The Parameter Variance Adjustment in Lognormal Linear Models for Loss Reserves: Bayesian versus Frequentist Analysis

by Frederick L. Klinker, FCAS

The Parameter Variance Adjustment in Lognormal Linear Models for Loss Reserves: Bayesian vs. Frequentist Analysis

by Fred Klinker

Abstract: In lognormal linear models for loss reserve estimation, losses are assumed to be lognormally distributed, where the expectations of the logarithms of losses are assumed linear in explanatory variables. A parameter variance term appears in the exponent of the estimator for expected losses. There is disagreement regarding the sign of this term. It will be argued in this note that the sign depends on whether one adopts a Bayesian or Frequentist viewpoint. Each sign is correct within the appropriate paradigm.

A number of actuarial papers have considered lognormal linear models for loss reserve estimation, among them Verrall [11], Verrall [12], Wright [14], and Zehnwirth [15]. This list is illustrative only and is far from exhaustive. In such models, losses (generally incremental, not cumulative) are assumed to be lognormally distributed, where the expectations of the logarithms of losses are assumed linear in explanatory variables. A parameter variance term appears in the exponent of the estimator for expected losses. There is disagreement regarding the sign of this term. The disagreement is implicit rather than explicit; none of the above referenced authors appears to acknowledge the different sign in other authors' works. However, Gary Venter, in his introduction to the papers on variability in reserves included in the Spring 1994 CAS Forum, specifically

in his comments on Verrall [12], notes that "...adjusting the maximum likelihood estimator of the lognormal mean for bias involves some controversy, with different authors advising upward or downward or no adjustment." (Venter [10], page 97.) It will be argued in the rest of this note that the sign of the adjustment depends on whether one adopts a Bayesian or Frequentist point of view. Each sign is correct within the appropriate paradigm.

Aside from its discussion of an admittedly technical fine point which may not interest many actuaries, this note may also serve to remind readers of the fundamental distinctions between Bayesian and Frequentist paradigms and the fact that the two do not always yield the same result. This last reminder is useful, since the statistical model most actuaries are most familiar with, the normal linear model, yields the same result whether from a Frequentist viewpoint or from a Bayesian (with uninformative prior), although the interpretation of the result differs somewhat according to viewpoint. (Regression and ANOVA are common examples of linear models. The normal linear model in a Frequentist setting assumes normally distributed errors. In a Bayesian setting, normal priors and normal errors are assumed, resulting in normal posteriors.) This happy coincidence of Bayesian and Frequentist results is not preserved in many other models, including lognormal linear models.

In subsequent sections of this note, first the general Bayesian and Frequentist paradigms are discussed, then the estimators that follow from these paradigms. Lastly, the special case of lognormal linear models is introduced.

Paradigms: Bayesian vs. Frequentist

The general problem is as follows. The state of nature is described by the parameter μ. There is a quantity of interest, y, whose expectation, conditional on μ, is a function of μ, $E[y|\mu]=f(\mu)$. If the state of nature were known to be μ, both Bayesians and Frequentists would probably agree that a good estimator for y is $f(\mu)$. However, the state of nature is not known. Data, x, either a single observation or a collection of observations, is collected in an attempt to determine μ and y. But Bayesians and Frequentists proceed differently.

First, the Frequentist approach: The state of nature, μ, is considered to be fixed, although unknown. The Frequentist refuses to quantify uncertainty surrounding μ via a probability distribution on μ. On the other hand, the Frequentist considers not only the actual outcome of the experiment, x, but also other outcomes that might have been, but weren't. The possible outcomes are described by a probability distribution on x, conditional on the fixed but unknown μ. Expectations and variances of functions of x are calculated over x, conditional on μ. The focus is on finding unbiased estimators $\hat{\mu}$ and $\hat{f}$ such that $E[\hat{\mu}(x)|\mu]=\mu$ and $E[\hat{f}(\hat{\mu}(x))|\mu]=f(\mu)=E[y|\mu]$.

Consider next the Bayesian point of view. Uncertainty surrounding the state of nature, μ, is quantified via a prior probability distribution on μ. This prior can be Objective Bayes (an uninformative prior), Subjective Bayes (based on personal estimates of probabilities), or Empirical Bayes (based on previous data from similar problems). Data,

x, is observed, and, based on this data and Bayes' Rule, a posterior distribution for μ follows. All inferences are conditioned on the observed data. There is no consideration given to other outcomes that might have come to pass but didn't. The focus is no longer on unbiased estimators. Unbiasedness is a Frequentist notion which requires taking expectations over actual and possible observed data, whereas the Bayesian does not consider the randomness of the data after the data has been observed and instead conditions on that observed data. Instead, the Bayesian desires an estimator which minimizes Bayes Risk across all states of nature still considered possible after observing x. Expectations and variances are calculated over μ via the posterior distribution for μ, conditioning on x. Adopting the standard loss function (quadratic), the minimum Bayes Risk estimator for y is its posterior expectation, $E[y|x] = E[E[y|\mu]|x] = E[f(\mu)|x]$.

To summarize the key distinctions between Bayesian and Frequentist, the Frequentist considers the data, x, to be a random variable, but not μ, which is considered fixed, although unknown. The Frequentist continues to worry, even after the data is observed, about observational outcomes that could have come to pass but didn't, and considers expectations and variances over x, conditional on μ. The Bayesian conditions all inferences on the observed data, x, and considers μ to be the random variable over which posterior expectations and variances are calculated. The Bayesian steadfastly refuses to be concerned about outcomes that could have come to pass but didn't. To clear up a common misconception, it is this conditioning on x which is the heart of the Bayesian paradigm, not the invocation of Bayes' Rule. Even some Frequentist methods invoke Bayes' Rule.

Before leaving this foundational section of this note, a few clarifying comments are in order.

1) In the above Frequentist discussion, I have focused on unbiased estimators. It should at least be noted that Frequentists do occasionally invoke considerations other than unbiasedness. However, it is certainly true that unbiasedness is one of the first characteristics that a new statistics student learns and one that is invoked often.

2) In the above Bayesian discussion, by mentioning Bayes Risk and loss functions, I have implicitly adopted a decision theoretic approach to Bayesian statistics. It should be noted that Bayesian theory and statistical decision theory are not synonymous. There are practicing Bayesians who are not decision theorists, at least not knowingly. And there are decision theorists who are not Bayesian, but rather quite decidedly Frequentist. On the other hand, of those discussions of Bayesian foundations with which I am familiar, all the best seem to adopt a decision theoretic viewpoint. Once one rejects the questionable Frequentist "objectivity", one seems driven naturally towards a decision theoretic viewpoint. Statistics appears to be less a method of discovering "truth" and more an aid to rational decision making. Any Bayesian can calculate the posterior expectation, $E[y|x]$. Only the Bayesian with a decision theoretic bent knows why this might be the appropriate quantity to calculate, because it minimizes posterior Bayes Risk under the most popular loss function, expected squared error.

The above has only scratched the surface. For those interested in more, Silvey [9] is a good introduction to Frequentist inference. There is a rapidly expanding literature on Bayesian foundations. Two good discussions are Berger [2] and Efron [6]. I particularly recommend the first of the two as an excellent discussion of Bayesian philosophy. (This is the source which first pounded into me the central role played by conditioning arguments.) The second of the two compares Bayesian and Frequentist paradigms. Although more applied, Gelman and others [8] and West and Harrison [13] also have interesting insights on Bayesian foundations.

Estimators: Bayesian vs. Frequentist

Suppose, first, that the function $f(\mu)$ of the previous section of this note is linear in μ. Then the following two operations commute: 1) taking expectations and 2) evaluating the function. In the Frequentist paradigm, $E[f(\hat{\mu}(x))|\mu]=f(E[\hat{\mu}(x)|\mu])=f(\mu)$ for $\hat{\mu}(x)$ an unbiased estimator of μ. In other words, $\hat{f}=f$ and $\hat{f}(\hat{\mu}(x))=f(\hat{\mu}(x))$ is an unbiased estimator for $E[y|\mu]=f(\mu)$. In the Bayesian paradigm, $E[y|x]=E[f(\mu)|x]=f(E[\mu|x])=f(\mu_x)$, where $\mu_x=E[\mu|x]$ is the posterior expectation of μ conditional on the observed x. Comparing the Bayesian and Frequentist estimators for y, they are of the same functional form as long as we identify the Bayesian μ_x with the Frequentist $\hat{\mu}(x)$. Why is the class of linear f so important? Because the normal linear model, already mentioned in the introduction, falls into this class.

Now assume that f is non-linear and take the Taylor series expansion to second order, about μ in the Frequentist case, and about μ_x in the Bayesian. It is not suggested that this calculation produces good estimators in all situations, but second order is the lowest order in which interesting phenomena arise, which are at least suggestive of the form of adjustments required for non-linear f. Considering first the Frequentist case,

$$E[f(\hat{\mu}(x)) \mid \mu] \approx E[f(\mu) + f'(\mu)(\hat{\mu}(x) - \mu) + \frac{1}{2} f''(\mu)(\hat{\mu}(x) - \mu)^2 \mid \mu]$$

$$= f(\mu) + \frac{1}{2} f''(\mu) \, Var[\hat{\mu}(x) \mid \mu]$$

(1)

where $\hat{\mu}(x)$ is an unbiased estimator for μ, and where the variance in the last line is the variance of the estimator $\hat{\mu}(x)$ conditional on μ. This equation suggests that $f(\hat{\mu}(x))$ would not in general be an unbiased estimator for $f(\mu)$ and, further, that the following might be approximately unbiased.

$$\hat{f}(\hat{\mu}(x)) = f(\hat{\mu}(x)) - \frac{1}{2} f''(\hat{\mu}(x)) \, Var[\hat{\mu}(x) \mid \mu] \qquad (2)$$

The unknown μ in the second derivative of f has been replaced by its unbiased estimator. The variance would also have to be estimated somehow. The unbiasedness would presumably be only approximate for a couple of reasons. First, higher order terms in the Taylor series expansion have been ignored. Second, both the μ (in the second derivative) and the variance in the variance adjustment term of equation (2) must be estimated, hence this variance adjustment term is itself a random variable, not a constant. There is no guarantee that the expectation of this random variable will be exactly numerically equal to the variance adjustment term of equation (1), barring a very judicious choice of variance estimator.

The Bayesian calculation is similar to the Frequentist.

$$E[f(\mu)\,|\,x] \approx E[f(\mu_x) + f'(\mu_x)(\mu-\mu_x) + \frac{1}{2}f''(\mu_x)(\mu-\mu_x)^2\,|\,x]$$

(3)

$$= f(\mu_x) + \frac{1}{2}f''(\mu_x)\,Var[\mu\,|\,x]$$

Equation (2) is the approximately unbiased estimator for y in the Frequentist case, equation (3) the approximate estimator for y in the Bayesian. Both have an adjustment for parameter variance. As before, upon identifying the Frequentist $\hat{\mu}(x)$ with the Bayesian μ_x, the functional forms would be identical, except that the signs of the parameter variance terms are opposite.

The Lognormal Linear Model

Consider first the lognormal distribution. A random variable z is said to be lognormally distributed with parameters μ and σ if and only if the natural log of z is normally distributed with expectation μ and standard deviation σ. μ and σ^2 are therefore the expectation and process variance in the log scale. Back in the original scale, the expectation of z, conditional on μ and σ, is $E[z|\mu,\sigma]=\exp(\mu+.5\sigma^2)$. For the actuarial reader unfamiliar with the lognormal distribution, past actuarial papers, such as Bickerstaff [3] and Finger [7], have made use of this distribution and include either a brief description or technical appendix on the lognormal. Those who desire considerably more detail on the lognormal distribution may consult Aitchison and Brown [1] or Crow and Shimizu [5].

Consider now the lognormal linear model. The data, x, and the quantity of interest, y, are assumed to be lognormally distributed, with expected logs that are linear in explanatory variables. The state of nature is characterized by the expectation of log(y), μ, and the process standard deviation of log(y), σ. μ will be linear in explanatory variables and their associated regression coefficients. The parameter variance of μ will depend on variances and covariances of the estimated regression coefficients via standard regression formulas involving the process variance and the structure matrix. In what follows, the process variance and the parameter variance will be assumed known. The fact that process and parameter variances must generally be estimated from the data is a technical complication which must be considered when designing exact estimators but which contributes nothing to the discussion at the present elementary level. So we will treat σ as a known rather than unknown descriptor of the state of nature and write $E[y|\mu]=\exp(\mu)\exp(.5\sigma^2)=f(\mu)$. (The additional problems introduced by unknown process and parameter variances, which must also be estimated, are treated in Verrall [11] and Verrall [12]. These two papers further reference Bradu and Mundlak [4], a highly educational paper in itself.)

The Frequentist now considers the problem to be one of estimating μ an $E[y|\mu]=f(\mu)$ from observed data x using unbiased estimators. Given the assumption that logs are normally distributed and linear in explanatory variables, standard regression analysis on the logs yields an unbiased linear estimator for μ, call it $\hat{\mu}(x)$, and an expression for the parameter variance of this estimator, $Var[\hat{\mu}(x)|\mu]$, in terms of the process variance, assumed known, and the structure matrix of the regression. Applying equation (2), an approximately unbiased estimator for $f(\mu)$ is:

$$\hat{f}(\hat{\mu}(x)) \approx f(\hat{\mu}(x))\left(1 - \frac{1}{2}Var[\hat{\mu}(x)|\mu]\right)$$

$$\approx f(\hat{\mu}(x))\exp\left(-\frac{1}{2}Var[\hat{\mu}(x)|\mu]\right) \quad (4)$$

$$= \exp\left(\hat{\mu}(x) - \frac{1}{2}Var[\hat{\mu}(x)|\mu] + \frac{1}{2}\sigma^2\right)$$

The first approximation follows from equation (2), because, given the present definition of $f(\mu)=\exp(\mu)\exp(.5\sigma^2)$, the second derivative of f with respect to μ is equal to $f(\mu)$ itself. The second approximation follows if the parameter variance is small (because $1-x \approx e^{-x}$ if x small), which is probably the regime in which equation (2) is valid anyway. (It should be noted that equation (4), derived under the above approximations, is an exact unbiased estimator for $f(\mu)$ if the variance terms are known, rather than estimated and the distribution of the data, x, is such that the estimator, $\hat{\mu}(x)$, is not only unbiased but normally distributed.) The second term in the exponent is the adjustment for parameter variance and appears with a negative sign.

Consider now the Bayesian estimator, E[y|x]. After observing the data x, μ has a posterior distribution with expectation μ_x and variance Var[μ|x]. Applying equation (3),

$$
\begin{aligned}
E[y|x] &= E[f(\mu)|x] \\
&\approx f(\mu_x)\left(1 + \frac{1}{2}Var[\mu|x]\right) \\
&\approx f(\mu_x)\exp\left(+\frac{1}{2}Var[\mu|x]\right) \\
&= \exp\left(\mu_x + \frac{1}{2}Var[\mu|x] + \frac{1}{2}\sigma^2\right)
\end{aligned}
\tag{5}
$$

This holds to the same level of approximation as equation (4). (Actually, if the posterior distribution for μ is normal with expectation μ_x and variance Var[μ|x], then equation (5) follows exactly, without approximation, because then exp(μ) is itself lognormally distributed.) Again, the second term in the exponent is the adjustment for parameter variance, but in the Bayesian setting it appears with a positive sign.

Note that equations (4) and (5) have the same functional form, except that the signs on the parameter variance term are reversed. Why? The Frequentist recognizes that his unbiased estimator for μ, $\hat{\mu}(x)$, has finite, non-zero variance. Because of the convex shape of the exponential function, excursions of $\hat{\mu}(x)$ above μ result in excursions of exp($\hat{\mu}(x)$) above exp(μ) of greater magnitude than excursions of exp($\hat{\mu}(x)$) below exp(μ) due to excursions of $\hat{\mu}(x)$ below μ. As an estimator of exp(μ), exp($\hat{\mu}(x)$) is therefore biased upward, and the bias is greater the greater the variance of the estimator $\hat{\mu}(x)$, the larger the excursions of $\hat{\mu}(x)$ from μ. The exp(-.5$Var[\hat{\mu}(x)|\mu]$) factor removes this bias (approximately).

The Bayesian, on the other hand, estimates E[y|x]=E[f(μ)|x]=E[exp(μ+.5σ^2)|x]. Again, because of the convex shape of the exponential function, excursions of μ

above μ_x=E[μ|x] have a larger impact on exp(μ) than excursions of μ below μ_x. Upward excursions of μ are more dangerous than downward excursions because of their greater impact on exp(μ), and the Bayes estimator, being a minimum risk estimator, augments the naive estimator exp(μ_x+.5σ^2) with the factor exp(+.5Var[μ|x]) to protect against the more dangerous upward excursions.

In closing this section of this note, what relation do the above results bear to those of other authors? I don't see an explicit parameter variance adjustment in Zehnwirth [15]. However, I know from the manual for his ICRFS loss reserving system and from private conversations with him that Zehnwirth is solidly in the Bayesian camp and advocates, or at least at one time advocated, the positive sign on the parameter variance adjustment. Verrall [12] actually appears to advocate both signs, depending on whether he is describing an unbiased Frequentist estimator or a Bayesian estimator, but he doesn't draw attention to the change in sign.

First, Verrall's equation (4.16) provides an unbiased Frequentist estimator. (Although he doesn't refer to this estimator as Frequentist, he notes its unbiasedness, which is a Frequentist notion. Furthermore, he invokes Bradu and Mundlak [4], which is a Frequentist paper.) To establish the connection between his notation and ours, note that $\underline{Z}$ is the vector of values of explanatory variables associated with our quantity of interest, y. $\underline{\beta}$ is the vector of regression coefficients associated with these explanatory variables, or rather the true but unknown values of these coefficients. $\underline{\hat{\beta}}$ is the vector of estimates of these regression coefficients derived from the regression. $\underline{Z}\underline{\hat{\beta}}$

and $Z\hat{\beta}$ are therefore inner products representing, respectively, our μ and our $\hat{\mu}$. From Verrall's equation (4.16), the unbiased estimator for $E[y|\mu]=\exp(\mu+.5\sigma^2)$ is

$$\exp(\hat{\mu})\, g_m\left[-\frac{1}{2}Z(X'X)^{-1}Z's^2 + \frac{1}{2}s^2\right] \tag{6}$$

where X is the regression structure matrix and s^2 is an unbiased estimator for σ^2. $g_m(t)$ is defined via power series expansion in Verrall's equation (4.5). It is clear from this definition that, as m becomes large, $g_m(t)$ tends to $\exp(t)$. m becomes large when the data base on which the regression is performed becomes large, without a corresponding increase in the number of explanatory variables. In this limit, the unbiased estimator for $E[y|\mu]$ of expression (6) above becomes

$$\exp\left(\hat{\mu} - \frac{1}{2}Z(X'X)^{-1}Z's^2 + \frac{1}{2}s^2\right) \tag{7}$$

From standard regression theory, the second term in the exponent is precisely -1/2 times the variance of the estimator $\hat{\mu}$. This estimator (7) therefore reproduces equation (4) above.

Lastly, Verrall provides, the middle of page 409, Bayesian estimators for posterior expected losses for lognormally distributed losses with parameters θ (our μ) and σ, where the posterior distribution of θ is normal with expectation m (our μ_x) and variance τ^2 (our $\mathrm{Var}[\mu|x]$). Verrall's estimator is

$$\exp\left(m + \frac{1}{2}\tau^2 + \frac{1}{2}\sigma^2\right) \tag{8}$$

which reproduces equation (5) above.

Concluding Remarks

So, which is the correct estimator in a lognormal linear model setting, equation (4) or equation (5)? Do we add or subtract a parameter variance adjustment? Each is correct, within its own paradigm, Bayesian or Frequentist. Unfortunately, for the lognormal linear model, unlike for the normal linear model, the result depends on the paradigm. It is up to the actuary to select the paradigm and, hence, the sign of the parameter variance adjustment. Unfortunately, there is no clear guidance as to which is appropriate for the loss reserving problem. Neither paradigm is without problems regarding its theoretical foundations, as Efron [6] is quick to point out.

A number of observations may be appropriate in closing, first some statistical ones, then some actuarial ones.

1) While calculating the Bayesian posterior mean, $E[y|x]$, it may be worthwhile to reflect on the fact that many Bayesians consider the greatest strength of the Bayesian paradigm to be its ability to produce readily interpretable posterior distributions and confidence intervals. (See, in particular, Gelman and others [8].) These Bayesians would consider someone who went to the trouble of constructing a Bayesian analysis only to extract posterior means and nothing else to have discarded most of the information revealed by their analysis. Yet, because of the very narrow focus of this note, I have ignored posterior variances, $Var[y|x]$, posterior predictive distributions for y, and posterior intervals resulting from those distributions.

2) A hard core Bayesian who wished to remain a Bayesian and yet was troubled by the above Bayesian/Frequentist discrepancy might be able to construct a valid Bayesian decision analysis that would reproduce the Frequentist unbiased result by considering loss functions other than quadratic, resulting in minimum Bayes Risk estimators other than the Bayesian posterior expectation, $E[y|x]$. I have not investigated what loss function might bring Bayesian and Frequentist analyses into agreement, but I might guess that such a loss function would appear quite ad hoc.

3) Both the unbiasedness of the Frequentist estimator and the minimum risk of the Bayesian estimator are predicated on the selected lognormal linear model being a reasonable approximation to reality. While we debate unbiasedness vs. minimum risk (tastes great vs. less filling), let us not forget that, if our model does not adequately approximate reality (incremental losses are not lognormally distributed, or expected logs are not linear in explanatory variables, or we have failed to include in the model important explanatory variables, etc), then, relative to a more adequate model, our Frequentist estimator is quite likely to be biased, and our Bayesian estimator is unlikely to be minimum risk.

Now, a few actuarial comments.

1) The Bayesian increases the indicated loss reserve for risk; the Frequentist reduces the indicated reserve to correct for presumed bias. The Bayesian indicated reserve is more conservative than the Frequentist. The Bayesian increase is, in effect, a kind of risk load. For those model parametrizations I have seen, the

greatest effect of parameter variance, percentagewise, tends to be out in the tail, at high development ages, because age tends to be selected as one of the explanatory variables and tends to be highly leveraged at high ages. Intuitively, out in the tail, at high development ages, is where an actuary would want the greatest risk load and conservativeness, because this is where the greatest uncertainty, percentagewise, lies.

2) The Frequentist loss reserver might believe the Bayesian indicated reserve to be redundant on average, because it fails to adjust for bias. Have you, or anyone you know, ever seen a truly redundant loss reserve (or Nessie, or Bigfoot)?

3) In the presence of controversy, with no clear indication as to how to resolve that controversy, perhaps we should employ the time-honored practice of practical actuaries everywhere: compromise. Ignore the parameter variance adjustment altogether. This produces indications intermediate between the bias adjusted Frequentist indication at the low end and the risk adjusted Bayesian indication at the high end.

My first preference would be for the Bayesian estimator because of its conservativeness, and because it is most conservative in the tail, where conservativeness is most appropriate. Upon failing to get my first preference, my second preference would be to ignore the parameter variance adjustment altogether. Why make any adjustment when we can't even agree on the sign of the adjustment? I would be very loathe to quote the Frequentist indication, to reflect the downward adjustment for bias, which is probably being

mis-estimated anyway because our selected lognormal linear model, on which the indicated bias is based, is likely to be an oversimplification of reality.

REFERENCES

[1] Aitchison, J. and J. A. C. Brown (1957), The Lognormal Distribution, Cambridge University Press, Cambridge.

[2] Berger, J. O. (1985), Statistical Decision Theory and Bayesian Analysis, Second Edition, Springer-Verlag, New York.

[3] Bickerstaff, D. R. (1972), "Automobile Collision Deductibles and Repair Cost Groups: The Lognormal Model," PCAS, Vol. 59, pps. 68-102 and discussions.

[4] Bradu, D. and Y. Mundlak (1970), "Estimation in Lognormal Linear Models," Journal of the American Statistical Association, Vol. 65, pps. 198-211.

[5] Crow, E. L. and K. Shimizu, editors (1988), Lognormal Distributions: Theory and Applications, Marcel Dekker, New York.

[6] Efron, B. (1978), "Controversies in the Foundations of Statistics," American Mathematical Monthly, Vol. 85, No. 4, pps. 231-246. (Reprinted in CAS Forum, Fall 1991 Edition.)

[7] Finger, R. J. (1976), "Estimating Pure Premium by Layer-An Approach," PCAS, Vol. 63, pps. 34-52 and discussions.

[8] Gelman, A., J. Carlin, H. Stern, and D. Rubin (1995), Bayesian Data Analysis, Chapman and Hall, New York.

[9] Silvey, S. D. (1975), Statistical Inference, Chapman and Hall, London.

[10] Venter, G. G. (1994), "Introduction," CAS Forum, Spring 1994, pps. 91-100.

[11] Verrall, R. J. (1991), "On the Estimation of Reserves from Loglinear Models," Insurance: Mathematics and Economics, Vol. 10, pps. 75-80.

[12] Verrall, R. J. (1994), "Statistical Methods for the Chain Ladder Technique," CAS Forum, Spring 1994, pps. 393-446.

[13] West, M. and J. Harrison (1989), Bayesian Forecasting and Dynamic Models, Springer-Verlag, New York.

[14] Wright, T. S. (1990), "A Stochastic Method for Claims Reserving in General Insurance," Journal of the Institute of Actuaries, Vol. 117, pps. 677-731.

[15] Zehnwirth, B. (1994), "Probabilistic Development Factor Models with Applications to Loss Reserve Variability, Prediction Intervals and Risk Based Capital," CAS Forum, Spring 1994, pps. 447-605.

An Introduction to Credibility
by Curtis Gary Dean, FCAS

This paper is derived from the presentation on basic credibility concepts that the author has given at the 1995 and 1996 CAS Seminars on Ratemaking.

AN INTRODUCTION TO CREDIBILITY

Credibility theory provides important tools to help the actuary deal with the randomness inherent in the data that he or she analyzes. Actuaries use past data to predict what can be expected in the future, but the data usually arises from a random process. In insurance, the loss process that generates claims is random. Both the number of claims and the size of individual claims can be expected to vary from one time period to another. If $1,500,000 in losses were paid by an insurer during the past year, one might estimate that $1,500,000 would likely be paid in the current year for the same group of policies. However, the expected accuracy of the estimate is a function of the variability in losses. Using credibility theory, the actuary estimates the randomness inherent in the data and then calculates a numeric weight to assign to the data.

Here is a dictionary definition of credible:

credible: Offering reasonable grounds for being believed

The actuary wants to know how much to believe the data that's being analyzed. To use the data to predict the future, this "belief in the data" must be quantified so that calculations can be made. This leads us to actuarial credibility:

actuarial credibility: the weight to be given to data relative to the weight to be given to other data

If we cannot fully believe our data, we may call on other information or data to supplement the data at hand. The data at hand and the supplemental data are each given an appropriate numeric weight in calculating an estimate.

The variability in insurance loss data can be seen in Table 1 which shows the loss experience for a group of policies covering contractor's pickup trucks. The last column shows that the average loss per truck varies widely from one year to the next. Any one year is a poor predictor of subsequent years.

TABLE 1

Contractor's Pickup Trucks

	(1) # of Insured Trucks	(2) Incurred Losses	Pure Premium (2)/(1)
1990	2,900	$2,030,000	$700
1991	3,000	1,470,000	490
1992	3,050	1,830,000	600
1993	3,050	1,250,500	410
1994	3,200	864,000	270

The variability in the average loss per pickup truck is depicted graphically in Figure 1. The expected average loss (pure premium) is $500 which we would observe if our body of data were infinite in size. But, for limited sample sizes, the observed average losses are randomly distributed. Note that as our sample size increases, the variability of the observed average loss decreases - the probability density curve becomes more concentrated around the $500 value. For a smaller sample size, the probability density curve flattens out. If our sample body of data consists of 50,000

trucks we can rely upon the observed average loss to estimate the true expected average loss to a much greater extent than if the data came from a smaller sample of only 3000 trucks.

FIGURE 1

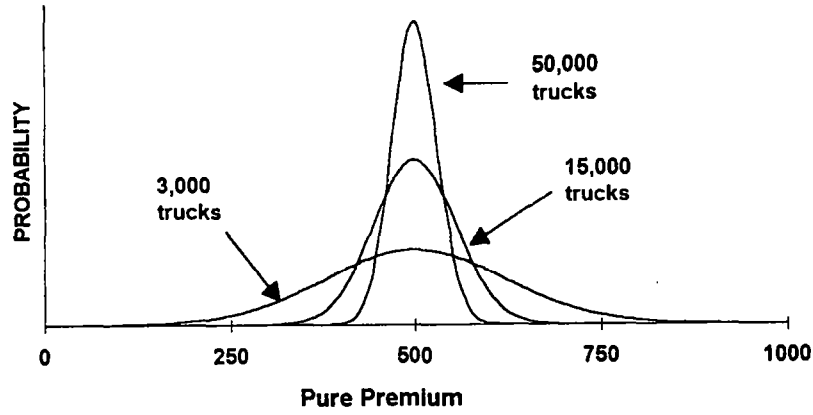

The actual distribution of pure premiums is not symmetric as shown in the prior graph, but is instead skewed to the right as shown in Figure 2. More of the observations would actually fall below the mean of $500 and the mode of the distribution is less than $500. The smaller the body of data, the greater the asymmetry in the graph. In an extreme case we could consider only one truck. In most years the truck would have no losses for an observed average loss of $0 in those loss-free years. But, every few years there would be a loss or, perhaps, several losses and the observed average loss would be substantial.

FIGURE 2

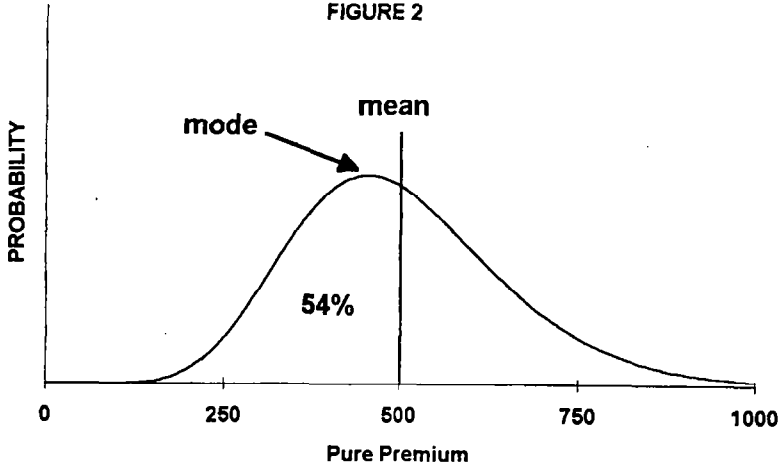

This leads us to a common problem that may occur when a group of non-actuaries is reviewing average losses or loss ratios for a series of years. The data may show, for example, four years with excellent loss ratios but a fifth year with a very high loss ratio. The five-year average may be close to some target loss ratio. Unfortunately, what frequently happens is that one of the reviewers will say that the one bad year is an anomaly that was caused by several severe claims and that the bad year should be thrown out of the data. This is a big mistake! For a small body of data, this pattern in the loss ratios is exactly what we expect to see. The majority of the loss ratios will look better than average, with a few being quite large. This doesn't mean that we should ignore the few high values; it usually means that our body of data is small.

The basic formula for calculating credibility weighted estimates is:

$$\text{Estimate} = Z \times [\text{Observation}] + (1-Z) \times [\text{Other Information}],$$

$$\text{and } 0 \leq Z \leq 1.$$

If our body of data is so large that we can give full weight to it in making our estimate, then we would set Z=1. If the data is not fully credible, then Z would be a number somewhere between 0 and 1. What is the "Other Information" that we might use in our formula? That depends on what we are trying to estimate. In Table 2, the left hand column shows our observed data and the right hand column may be the "Other Information" that we might use in the above formula.

TABLE 2

Observation		Other Information
Pure premium for a class	←→	Pure Premium for all classes
Loss ratio for an individual risk	←→	Loss ratio for entire class
Indicated rate change for a territory	←→	Indicated rate change for entire state
Indicated rate change for entire state	←→	Trend in loss ratio

Suppose you are trying to estimate the indicated rate change for a territory within a state, but your company has a limited volume of business in the territory. An option may be to weight the indicated change from territorial data alone with the indicated change for the entire state. This way you have reflected territorial experience in your rate change to the extent that it is credible.

The loss ratios shown below in Table 3 were produced in a computer simulation that modeled the insurance random loss process. The expected loss ratio is 60 for both the small and big states, but the observed (simulated) loss ratios will randomly vary around this value. As we would expect, the variation is much larger for the small state. In the larger state the loss ratio hovers around 60 in each year. Five-year average loss

ratios were calculated and then state indicated rate changes were calculated using the expected loss ratio of 60 as the permissible loss ratio. For example, in the small state -28.3% = (43/60 - 1.000). Using one of the formulas that we will discuss in a moment, credibility values Z were calculated for each state.

TABLE 3

	Small State		Large State	
	Earned ($000)	Loss Ratio	Earned ($000)	Loss Ratio
1990	69	17	7,100	58
1991	71	109	7,120	58
1992	72	62	7,180	60
1993	74	7	7,200	58
1994	74	19	7,400	61
Total	360	43	36,000	59
Permissible Loss Ratio		60		60
State Indication		-28.3%		-1.7
Credibility		10%		100%

Perhaps this data comes from a line of insurance that has an aggressive insurance to value program such that the inflationary trend in losses is exactly offset by the annual increases in the amount of insurance. In this case the trend in our loss ratio would be 0%. (For our data, we know that the trend in the loss ratio is 0% because each year has an expected loss ratio of 60.) We will apply our complement of credibility factor (1-Z) to this information. So, we would get the following two indications:

small state: .10 X [-28.3%] + (1 - .10) X [0.0%] = -2.8%

large state: 1.00 X [-1.7%] + (1 - 1.00) X [0.0%] = -1.7%

In both cases we know the right answer! We should take a 0.0% rate change in each state because our expected loss ratios are what we used for the permissible loss ratios. But, because of the randomness inherent in our data, our indications are slightly off the mark.

The important thing in the prior example is that we greatly improved the accuracy of our rate indication in the small state by incorporating credibility. We gave only a 10% weight to the raw indication arising from the small state's loss ratio. This had the result of dampening the effect of the randomness. To the extent possible we would like to use our observed data to calculate our estimate rather than rely on supplementary data, but given the randomness present in our observations, we need to temper the data. Using credibility theory we weight an estimate based on limited data with data from other sources. We want to find a weight Z that allows us to rely on our limited data to the extent reasonable, but which also recognizes that our limited data is variable.

There are two widely used formulas for the credibility Z as shown side by side in Table 4. For the classical credibility formula, if n > N then Z is set equal to 1.00. In the case of Bühlmann credibility, Z asymptotically approaches 1.00 as n goes to infinity.

TABLE 4

Classical Credibility	Bühlmann credibility
$Z = \sqrt{\frac{n}{N}}$	$Z = \frac{n}{n+K}$
Also called:	Also Called:
(1) Limited Fluctuation Credibility	(1) Least Squares Credibility (2) Empirical Bayesian Credibility (3) Bayesian Credibility

In both formulas n is a measure of the size of the body of data and is an indicator of the variability of the loss ratio or pure premium calculated from the data. n can be any of the following:

- number of claims
- amount of incurred losses
- number of policies
- earned premium
- number of insured unit-years.

These are not the only possibilities for n, but n needs to be some measure that grows directly with the size of the body of data that we have collected.

In practice both of the formulas can give about the same answer if N and K are chosen appropriately as displayed in Figure 3. Note that in the classical credibility case, when n is greater than or equal to 10,000, Z is identically 1.00.

FIGURE 3

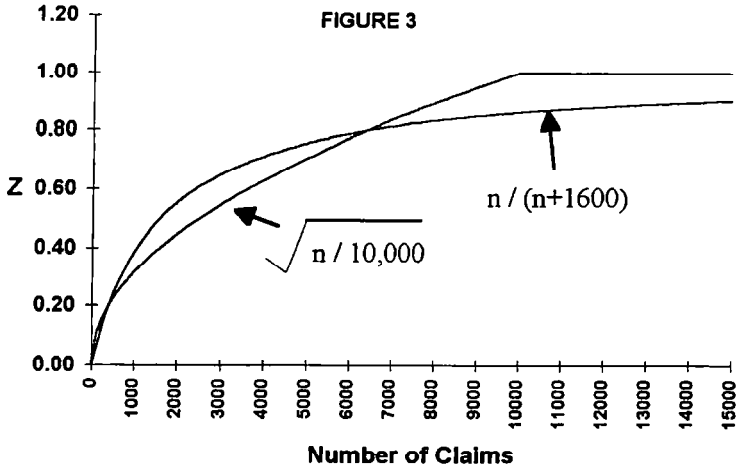

Classical Credibility

First we will discuss the classical credibility formula. Classical credibility attempts to restrict the fluctuation in the estimate to a certain range. N is calculated such that for fully credible data with n=N and Z=1.00, the observed pure premium or loss ratio will fall within a band about the expected value a specified percentage of the time. This is illustrated in Figure 4.

FIGURE 4

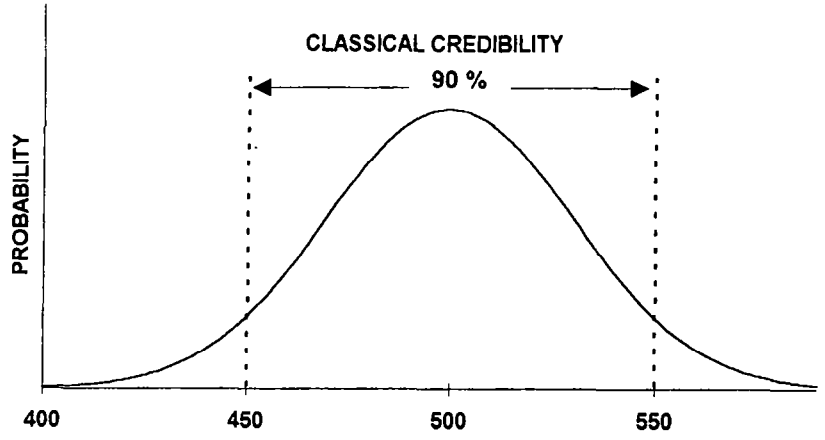

If N=5,200 claims, then the observed Pure Premium is within 10% of the "true" value 90% of the time.

In this example the measure of the size of the body of data is the expected number of claims. When our body of data is large enough so that we expect 5,200 claims in our observation period, the observed pure premium will fall within k=10% of the true value P=90% of the time; that is, 90% of the time our pure premium calculated from our body of data will fall into the interval [450,550]. Both the 90% probability and the 10% width of the range must be selected by the ratemaker. If you wanted much less variance in your estimate you might select a P=99% probability and a k=2.5% error in your estimate. Of course, it would require a much larger body of data in the observation period to achieve this level of certainty.

The full credibility standard N is a function of the selected P and k values. A larger P value results in a larger N and a smaller k also produces a larger N. In order to calculate the N that corresponds to the selected P and k, one needs to make certain assumptions and also know something about the loss process. In classical credibility one assumes that the frequency of claims can be modeled by a Poisson distribution. Also, one needs an estimate of the average claim size and the variance in claim sizes. Using these an estimate of the variance in total losses can be computed. The next assumption is that the distribution of the total losses is normal, i.e. bell-shaped. Then, the N value can be calculated. This is all covered in much detail in the syllabus material for the actuarial exam that tests credibility theory.

One does not have to use the number of claims in the classical credibility formula, but instead can use earned premium, number of policies, or some other basis. We could convert our formula developed above to an earned premium basis. Suppose that in reviewing our data we calculate that on average there is approximately $2,500 in earned premium for each claim; that is, the ratio of earned premium to the number of claims is $2,500. A full credibility standard of (2,500 dollars/claim) x (5,200 claims) = $13,000,000 could be used in place of the 5,200 claims. Then, the credibility assigned to any data could be calculated from the earned premium of the data.

To calculate the full credibility standard, the denominator in the formula, the amount of variability acceptable in fully credible data must be defined by the selection of P and k values. For less than fully credible data the square-root formula determines the credibility Z. Figure 5 displays graphically the calculation of partial credibility.

FIGURE 5

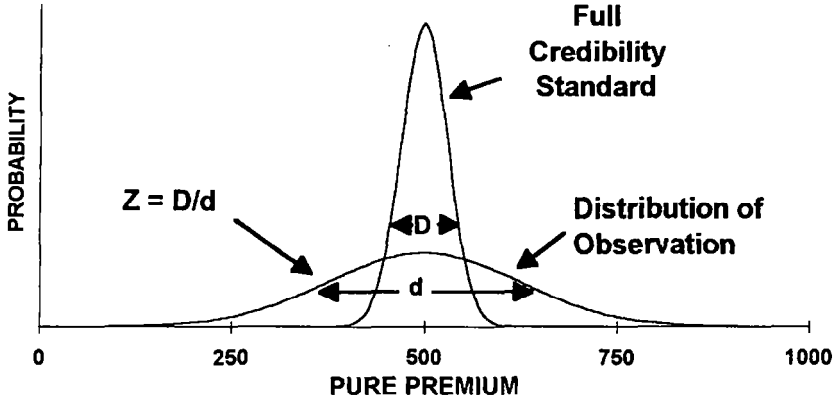

In the graph the width of the curve representing the variability of data which just meets the standard for full credibility is represented by D. D can be considered the standard deviation of the curve. (If you prefer, D can be two standard deviations.) Likewise, d is the width corresponding to a smaller body of data that is less credible. It turns out that the credibility that should be assigned to the smaller body of data in this model is Z = D/d, the ratio of the standard deviation of the pure premium of the fully credible data to the standard deviation of the pure premium of the partially credible data. We will allow a standard deviation of size D, but if our body of data has a standard deviation of d, then we apply a weight of D/d to the data. If the pure premium (p.p.) calculated from the data is expected to have a standard deviation of d, then the quantity Z x (p.p) has a standard deviation of D, which is our target.

Bühlmann Credibility

The least-squares credibility model uses the credibility formula:

$$Z = n/(n + K)$$

K is defined by the following intimidating expression:

$$K = \frac{\textit{Expected Value of the Process Variance}}{\textit{Variance of the Hypothetical Means}}$$

A good way to think about least-squares credibility is in the context of experience rating where the rate charged to an insured is a manual rate modified to reflect the experience of the individual insured. The losses incurred by an insured are random, so an insured's loss ratio will fluctuate. The term "process variance" is the variance in the loss ratio of the risk. The "expected value of the process variance" is the average value of the variance across the risks within the population. Since each risk is unique, the expected loss ratios of the individual risks at the manual rates will vary across the population because the manual rates are based on averages calculated for groups of risks who are classified alike in the rating plan. Each risk has it's own "hypothetical mean" loss ratio. The "variance of the hypothetical means" is the variance across the population of risks of their individual hypothetical mean loss ratios.

In Figure 6 there are two risks, risk #1 and risk #2, each with its own loss ratio distribution curve. The process variance is a function of the width of the curve indicated by the [1] in the figure. As mentioned above the width of the curve can be thought of as some multiple of the standard deviation. The process variance is the square of the standard deviation. So the wider the curve, the larger the process variance. [2] marks the difference in the hypothetical means between the risks. The variance in the hypothetical means across the population is a function of the differences in the hypothetical means between the risks.

When the process variance of the risks is large in relation to the difference in the means of the risks, K is large. A large K means that the credibility $Z = n/(n+K)$ is small. Looking at the second graph in Figure 6, we see that there is a broad band where the two risks' loss ratios overlap. Since the loss ratio of each risk is so variable, it makes sense to give more weight to the manual rate calculated from the average experience of a large group of similar risks and less weight to the experience of the individual risk.

Small process variances in relation to the differences in the means of the risks results in a small K value and a larger credibility Z. This scenario is represented by the bottom graph in Figure 6. The distributions of the two risks do not overlap. The larger credibility Z means more weight is assigned to the experience of the individual risk and less, (1-Z), to the experience of the population.

Several Examples

Examples of credibility formulas developed by the Insurance Services Office are displayed in Table 5. The first set of formulas are used in Homeowners ratemaking and are based on the classical credibility model. The measure of the size of the body of data and its consequent variability is in the units of house-years; that is, one house insured for one year contributes one unit. In making a statewide change 240,000 house-years are required for full credibility, and with that large of a body of data, the observed experience should be within 5% of the actual value 90% of the time. In computing territorial changes within the state, 60,000 house-years are assigned full

FIGURE 6

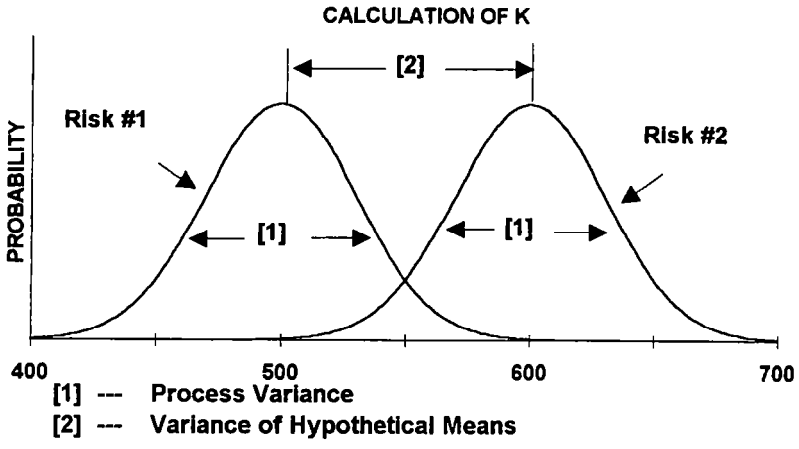
CALCULATION OF K
[2]
Risk #1
Risk #2
PROBABILITY
[1]
[1]
400
500
600
700
[1] --- Process Variance
[2] --- Variance of Hypothetical Means

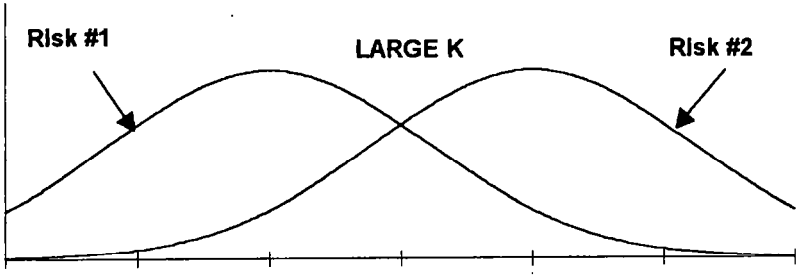
Risk #1
LARGE K
Risk #2

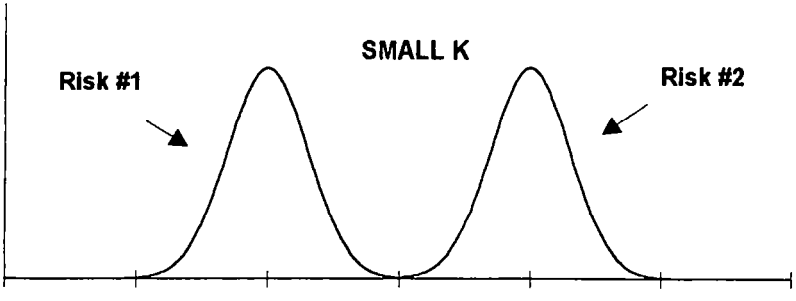
SMALL K
Risk #1
Risk #2

credibility and the observed territorial experience is expected to be within 10% of the expected value of 90% of the time. As stated previously, the actuary needs to decide on the units for n, the size of the P value, and the size of the k value.

TABLE 5

Credibility Formulas
Insurance Services Office

Homeowners: Owners Forms	
Statewide Changes	*Territorial Changes*
$Z = \sqrt{\frac{n}{240,000}}$	$Z = \sqrt{\frac{n}{60,000}}$
n = house years 90% confident within 5% of actual value	n = house years 90% confident within 10% of actual values
Manufacturers & Contractors	
Statewide Changes	*Relativities*
$Z = \sqrt{\frac{n}{8,000}}$	$Z = \sqrt{\frac{n}{25,000}}$
n = number of occurrences in three years 90% confident within 7% of actual value	n = number of occurrences in five years 95% confident within 5% of actual values
General Liability Experience Rating	
$Z = \frac{L}{L + \$177,000}$	
L = expected loss costs (including ALAE) at $100,000 basic limits	

The next set of formulas in Table 5 are used by ISO in Manufacturers & Contractors ratemaking. Statewide changes require 8,000 claims (occurrences) in a three-year period, and with this many expected claims, the experience of the body of data should be within 7% of the expected value 90% of the time. The full credibility standard for relativities within M&C, such as class relativities, is much tougher with 25,000 claims required for a P=95% and k=5%.

The selection of P and k is probably more art than science. If the body of data that the actuary is working with is of limited size and there is no good surrogate for the data to which to assign the complement of credibility, then the actuary may select a smaller P and larger k to produce a smaller requirement for full credibility. If the actuary wants to make the rates more responsive to current experience he or she may also select a smaller P and a larger k. If rate stability is the most important goal then larger P and smaller k may be selected.

The last formula in Table 5 is the credibility to be assigned to an individual insured's data in General Liability experience rating and it is based on the Bühlmann model. In a loss cost environment, L represents the expected loss costs (expected incurred losses and allocated loss adjustment expenses) for the individual risk. Before the advent of loss costs, premium designated by E was used instead of L. The expected loss costs included in L are $100,000 basic limits losses. ISO has recently converted from $25,000 basic limits to $100,000. At $100,000 basic limits it was necessary to increase the K value in the denominator to $177,000 from its previously smaller value that applied when $25,000 basic limits losses were used in computing the experience rating adjustment. If unlimited losses were used in the experience rating formula, then an even larger K value would be necessary because the expected value of the process variance would become even larger.

Reducing Variability of the Data

The data used by ratemakers in the insurance business arises from a random process; in fact, it is this randomness that makes insurance necessary. The ratemaker is confronted with the task of finding the proper premiums to charge insureds without knowing for sure what the cost will be to the company to provide the insurance. The ratemaker estimates the cost of future payments in insurance claims by his or her company by analyzing past costs. The ratemaker wants to use the most relevant data to estimate future costs, but he or she must also deal with the variability inherent in the data.

One way to decrease the variability in ratemaking data is to use a larger body of data. Here are several ways to do this:

- include more years in the experience period
- use Bureau data
- combine data into fewer, but larger groups

Each of these involves a tradeoff. If more years are included in the experience period then it becomes necessary to apply larger trend factors to the older data and trend can be tough to estimate. Also, the book of business to which new rates will apply may be different from the business that produced the experience years ago. The same goes for Bureau data. The insureds included in Bureau data may be very different from the average insured in the ratemaker's data. Combining the data into fewer, but larger groups, may limit a company's ability to effectively compete against competitors who can better identify the proper price to charge an insured.

Another approach to decreasing the variability in losses used in ratemaking is to:

- cap large losses
- remove catastrophes

Of course, if we do either of the above we must put something back to make up for the losses we removed. One method to cap large losses is to do basic limits ratemaking by state, territory, class, etc., and calculate basic limits rates. Then, rates for higher limits are computed using increased limits factors calculated based on the aggregate data for many states and classes. Another approach is to limit all losses at some set amount, for example $150,000, and then to prorate the excess losses amount back by state, territory, class, etc. Catastrophe losses can be removed from the data and a catastrophe load substituted in its place. This load can be computed from a very long observation period, thirty years or more for weather losses, or a computer model that attempts to model the catastrophe loss process.