

The Path of the Ultimate Loss Ratio Estimate

by Michael G. Wacek

ABSTRACT

This paper presents a framework for stochastically modeling the path of the ultimate loss ratio estimate through time from the inception of exposure to the payment of all claims. The framework is illustrated using Hayne's log-normal loss development model, but the approach can be used with other stochastic loss development models. The behavior of chain ladder and Bornhuetter-Ferguson estimates consistent with the assumptions of Hayne's model is examined. The general framework has application to the quantification of the uncertainty in loss ratio estimates used in reserving and pricing as well as to the evaluation of risk-based capital requirements for solvency and underwriting analysis.

KEYWORDS

Stochastic model, diffusion process, loss development, loss ratio estimation, lognormal, Student's t, log t, parameter uncertainty

1. Introduction

Ultimate loss ratio estimates change over time. The initial loss ratio estimate that emerges from the pricing analysis for a tranche of policies soon gives way to a new estimate as time passes and claims begin to emerge (or not). By the time all claims have been paid, the loss ratio is likely to have been re-estimated many times. The focus of this paper is on how to model the future revisions of these ultimate loss ratio estimates. We illustrate the approach using loss ratio estimates based on chain ladder and Bornhuetter-Ferguson methods underpinned by a simple stochastic model described by Hayne [4].

There appears to be little, if any, actuarial literature on the subject of behavior of an ultimate loss ratio estimate *between* the time when it is made and the time when its final value becomes known, i.e., the point at which all claims have been paid. Various authors have sought to address uncertainty in the ultimate loss ratio estimate, but generally from the perspective of a single point in time.

For example, Hayne [4] proposed a lognormal model of loss development that supports the construction of confidence intervals around the ultimate loss ratio estimate.¹ Kelly [7] and Kreps [9] also used a lognormal framework to explore issues of parameter estimation and parameter uncertainty, respectively. Hodes, Feldblum, and Blumsohn [5] used a slightly different lognormal development model to quantify the uncertainty in workers compensation reserves. Mack, Venter, and Zehn-wirth have all written extensively about stochastic modeling of the loss development process.² Others, including Van Kampen [12], Wacek [15], and the American Academy of Actuaries Property and Casualty Risk-Based

Capital Task Force [1], have sought to quantify the uncertainty in the ultimate loss ratio estimate used in pricing and reserving applications directly, without reference to the loss development process. The question on which all of these authors focused their attention is the potential variation in the final loss ratio at ultimate compared to the current ultimate loss ratio estimate, with no reference to how the ultimate loss ratio estimate might vary at intermediate points in time.

In contrast, in his acclaimed paper on solvency measurement, Butsic [3] observed that loss estimates change in their march through time. He recognized that they, like stock prices, are governed by a diffusion process, a type of continuous stochastic process with a time-dependent probability structure. However, he did not propose a model of this stochastic process.

How ultimate loss ratio estimates change in the future depends in part on the method used to make the estimates. In this paper we assume that loss ratio estimates are derived from a consistently applied estimation process with minimal subjective overriding of the indicated result. We model the behavior of loss ratio estimates using stochastic versions of two loss development methods: the chain ladder method and the Bornhuetter-Ferguson method, both using paid development data. To model chain ladder estimates, we combine Hayne's and Butsic's ideas to synthesize a lognormal diffusion model for the path of the ultimate loss ratio. Then we adapt that model to the Bornhuetter-Ferguson method.

This conceptual framework, which could easily be adapted to handle other loss development models, provides actuaries with the means to give their clients more information about how much their loss ratio or reserve estimates may fluctuate from period to period. As such, it can be a useful tool for managing expectations about the variability of loss reserve estimates. It also has potential application in a number of other areas of actuarial analysis, as we will discuss later.

¹Conscious that the confidence intervals he derived were dependent on the lognormal model being the correct choice, he cautiously described his results as providing a "range of reasonableness."

²For example, see Mack [10], [11], Venter [13], [14], and Zehn-wirth [16], [2] (the last co-authored with Barnett).

The paper comprises six sections, the first being this introduction. In Section 2 we outline Hayne's lognormal model of chain ladder loss development and illustrate its application using industry Private Passenger Auto Liability data from the 2004 Schedule P. We illustrate the main benefit of a stochastic model for loss development, namely, the ability to measure the uncertainty in loss development factors and in the ultimate loss ratio estimate.

In Section 3 we discuss the effect of future loss emergence on future ultimate loss ratio estimates. We show how to use information implicit in Hayne's model to determine the distribution of future estimates derived from our stochastic versions of the chain ladder and Bornhuetter-Ferguson methods, with particular attention to the loss ratio estimate one year out. We again use industry Private Passenger Auto Liability data to illustrate the process.

In Section 4 we adjust Hayne's model to allow for parameter uncertainty and illustrate the effect. Because the adjusted distribution does not have the multiplicative properties of the lognormal, we illustrate the use of Monte Carlo simulation to model the distribution of future ultimate loss ratio estimates.

In Section 5 we conclude with an outline of potential applications of the framework for future ultimate loss ratio estimates in loss reserving and risk-based capital applications.

2. Hayne's lognormal loss development model

Hayne presented two models of chain ladder loss development: one that assumed that development is independent from one period to the next, and a second one that relaxed the independence assumption. We will adopt the first model (and henceforth refer to it simply as "Hayne's model"). Kelly [7] argued that independence is more plausible for paid loss development than for case-incurred development. We will use paid

losses to illustrate our framework, but a good case can also be made for its application to ultimate loss ratio estimates based on case-incurred losses, even if all of the underlying assumptions are not always met.

Hayne's model is quite simple. He assumed that age-to-age development factors are lognormally distributed. The product of independent lognormal random variables is also lognormal, which implies that age-to-ultimate loss development factors are lognormal. Because the product of a constant and a lognormal random variable is lognormal, if we are given the cumulative paid loss ratio at any age and the estimated parameters of the matching age-to-ultimate factor, we can determine the parameter estimates of the ultimate loss ratio. Using these parameters we can estimate the expected loss ratio (which we will take as the "best" estimate) as well as confidence intervals around that estimate.

The lognormal parameters μ and σ of the age-to-age factors can be estimated by a variety of methods. Hayne used (and we also prefer) the unbiased estimators

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i = \frac{1}{n} \sum_{i=1}^n \ln(x_i) \quad \text{and}$$

$$s^2 = \sum_{i=1}^n \frac{(y_i - \bar{y})^2}{n - 1}$$

for μ and σ^2 , respectively, where $x_1, x_2, x_3, \dots, x_n$ are the observed age-to-age factors.³ \bar{y} is also a maximum likelihood estimator.

2.1. Illustration of model parameter estimation

We illustrate the parameter estimation for Hayne's model using the real loss development data presented in Tables 1 and 2. Table 1 shows

³We used unweighted unbiased estimators for μ and σ throughout this paper. For formulas for estimators using unequal weights for the observations, see Section 5.5 of [15].

Table 1. Annual Statement for the year 2004 of the industry aggregate Schedule P, Part 3B, private passenger auto liability/medical

Cumulative paid net losses and defense and cost containment expenses reported at annual intervals (\$000,000 omitted)										
AY	1	2	3	4	5	6	7	8	9	10
1995	17,674	32,062	38,619	42,035	43,829	44,723	45,162	45,375	45,483	45,540
1996	18,315	32,791	39,271	42,933	44,950	45,917	46,392	46,600	46,753	
1997	18,606	32,942	39,634	43,411	45,428	46,357	46,681	46,921		
1998	18,816	33,667	40,575	44,446	46,476	47,350	47,809			
1999	20,649	36,515	43,724	47,684	49,753	50,716				
2000	22,327	39,312	46,848	51,065	53,242					
2001	23,141	40,527	48,284	52,661						
2002	24,301	42,168	50,356							
2003	24,210	41,640								
2004	24,468									

Paid age-to-age loss development factors										
AY	1 Loss Ratio	1-2 LDF	2-3 LDF	3-4 LDF	4-5 LDF	5-6 LDF	6-7 LDF	7-8 LDF	8-9 LDF	9-10 LDF
1995	28.0%	1.814	1.205	1.088	1.043	1.020	1.010	1.005	1.002	1.001
1996	27.7%	1.790	1.198	1.093	1.047	1.022	1.010	1.004	1.003	
1997	27.1%	1.771	1.203	1.095	1.046	1.020	1.007	1.005		
1998	27.1%	1.789	1.205	1.095	1.046	1.019	1.010			
1999	29.8%	1.768	1.197	1.091	1.043	1.019				
2000	32.2%	1.761	1.192	1.090	1.043					
2001	31.6%	1.751	1.191	1.091						
2002	30.4%	1.735	1.194							
2003	27.7%	1.720								
2004	26.6%									
Mean	28.8%	1.767	1.198	1.092	1.045	1.020	1.009	1.005	1.003	1.001
S.D.	2.0%	0.029	0.006	0.003	0.002	0.001	0.002	0.000	0.001	0.000
C.V.	7.0%	0.016	0.005	0.002	0.002	0.001	0.002	0.000	0.001	0.000
Cum. Mean		2.508	1.420	1.185	1.085	1.039	1.018	1.009	1.004	1.001

Source: Highline Data LLC as reported in the statutory filings (OneSource)

industry aggregate Schedule P net paid loss development data for Private Passenger Auto Liability for accident years 1995 through 2004 from the 2004 Annual Statement⁴ together with the associated paid loss age-to-age development factors. The paid loss ratios at age one year are also included in the development factor table. Table 2 shows the natural logarithms of the age-to-age factors and the age one year paid loss ratios. The rows labeled “Mean” and “S.D.” in Table 2 show the unbiased estimators for μ and σ , respectively, given the data in the body of the column.⁵

⁴Source: Highline Data LLC as reported in the statutory filings (OneSource).

⁵Note that the standard deviation for the age 9 to 10 development factor, which is undefined, has been selected to be equal to that of the age 8 to 9 development factor in both Table 1 and Table 2.

For example, in Table 2 the mean and standard deviation of the natural logarithms of the observed age 1 to 2 development factors are 0.569 and 0.016, respectively. If we set $\mu = 0.569$ and $\sigma = 0.016$,⁶ these parameter estimates for prospective age 1 to 2 development imply a lognormal mean, defined as $E(x) = \exp[\mu + 0.5\sigma^2]$, of 1.767, which matches the mean loss development factor calculated by the traditional method in Table 1. The same is true for all of the other age-to-age factors. Similarly, the parameter estimates for the age one paid loss ratio are -1.246 and 0.069 for μ and σ , respectively, which imply a lognormal mean of 28.8%. This, too, matches

⁶These parameters define the lognormal distribution that best fits the data, using unbiasedness as the criterion for “best.” However, there is uncertainty about whether those parameters are correct. We address the issue of parameter uncertainty later in the paper.

Table 2. Annual Statement for the year 2004 of the industry aggregate Schedule P, Part 3B, private passenger auto liability/medical

AY	Natural logarithms of paid age-to-age loss development factors in Table 1									
	Loss Ratio	Loss Development Factors								
	1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10
1995	-1.274	0.596	0.186	0.085	0.042	0.020	0.010	0.005	0.002	0.001
1996	-1.282	0.582	0.180	0.089	0.046	0.021	0.010	0.004	0.003	
1997	-1.307	0.571	0.185	0.091	0.045	0.020	0.007	0.005		
1998	-1.304	0.582	0.187	0.091	0.045	0.019	0.010			
1999	-1.210	0.570	0.180	0.087	0.042	0.019				
2000	-1.135	0.566	0.175	0.086	0.042					
2001	-1.151	0.560	0.175	0.087						
2002	-1.191	0.551	0.177							
2003	-1.282	0.542								
2004	-1.325									
Mean	-1.246	0.569	0.181	0.088	0.044	0.020	0.009	0.005	0.003	0.001
S.D.	0.069	0.016	0.005	0.002	0.002	0.001	0.002	0.000	0.001	0.001
LN Fit LDFs	28.8%	1.767	1.198	1.092	1.045	1.020	1.009	1.005	1.003	1.001
Cum. Mean	-0.327	0.919	0.350	0.170	0.082	0.038	0.018	0.009	0.004	0.001
Cum. S.D.	0.071	0.018	0.006	0.004	0.003	0.002	0.002	0.001	0.001	0.001
LN Fit LDFs	72.3%	2.508	1.420	1.185	1.085	1.039	1.018	1.009	1.004	1.001

Source: Highline Data LLC as reported in the statutory filings (OneSource)

the mean age one paid loss ratio shown in Table 1.⁷

The parameter estimates for the prospective age-to-age factors can be combined using the multiplicative property of lognormal distributions to determine parameter estimates for prospective age-to-ultimate factors. The product of n lognormal random variables with respective parameter sets $(\mu_1, \sigma_1), (\mu_2, \sigma_2), (\mu_3, \sigma_3), \dots, (\mu_n, \sigma_n)$ is a lognormal random variable with parameters

$$\mu = \sum_{i=1}^n \mu_i \quad \text{and}$$

$$\sigma = \left(\sum_{i=1}^n \sigma_i^2 \right)^{1/2}.$$

For example, treating age 10 as ultimate, in Table 2 the μ parameter estimate for the age 7 to ultimate development factor is the sum of the mean age-to-age factor natural logarithms for ages 7 to 8, 8 to 9, and 9 to 10: $0.005 + 0.003 + 0.001 = 0.009$. The corresponding σ parameter estimate

⁷We point this out to emphasize that the lognormal model appears to fit this data well. However, it is important to note that the mean of the fitted lognormal does not necessarily equal the mean of the sample.

Table 3. Summary of estimated ultimate loss ratios from paid loss development analysis: private passenger auto liability, based on industry aggregate experience as of December 2004

Accident Year	Net Earned Premiums	Net Paid Losses	Net Paid Loss Ratio	Age-to-Ult Factor	Estimated Ultimate Loss Ratio
1995	63,183	45,540	72.1%	1.000	72.1%
1996	66,006	46,753	70.8%	1.001	70.9%
1997	68,764	46,921	68.2%	1.004	68.5%
1998	69,343	47,809	68.9%	1.009	69.6%
1999	69,231	50,716	73.3%	1.018	74.6%
2000	69,444	53,242	76.7%	1.039	79.6%
2001	73,143	52,661	72.0%	1.085	78.1%
2002	79,922	50,356	63.0%	1.185	74.6%
2003	87,242	41,640	47.7%	1.420	67.8%
2004	92,064	24,468	26.6%	2.508	66.7%

is the square root of the sum of the variances of the natural logarithms of the same age-to-age factors: $\sqrt{0.000^2 + 0.001^2 + 0.001^2} = 0.001$. Note that the lognormal means (labeled “LN Fit LDFs” in Table 2) implied by these age-to-ultimate parameters match the age-to-ultimate development factors shown in Table 1.

The ultimate chain ladder loss ratio estimates indicated by this analysis as of the end of 2004 for accident years 1995 through 2004 are summarized in Table 3. In this example, the lognor-

Table 4. Summary of paid loss development factors with associated lognormal 95% confidence intervals: private passenger auto liability based on industry aggregate experience through December 2004

Age-to-age factors					
Paid Loss Development Period	Est μ for LDF	Est σ for LDF	Mean LDF	Lognormal 95% Confidence	
				LDF Lower Bound	LDF Upper Bound
9-Ult*	0.001	0.001	1.001	1.000	1.002
8-9	0.003	0.001	1.003	1.002	1.004
7-8	0.005	0.000	1.005	1.004	1.005
6-7	0.009	0.002	1.009	1.006	1.012
5-6	0.020	0.001	1.020	1.018	1.022
4-5	0.044	0.002	1.045	1.041	1.048
3-4	0.088	0.002	1.092	1.087	1.097
2-3	0.181	0.005	1.198	1.187	1.209
1-2	0.569	0.016	1.767	1.710	1.824

Age-to-ultimate factors					
Paid Loss Development Period	Est μ for LDF	Est σ for LDF	Mean LDF	Lognormal 95% Confidence	
				LDF Lower Bound	LDF Upper Bound
9-Ult*	0.001	0.001	1.001	1.000	1.002
8-Ult	0.004	0.001	1.004	1.002	1.006
7-Ult	0.009	0.001	1.009	1.007	1.011
6-Ult	0.018	0.002	1.018	1.015	1.022
5-Ult	0.038	0.002	1.039	1.034	1.043
4-Ult	0.082	0.003	1.085	1.079	1.091
3-Ult	0.170	0.004	1.185	1.176	1.193
2-Ult	0.350	0.006	1.420	1.403	1.436
1-Ult	0.919	0.018	2.508	2.423	2.595

*Age 10 deemed to be ultimate

mal loss development model produces the same loss ratio estimates as the traditional deterministic chain ladder loss development method. If we were interested only in these mean estimates, the traditional approach would suffice. However, we also want to measure the uncertainty in the loss ratio estimates, and for that purpose the richer lognormal model is superior.

2.2. Measurement of loss development uncertainty

If we assume $\mu = \bar{y}$ and $\sigma = s$ based on the data for each age-to-age development period, we can calculate the lower and upper bounds of a two-sided 95% confidence interval for prospective age-to-age factors as $\exp[\bar{y} - N^{-1}(97.5\%) \cdot s]$

and $\exp[\bar{y} + N^{-1}(97.5\%) \cdot s]$, respectively, where $N^{-1}(97.5\%)$ is the value of the standard normal cdf corresponding to a cumulative probability of 97.5%.⁸ Similarly, using the parameter estimates for the age-to-ultimate factors, we can also determine confidence intervals for age-to-ultimate factors. We have tabulated these 95% confidence intervals based on the industry Private Passenger Auto Liability Schedule P data as of the end of 2004 in Table 4.⁹

Table 4 indicates that the age 1 to 2 development factor, which has an estimated mean of 1.767, should fall within a range of 1.710 to 1.824 95% of the time. The age 1 to ultimate development factor, which has an estimated mean of 2.508, can be expected to fall within a range of 2.423 to 2.595 95% of the time. Given the accident year 2004 paid loss ratio of 26.6% at age 1, these confidence intervals imply a paid loss ratio range at age 2 of 45.5% to 48.5% (47.0% \pm 1.5%) and an ultimate loss ratio range of 64.4% to 69.0% (66.7% \pm 2.3%).¹⁰

As we would expect, the development factors for more mature accident years have tighter confidence intervals. For example, the age 5 to 6 factor, which in a year end 2004 analysis would be applicable to accident year 2000, has an estimated mean of 1.020 and a 95% confidence range of 1.018 to 1.022, implying that 95% of the time the accident year 2000 paid loss ratio of 76.7% as of the end of 2004 will develop to a paid loss ratio of 78.1% to 78.4% by the end of 2005, a range of 0.3 points. The 95% confidence interval for the age 5 to ultimate factor,

⁸ $N^{-1}(97.5\%)$ is replicated in Excel by *NORMSINV(0.975)*.

⁹Bear in mind that these confidence intervals are premised on the parameter estimates being correct and are narrower than confidence intervals that incorporate parameter uncertainty.

¹⁰While the lognormal is a skewed distribution, the skewness is imperceptible for small values of σ , and in those cases the confidence intervals are, for most practical purposes, symmetrical. In this example where the age 1 to 2 factor has an estimated $\sigma = 0.016$, the skewness coefficient is 0.05. In contrast, in the case of $\sigma = 1$ the skewness coefficient is 6.18 and the confidence interval is highly asymmetrical!

which has an estimated mean of 1.039, is a range of 1.034 to 1.043. That implies an ultimate loss ratio range of 79.3% to 80.0%, or 0.7 points.

All of these development factor, loss ratio, and confidence interval estimates are as of the end of 2004. They are all subject to change as new information in the form of actual future loss emergence becomes available. In the next section we will show how to use information implicit in Hayne's approach to model the effect of future loss emergence on these estimates.

3. A model for future ultimate loss ratio estimates

Any estimate of the ultimate loss ratio for a particular accident year is quickly made obsolete by subsequent actual loss emergence. Because of this rapid obsolescence, the ultimate loss ratio must be re-estimated periodically in light of the loss development in the period since the previous evaluation. That loss development affects the new estimate in two ways.

3.1. Sources of variation in future loss ratio estimates

First, the actual accident year loss emergence replaces the expected emergence in the loss ratio projection. For example, in Table 3 the Private Passenger Auto Liability accident year 2004 ultimate loss ratio of 66.7%, estimated as of the end of 2004, was determined by applying an age-to-ultimate factor of 2.508 to the paid loss ratio of 26.6%. That age-to-ultimate factor reflected an *expected* age 1 to 2 development factor of 1.767 combined with an age 2 to ultimate factor of 1.420.

It is likely that *actual* age 1 to 2 loss development will vary from the expected. If, for example, the actual accident year 2004 emergence during 2005 (from age 1 to 2) corresponds to a development factor of 1.75, then in the ultimate loss ratio analysis conducted at the end of 2005 this actual development factor will replace the

expected development factor of 1.767. If the age 2 to ultimate factor remains unchanged at 1.420, the chain ladder estimate of the ultimate loss ratio will become $26.6\% \times 1.75 \times 1.42 = 66.1\%$. Assuming an expected ultimate loss ratio of 66.7%, the Bornhuetter-Ferguson loss ratio estimate will become $26.6\% \times (1.75 - 1.767) + 26.6\% \times 1.767 \times 1.42 = 66.1\%$.^{11,12}

Of course, loss emergence with respect to older accident years might cause a revision in the prospective age 2 to ultimate factor. This potential for tail factor revision is a second source of uncertainty. For example, suppose the actual age 2 to 3 development on accident year 2003 during 2005 corresponds to a factor of 1.210. If that factor is averaged with the previous eight-point mean of 1.198 determined in Table 1 (using loss development data through 2004), the result is a revised age 2 to 3 development factor of 1.199. Assuming the same process is repeated for the other development periods, a revised age 2 to ultimate factor will be obtained. If the resulting age 2 to ultimate factor is 1.425, the revised chain ladder ultimate loss ratio estimate is given by $26.6\% \times 1.75 \times 1.425 = 66.3\%$, a reduction of 0.4% from the year end 2004 ultimate loss ratio estimate of 66.7%. The revised Bornhuetter-Ferguson estimate in this case is given by $26.6\% \times (1.75 - 1.767) + 26.6\% \times 1.767 \times 1.425 = 66.5\%$.

The foregoing is an illustration of just one scenario of the loss development that might occur in 2005 and its effect on the ultimate loss ratio estimate. We can use information developed in

¹¹This is algebraically equivalent to the conventional statement of the B-F estimate as Emerged LR + Expected Ultimate LR $\times (1 - 1/\text{Ultimate LDF})$, which in this example would be expressed as $26.6\% \times 1.75 + 26.6\% \times 1.767 \times 1.42 \times (1 - 1/1.42)$. For the December 2005 valuation we will assume a B-F Expected Ultimate LR for each accident year equal to the corresponding chain ladder ultimate estimate as of December 2004 shown in Table 3. There are various ways to choose B-F Expected Ultimate LRs and ours may not be the one most commonly used in practice.

¹²Note that it is merely a coincidence that the chain ladder and B-F estimates in this example are both 66.1%.

Table 5. Summary of revised mean paid loss development factors one year out: private passenger auto liability, based on industry aggregate experience through December 2004

Paid Loss Development Period	Est σ for Actual LDF	Actual LDF Weight	Est σ for Revised Mean LDF	Est μ for Revised Mean LDF	Revised Mean LDF	Lognormal 95% Confidence	
						Revised Mean LDF (Lower Bound)	Revised Mean LDF (Upper Bound)
<i>Mean age-to-age factors one year out</i>							
9-Ult*	0.001	1/2	0.000	0.001	1.001	1.001	1.002
8-9	0.001	1/3	0.000	0.003	1.003	1.002	1.003
7-8	0.000	1/4	0.000	0.005	1.005	1.005	1.005
6-7	0.002	1/5	0.000	0.009	1.009	1.009	1.010
5-6	0.001	1/6	0.000	0.020	1.020	1.020	1.020
4-5	0.002	1/7	0.000	0.044	1.045	1.044	1.045
3-4	0.002	1/8	0.000	0.088	1.092	1.091	1.093
2-3	0.005	1/9	0.001	0.181	1.198	1.197	1.199
1-2	0.016	1/10	0.002	0.569	1.767	1.761	1.772
<i>Mean age-to-ultimate factors one year out</i>							
9-Ult*			0.000	0.001	1.001	1.001	1.002
8-Ult			0.000	0.004	1.004	1.003	1.005
7-Ult			0.000	0.009	1.009	1.008	1.010
6-Ult			0.000	0.018	1.018	1.017	1.019
5-Ult			0.001	0.038	1.039	1.038	1.040
4-Ult			0.001	0.082	1.085	1.084	1.086
3-Ult			0.001	0.170	1.185	1.183	1.186
2-Ult			0.001	0.350	1.420	1.417	1.422
1-Ult			0.002	0.919	2.508	2.499	2.517

*Age 10 deemed to be ultimate

Hayne’s framework to model these two effects generally.

3.2. Modeling the first source of variation—Accident year development

The first effect is captured by the lognormal random variable estimated for the next year of development with respect to the accident year under review. For example, for accident year 2004, which at the end of 2004 is age 1, the lognormal distribution with $\mu = 0.569$ and $\sigma = 0.016$ models age 1 to 2 paid development. Then, since the age 1 paid loss ratio is 26.6%, the paid loss ratio distribution at age 2 is lognormal with parameters $\mu = \ln 26.6\% + 0.569 = -0.756$ and $\sigma = 0.016$, implying a mean of 47.0%.

If the mean age 2 to ultimate factor (the tail factor) of 1.42 does not change, then the distribution of the revised chain ladder ultimate loss ratio estimate at age 2 (i.e., one year out) has lognor-

mal parameters $\mu = \ln 26.6\% + 0.569 + \ln 1.42 = -0.406$ and $\sigma = 0.016$. The random variable for this chain ladder estimate x_{CL} can be expressed as a function of the paid loss ratio random variable x_P and the expected value of the mean tail factor:

$$x_{CL} = x_P \cdot E(\text{tail}) \quad (3.1)$$

The random variable x_{BF} for the comparable Bornhuetter-Ferguson estimate is a shifted version of the random variable for the age 2 paid loss ratio:

$$x_{BF} = x_P - E(x_P) + E(x_P) \cdot E(\text{tail}) \quad (3.2)$$

As defined by Formulas 3.1 and 3.2, both x_{CL} and x_{BF} reflect the uncertain impact of accident year 2004 development during 2005 on the updated ultimate loss ratio estimate that will be made at the end of 2005, but do not reflect the potential impact of tail factor revision.

3.3. Modeling the second source of variation—Tail factor revision

The second effect, due to tail factor revision, is captured by measuring the effect of the lognormal loss development modeled for the next year on the existing mean age-to-age and age-to-ultimate factors. For example, the mean age 2 to 3 development factor shown in Table 1 is 1.198. This is a mean of eight data points. What will be the effect on the mean of adding a ninth data point (representing 2005 development on accident year 2003), given that it will arise from a lognormal distribution with parameters $\mu = 0.181$ and $\sigma = 0.005$ (and mean of 1.198)? The uncertain ninth data point will contribute one-ninth weight to the revised mean age-to-age factor. There is no uncertainty about the existing mean age 2 to 3 factor—it is a constant. Therefore, the σ parameter of the distribution of the revised mean age 2 to 3 factor one year out, given an additional year of actual development, is given by $\sqrt{(\frac{8}{9} \cdot 0)^2 + (\frac{1}{9} \cdot 0.005)^2} = 0.001$. The μ parameter is given by $\ln 1.198 - 0.5 \cdot 0.001^2 = 0.181$. We can use the same process to estimate μ and σ parameters for the comparable distributions of mean age-to-age factors one year out for all such factors comprising the development tail.¹³ We can then combine the revised mean age-to-age factor parameters to determine the parameters of the revised mean age-to-ultimate factor distributions. See Table 5 for a tabulation of the parameters of these revised mean age-to-age and age-to-ultimate distributions for all ages. The σ of the distributions of revised factors for age 3 to 4 and beyond is less than 0.0005 (and thus displayed as 0.000 in Table 5), indicating that for Private Passenger Auto Liability, the uncertainty arising from the potential for tail factor revision is very

¹³Bear in mind that these parameters refer to distributions of the mean age-to-age development factor one year out and not to distributions of the development factor itself. We are interested in the distribution of the mean development factor because changes in the mean directly affect the ultimate loss ratio estimate (which is also a mean).

small. This is confirmed by the very narrow confidence intervals.

3.4. Modeling the revised loss ratio estimate one year out

We can now combine these two effects to determine the distribution of the revised ultimate loss ratio estimate that will be determined in one year's time based on the updated loss development experience that will then be available.

To determine the distribution of the revised chain ladder estimate, we start with the actual accident year paid loss ratio, which we then multiply by the lognormal random variables for (1) the age-to-age factor for the next year of development (obtaining the random variable x_P of the paid loss ratio one year out) and (2) the revised age-to-ultimate factor beyond the next year of development. Using accident year 2004 as an example, as of the end of 2004 the ultimate loss ratio estimate is 66.7%, which has been determined by multiplying the paid loss ratio of 26.6% first by an age 1 to 2 factor of 1.767 and then by an age 2 to ultimate factor of 1.420. In order to model the ultimate loss ratio estimate one year later, at the end of 2005, we replace the constant age 1 to 2 factor of 1.767 with the lognormal random variable with parameters $\mu_1 = 0.569$ and $\sigma_1 = 0.016$. In addition, we replace the constant age-to-ultimate factor of 1.420 with the lognormal random variable with parameters $\mu_2 = 0.350$ and $\sigma_2 = 0.001$. The expected values of these two lognormal random variables are 1.767 and 1.420, respectively. The product of the paid loss ratio (a constant) and these two lognormal random variables is lognormal with parameters $\mu = \ln P + \mu_1 + \mu_2$ and $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$, where P represents the actual paid loss ratio at the end of 2004, which, in this example, implies $\mu = -1.325 + 0.569 + 0.350 = -0.406$ and $\sigma = \sqrt{1.016^2 + 0.001^2} = 0.017$.

Table 6. Analysis of estimated ultimate loss ratios (chain ladder) one year out: private passenger auto liability, based on industry aggregate experience through December 2004

Accident Year	Dev't Age	Net Paid Loss Ratio	First Effect Accident Year Dev't		Second Effect Tail Factor Revision				Estimated Ultimate Loss Ratio 1 Yr Out	Lognormal 95% Confidence	
			Est μ for Actual 1-Yr LDF	Est σ for Actual 1-Yr LDF	Est μ for Revised Mean LDF	Est σ for Revised Mean LDF	Est μ for Est Ult L/R 1 Yr Out	Est σ for Est Ult L/R 1 Yr Out		Est L/R 1 Yr Out (Lower Bound)	Est L/R 1 Yr Out (Upper Bound)
			1995	10	72.1%	0.000	0.000	0.000		0.000	-0.327
1996	9	70.8%	0.001	0.001	0.000	0.000	-0.344	0.001	70.9%	70.8%	71.0%
1997	8	68.2%	0.003	0.001	0.001	0.000	-0.378	0.001	68.5%	68.4%	68.6%
1998	7	68.9%	0.005	0.000	0.004	0.000	-0.363	0.001	69.6%	69.5%	69.6%
1999	6	73.3%	0.009	0.002	0.009	0.000	-0.293	0.002	74.6%	74.4%	74.8%
2000	5	76.7%	0.020	0.001	0.018	0.000	-0.228	0.001	79.6%	79.5%	79.8%
2001	4	72.0%	0.044	0.002	0.038	0.001	-0.247	0.002	78.1%	77.8%	78.4%
2002	3	63.0%	0.088	0.002	0.082	0.001	-0.292	0.003	74.6%	74.3%	75.0%
2003	2	47.7%	0.181	0.005	0.170	0.001	-0.389	0.005	67.8%	67.1%	68.4%
2004	1	26.6%	0.569	0.016	0.350	0.001	-0.406	0.017	66.7%	64.5%	68.8%
2005	0	0.0%	-1.246	0.069	0.919	0.002	-0.327	0.069	72.3%	63.0%	82.6%

Generally, we can express the random variable x_{CL} as the product of the two lognormal random variables x_P and $tail$, representing the paid loss ratio one year out and the mean tail factor:

$$x_{CL} = x_P \cdot tail \tag{3.3}$$

Now we are in a position to determine confidence intervals for the revised chain ladder ultimate loss ratio estimate at the end of 2005. The endpoints of the two-sided 95% confidence interval are given by $\exp[\mu - N^{-1}(97.5\%) \cdot \sigma]$ and $\exp[\mu + N^{-1}(97.5\%) \cdot \sigma]$, which imply an estimated loss ratio range one year out for accident year 2004 of 64.5% to 68.8%, or approximately $66.7\% \pm 2.1\%$. Confidence intervals for ultimate loss ratio estimates one year out for the other accident years can be estimated in the same way and are tabulated together with those for accident year 2004 in Table 6.

To determine the distribution of the comparable revised Bornhuetter-Ferguson estimate, we replace the constant $E(tail)$ in Formula 3.2 with the random variable $tail$:

$$x_{BF} = x_P - E(x_P) + E(x_P) \cdot tail \tag{3.4}$$

We can also determine confidence intervals for the revised Bornhuetter-Ferguson loss ratio estimate at the end of 2005. However, because the

sum of two lognormal random variables, in this case x_P and $tail$, is not expressible in closed distributional form, the confidence intervals must be estimated using Monte Carlo simulation. The results of a simulation involving 10,000 trials are shown in Table 7. For each of the trials we randomly selected observations from the distributions of x_P and $tail$, assuming independence, and combined them according to Formula 3.4 to arrive at a simulated Bornhuetter-Ferguson estimate. After tabulating the results of 10,000 such trials, we determined the lower and upper bounds of the 95% confidence interval of the loss ratio estimate by identifying the 2.5 percentile and the 97.5 percentile of the trial values. Not surprisingly, the 95% confidence intervals for the revised Bornhuetter-Ferguson estimates are narrower in every case than the revised chain ladder estimates.

3.5. Modeling the revised loss ratio estimate—Other time horizons

We can extend this process to longer time horizons and determine the distribution of the ultimate loss ratio estimate two years out, three years out, and so on, until the time horizon encom-

Table 7. Analysis of estimated ultimate loss ratios (Bornhuetter-Ferguson) one year out: private passenger auto liability, based on industry aggregate experience through December 2004

Accident Year	Devt Age	Net Paid Loss Ratio	First Effect Accident Year Devt		Expected Paid L/R 1 Yr Out	Second Effect Tail Factor Revision			Estimated Ultimate Loss Ratio 1 Yr Out	Lognormal 95% Confidence *	
			Est μ for Actual 1-Yr LDF	Est σ for Actual 1-Yr LDF		Est μ for Revised Mean LDF	Est σ for Revised Mean LDF	Expected Mean Tail 1 Yr Out		Est L/R 1 Yr Out (Lower Bound)	Est L/R 1 Yr Out (Upper Bound)
			1995	10		72.1%	0.000	0.000		72.1%	0.000
1996	9	70.8%	0.001	0.001	70.9%	0.000	0.000	1.000	70.9%	70.8%	71.0%
1997	8	68.2%	0.003	0.001	68.4%	0.001	0.000	1.001	68.5%	68.4%	68.6%
1998	7	68.9%	0.005	0.000	69.3%	0.004	0.000	1.004	69.6%	69.5%	69.6%
1999	6	73.3%	0.009	0.002	73.9%	0.009	0.000	1.009	74.6%	74.4%	74.8%
2000	5	76.7%	0.020	0.001	78.2%	0.018	0.000	1.018	79.6%	79.5%	79.8%
2001	4	72.0%	0.044	0.002	75.2%	0.038	0.001	1.039	78.1%	77.8%	78.4%
2002	3	63.0%	0.088	0.002	68.8%	0.082	0.001	1.085	74.6%	74.3%	75.0%
2003	2	47.7%	0.181	0.005	57.2%	0.170	0.001	1.185	67.8%	67.2%	68.3%
2004	1	26.6%	0.569	0.016	47.0%	0.350	0.001	1.420	66.7%	65.1%	68.2%
2005	0	0.0%	-1.246	0.069	28.8%	0.919	0.002	2.508	72.3%	68.6%	76.3%

*Based on Monte Carlo simulation of $x_{BF} = x_p - E(x_p) + E(x_p) \cdot tail$

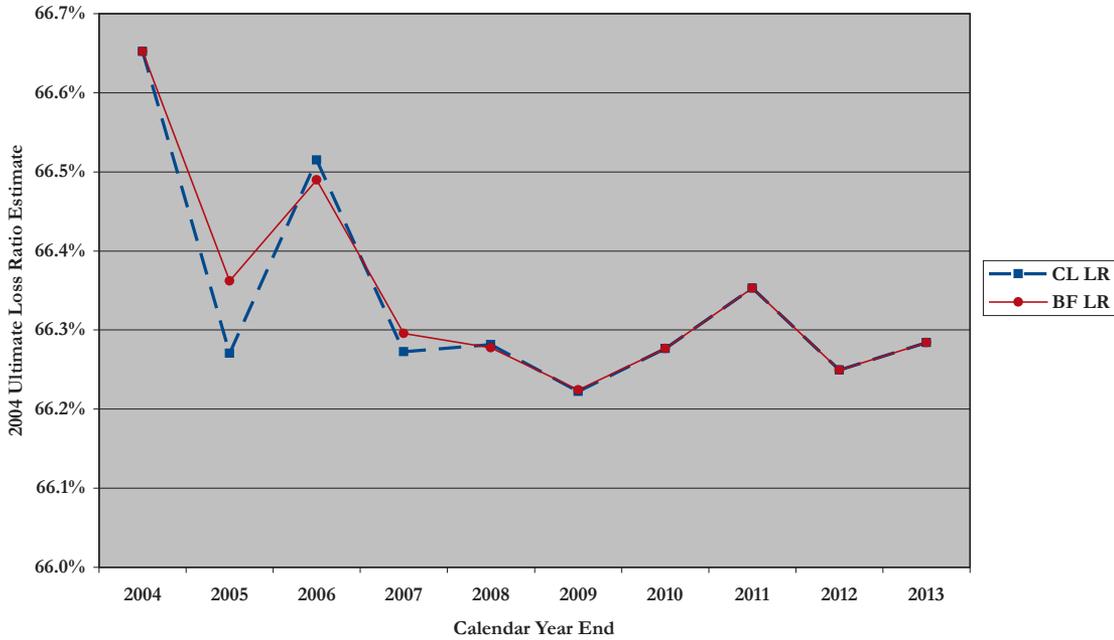
passes the point when all claims are expected to have been settled. The modeling is conducted in essentially the same way as for the one-year time horizon. For example, in the case of a two-year horizon, the first source of uncertainty (accident year development) is modeled using the distribution of the age j to $j + 2$ development factor, where j is the age in years of the accident year under review. The second source of uncertainty (potential tail factor revision) is modeled by reference to the potential effect of two additional development data points on the mean tail factor for age $j + 2$ to ultimate development. The analysis of a three-year time horizon focuses on accident year development from age j to $j + 3$ and the tail factor from $j + 3$ to ultimate, but is otherwise identical to that for the one-year and two-year time horizons. The analysis of the ultimate loss ratio estimate at points further in the future proceeds in the same way.

Alternatively, we can model the path of the ultimate loss ratio estimate as a succession of annual revaluations. Figure 1 illustrates this by plotting the results of one simulation of the path of the accident year 2004 loss ratio estimates through time for estimates determined from both

chain ladder and Bornhuetter-Ferguson methods. It represents just one path among many possibilities. The simulation was performed from the vantage point of the end of 2004. As such, it incorporates everything we know about actual loss development through that time as well as what we can infer about the structure of future development. We started with the accident year 2004 loss ratio estimate as of the end of 2004, which was 66.7%. Then, based on one random simulation of loss development during calendar year 2005, we made new chain ladder and Bornhuetter-Ferguson estimates of the ultimate loss ratio as of the end of 2005. We repeated the process for calendar years 2006 through 2013, using the simulated cumulative loss development through each valuation date. Figure 1 is a plot of the results. A complete description of the probability structure of the path can be built up from a simulation involving a large number of random trials, or, in the chain ladder case, directly from the properties of the lognormal distribution.

In practice, there might not be much benefit in determining the distribution of the chain ladder ultimate loss ratio estimate for time horizons between one year and the ultimate horizon (when

Figure 1. One path of the accident year 2004 ultimate loss ratio estimate.



all claims have been settled), at least for Private Passenger Auto Liability.¹⁴ We see this in Table 8, the top half of which compares the 95% confidence intervals for the accident year 1995 through 2004 chain ladder loss ratio estimates one year out with confidence intervals for the accident year loss ratio estimates over the ultimate time horizon. If we contrast the 95% confidence interval for accident year 2004 for the one-year horizon with the 95% confidence interval for the chain ladder loss ratio estimate over the ultimate time horizon, we can see that the contribution from the out years is dwarfed by the contribution from the next 12 months. The 95% confidence interval for the ultimate time horizon indicates a range for the accident year 2004 loss ratio of $66.7\% \pm 2.3\%$, which is barely wider than the range for just one year out. This is true not only for accident year 2004, but also holds for accident years 1995 through 2003.

For example, the accident year 2003 confidence interval of approximately $67.8\% \pm 0.7\%$

Table 8. Lognormal confidence intervals—ultimate loss ratios for one year vs. ultimate time horizons: private passenger auto liability, based on industry aggregate experience through December 2004

Accident Year	Dec 2004	95% Confidence Intervals—Chain Ladder Estimates			
		One-Year Horizon		Ultimate Horizon	
1995	72.1%	72.1%	72.1%	72.1%	72.1%
1996	70.9%	70.8%	71.0%	70.8%	71.0%
1997	68.5%	68.4%	68.6%	68.4%	68.6%
1998	69.6%	69.5%	69.6%	69.4%	69.7%
1999	74.6%	74.4%	74.8%	74.3%	74.8%
2000	79.6%	79.5%	79.8%	79.3%	80.0%
2001	78.1%	77.8%	78.4%	77.7%	78.5%
2002	74.6%	74.3%	75.0%	74.1%	75.2%
2003	67.8%	67.1%	68.4%	67.0%	68.6%
2004	66.7%	64.5%	68.8%	64.4%	69.0%

Accident Year	Dec 2004	95% Confidence Intervals—B-F Estimates			
		One-Year Horizon		Ultimate Horizon	
1995	72.1%	72.1%	72.1%	72.1%	72.1%
1996	70.9%	70.8%	71.0%	70.8%	71.0%
1997	68.5%	68.4%	68.6%	68.4%	68.6%
1998	69.6%	69.5%	69.6%	69.4%	69.7%
1999	74.6%	74.4%	74.8%	74.3%	74.8%
2000	79.6%	79.5%	79.8%	79.3%	80.0%
2001	78.1%	77.8%	78.4%	77.7%	78.5%
2002	74.6%	74.3%	75.0%	74.1%	75.2%
2003	67.8%	67.2%	68.3%	67.0%	68.6%
2004	66.7%	65.1%	68.2%	64.4%	69.0%

¹⁴There might be value in doing so for other lines that display more loss development variability.

for a one-year time horizon is almost as wide as that for the time horizon to ultimate of $67.8\% \pm 0.8\%$. For all of the older accident years, the first year of future development accounts for more than half of the variation associated with the ultimate time horizon.

This phenomenon is not confined to loss ratio estimates over short vs. longer time horizons. The same effect is also seen in other situations not related to insurance, where variability is a function of time. For example, given the common assumption that future stock price movements are lognormally distributed and independent, the 95% confidence interval for a stock price one year out, given constant annualized volatility of $\sigma = 20\%$ and an expected value of \$66.70, is \$45.07 to \$98.71, a range of \$53.64. Assuming the same expected value of \$66.70, the 95% confidence interval for the stock price two years out is \$38.22 to \$116.11, a range of \$77.80. The confidence interval range for the one-year time horizon stock price is 69% of the price range for the two-year time horizon. The reason for the disproportionate impact of the first period is that the confidence interval is not a linear function of σ but rather of $\sigma\sqrt{t}$, where t represents the time lag in years. In the case of chain ladder ultimate loss ratio estimation, where the age-to-age σ typically declines as the accident year ages, this effect can be even more pronounced.

Turning now to the Bornhuetter-Ferguson estimates, which are inherently less variable, the effect is smaller but still evident. The bottom half of Table 8 compares the 95% confidence intervals for accident year 1995 through 2005 loss ratio estimates one year out with the confidence intervals for the loss ratio estimates over the ultimate time horizon. In the Private Passenger Auto Liability example considered here, the 95% confidence interval for the accident year 2004 loss ratio estimate is approximately $66.7\% \pm 1.6\%$, which is about two-thirds of the range of the confidence interval for estimates at the ultimate time

horizon. For all of the older accident years, as in the case of the chain ladder estimates, the first year of future development accounts for more than half of the variation associated with the ultimate time horizon.

3.6. Modeling the loss ratio estimate at inception

Up to this point we have focused on modeling the distribution of the ultimate loss ratio after losses have begun to emerge. However, there is no reason why we cannot extend essentially the same procedure backward to the inception of loss exposure at age 0. Indeed, the benefit of doing so is that we can obtain a complete model of the path of the ultimate loss ratio from inception to ultimate.

The main difference in the procedure is that the lognormal model for loss emergence between age 0 and 1 describes the behavior of the paid loss ratio rather than an age-to-age factor. The rest of the analysis is merely an application of Formula 3.3.

For example, assume for the sake of illustration that the age 1 paid loss ratios in Table 1 are lognormally distributed and reflect “on level” adjustments to the accident year 2005 level. The mean age 1 paid loss ratio is 28.8%, which we can take as an estimate of the 2005 “on level” age 1 paid loss ratio. The unbiased estimates of the parameters of the lognormal distribution representing the paid loss ratio at age 1 are $\mu = -1.246$ and $\sigma = 0.069$. These parameters imply a lognormal mean paid loss ratio of 28.8% that matches the sample mean. The age 1 to ultimate development factor of 2.508 implies an ultimate loss ratio estimate at inception of 72.3%.

Applying the lognormal multiplicative rule described in Section 2, the parameters of the lognormally distributed ultimate loss ratio (at the ultimate time horizon) are $\mu = -1.246 + 0.919 = -0.327$ and $\sigma = \sqrt{0.069^2 + 0.018^2} = 0.071$, implying a 95% confidence interval of 62.8% to 82.9%, a range of 20.1%. The parameters of the

Table 9. Log t confidence intervals for paid loss development factors reflecting parameter uncertainty: private passenger auto liability, based on industry aggregate experience through December 2004

Paid Loss Development Period	Supporting Information for Confidence Interval Determination			Log t 95% Confidence			Log t Ratio to Lognormal	
	Sample Size n	Degrees of Freedom $n - 1^{**}$	$T_{n-1}^{-1}(97.5\%)$	LDF	LDF	LDF	At	At
				Lower Bound	Mean	Upper Bound	Lower Bound	Upper Bound
<i>Age-to-age factors</i>								
9-Ult*	1	3	3.182	0.998	1.001	1.004	0.998	1.002
8-9	2	3	3.182	1.000	1.003	1.005	0.999	1.001
7-8	3	3	3.182	1.004	1.005	1.006	0.999	1.001
6-7	4	3	3.182	1.004	1.009	1.015	0.998	1.002
5-6	5	4	2.776	1.017	1.020	1.023	0.999	1.001
4-5	6	5	2.571	1.039	1.045	1.050	0.998	1.002
3-4	7	6	2.447	1.085	1.092	1.099	0.998	1.002
2-3	8	7	2.365	1.184	1.198	1.212	0.997	1.003
1-2	9	8	2.306	1.697	1.767	1.839	0.992	1.008
<i>Age-to-ultimate factors***</i>								
9-Ult*				0.998	1.001	1.004	0.998	1.002
8-Ult				1.000	1.004	1.008	0.998	1.002
7-Ult				1.005	1.009	1.013	0.998	1.002
6-Ult				1.011	1.018	1.025	0.997	1.003
5-Ult				1.031	1.039	1.047	0.997	1.004
4-Ult				1.075	1.085	1.095	0.996	1.004
3-Ult				1.171	1.185	1.198	0.996	1.004
2-Ult				1.397	1.420	1.443	0.996	1.005
1-Ult				2.401	2.508	2.619	0.991	1.009

* Age 10 deemed to be ultimate
 ** Judgmentally limited to a minimum of 3. (Variance not defined, if d.f. < 3.)
 ***From Monte Carlo simulation (10,000 trials)

ultimate loss ratio one year out are $\mu = -1.246 + 0.919 = -0.327$ and $\sigma = \sqrt{0.069^2 + 0.002^2} = 0.069$. The indicated 95% confidence interval is 63.0% to 82.6%, a range of 19.6%. These calculations are summarized in Table 6.

The comparable Bornhuetter-Ferguson estimate can be determined by applying Formula 3.4. Table 7 shows that the 95% confidence interval for the revised Bornhuetter-Ferguson estimate of the accident year 2005 loss ratio one year out is 68.6% to 76.4%, a range of 7.8%.

4. Adjusting the model for parameter uncertainty

In Section 2 we explained that, given the observations $x_1, x_2, x_3, \dots, x_n$ arising from a log-normal process and the natural logarithms of the same observations $y_1, y_2, y_3, \dots, y_n$ (where $y_i =$

$\ln x_i$), the mean \bar{y} and standard deviation s of the log-transformed sample are unbiased estimators of the lognormal process parameters μ and σ , respectively. The parameter selections $\mu = \bar{y}$ and $\sigma = s$ define the lognormal distribution $f(x | \mu, \sigma)$ that best fits the data, using unbiasedness as the criterion for “best.”

However, while these are good estimates of the parameters, there is uncertainty about their true values. Fortunately, by combining information contained in the sample with results from sampling theory, it is possible to determine the mixed distribution $f(x)$ that reflects the probability weighted contribution of all of the potential parameter values.¹⁵ Wacek [15] showed that $f(x)$

¹⁵This assumes that the historical data and future observations are samples from the same distribution. To the extent that future economic, regulatory, and market conditions are different from those of the past, there will be additional sources of parameter uncertainty beyond those measured here.

defines a “log t ” distribution¹⁶ and in particular that the random variable $y = \ln x$ is Student’s t with $n - 1$ degrees of freedom, mean \bar{y} and variance $s^2 \cdot (n + 1)/n \cdot (n - 1)/(n - 3)$.¹⁷

4.1. Log t confidence intervals

The bounds of the two-sided log t 95% confidence interval are given by $\exp[\bar{y} - T_{n-1}^{-1}(97.5\%) \cdot s \cdot \sqrt{(n+1)/n}]$ and $\exp[\bar{y} + T_{n-1}^{-1}(97.5\%) \cdot s \cdot \sqrt{(n+1)/n}]$, respectively, where $T_{n-1}^{-1}(97.5\%)$ is the value of the standard Student’s t cdf with $n - 1$ degrees of freedom corresponding to a cumulative probability of 97.5%.¹⁸ Two-sided 95% confidence intervals for Private Passenger Auto Liability age-to-age factors, based on the log t distribution, are shown in Table 9. Unfortunately, the log t distribution does not share the multiplicative property of the lognormal. As a result, we cannot specify the distribution of age-to-ultimate development factors in closed form. Instead, the age-to-ultimate factor distributions and related confidence intervals must be estimated using a Monte Carlo simulation procedure that determines the age-to-ultimate factor from the underlying age-to-age factors for each random trial.

In the top section of Table 9, we have tabulated the indicated log t 95% confidence intervals for age-to-age factors based on the industry Private Passenger Auto Liability 2004 Schedule P data, together with the ratios of these confidence interval bounds to the lognormal confidence interval bounds given in Table 4. In addition, we have tabulated the sample size for each development period as well as $T_{n-1}^{-1}(97.5\%)$ and the degrees of freedom used in the calculations. At the risk of being seen as statistically less than rigorous, we set a minimum degrees of freedom value of

3 for purposes of calculating the confidence intervals to avoid using log t distributions with an undefined variance.

The log t confidence intervals shown in Table 9 for age-to-age factors are very close to the lognormal confidence intervals given in Table 4. The largest difference is in the age 1 to 2 factor, where the upper bound of the log t interval is 1.839, which is only 0.8% larger than the lognormal upper bound of 1.824. The percentage differences for the other age-to-age factors are smaller.

In the lower section of Table 9, we have tabulated the 95% confidence intervals for age-to-ultimate factors indicated by a Monte Carlo simulation involving 10,000 trials. As was the case with the age-to-age factors, the differences between the log t confidence intervals and lognormal confidence intervals for the age-to-ultimate factors are quite small. For example, the largest difference is in the age 1 to ultimate confidence interval, where the upper bound of the log t interval is 2.619. This is only 0.9% larger than the lognormal upper bound of 2.595. The percentage differences for the other age-to-ultimate factors are smaller. This suggests that, at least for Private Passenger Auto Liability at the industry level, the effect of parameter uncertainty is small enough that it can be ignored. However, that is probably not the case for individual insurers, particularly small ones, writing Private Passenger Auto Liability or for other lines of business.

4.2. Log t simulation of development factors

In the Monte Carlo simulation of age-to-ultimate factors, for each trial we randomly selected one age-to-age factor from each of the log t distributions representing development from age 1 to 2, age 2 to 3, ..., age 9 to 10. Treating age 10 as ultimate, we then multiplied these age-to-age factors in the usual way to determine a set of age-to-ultimate factors for that trial. After the results of the 10,000 trials were tabulated, we de-

¹⁶The log t bears the same relationship to the Student’s t distribution that the lognormal bears to the normal.

¹⁷Note that if we used the maximum likelihood estimator $s_{ml}^2 = \sum_{i=1}^n ((y_i - \bar{y})^2/n)$, the variance of this Student’s t distribution could be expressed as $s_{ml}^2 \cdot (n + 1)/(n - 3)$.

¹⁸ $T_{n-1}^{-1}(97.5\%)$ is replicated in Excel by $TINV(0.05, n - 1)$.

Table 10. Monte Carlo simulation of estimated ultimate loss ratio for accident year 2004 one year out: illustration of one random trial reflecting parameter uncertainty, private passenger auto liability, based on industry aggregate experience through December 2004

Dev't Period	Expected LDF	Sample Size n	Degrees of Freedom $n - 1^{**}$	Uniform Random Number R	$T_{n-1}^{-1}(R)$	\bar{y}	s	$\sqrt{\frac{n+1}{n}}$	Random LDF: Accident Yr Dev't*	Random LDF: Revised Tail*
9-Ult***	1.001	2	3	0.561	0.167	0.001	0.000	1.225		1.001
8-9	1.003	3	3	0.074	-1.938	0.003	0.000	1.155		1.002
7-8	1.005	4	3	0.084	-1.810	0.005	0.000	1.118		1.005
6-7	1.009	5	4	0.484	-0.043	0.009	0.000	1.095		1.009
5-6	1.020	6	5	0.899	1.468	0.020	0.000	1.080		1.020
4-5	1.045	7	6	0.128	-1.255	0.044	0.000	1.069		1.044
3-4	1.092	8	7	0.131	-1.220	0.088	0.000	1.061		1.092
2-3	1.198	9	8	0.396	-0.273	0.181	0.001	1.054		1.198
1-2	1.767	9	8	0.116	-1.293	0.569	0.016	1.054	1.727	
									1.727	1.418

Revised Chain Ladder Loss Ratio Estimate One Year Out = Paid Loss Ratio x Actual Acc Year Dev't x Revised Tail Factor
 = $26.6\% \times 1.727 \times 1.418 = 65.1\%$

Revised B-F L/R Loss Ratio Estimate One Year Out = Actual Paid L/R - Expected Paid L/R + Expected Paid L/R x Revised Tail Factor
 = $(26.6\% \times 1.727) - (26.6\% \times 1.767) + (26.6\% \times 1.767) \times 1.418 = 65.6\%$

* = $\exp(\bar{y} + T_{n-1}^{-1}(R) \cdot s\sqrt{(n+1)/n})$

** Judgmentally limited to a minimum of 3. (Variance not defined, if d.f. < 3.)

***Age 10 deemed to be ultimate

terminated the lower and upper bounds of the 95% confidence interval for each age-to-ultimate factor (age 1 to ultimate, age 2 to ultimate, etc.) by identifying the 2.5 percentile and the 97.5 percentile of the 10,000 trial values.

To make the random age-to-age factor selections, we started with a random draw R from the uniform distribution defined on the interval $[0, 1]$. Because R has a value between 0 and 1, it can be treated as though it is a cumulative probability. The number $T_{n-1}^{-1}(R)$ that corresponds to a standard Student's t cumulative probability of R is a random number from the standard Student's t distribution with $n - 1$ degrees of freedom, which has a mean of zero and a variance of $(n - 1)/(n - 3)$. More generally, the corresponding random number from the Student's t distribution with $n - 1$ degrees of freedom, mean M and variance $C^2 \cdot (n - 1)/(n - 3)$ is given by $M + T_{n-1}^{-1}(R) \cdot C$, which corresponds to a random number of $\exp[M + T_{n-1}^{-1}(R) \cdot C]$ from the related $\log t$ distribution. Substituting the appropriate values of \bar{y} for M and $s\sqrt{(n+1)/n}$ for C , we obtain

$\exp[\bar{y} + T_{n-1}^{-1}(R) \cdot s\sqrt{(n+1)/n}]$ as the value of a randomly selected age-to-age factor.

Putting some numbers to it, a draw of $R = 0.873$ implies a random age 1 to 2 development factor from the corresponding $\log t$ with 8 degrees of freedom of $\exp(0.569 + 1.229 \cdot 0.016 \cdot \sqrt{10/9}) = 1.803$.¹⁹ If the next draw is $R = 0.239$, then the random age 2 to 3 factor, drawn from the corresponding $\log t$ with 7 degrees of freedom, is $\exp(0.181 + (-0.749) \cdot 0.005\sqrt{9/8}) = 1.194$. Random numbers corresponding to the other development periods are similarly obtained. Then the age 1 to ultimate factor, the age 2 to ultimate factor, age 3 to ultimate factor, and so on, are obtained by multiplication. Tabulation of these results completes the first trial. The process is repeated in the same way for 10,000 trials.

¹⁹ $T_{n-1}^{-1}(R)$ is replicated in Excel by $TINV(2(1-R), n-1)$ if $R > 0.5$, and $-TINV(2R, n-1)$, if $R \leq 0.5$. $TINV$ assumes users are interested in two-tailed applications and therefore takes as its first argument the total two-tail probability. It returns values only from the right half of the distribution.

4.3. Log *t* simulation of future loss ratio estimates

Under conditions of parameter uncertainty, the distribution of future loss ratio estimates must also be modeled using Monte Carlo simulation. Each of the lognormal age-to-age development components identified in Section 3 must be replaced with corresponding log *t* components.

For example, to estimate the distribution of the updated chain ladder estimate of the accident year 2004 ultimate loss ratio at the end of 2005, given the year-end 2004 estimate of 66.7%, we tabulated 10,000 randomly obtained year-end 2005 loss ratio estimates. To determine each loss ratio estimate, we randomly selected from the log *t* distributions that represent the factors that contribute to the uncertainty in that estimate. For each trial we randomly selected one factor from the distribution of accident year 2004 development during 2005 and one factor from each of the age-to-age factor distributions that contribute to the revised tail factor. Then we multiplied all of these factors and the paid loss ratio as of year end 2004 to arrive at the ultimate loss ratio estimate for that trial.

This is illustrated in detail in Table 10 for one trial, where the simulated actual accident year 2004 age 1 to 2 development factor is 1.727 (compared to an expected factor of 1.767) and the revised tail factor is 1.418 (compared to an expected 1.420). The product of the year-end 2004 paid loss ratio and these two factors is the revised estimated ultimate loss ratio for accident year 2004 as of the end of 2005.

To arrive at approximate distributions of revised chain ladder ultimate loss ratio estimates for all of the accident years 1995 through 2004 as of the end of 2005, the process described in the preceding paragraph was repeated 10,000 times for each accident year.²⁰ The results of this pro-

²⁰Note that largely because their tail factors overlap, the accident year 1995 through 2004 ultimate loss ratio estimates are not independent, and for that reason their distributions were modeled simultaneously. To give one example of the tail factor overlap, the

Table 11. Log *t* confidence intervals—ultimate loss ratios, for one year vs. ultimate time horizons: private passenger auto liability, based on industry aggregate experience through December 2004

Accident Year		95% Confidence Intervals—Chain Ladder Estimates			
		Dec 2004	One-Year Horizon		Ultimate Horizon
1995	72.1%	72.1%	72.1%	72.1%	72.1%
1996	70.9%	70.7%	71.1%	70.7%	71.1%
1997	68.5%	68.3%	68.7%	68.3%	68.8%
1998	69.6%	69.4%	69.7%	69.3%	69.8%
1999	74.6%	74.2%	75.0%	74.1%	75.1%
2000	79.6%	79.3%	79.9%	79.0%	80.3%
2001	78.1%	77.7%	78.5%	77.4%	78.9%
2002	74.6%	74.1%	75.1%	73.8%	75.5%
2003	67.8%	66.9%	68.6%	66.7%	68.9%
2004	66.7%	64.0%	69.4%	63.8%	69.6%

Accident Year		95% Confidence Intervals—B-F Estimates			
		Dec 2004	One-Year Horizon		Ultimate Horizon
1995	72.1%	72.1%	72.1%	72.1%	72.1%
1996	70.9%	70.7%	71.1%	70.7%	71.1%
1997	68.5%	68.3%	68.7%	68.3%	68.8%
1998	69.6%	69.4%	69.7%	69.3%	69.8%
1999	74.6%	74.2%	75.0%	74.1%	75.1%
2000	79.6%	79.3%	79.9%	79.0%	80.3%
2001	78.1%	77.7%	78.5%	77.4%	78.9%
2002	74.6%	74.2%	75.1%	73.8%	75.5%
2003	67.8%	67.1%	68.4%	66.7%	68.9%
2004	66.7%	64.8%	68.5%	63.8%	69.6%

cess are summarized in Table 11, which, as the log *t* version of Table 8, compares the 95% confidence intervals for the accident year 1995–2004 loss ratio estimates one year out with the confidence intervals for the estimates over the ultimate time horizon. The chain ladder estimates are summarized in the top half of the table and the Bornhuetter-Ferguson estimates in the bottom half. As we observed in the lognormal case, much of the potential variation in the ultimate loss ratio estimates that is expected over the time horizon to ultimate is encompassed in the variation expected over a one-year time horizon. For example, the log *t* 95% confidence interval for the chain ladder estimate of the accident year 2004 loss ratio one year out of 66.7% ± 2.7% is

mean age 8 to 9 factor of 1.002 shown in Table 10, which illustrates the calculation of the accident year 2004 ultimate loss ratio estimate for one Monte Carlo trial, was also used in the calculation of ultimate loss ratio estimates for accident years 1996 through 2003.

nearly as wide as the 95% confidence interval of $66.7\% \pm 2.9\%$ for the same loss ratio over the ultimate time horizon. Similarly, the accident year 2003 confidence interval for the chain ladder estimate of approximately $66.7\% \pm 0.9\%$ for a one-year time horizon is also nearly as wide as that for the time horizon to ultimate of $67.8\% \pm 1.1\%$. For the older accident years, the proportion of the variation associated with the ultimate time horizon accounted for by the first year of future development is somewhat smaller, but the absolute size and significance of the confidence intervals for those years is much smaller.

Note that the $\log t$ confidence intervals are at least as wide in every case as the comparable lognormal confidence intervals shown in Table 7. In fact, in the case of the chain ladder estimates, for every accident year 1995–2004 the $\log t$ confidence intervals for the one-year time horizon are at least as wide as the lognormal confidence intervals for the ultimate time horizon!

5. Conclusions

There are a number of potential applications of the framework we have described for modeling future estimates of the ultimate loss ratio, ranging from loss reserving to pricing to analysis of risk-based capital. While a detailed discussion of these applications is beyond the scope of this paper, we will touch briefly on some examples.

5.1. Loss reserving

The framework presented in this paper gives reserve actuaries a way to manage their clients' expectations. Reserve clients don't like surprises and often express frustration that loss ratio or reserve estimates change significantly from one period to the next. We have shown in this paper that a large proportion of the potential variation in ultimate estimates can be present in the first year of future development. As we saw in the Private Passenger Auto Liability example we

presented, this phenomenon is particularly pronounced when the estimates are determined using the chain ladder method, but it can also be present if the estimates are derived from the Bornhuetter-Ferguson approach. It seems likely that most reserve clients do not understand this phenomenon. Actuaries have done a good job in getting clients to understand that ultimate loss estimates are subject to large potential variation, but many clients seem to expect that variation to emerge only in the distant future, if at all.

We suggest that the uncertainty in loss ratio and reserve estimates be framed in terms of how these estimates might change at the *next* valuation by presenting the ultimate estimates together with confidence intervals consistent with the valuation time horizon. For example, if the next valuation will be in one year, then the results would be presented with one-year time horizon confidence intervals. Then, because the potential variation has been explained to them in advance, clients might be better able to accept the revised estimates produced at the next valuation. This framework also naturally facilitates the explanation of the reasons for estimate revisions in terms of the sources of variation. For example, how much of the revision is due to actual accident year development and how much is due to a tail factor revision caused by loss emergence on the older accident years?

While we have focused much of our discussion on historical accident years and thus implicitly on reserving, we can easily extend this framework to encompass certain aspects of the pricing and underwriting, which can be used to assess risk load requirements and reinsurance risk transfer characteristics as well as to establish expectations for paid loss emergence during the first year after inception.

5.2. Risk-based capital

The framework described can also be applied to analysis of the issues outlined by Butsic [3]

in his paper on solvency measurement in risk-based capital applications. He advocated the use of a common time horizon for measurement of all kinds of risks on both sides of the balance sheet. He showed how long-term solvency protection could be achieved by periodic assessment and adjustment of risk-based capital using a short time horizon, e.g., one year. In particular, Butsic proposed that the risk-based capital charge at the beginning of each period be calibrated to a suitably small Expected Policyholder Deficit (EPD)²¹ expressed as a ratio to expected unpaid losses. The capital charge would be reset at the beginning of each new period based on asset and/or liability changes during the period just ended. While he illustrated his approach with numerical examples, he did not describe a model for how claim liabilities change from one period to the next. The model presented in this paper, using parameters determined from Schedule P data, could be used together with Butsic's approach to test and refine the capital charges employed in the NAIC and rating agency risk-based capital models.²² Moreover, to the extent that these risk-based capital charges imply the minimum amount of capital needed by an underwriter to assume risk, the model potentially has application to the problem of capital allocation for pricing applications as well.

5.3. Other stochastic loss development models

We have used Hayne's simple lognormal model to illustrate how to model the future behavior of loss ratio estimates. However, the same

²¹The EPD is defined as the expectation of losses exceeding available assets. It can be viewed as the expected value of the proportion of policyholder claims that will be unrecoverable because of insurer insolvency.

²²For stress testing these solvency models it may make sense to use the chain ladder model, which produces more variable loss ratio estimates, rather than the Bornhuetter-Ferguson model.

conceptual approach can be used with other stochastic models. If ultimate loss ratios are estimated using a different stochastic model, the path of future revisions to those ultimate loss ratio estimates can be determined using the ideas presented in this paper.

Abbreviations and notations

- μ , first parameter of lognormal, $E(\bar{y}) = \mu$
- σ , second parameter of lognormal, $E(s^2) = \sigma^2$
- EPD, expected policyholder deficit
- $f(x | \mu, \sigma)$, distribution of x , given known parameters μ, σ
- $f(x)$, distribution of x (unknown parameters)
- n , number of points in sample
- $N^{-1}(\text{prob})$, standard normal inverse distribution function
- P , actual paid loss ratio
- R , random number from unit uniform distribution
- s , standard deviation of log-transformed sample
- $T_{n-1}^{-1}(\text{prob})$, standard student's t inverse distribution function
- tail*, random variable for mean tail factor one year out
- $x_1, x_2, x_3, \dots, x_n$, lognormal sample
- x_{BF} , Bornhuetter-Ferguson estimate of ultimate loss ratio
- x_{CL} , chain ladder estimate of ultimate loss ratio
- x_P , cumulative paid loss ratio
- $y_1, y_2, y_3, \dots, y_n$ log-transformed sample
- \bar{y} , mean of log-transformed sample

References

- [1] American Academy of Actuaries Property/Casualty Risk-Based Capital Task Force, "Report on Reserve and Underwriting Risk Factors," *Casualty Actuarial Society Forum*, Summer 1993, pp. 105–171, <http://casact.org/pubs/forum/93sforum/93sf105.pdf>.
- [2] Barnett, G., and B. Zehnirith, "Best Estimates for Reserves," *Proceedings of the Casualty Actuarial Society* 87, 2000, pp. 245–321, <http://www.casact.org/pubs/proceed/proceed00/00245.pdf>.

- [3] Butsic, R. P., “Solvency Measurement for Property-Liability Risk-Based Capital Applications,” *Casualty Actuarial Society Discussion Paper Program*, 1992, pp. 311–354, <http://www.casact.org/pubs/dpp/dpp92/92dpp311.pdf>.
- [4] Hayne, R. M., “An Estimate of Statistical Variation in Development Factor Methods,” *Proceedings of the Casualty Actuarial Society* 72, 1985, pp. 25–43, <http://www.casact.org/pubs/proceed/proceed85/85025.pdf>.
- [5] Hodes, D. M., S. Feldblum, and G. Blumsohn, “Workers Compensation Reserve Uncertainty,” *Casualty Actuarial Society Forum*, Summer 1996, pp. 61–150, <http://www.casact.org/pubs/forum/96sforum/96sf061.pdf>.
- [6] Johnson, N. L., S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions 2* (2nd edition), New York: Wiley, 1995.
- [7] Kelly, M. V., “Practical Loss Reserving Method with Stochastic Development Factors,” *Casualty Actuarial Society Discussion Paper Program*, 1992, pp. 355–381, <http://www.casact.org/pubs/dpp/dpp92/92dpp355.pdf>.
- [8] Klugman, S. A., H. H. Panjer, and G. E. Willmot, *Loss Models: From Data to Decisions* (2nd edition), New York: Wiley, 1994.
- [9] 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>.
- [10] Mack, T., “Distribution-Free Calculation of the Standard Error of Chain Ladder Reserve Estimates,” *ASTIN Bulletin* 23, 1993, pp. 213–225, <http://www.casact.org/library/astin/vol23no2/213.pdf>.
- [11] Mack, T., “Measuring the Variability of Chain Ladder Reserve Estimates,” *Casualty Actuarial Society Forum*, Spring (1) 1994, pp. 101–182, <http://www.casact.org/pubs/forum/94spforum/94spf101.pdf>.
- [12] 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>.
- [13] Venter, G. G., “Introduction to Selected Papers from the Variability in Reserves Prize Program,” *Casualty Actuarial Society Forum*, Spring (1) 1994, pp. 91–100, <http://www.casact.org/pubs/forum/94spforum/94spf091.pdf>.
- [14] Venter, G. G., “Testing the Assumptions of Age-to-Age Factors,” *Proceedings of the Casualty Actuarial Society* 85, 1998, pp. 807–847, <http://www.casact.org/pubs/proceed/proceed98/980807.pdf>.
- [15] Wacek, M. G., “Parameter Uncertainty in Loss Ratio Distributions and its Implications,” *Casualty Actuarial Society Forum*, Fall 2005, pp. 165–202, <http://casact.org/pubs/forum/05fforum/05f165.pdf>.
- [16] Zehnwirth, B., “Probabilistic Development Factor Models with Applications to Loss Reserve Variability, Prediction Intervals and Risk Based Capital,” *Casualty Actuarial Society Forum*, Spring (2) 1994, pp. 447–606, <http://www.casact.org/pubs/forum/94spforum/94spf447.pdf>.