameters and these can best be derived from individual claim data. When these parameters are selected based on aggregate data or judgment, does the aggregate probability model reflect the additional uncertainty contributed by these less rigorous parameter estimates?
2. More work is needed on the “tail” problem. How does the actuary quantify the variability for development ages beyond the last observation in the data? The uncertainty associated with the tail can be substantial.

3. The impact of correlation needs further analysis. This includes correlation within the development triangle, among lines of business, with inflation, and with interest rates, especially for discounted reserved.
4. How does the actuary realistically reflect the uncertainty in reserve estimates for companies or lines of business with little or no data, or with recent changes in the data? It is reasonable to assume that the variability of such reserve estimates should be higher than for a company or line of business with abundant data. What about when different data sets are combined (company/industry, external indexes, etc.)?
5. What kind of testing is needed to truly validate the use of these models? Tests based on the triangle and fitted data can invalidate models, but failure to invalidate does not necessarily give much comfort for forecasting. An understanding of the assumptions used, and reflecting on their reasonableness may always be necessary, regardless of the fit provided.
6. How can the regression models be enhanced to incorporate a finite probability of no losses paid in a future period for given accident years? For small companies this is a realistic possibility, and should be reflected in prediction intervals.

Gary G. Venter