A Top-Down Approach to Understanding Uncertainty in Loss Ratio Estimation

by Alice Underwood and Jian-An Zhu

ABSTRACT

In this paper we define a specific measure of error in the estimation of loss ratios; specifically, we focus on the discrepancy between the original estimate of the loss ratio and the ultimate value of the loss ratio. We also investigate what publicly available data can tell us about this measure. Using Other Liability Occurrence data as reported in Schedule P, we find that in a given accident year the values of this "estimation error ratio" for different companies are lognormally distributed. Furthermore, we find that the average accident year estimation error ratio is amenable to time series analysis. Using the time series analysis and the lognormal accident year model, we can estimate the distribution of possible estimation error ratios for the industry in a future year.

KEYWORDS

Schedule P, loss ratio, time series, ARIMA, Other Liability, estimation error, parameter risk

1. Introduction

Many factors contribute to uncertainty in actuarial loss ratio estimates, including

- data issues (flawed data; finite sample size)
- projection issues (selection of trend, development, and on-level factors)
- judgmental adjustments for factors which cannot be directly quantified
- unforeseen external influences (law changes; coverage changes)

Most research on the topic of estimation uncertainty has focused on quantifying the effects of finite sample size [see, e.g., Kreps (1997); Van Kampen (2003); Wacek (2005)]. But this is far from the only hurdle that practicing actuaries must face—and often not the most important one.

A review of insurance industry experience over the past several years indicates that while for the most part companies do a fairly good job of estimating loss ratios, there are times when the industry as a whole gets it very wrong. This supports the argument that finite sample size is *not* the key driver of estimation error. Furthermore, for long-tail casualty lines of business, the error in initial loss ratio estimate is strongly correlated to the insurance market cycle and persists over a number of years. This is not surprising, since in these lines it takes a number of years for the effects of changing influences to become fully evident.

In this paper we will not attempt to identify, much less quantify, all the various "known unknowns" and "unknown unknowns" that can contribute to uncertainty in the estimation of loss ratios. Instead, we will take a top-down approach, looking at the loss ratio data itself.

2. Definitions and notation

Let's start by establishing some notation for the different quantities to be examined; using this notation, we will define a measure of estimation error that can be studied using available data. Given a set C of companies and a set T of accident years, denote

OLR(c,t) = original loss ratio for company cand accident year t at age 12 months,

ULR(c,t) = ultimate loss ratio for company cand accident year t.

One measure of the discrepancy between initial loss ratio estimates and the ultimate loss ratio is the quotient

$$\mathbf{R}(c,t) = \mathbf{ULR}(c,t) / \mathbf{OLR}(c,t).$$

Others [see Wacek (2007)] have sought to understand the trajectory from OLR to ULR; the main goal of this paper is to quantify and explore the behavior of R(c,t).

2.1. Available data

Values for OLR can readily be found as the Schedule P estimate booked 12 months after the start of the accident year. While the final value of ULR may take many years to be precisely determined, we may look to the most recent Schedule P estimate of the ultimate loss ratio. Schedule P provides only 10 development years, but ISO data indicate that at this age unlimited General Liability losses are more than 93% reported (i.e., the 120-to-Ultimate development factor is less than 1.07); it seems reasonable to assume that companies' estimates of ultimate have largely stabilized by 120 months. We consider this to be a reasonable proxy for the true ultimate loss ratio.

Of course, the latest available Schedule P loss ratio for more recent accident years reflects fewer than 10 years of development. We began this study using data available as at 12/31/2005. On this basis, our Schedule P proxy for the Accident Year 2005 ULR is the same as our proxy for the OLR—in other words, no information useful for our study can be obtained for Accident Year 2005. Similarly, the data for Accident Year 2004 would be of very little value. For purposes of this study, we selected Accident Year 2003 as the most recent accident year to include; even so, as our estimate of ULR is only 36 months old, one might question whether this cutoff is sufficient. ISO data indicate that approximately 90% of General Liability claims and 63% of the associated loss dollars are reported at 36 months, which should provide some measure of stability to companies' ULR estimates—but certainly less than that achieved at 120 months. Because we wished to use the results of this study to project behavior for future accident years, it was desirable to include the most recent data to the extent it would be meaningful; and therefore we decided to include 2003.

The loss ratio data used in this study is "as was" reserving-type data, not "as-if" trended and on-leveled prospective data. This is crucial, because we specifically want to understand how the estimation error behaves over a span of successive accident years, given changing market conditions. This is not to say that the effects of trend and onlevel are being ignored: they are, in fact, incorporated in the loss ratio estimates which are the subject of our analysis, and their estimation contributes to the estimation error. As mentioned in the Introduction, most published research on estimation error focuses on the portion of estimation error which is due to finite sample size. Implicit in our current analysis of estimation error are *all* the factors that go into an actuary's original estimated loss ratio: uncertainty about the form of the distribution, the trend factors, the development factors, application of professional judgment, and other known and unknown sources of uncertainty, as well as uncertainty due to finite sample size.

2.2. Remarks

For a fixed company c_0 , consider the average of $R(c_0, t)$ over a large set *T* of years:

Average[
$$R(c_0, t) | t \in T$$
] := $R(c_0)$

where $R(c_0)$ is the long-term company-specific estimation error. If the company's loss ratio estimates are unbiased, $R(c_0)$ should be close to 1; but it could be the case that a given company has a particular estimation bias, either upward or downward [see Kahneman and Lovello (1993)].

For a fixed accident year t_0 , consider the average of $R(c,t_0)$ over a large set *C* of companies:

Average[$\mathbf{R}(c,t_0) \mid c \in C$] := $\mathbf{R}(t_0)$

where $R(t_0)$ is the accident-year specific average estimation error across all companies. In this case, we would not necessarily expect $R(t_0)$ to equal 1. We can think of $R(t_0)$ as the "industry delusion factor"—the ratio of the *actual* loss potential faced by the insurance industry in accident year t_0 to the industry's *initial view* of their loss potential. Ratios greater than 1 correspond to an initial under-estimation of loss and subsequent adverse development; ratios less than 1 correspond to an initial over-estimation of loss and subsequent favorable development.

The interested reader may consult Meyers (1999) for an investigation of the effects of company-specific and industry-wide uncertainty effects.

3. Data and analysis

We turned to publicly available data to see what we could learn about the behavior of the estimation error ratio R.

3.1. Description of data set and overview of analysis

For 63 of the largest casualty writers, we used Schedule P data to compile Other Liability Occurrence booked loss ratios at 12 months, and the most recent booked loss ratios as of 12/31/05, for accident years 1980–2003. No specific adjustments were made for mergers and acquisitions, though a few companies for which this was considered to be a potentially problematic

| Accident Year | Mean of ULR/OLR | Stdev of ULR/OLR | μ := Mean of In(ULR/OLR) | σ := Stdev of In(ULR/OLR) | Mean of Lognormal Distribution | Stdev of Lognormal Distribution |
|------------------|--------------------|---------------------|------------------------------|----------------------------------|--------------------------------------|---------------------------------------|
| 1980 | 1.067 | 0.244 | 0.039 | 0.237 | 1.069 | 0.256 |
| 1981 | 1.206 | 0.335 | 0.152 | 0.263 | 1.206 | 0.322 |
| 1982 | 1.447 | 0.624 | 0.299 | 0.361 | 1.440 | 0.538 |
| 1983 | 1.539 | 0.887 | 0.331 | 0.416 | 1.519 | 0.661 |
| 1984 | 1.714 | 0.881 | 0.427 | 0.469 | 1.712 | 0.849 |
| 1985 | 1.325 | 0.518 | 0.224 | 0.323 | 1.319 | 0.437 |
| 1986 | 0.900 | 0.327 | (0.156) | 0.314 | 0.898 | 0.289 |
| 1987 | 0.768 | 0.263 | (0.336) | 0.424 | 0.782 | 0.347 |
| 1988 | 0.870 | 0.273 | (0.187) | 0.324 | 0.874 | 0.291 |
| 1989 | 0.947 | 0.356 | (0.117) | 0.359 | 0.949 | 0.352 |
| 1990 | 0.935 | 0.300 | (0.109) | 0.286 | 0.934 | 0.273 |
| 1991 | 0.919 | 0.349 | (0.140) | 0.323 | 0.916 | 0.304 |
| 1992 | 0.919 | 0.313 | (0.133) | 0.309 | 0.918 | 0.291 |
| 1993 | 0.868 | 0.224 | (0.175) | 0.266 | 0.869 | 0.235 |
| 1994 | 0.841 | 0.249 | (0.215) | 0.292 | 0.842 | 0.251 |
| 1995 | 0.929 | 0.265 | (0.114) | 0.297 | 0.932 | 0.283 |
| 1996 | 0.962 | 0.235 | (0.073) | 0.278 | 0.966 | 0.274 |
| 1997 | 1.049 | 0.206 | 0.028 | 0.206 | 1.050 | 0.219 |
| 1998 | 1.152 | 0.283 | 0.110 | 0.261 | 1.155 | 0.307 |
| 1999 | 1.301 | 0.394 | 0.220 | 0.296 | 1.302 | 0.394 |
| 2000 | 1.281 | 0.416 | 0.201 | 0.303 | 1.280 | 0.397 |
| 2001 | 1.130 | 0.337 | 0.086 | 0.260 | 1.128 | 0.298 |
| 2002 | 1.042 | 0.223 | 0.018 | 0.222 | 1.044 | 0.235 |
| 2003 | 0.949 | 0.152 | (0.067) | 0.180 | 0.950 | 0.173 |

Table 1. Observed vs. fitted mean and standard deviation

issue were excluded from the sample. Not every company had data available for each year; the number of companies for a given year varied from 38 for accident year 1980 to 47 for accident years 1993–2002. It must also be noted that the definitions of data to be included in specific exhibits of the statutory blank do change over time.

For each company and accident year where OLR and ULR data were available we constructed the ratio R = ULR/OLR. As described in the remainder of Section 3, our analysis of the data indicates that

- For each accident year *t_i* the values R(*c*,*t_i*) are lognormally distributed across companies *c*
- The μ and σ parameters of these lognormal distributions are strongly linked
- The σ parameter can be approximated by a linear function of $|\mu|$
- The μ parameter can be analyzed using time series methods

3.2. Distribution of company estimation error ratios in a fixed accident year

Observation of the data points in each accident year suggested that the company estimation errors might be lognormally distributed. Therefore, for each accident year t_i , we fitted a normal distribution to the values $\ln(R(c,t_i))$ using the method of moments, which is also the maximum likelihood fit. This is shown in Table 1.

We then applied the Kolmogorov-Smirnov test (K-S) test to compare the observed values of $R(c,t_i)$ to those indicated by the lognormal approximation. For no accident year did the K-S test imply rejection of the lognormal model. A sample fit is shown in Figure 1.

Further examination of the accident year distributions showed a strong linkage between each accident year's mean and its standard deviation; one might also formulate this relationship as a linkage between the μ and σ parameters of the fitted lognormal distribution.





In Figure 2, each data point represents a single accident year. In the left-hand chart, year t_i is represented by the point (Average[$\mathbf{R}(c,t_i) \mid c \in C$], Stdev[$\mathbf{R}(c,t_i) \mid c \in C$]) and in the right-hand chart year t_i is represented by the point $(\mu(t_i), \sigma(t_i))$. We examined various ways to express this relationship. The linear-absolute model offered a better fit than the linear model, with the simplicity of

Figure 2. Mean vs. Stdev and $|\mu|$ vs. σ for fitted lognormals

fewer parameters and greater symmetry than the polynomial fit. Therefore, we selected this linearabsolute model.

In essence this model says that the more pronounced the industry average error in a given year (whether this error is favorable or adverse), the greater the spread of the distribution across companies. The business interpretation of this could be as follows: for years when the industry as a whole does a good job of estimation, most companies will fall fairly close to the average in the accuracy of their estimates. On the other hand, for years in which the industry as a whole does a poor job of estimation, the spread of company errors is larger.

While certainly some of the spread of individual company ratios is attributable to random variation (which we assume has been largely "smoothed out" when looking at the industry average), the strength of the relationship between the μ and σ parameters of the fitted lognormal distribution suggests that the degree of spread in

|mu| vs. sigma







Figure 3. Stability of fitted linear-absolute parameters relating σ and μ

company values for a given accident year may be driven largely by estimation error contagion—in other words, by the level of industry delusion.

It seems reasonable to ask whether the relationship between μ and σ has remained stable over time. We tested this by fitting the slope and intercept for the linear-absolute relationship to restricted data sets, starting with accident years 1980–1989 and then successively including the next accident year to consider 1980-1990, 1980-1991, and so on. This is shown in Figure 3. We found that the linear-absolute fit remained good, with least-squares slope and intercept parameters changing only slightly over time. This consistency supports the idea that the degree of spread in company delusion around the average industry delusion in a given year may be driven more by estimation error contagion than by purely random variation.

3.3. Time series analysis of the μ parameter

Previous sections have largely focused on the dispersion of company estimation error around the industry average estimation error for a single given accident year. Now we turn our attention to a different question: what can we say about how the value of the industry average estimation error





in one accident year relates to the value of the industry average estimation error in a different accident year? In particular, we investigate the behavior of the underlying parameter, μ , across time. As can be observed from Figure 4, the behavior of the μ values indicates autocorrelation.

Accident Year

We examined the time series $\mu(t)$ using ARIMA methodology and identified the AR(2) model as the best fit for the data set. Using the maximum likelihood method to fit the AR(2) parameters, we obtained the following model:

$$\mu(t) = 1.33 * \mu(t-1) - 0.66 * \mu(t-2) + e(t).$$

This can be interpreted as saying that μ tends to "overshoot" its immediate prior value, with an offsetting effect determined by the second prior value. This is depicted in Figure 5.

Figure 5. $\mu(t)$ vs. AR(2) model



Figure 6. Approximation of e(t) by normal distribution

Analysis indicates that the residuals e(t) are stationary and lack autocorrelation. Furthermore, the K-S test indicates that e(t) may be approximated with a normal distribution having mean 0 and standard deviation 0.09, as shown in Figure 6.

3.4. Applicability in practice

The AR(2) model enables the practitioner to develop a view on future years' behavior. However, there is one significant obstacle to applicability. Our data set consists of information through 12/31/2005; consider the perspective of the actuary who, during the course of 2006, wishes to forecast the likely level of "industry delusion" for Accident Year 2007. As discussed in Section 2.1, at this point in time, the data for Accident Year 2005 is useless for this purpose, and Accident Year 2004 remains extremely green. In other words, the values of $\mu(t-1)$ and $\mu(t-2)$ are not yet "ripe" enough for use in projecting $\mu(t)$. If our practitioner makes the same assumption that we used for purposes of this study, namely that 2003 is the most recent accident year for which the ULR/OLR value has predictive value, in order to forecast 2007 the AR(2) formula must be applied recursively.

Alternatively, since in effect the actuary is forced to estimate accident year *t* using data from accident years t-4 and t-5, we could seek to directly fit a model that estimates $\mu(t)$ as a linear

combination of $\mu(t-4)$ and $\mu(t-5)$. Using our 1980–2003 data set and applying the method of least squares yields

Alternate model:

 $\mu(t) = 0.07 * \mu(t-4) - 0.33 * \mu(t-5) + e_{alt}(t).$

For comparison, recursive application of the AR(2) model yields

Recursive AR(2) model:

$$\mu(t) = 0.06 * \mu(t-4) - 0.39 * \mu(t-5) + e_{\text{cumul}}t).$$

We can see that the coefficients are similar, as are the resulting projections, shown in Table 2.

Clearly, neither recursive application of the AR(2) model to project four steps ahead nor the alternative four-step-ahead model is as accurate as the AR(2) model applied to ripe data illustrated in Figure 5; and clearly the residuals as seen in Table 2 no longer lack autocorrelation. However, a four-step-ahead projection may still provide some insight as to the potential range of μ values in a future year. In the next section we will investigate such a projection. Given the similarity of the two models described above, we will continue with the recursive application of the AR(2) model; this offers greater theoretical simplicity and enables the development of a trajectory for μ that may provide additional insight.

3.5. Simulation results

In order to gain additional perspective on the "probability cone" for the future trajectory of μ , we randomly generated residual values e(t) for recursive application of the AR(2) formula. The results of 10,000 Monte Carlo simulations form the probability cone depicted in Figure 7.

In Figure 8 we graph the simulated probability density function and cumulative distribution function of the forecast industry mean error R(2007).

The median of the forecast values for R(2007) is approximately 1.0, while the mean is approximately 1.05. So, while as of 2003 the industry



Figure 7. Percentiles of projection of μ for 2004–2007

Table 2. Comparison of alternative model with recursive AR(2)

| | | Alt Model | | Recursive AR(2) | | |
|---------------|----------|-----------|--------------|-----------------|----------------|--|
| Accident Year | $\mu(t)$ | Forecast | $e_{alt}(t)$ | Forecast | $e_{cumul}(t)$ | |
| 1985 | 0.224 | (0.001) | 0.226 | (0.006) | 0.230 | |
| 1986 | (0.156) | (0.027) | (0.129) | (0.042) | (0.115) | |
| 1987 | (0.336) | (0.073) | (0.263) | (0.097) | (0.238) | |
| 1988 | (0.187) | (0.076) | (0.111) | (0.104) | (0.083) | |
| 1989 | (0.117) | (0.122) | 0.005 | (0.154) | 0.037 | |
| 1990 | (0.109) | (0.085) | (0.024) | (0.098) | (0.012) | |
| 1991 | (0.140) | 0.026 | (0.165) | 0.041 | (0.180) | |
| 1992 | (0.133) | 0.095 | (0.228) | 0.120 | (0.253) | |
| 1993 | (0.175) | 0.052 | (0.227) | 0.066 | (0.241) | |
| 1994 | (0.215) | 0.030 | (0.245) | 0.039 | (0.254) | |
| 1995 | (0.114) | 0.025 | (0.139) | 0.034 | (0.148) | |
| 1996 | (0.073) | 0.035 | (0.108) | 0.047 | (0.120) | |
| 1997 | 0.028 | 0.030 | (0.002) | 0.042 | (0.014) | |
| 1998 | 0.110 | 0.041 | 0.069 | 0.056 | 0.054 | |
| 1999 | 0.220 | 0.061 | 0.158 | 0.077 | 0.142 | |
| 2000 | 0.201 | 0.032 | 0.169 | 0.040 | 0.161 | |
| 2001 | 0.086 | 0.026 | 0.060 | 0.030 | 0.056 | |
| 2002 | 0.018 | (0.001) | 0.019 | (0.004) | 0.0222 | |
| 2003 | (0.067) | (0.019) | (0.048) | (0.030) | (0.038) | |

seems to have moved into the over-estimation part of the cycle, where R = ULR/OLR is less than 1, the model suggests than by 2007 the industry could well have shifted into under-estimation behavior. This can also be seen in terms of the trajectory of the projection illustrated in Figure 7; given the lognormal model, μ values greater than 0 correspond to ULR/OLR ratios greater than 1.

The 80th percentile of the simulated distribution for R(2007) is approximately 1.22, i.e., there is approximately a 20% chance that the average company's ultimate 2007 loss ratio will be 22% higher than the initial estimate. The 95th percentile of the simulated distribution is approximately 1.47, i.e., there is approximately a 5% chance that the industry average ultimate loss ratio will be approximately 47% higher than the initial estimate. The 99th percentile is approximately 1.74 and the 99.9th percentile is approximately 2.10.

3.6. Back-testing

To investigate the reasonableness of this "four years forward" projection using recursive appli-





Table 3. Actual μ as percentile of forecast distribution

| Accident Year | $\begin{array}{l} \mbox{Mean of Back-} \\ \mbox{Testing} \\ \mbox{Forecast Distr.} \\ \mbox{For } \mu \end{array}$ | Actual μ | Actual as %ile of Forecast Distr. | Quintile of Forecast Distr. |
|------------------|--|--------------|---|--------------------------------|
| 1989 | (0.153) | (0.117) | 58% | 3 |
| 1990 | (0.096) | (0.109) | 47% | 3 |
| 1991 | 0.039 | (0.140) | 18% | 1 |
| 1992 | 0.116 | (0.133) | 9% | 1 |
| 1993 | 0.065 | (0.175) | 11% | 1 |
| 1994 | 0.038 | (0.215) | 10% | 1 |
| 1995 | 0.033 | (0.114) | 22% | 2 |
| 1996 | 0.051 | (0.073) | 26% | 2 |
| 1997 | 0.042 | 0.028 | 47% | 3 |
| 1998 | 0.057 | 0.110 | 61% | 4 |
| 1999 | 0.077 | 0.220 | 77% | 4 |
| 2000 | 0.042 | 0.201 | 80% | 4 |
| 2001 | 0.030 | 0.086 | 62% | 4 |
| 2002 | (0.003) | 0.018 | 54% | 3 |
| 2003 | (0.029) | (0.067) | 42% | 3 |

cation of the AR(2) model, we back-tested it against our data set. Using the AR(2) parameters derived from the full 1980–2003 data set, we created simulated forecast distributions for μ in years for which we could actually calculate the value of μ from the data. We felt that the back-testing time frame of 1989–2003 could be viewed as fairly representative of a full market cycle (see Figure 4).

While we would not expect each year's observed μ value to fall right at the center of the



Table 4. Observed vs. theoretical quintiles

| Quintile | Observed Count | Observed F Incremental | Percentage Cumulative | Theoretical Cumul. % |
|----------|-------------------|---------------------------|--------------------------|-------------------------|
| 1 | 4 | 27% | 27% | 20% |
| 2 | 2 | 13% | 40% | 40% |
| 3 | 5 | 33% | 73% | 60% |
| 4 | 4 | 27% | 100% | 80% |
| 5 | 0 | 0% | 100% | 100% |
| Total | 15 | 100% | 100% | 100% |

forecast distribution, we can compare where the observed μ values fell within the percentiles of the forecast distributions and see if the observations were distributed evenly across the distribution percentiles, or clustered inappropriately, as shown in Table 3.

We checked the percentage of μ values falling into the quintile groupings of their respective forecast distributions and found reasonably good agreement with theoretical quintiles, shown in Table 4.

As noted above, we would not expect the observed μ values to consistently fall near the center of the forecast distributions. The method cannot provide a precise point estimate of μ in future years—it's not a crystal ball to tell us when a major shift in the industry delusion factor is on the horizon. Such shifts are generally driven by coverage changes, law changes, and so on: external variables that affect market behavior. However, as shown in the quintile analysis, we do find that the method is useful in projecting the likely *range* of possible μ values. In other words, the method does appear to provide a reasonable answer to the question "how likely are we as an industry to get it wrong-and if we do get it wrong, how wrong might we be?" As observed above, the method indicates that for the 2007 accident year, there is approximately a 20% chance that the industry average ultimate loss ratio will be at least 22% higher than the initial estimate. We believe this type of estimate is the best application of the analysis, and can be used in helping companies to understand the potential for and potential magnitude of estimation error. This in turn can help companies in stress testing, making reinsurance decisions, conducting dynamic financial analysis, and applying enterprise risk management techniques.

4. Caveats

Throughout the preceding sections we have noted various limitations to this methodology. To recapitulate:

- *Imperfect data*. As noted earlier, definitions of data to be included in specific exhibits of the statutory blank do change over time. And, of course, companies enter and leave the business over time and/or go through mergers and acquisitions.
- ULR approximation. To the extent that reported loss ratios after several years may still be subject to additional development (either adverse or favorable) before reaching the true ultimate loss ratio, the calculated estimation error ratio may tend to be too close to 1 (i.e., μ too close to 0) and the forecast distribution may show insufficient variability.
- *External "shock" influences*. There have been in the past (and there surely will be in the future) external factors which influence the loss

ratio estimation error. Examples include unforeseen losses such as asbestos and regulatory changes such as Sarbanes-Oxley.

• *Intrinsic "component" influences.* As described in the Introduction, this study was purposely created from a top-down perspective. We did not attempt to quantify or even identify all of the component influences that can contribute to uncertainty in the estimation of loss ratios. While the resulting model is not perfect, we believe that there is value in developing a model based solely on loss ratio behavior over time—one that does not explicitly rely on parameters such as loss trend, rate change, etc., which are often unreliable and/or unavailable.

The results of the back-testing shown in Table 3 suggest that overall the tail of our forecast may be somewhat conservative, in that we had no observation that fell into the highest quintile of the corresponding forecast distribution. However, with only 15 data points it is difficult to draw a firm conclusion in that regard.

5. Conclusions

The top-down approach to loss ratio estimation error we have taken in this paper-looking at the loss ratio data itself, rather than trying to quantify the individual factors that might contribute to estimation error-has two important advantages. By its nature, the method incorporates all sources of estimation error. Furthermore, unlike a bottom-up approach, this method does not require contemplation of the complicated and difficult-to-quantify relationships and codependencies among all these various contributors. Using the time series analysis and the lognormal accident year model described in this paper provides a way to help quantify the likelihood and magnitude of estimation error for current and future accident years at the industry level. This information can help companies in stress testing, making reinsurance decisions, conducting dynamic financial analysis, and applying enterprise risk management techniques.

References

- Kahneman, D., and D. Lovello, "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking," *Management Science* 39, 1993, pp. 17–31.
- Kreps, R. E., "Parameter Uncertainty in (Log)Normal Distributions," *Proceedings of the Casualty Actuarial Society* 84, 1997, pp. 553–580, http://www.casact.org/pubs/ proceed/proceed97/97553.pdf.
- Meyers, G. G., "Estimating Between Line Correlations Generated by Parameter Uncertainty," Casualty Actuarial Society *Forum*, Summer 1999, pp. 197–222, http://www. casact.org/pubs/forum/99sforum/99sf197.pdf.

- Van Kampen, C. E., "Estimating the Parameter Risk of a Loss Ratio Distribution," Casualty Actuarial Society *Forum*, Spring 2003, pp. 177–213, http://www.casact.org/ pubs/forum/03spforum/03spf177.pdf.
- Wacek, M. G., "Parameter Uncertainty in Loss Ratio Distributions and its Implications," Casualty Actuarial Society *Forum*, Fall 2005, pp. 165–202, http://www.casact.org/ pubs/forum/05fforum/05f165.pdf.
- Wacek, M. G., "The Path of the Ultimate Loss Ratio Estimate," *Variance* 1, 2007, pp. 173–192, http://www. variancejournal.org/issues/?fa=article&abstrID=6414.