

A Bayesian Approach to Excess of Loss Pure Premium Rating

by Jack Barnett

ABSTRACT

This paper demonstrates a Bayesian approach for estimating loss costs associated with Excess of Loss Reinsurance programs. The main features of this approach are (1) prior severity distributions are adjusted for historical emergence patterns underlying the experience data; (2) Maximum Likelihood Estimation is used to estimate a ground up loss ratio for each prior severity distribution; (3) a posterior severity distribution is derived using a Bayesian approach; and (4) a posterior ground up loss ratio is derived using a Bayesian approach. This paper illustrates a simple implementation of the approach. We test the model by simulating from known frequency and severity distributions and fitting the model to the simulated “data”.

KEYWORDS

Bayesian estimation, Maximum Likelihood Estimation, Emergence Patterns, Exposure Curves

1. Introduction

When pricing excess of loss programs, many reinsurance and large account pricing actuaries may rely on assumptions which are difficult to support. This is particularly true when the actuary is uncertain of the underlying exposure, or which exposure curve(s) to use. An additional complication arises when the reinsurer does not receive all individual claims in the reinsurance submission. Claims are typically reported above some truncation point. Finally, Maximum likelihood estimation and, by extension, Bayesian analysis is generally not amenable to modeling long tail lines of business where claims tend to develop upward or downward over time. This paper attempts to solve all these problems within a single coherent framework.

Meyers (2005) provides a Bayesian methodology for estimating loss costs associated with high layers. An introduction to the methodology can be found as a solution to the Cotor Challenge. It's been many years in the making, but his solution to the Cotor Challenge was a primary motivation for much of the work I have done here. Following the Cotor Challenge, Meyers offered solutions to address what he felt were two serious shortcomings of the methodology: (1) the way in which prior severity models were developed, and (2) the issue of settlement time on ultimate value. These are major issues for which this paper offers two more solutions. Finally, the framework presented here includes a Bayesian estimate (mean of the posterior loss ratio distribution) for the ground up loss ratio. Actuarial literature is relatively absent any discussion of Bayesian loss ratios for Excess of Loss (XOL) rating.

Gelman (2006) distinguishes three categories of priors: (1) Prior distributions giving numerical information that is crucial to estimation of the model. This would be a traditional informative prior, which might come from a literature review or explicitly from an earlier data analysis, (2) Prior distributions that are not supplying any controversial information but are strong enough to pull the data away from inappropriate inferences that are consistent with the likelihood. This might be called a weakly informative prior, and (3) Prior distributions that are uniform, or nearly so, and basically allow the information from the likelihood to be interpreted probabilistically. These are noninformative priors, or maybe, in some cases, weakly informative. For approach taken here, a weakly informative prior is utilized for the prior severity distributions. The prior loss ratio distribution is implied by the prior severity distributions (as described in Section 5).

When significant claim volume is available, accurate severity distributions can be estimated by using empirical methods, nonparametric statistics, or maximum likelihood methods (Klugman 2009). When little or no data is available, actuaries may rely on an exposure curve. There are numerous approaches taken to decide upon an exposure curve. For example, one can use an external curve acquired from a

statistical agency, or an internal curve developed from similar business. In the limiting case, where there is no data, the Bayesian approach defaults to the average of the prior distributions. In such cases with little or no data, it becomes useful to calibrate prior severity distributions to agree with the exposure curve selected for the underlying insured. In such a case, we prefer an informative prior. As the volume of claim data grows, it becomes less critical for the average of the priors to agree with the exposure curve selected for the underlying insured. In this paper, we do not assume an exposure curve underlying the Insured risk is available. A Bayesian approach is most useful for when we are between the two extremes of no data and a significant volume of data. The posterior severity distribution can be viewed as a nonlinear weighing of data and prior curves.

Prior to the application of the model, it will be assumed that premium is on-level and losses have been properly adjusted for trend and changes in attachment/limit. Improper adjustments may lead to estimates with greater error.

Section 2 will outline how the severity distribution and the likelihood function can be adjusted for age of claim. This is accomplished through estimating claim count emergence patterns from various size of loss triangles and applying these patterns directly to the likelihood function.

Section 3 examines 24 severity distributions used as priors for Bayesian analyses purposes. Using “only” 24 curves is a relatively simple implementation of the approach. The curves are not meant to be prescriptive, but are selected for illustrative purposes.

Section 4 describes assumptions and sample data (one simulation) underlying the model.

Section 5 describes how we estimate an underlying loss ratio for each curve using maximum likelihood estimation.

Section 6 describes the calculation of the posterior weight for each curve.

Section 7 describes a Bayesian approach for estimating the ground up loss ratio weighting over all maximum likelihood estimates, also called the mean of the posterior loss ratio distribution.

Section 8 compares Bayesian claim count estimates to experience and maximum likelihood estimates for 200 sets of data (simulations).

Section 9 is summary and conclusions.

2. Adjusting the Likelihood Function

The likelihood function is the probability of observing the data given a hypothetical value of the parameter. In the example presented here, the data is the number of reported claims in each layer and the hypothetical parameter is a hypothetical severity distribution. The maximum likelihood estimate is the distribution yielding the highest probability of obtaining the observed reported claim counts weighted over all layers. However, in the case of liability business, data is immature. Ultimate claim counts by layer will generally not be the same as reported claim counts by layer. In general, curves fit to reported and ultimate claim counts utilizing maximum likelihood estimation will not produce the same parameter estimates.

Reported claims may be available, but can we get the severity distribution for ultimate claims? The form of the likelihood function used in this paper is based on grouped data. Let a range of possible values be partitioned as $c_0 < c_1 < \dots < c_k$ where c_0 is the smallest possible value in the model (i.e., the truncation point) and c_k is the largest possible value (i.e, maximum policy limit). Let $[F(c_j; \theta) - F(c_{j-1}; \theta)]$ represent the proportion of ultimate claims in the band c_j to c_{j-1} given the distribution F and parameter θ . The grouped likelihood function is given by the expression:

$$L(\theta) = \prod_j [F(c_j; \theta) - F(c_{j-1}; \theta)]^{n_j} \tag{1}$$

Where n_j is the number of claims in the band. Note, the proportion of claims for the reported data will differ from ultimate by the emergence pattern. In other words, the relative proportion of claims reported in each band is:

$$\% \text{ Reported}_j \propto [F(c_j; \theta) - F(c_{j-1}; \theta)] \tag{2}$$

For reported claims, the proportion of claims in each band is:

$$\% \text{ Reported}_j \propto [F(c_j; \theta) - F(c_{j-1}; \theta)] / \sum_j \% \text{ Reported}_j \times [F(c_j; \theta) - F(c_{j-1}; \theta)] \tag{3}$$

The denominator is simply a normalizing constant so that the probabilities sum to unity.

The **adjusted** likelihood function for reported claims is:

$$L(\theta) = \prod_j \{ \% \text{ Reported}_j \times [F(c_j; \theta) - F(c_{j-1}; \theta)] / \sum_j \% \text{ Reported}_j \times [F(c_j; \theta) - F(c_{j-1}; \theta)] \}^{n_j} \tag{4}$$

Exhibit 1.1 illustrates one way we might simulate reported claim counts from a lognormal distribution with an underlying emergence pattern. Only 10 of 30,000 simulations are displayed.

Column (2) represents an ultimate claim value. At any point in time, the estimated claim value may be more, or less, than the ultimate value. We are assuming an emergence pattern as shown in column (9). Smaller claims, from \$0 to \$100K are assumed to be reported more quickly. We assume that 90% of all claims less than \$100K have been reported. For the largest claims, only 60% have been reported. It is common for reinsurers to separately examine incurred and/or paid claim count triangles by size of loss. To convert ultimate counts to reported counts, we first augment the emergence pattern, column (3). A unit uniform random variable is generated in column (4). If the random variable is less than the emergence pattern $[(4) < (3)]$, the claim is considered reported [column (5) equals 1]. If the claim is reported, it is shown in column (6). At this point, roughly 90% of the claims less than \$100K show up as reported claims in column (6) while only 60% of the claims from \$2M to \$4M show up as reported.

Exhibit 1.2 illustrates the application of equation (4) to the reported claim counts simulated in Exhibit 1.1. In practice, it is typical to minimize the negative loglikelihood function, or: $-\ln[L(\theta)] = \sum_j n_j \ln[(3)]$. Minimizing column (9) produces the parameters shown in the upper left of the exhibit. You can see the fit is in close agreement to the underlying severity parameters, or “actual.”

So, what does this tell us? It tells us that if we have reported claim amounts (counts) and an estimate of the emergence pattern, we can replicate the parameters of the underlying severity distribution, and we are well on our way to implementing a Bayesian approach. Why not just use adjusted MLE rather than a Bayesian approach? Because we usually don't see 25,655 reported claims [total of column (8) in Exhibit 1.2]. Fewer claims leads to greater parameter uncertainty underlying the maximum likelihood estimates. For smaller portfolios, Bayesian estimates utilizing the mean of the posterior distribution are less impacted by random variation than MLE estimates. Boswell (2011) provides an illustration and explanation for this case. However, the adjusted MLE does seem especially well suited for pricing very large portfolios, or creating exposure curves for individual lines of business.

Exhibit 1.1

Exhibit 1A									
Actual									
μ : 10.42 σ : 2.12									
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Simulation	Rand()	Loss	Emergence Pattern	Random Variable	Reported (1=Yes/0=No)	Reported	c_{j-1}	c_j	Emergence Pattern
1	0.379	17,451	0.90	0.295	1	17,451	0	100,000	0.90
2	0.942	930,844	0.70	0.943	0		100,000	250,000	0.80
3	0.068	1,412	0.90	0.837	1	1,412	250,000	500,000	0.75
4	0.766	155,793	0.80	0.303	1	155,793	500,000	1,000,000	0.70
5	0.566	47,743	0.90	0.546	1	47,743	1,000,000	2,000,000	0.65
6	0.344	14,327	0.90	0.638	1	14,327	2,000,000	4,000,000	0.60
7	0.226	6,789	0.90	0.397	1	6,789			
8	0.247	7,843	0.90	0.381	1	7,843	Maximum Policy Limit		
9	0.763	152,775	0.80	0.955	0				
10	0.967	1,656,936	0.65	0.939	0				

Exhibit 1.2

Exhibit 1B											
Fit (Estimate)											
$\hat{\mu}$ 10.43 $\hat{\sigma}$ 2.12											
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)			
c_{j-1}	c_j	Emergence Pattern	$F(c_{j-1};\theta)$ lower	$F(c_j;\theta)$ upper	layer probabilities $F(c_j;\theta) - F(c_{j-1};\theta)$	(4)x(1)	(5)/ constant	ln(6)	Reported Claims n_j	negative-log likelihood $-(7)*(8)$	
0	100,000	0.90	0.000	0.694	0.694	0.625	0.734	-0.309	18831	5823.933	
100,000	250,000	0.80	0.694	0.826	0.132	0.106	0.124	-2.087	3187	6651.261	
250,000	500,000	0.75	0.826	0.897	0.071	0.053	0.062	-2.773	1593	4416.801	
500,000	1,000,000	0.70	0.897	0.944	0.047	0.033	0.039	-3.252	982	3193.341	
1,000,000	2,000,000	0.65	0.944	0.973	0.028	0.018	0.021	-3.842	574	2205.323	
2,000,000	4,000,000	0.60	0.973	1.000	0.027	0.016	0.019	-3.947	488	1926.150	
Maximum Policy Limit				Totals:	1.000	0.851	1.000		25655	24216.808	
						Normalizing Constant		Minimize this			

3. Prior Distributions

For this paper, we implement just 24 severity distributions as priors. Exhibit 2 below displays the parameters and some statistics for the 24 curves. The mean of the prior distribution is an equal weighted average of the curves, labelled “PRIOR (AVG)” (the excess severity is also weighted by the probability of a claim exceeding the truncation, \$100K). “ACTUAL” are the statistics associated with the lognormal curve in section 2. It should be noted that statistics and probabilities for the prior curves are net of any insurer attachment/limit profile. These figures can be easily calculated in most exposure rating models. For the implementation here, I have selected lognormal curves for prior distributions. Although we may be unsure of the underlying stochastic process, I have found that in *most* cases this choice of prior produces a reasonable estimate of the underlying severity distribution even when the underlying process is not lognormal. However, using priors that better match the distributional form generally produces estimates with lower variance and bias. I suggest the reader identify the most appropriate distribution underlying their data and use priors that match that distribution. For this exposition, expressions and calculations for the lognormal distribution will be shown/used. Maximum likelihood estimation is one possible approach for identifying the distribution. Given the same number of parameters, one could select the distribution that produces the greatest value of the likelihood function.

How might prior severity distributions be constructed? Assume we are pricing a \$5M excess \$5M layer on a per occurrence basis and receive only claims above some truncation, say \$100K.¹ To calibrate the curves as shown in Exhibit 2, and in general, the following procedure can be utilized: Examine your company’s exposure curves. What are the lowest and highest probabilities indicated by your curves for the highest layer, \$5M excess \$5M. The prior severity distributions should span the low and high probabilities indicated by your curves. For the exposure curves available to me, a range from 1% to 9% was sufficient. These figures represent the proportion of claim counts above \$5M given that a claim is above \$100K. Next, I’ve judgmentally selected three scale parameters: 9, 10 and 11. With the given scale parameters, I use Excel’s solver to find the shape parameters to match the target probabilities shown in column (14). In practice, I suggest using more than three scale parameters and for readers to extend the range of probabilities beyond the highest and lowest indicated by your exposure curves. Extending the ranges reduces the bias that may result when rating an account that tests the extremes of your prior severity curves.

¹ For reinsurance, it is uncommon to receive claims at such a low threshold (I have personally seen all claims for medical malpractice and legal liability lines). However, for large account pricing, it is common to receive all claims. The truncation here is only for illustration purposes.

Exhibit 2: Prior Severity Distributions

Curve	Lognormal Parameters		Excess Severity	Severity Ground Up	Probability >\$100K	% Loss>T [3]X[5]/[4]	[F(c;θ) - F(c _s ;θ)]							
	μ	σ					35,000 x	50,000 x	65,000 x	250,000 x	500,000 x	1,000,000 x	3,000,000 x	5,000,000 x
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
1	9.0	2.111	393,605	71,166	0.117	64.7%	0.219	0.189	0.146	0.228	0.121	0.058	0.029	0.010
2	9.0	2.198	453,173	83,162	0.126	68.9%	0.207	0.182	0.142	0.229	0.128	0.064	0.035	0.014
3	9.0	2.295	527,510	98,833	0.137	73.0%	0.195	0.173	0.138	0.229	0.134	0.071	0.041	0.019
4	9.0	2.407	621,171	119,601	0.148	77.0%	0.182	0.165	0.133	0.228	0.139	0.079	0.049	0.026
5	9.0	2.536	740,297	147,523	0.161	80.7%	0.169	0.155	0.128	0.225	0.144	0.086	0.058	0.035
6	9.0	2.687	893,192	185,605	0.175	84.1%	0.156	0.145	0.121	0.220	0.148	0.094	0.067	0.048
7	9.0	2.867	1,091,159	238,286	0.190	87.2%	0.142	0.135	0.114	0.213	0.151	0.101	0.078	0.066
8	9.0	3.087	1,349,634	312,195	0.208	89.8%	0.129	0.123	0.107	0.204	0.152	0.107	0.089	0.090
9	10.0	1.898	413,781	127,332	0.213	69.1%	0.202	0.182	0.145	0.237	0.131	0.063	0.031	0.010
10	10.0	1.977	473,573	144,931	0.222	72.6%	0.191	0.174	0.141	0.236	0.137	0.070	0.037	0.014
11	10.0	2.068	547,988	167,404	0.232	76.0%	0.180	0.166	0.136	0.235	0.142	0.077	0.044	0.019
12	10.0	2.171	641,529	196,470	0.243	79.3%	0.169	0.158	0.131	0.232	0.147	0.084	0.052	0.026
13	10.0	2.290	760,281	234,553	0.254	82.5%	0.158	0.149	0.126	0.228	0.151	0.092	0.061	0.035
14	10.0	2.431	912,485	285,091	0.267	85.4%	0.146	0.139	0.119	0.222	0.155	0.099	0.071	0.048
15	10.0	2.599	1,109,379	353,015	0.280	88.1%	0.134	0.129	0.112	0.215	0.156	0.106	0.082	0.066
16	10.0	2.804	1,366,336	445,459	0.295	90.4%	0.121	0.119	0.104	0.205	0.156	0.112	0.093	0.090
17	11.0	1.657	447,754	228,304	0.378	74.2%	0.176	0.169	0.142	0.249	0.147	0.073	0.035	0.010
18	11.0	1.729	507,550	253,208	0.383	76.8%	0.168	0.162	0.138	0.246	0.152	0.080	0.042	0.014
19	11.0	1.811	581,709	284,379	0.389	79.5%	0.159	0.155	0.133	0.243	0.156	0.087	0.049	0.019
20	11.0	1.905	674,658	323,842	0.394	82.1%	0.150	0.147	0.128	0.238	0.160	0.094	0.058	0.026
21	11.0	2.014	792,399	374,379	0.400	84.6%	0.141	0.139	0.122	0.233	0.163	0.101	0.067	0.035
22	11.0	2.143	943,086	439,834	0.405	86.9%	0.131	0.130	0.116	0.225	0.164	0.108	0.077	0.048
23	11.0	2.298	1,137,868	525,558	0.412	89.1%	0.121	0.121	0.109	0.217	0.164	0.114	0.088	0.066
24	11.0	2.488	1,392,048	639,044	0.418	91.1%	0.111	0.112	0.101	0.206	0.162	0.119	0.099	0.090
PRIOR (AVG):			651,163	261,632	0.269	66.9%	0.154	0.146	0.124	0.227	0.152	0.093	0.063	0.041
ACTUAL:	10.42	2.12	709,952	262,713	0.303	81.9%	0.157	0.150	0.127	0.232	0.154	0.091	0.059	0.030

PRIOR (averages) for columns 7 thru 14 are weighted by probabilities in column (5).

(3)=(EXP((μ+0.5*σ²)*NORMSDIST((LN(\$10M)-μ)/σ-σ)+\$10M*(1-NORMSDIST((LN(\$10M)-μ)/σ)))-(EXP((μ+0.5*σ²)*NORMSDIST((LN(\$100K)-μ)/σ-σ)+\$100K*(1-NORMSDIST((LN(\$100K)-μ)/σ)))/(5)
 (4)=EXP((m+0.5*s²)*NORMSDIST((LN(\$10M)-m)/s-s)+\$10M*(1-NORMSDIST((LN(\$10M)-m)/s))
 (5)=1 - NORMSDIST((LN(\$100K) - μ)/σ)
 (7)=(NORMSDIST((LN(\$135k) - μ) / σ)-NORMSDIST((LN(\$100K) - μ) / σ))/(5)

4. Assumptions and “Data”

In the discussion that follows, an identical severity distribution for each risk is assumed. This is an appropriate assumption for testing the efficacy of the model. The underlying ground up loss ratio for each risk will vary. Since each risk is of identical severity, the impact of loss ratio is felt through the frequency of claims.

In the example that follows, “data” will be the result of one simulation to which we fit the model. We simulate an emergence pattern, a ground up loss ratio and the size of claim(s). The Bayesian approach will simultaneously fit both the severity distribution and ground up loss ratio. We examine 10 years of losses with following additional assumptions: \$6M premium is written each year and the loss ratio drawn from a lognormal distribution with mean 100% and CV 20.2% ($\sigma=2$)². Size of loss is drawn from a lognormal distribution with parameters as shown in Exhibit 1.1 above, $\mu=10.42$ and $\sigma=2.12$. Maximum loss size (policy limit) is \$10M. The expected annual number of ground up claims over all risks is 22.8 [\$6M divided by an expected severity of \$262,713]. However, for each risk the

² Since the underlying loss ratio is 100%, the premium in this example is equal to pure premium (i.e., the amount necessary to pay losses).

expected frequency varies as a function of the underlying loss ratio. We assume that the reinsurer receives only those claims above \$100K (Truncation: $T=\$100K$). Given this lognormal assumption, roughly 6.9 claims per year are excess of \$100K [69 claims over 10 years]. Based on simulated emergence patterns, roughly 78% of these claims are reported. The percentage reported declines as the size of loss increases. Frequency is simulated from a negative binomial distribution with a variance-to-mean ratio of $1+cL$, with $c=.01^3$ where L is the expected number of ground up claims. Exhibit 3 below displays the expected emergence pattern for each of the layers [bands] of loss. The emergence patterns are unknown parameters which are estimated. This paper suggests an approach to estimate the patterns, but this is not a goal of the model. For lower layers, losses are reported more quickly than higher layers. The variation around the expected pattern for each age is modeled as lognormal with a sigma of 50%. After emergence patterns are simulated for each age, they are scaled so that the patterns add up to unity for each accident period. Exhibit 4.1 displays the reported claims, ultimate claims, estimated emergence pattern and actual emergence pattern for the 3,000,000 x 2,000,000 layer. Since each year has identical exposure, the incremental claim count pattern can be simply calculated as the incremental number of claims reported in the lag period [development period] divided by the years of data for that age. For example, in the triangle shown in Exhibit 4.1, one claim is reported from age 1 to age 2 (over accident periods 2007 through 2015). This results in .111 claims expected to emerge from age 1 to age 2 (1 claim over 9 years). For 10 years of data, 5 claims have been reported to date, 7 ultimate claims have been simulated, while we estimate 2.46 IBNR claims. The estimated percentage reported for \$3M x \$2M layer is 67.1% [$5/(5+2.46)$]. The 67.1% will be the \$3M x \$2M contribution to adjust the 24 severity distributions to a reported level. Exhibit 4.2 displays the simulated loss ratio, reported claims, IBNR claims, ultimate claims and estimated emergence pattern (percentage reported) for all the layers. The simulated loss ratio is an unknown parameter which the Bayesian framework will estimate. The reported claims and the ultimate claims are totals over 10 years. Column (4), “Ultimate”, represents an experience or emergence estimate. The Bayesian estimate will be compared to this estimate.

³ In my experience, I have found $c=.01$ to be a reasonable estimate of “contagion”. The model is robust to the choice of “c”.

Exhibit-3: Expected Emergence Patterns

Expected Emergence Patterns										
Layer	Age 1	Age 2	Age 3	Age 4	Age 5	Age 6	Age 7	Age 8	Age 9	Age 10
35,000 x 100,000	0.501	0.177	0.096	0.063	0.045	0.034	0.027	0.022	0.019	0.016
50,000 x 135,000	0.454	0.178	0.103	0.070	0.052	0.040	0.033	0.027	0.023	0.020
65,000 x 185,000	0.437	0.178	0.105	0.072	0.054	0.043	0.035	0.029	0.025	0.022
250,000 x 250,000	0.421	0.177	0.107	0.074	0.056	0.045	0.037	0.031	0.027	0.024
500,000 x 500,000	0.341	0.171	0.114	0.085	0.068	0.057	0.049	0.043	0.038	0.034
1,000,000 x 1,000,000	0.266	0.158	0.117	0.094	0.080	0.069	0.062	0.056	0.051	0.047
3,000,000 x 2,000,000	0.199	0.141	0.115	0.100	0.089	0.081	0.075	0.070	0.066	0.063
5,000,000 x 5,000,000	0.134	0.117	0.107	0.101	0.097	0.094	0.091	0.088	0.086	0.084

Exhibit-4.1

cumulative to date: 3,000,000 x 2,000,000												(1)	(2)	(3)	(4)
Year	OLP	age 1	age 2	age 3	age 4	age 5	age 6	age 7	age 8	age 9	age 10	Reported to date	IBNR Claims	Expected Ultimate	Simulated Ultimate
2007	6,000,000	0	0	0	0	0	0	0	0	0	0	0	0.000	0.000	0
2008	6,000,000	0	1	1	1	1	2	2	2	2		2	0.000	2.000	2
2009	6,000,000	0	0	0	0	1	1	1	1			1	0.000	1.000	1
2010	6,000,000	0	0	0	1	1	1	1				1	0.000	1.000	1
2011	6,000,000	0	0	0	0	0	0					0	0.000	0.000	0
2012	6,000,000	0	0	0	0	0						0	0.200	0.200	0
2013	6,000,000	0	0	1	1							1	0.367	1.367	1
2014	6,000,000	0	0	0								0	0.510	0.510	0
2015	6,000,000	0	0									0	0.635	0.635	1
2016	6,000,000	0										0	0.746	0.746	1
Expected Incremental Emergence:			0.111	0.125	0.143	0.167	0.200	0.000	0.000	0.000	0.000	5	2.456	7.456	7
Expected Cumulative Emergence:			0.746	0.635	0.510	0.367	0.200	0.000	0.000	0.000	0.000				
												Estimated Emerged Claims: Total (1)/Total (3)		0.671	
												Actual Emerged Claims: Total (1)/Total (4)		0.714	

Exhibit 4.2

Simulated Loss Ratio: 133.4%					
(1)	(2)	(3)	(4)	(5)	(6)
Layer	Reported	IBNR	Ultimate, or "Experience Estimate" (2)+(3)	Percentage Reported, or "Emergence Pattern" (2)/(4)	Underlying Frequency [Parameter] "ACTUAL"
35,000 x 100,000	17	1.98	18.98	89.6%	14.48
50,000 x 135,000	13	1.76	14.76	88.1%	13.82
65,000 x 185,000	10	1.57	11.57	86.4%	11.75
250,000 x 250,000	15	2.79	17.79	84.3%	21.46
500,000 x 500,000	9	2.08	11.08	81.2%	14.20
1,000,000 x 1,000,000	11	4.01	15.01	73.3%	8.45
3,000,000 x 2,000,000	5	2.46	7.46	67.1%	5.41
5,000,000 x 5,000,000	1	0.67	1.67	60.0%	2.78
Total:	81	17.32	98.32	82.4%	92.34

The adjustment of the severity distributions to a reported basis from an ultimate basis is shown in Exhibit 5 below. In the top left of the exhibit, you will see the emergence pattern, or percentage reported, in each layer (see Exhibit 4.2). By examining layer probabilities, you can see that adjusting the distributions shifts some probability from the higher attaching layers to the lower attaching layers.

Exhibit 5: Ultimate and Reported Probabilities

								% Reported									
89.6%		88.1%		86.4%		84.3%		81.2%		73.3%		67.1%		60.0%			
(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
$[F(c_i; \theta) - F(c_{i+1}; \theta)]$								Normalizing Constant	% Reported x $[F(c_i; \theta) - F(c_{i+1}; \theta)] / \text{Normalizing Constant}$								Total Probability
35,000 x 100,000	50,000 x 135,000	65,000 x 185,000	250,000 x 250,000	500,000 x 500,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000		35,000 x 100,000	50,000 x 135,000	65,000 x 185,000	250,000 x 250,000	500,000 x 500,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000	
0.219	0.189	0.146	0.228	0.121	0.058	0.029	0.010	0.847	0.231	0.197	0.149	0.227	0.116	0.050	0.023	0.007	1.000
0.207	0.182	0.142	0.229	0.128	0.064	0.035	0.014	0.844	0.220	0.190	0.146	0.229	0.123	0.056	0.027	0.010	1.000
0.195	0.173	0.138	0.229	0.134	0.071	0.041	0.019	0.839	0.208	0.182	0.142	0.230	0.129	0.062	0.033	0.013	1.000
0.182	0.165	0.133	0.228	0.139	0.079	0.049	0.026	0.834	0.195	0.174	0.138	0.230	0.136	0.069	0.039	0.018	1.000
0.169	0.155	0.128	0.225	0.144	0.086	0.058	0.035	0.828	0.183	0.165	0.133	0.229	0.142	0.076	0.047	0.025	1.000
0.156	0.145	0.121	0.220	0.148	0.094	0.067	0.048	0.821	0.170	0.156	0.128	0.226	0.147	0.084	0.055	0.035	1.000
0.142	0.135	0.114	0.213	0.151	0.101	0.078	0.066	0.813	0.157	0.146	0.122	0.221	0.151	0.091	0.064	0.049	1.000
0.129	0.123	0.107	0.204	0.152	0.107	0.089	0.090	0.803	0.143	0.135	0.115	0.214	0.153	0.098	0.074	0.067	1.000
0.202	0.182	0.145	0.237	0.131	0.063	0.031	0.010	0.845	0.214	0.189	0.148	0.236	0.125	0.055	0.025	0.007	1.000
0.191	0.174	0.141	0.236	0.137	0.070	0.037	0.014	0.841	0.204	0.182	0.145	0.237	0.132	0.061	0.030	0.010	1.000
0.180	0.166	0.136	0.235	0.142	0.077	0.044	0.019	0.837	0.193	0.175	0.141	0.237	0.138	0.067	0.035	0.013	1.000
0.169	0.158	0.131	0.232	0.147	0.084	0.052	0.026	0.832	0.182	0.167	0.136	0.235	0.144	0.074	0.042	0.019	1.000
0.158	0.149	0.126	0.228	0.151	0.092	0.061	0.035	0.826	0.171	0.159	0.132	0.233	0.149	0.081	0.050	0.026	1.000
0.146	0.139	0.119	0.222	0.155	0.099	0.071	0.048	0.819	0.160	0.150	0.126	0.229	0.153	0.089	0.058	0.035	1.000
0.134	0.129	0.112	0.215	0.156	0.106	0.082	0.066	0.811	0.148	0.141	0.120	0.223	0.156	0.096	0.068	0.049	1.000
0.121	0.119	0.104	0.205	0.156	0.112	0.093	0.090	0.801	0.136	0.131	0.113	0.216	0.158	0.102	0.078	0.067	1.000
0.176	0.169	0.142	0.249	0.147	0.073	0.035	0.010	0.841	0.188	0.177	0.146	0.249	0.142	0.063	0.028	0.007	1.000
0.168	0.162	0.138	0.246	0.152	0.080	0.042	0.014	0.837	0.179	0.170	0.142	0.248	0.147	0.070	0.033	0.010	1.000
0.159	0.155	0.133	0.243	0.156	0.087	0.049	0.019	0.833	0.171	0.164	0.138	0.246	0.152	0.076	0.040	0.013	1.000
0.150	0.147	0.128	0.238	0.160	0.094	0.058	0.026	0.828	0.162	0.156	0.133	0.243	0.157	0.083	0.047	0.019	1.000
0.141	0.139	0.122	0.233	0.163	0.101	0.067	0.035	0.822	0.153	0.149	0.128	0.239	0.161	0.090	0.055	0.026	1.000
0.131	0.130	0.116	0.225	0.164	0.108	0.077	0.048	0.815	0.144	0.141	0.123	0.233	0.163	0.097	0.064	0.035	1.000
0.121	0.121	0.109	0.217	0.164	0.114	0.088	0.066	0.808	0.135	0.132	0.116	0.226	0.165	0.103	0.073	0.049	1.000
0.111	0.112	0.101	0.206	0.162	0.119	0.099	0.090	0.798	0.125	0.123	0.109	0.217	0.165	0.109	0.083	0.068	1.000

5. Maximum Likelihood Estimate for the Ground Up Loss Ratio

For each curve, we can calculate the probability of a claim exceeding the truncation ($T=\$100K$), along with the ground up and the excess claim severities. Given a ground up loss ratio, it is possible to estimate the ground up claim count and the number of claims excess a given truncation point. Recall, we start our example (or analysis) with an estimate of the total on-level premium for all years. Applying a loss ratio produces total loss, total ground up claim counts and the number of claims above the truncation.

For each curve, given a loss ratio we can estimate the expected number of claims in each layer. Applying the emergence pattern (Exhibit 4.2) to these counts gives us the expected number of reported claims in each layer. We can assume a frequency distribution for each layer and solve for the loss ratio that maximizes the likelihood of observing the reported claim counts in each layer given the expected number of reported claims in each layer. Under this framework, we can view the loss ratio as a parameter of the frequency distribution. For the ground up claim counts, I assume a negative binomial distribution with a variance-to-mean ratio (VTM) of: $1+cL$, with $c=.01$ and L is the number of ground up claims as a function of the MLE. The model is robust to the choice of “ c ”. The maximum likelihood function is the product of eight frequency distributions normalized for probability in each band (the probability in each band can be found in Exhibit (5) columns 16 through 23). Normalizing assures us that the total probability summed over the eight frequency distributions is 100%. Normalizing doesn’t affect the MLE for each curve, but does affect the relative likelihood between curves. This influences the ground up loss ratio (mean of the posterior loss ratio distribution) weighted over all curves (more in section 7).

The VTM ratio generally declines as we move up successively into higher layers. Patrik and Mashitz (1990) describe expressions that modify the ground up variance-to-mean ratio as we move into higher layers. We can view the dispersion in the ground up claim count as the result of independent Poisson random variables each with different underlying mean. We assume the mean is gamma distributed with parameters A and B calibrated to produce the observed ground up claim count and VTM ratio. Expressions are as follows:

$$\text{Ground up claim counts} = A/B$$

$$\text{VTM Ratio} = (1+ B)/B$$

$$\text{Ground up parameters are: } B=1/(\text{VTM}-1)$$

$$A=[\text{Ground up counts}]x[\text{VTM}-1]$$

As you move into higher layers, Patrik & Mashitz proceed to show that the scale parameter $B \rightarrow B/(\text{probability of attaching})$. The expression for the VTM above gets updated as B gets transformed. The probability of attaching is calculated directly from the prior severity distribution.

Exhibit 6 takes us through the calculations for curve #1 (see Exhibit 2 for parameters). The Maximum Likelihood Estimate shown on line (2), 97.8%, minimizes the negative loglikelihood function on line (22). The negative binomial density function on line (19) evaluates the probability of observing the actual number of reported claims given the number of expected reported claims and the VTM ratio. The expected number of claims is a function of the Ground Up Loss Ratio. An identical calculation is performed for each curve.

Exhibit 6 – Maximum Likelihood Estimate of the Ground Up Loss Ratio for Curve #1.

Estimated on-level premium ALL YEARS	60,000,000	(1)								
Ground up Loss Ratio (MLE):	97.8%	(2)								
Excess Loss Percentage:	64.7%	(3); See Exhibit 2, Column (6)								
Excess Loss Rate :	63.3%	(4); (2)x(3)								
Excess Loss:	37,958,445	(5); (4)x(1)								
Excess Severity:	393,605	(6); See Exhibit 2, Column (3)								
Expected Excess Claims:	96.44	(7); (5)/(6)								
Probability of attaching:	11.7%	(8); See Exhibit 2, Column (5)								
Expected Ground Up Claims (over 10 year period):	824.6	(9); (7)/(8)								
Assumed Contagion:	1.0%	(10)								
Assumed VTM Ground Up:	1.825	(11); $1+(10)x(9)/10$								
B (scale parameter gamma distribution):	1.213	(12); $1/[(11)-1]$								
A (scale parameter gamma distribution):	100	(13); $(9)x(12)/10$								
			Layer1	Layer2	Layer3	Layer4	Layer5	Layer6	Layer7	Layer8
Curve #1_Probability of attaching [q]:	11.7%	9.1%	6.9%	5.2%	2.5%	1.1%	0.5%	0.1%	(14)	
B' = B/q	10.37	13.28	17.52	23.26	47.70	107.58	267.40	1036.94	(15); (12)/(14) [by layer]	
VTM:	1.096	1.075	1.057	1.043	1.021	1.009	1.004	1.001	(16); $[1+(15)]/(15)$	
Layer Probability:	0.219	0.189	0.146	0.228	0.121	0.058	0.029	0.010	(17)	
Emergence Pattern [% Reported]:	89.6%	88.1%	86.4%	84.3%	81.2%	73.3%	67.1%	60.0%	(18)	
Expected Reported:	18.9	16.1	12.2	18.6	9.5	4.1	1.9	0.6	(19); $(7)x(17)x(18)$	
Actual Reported Claims (simulated):	17.0	13.0	10.0	15.0	9.0	11.0	5.0	1.0	(20)	
Negative Binomial pdf:	0.0843	0.0795	0.1007	0.0706	0.1289	0.0023	0.0291	0.3242	(21)	
Normalizing constants (adjusted layer probs):	0.231	0.197	0.149	0.227	0.116	0.050	0.023	0.007	(22)	
log likelihood :	-3.94	-4.16	-4.20	-4.13	-4.20	-9.09	-7.32	-6.08	(23); $\ln[(21)x(22)]$	
Negative loglikelihood:	43.1159								(24); $-\text{sum}[(23)]$	

6. Posterior weights for the severity distribution

The severity likelihood function depends on two expressions. The first expression is equation (4) above, the adjusted likelihood function. The second expression will be described below.

It is quite common that only claims excess of some given truncation are submitted to a reinsurer. In many instances, a severity distribution that appears to be a good candidate (fit) from an excess

perspective may not be a good candidate from a ground up perspective. Why does this matter? Because, in many instances curves that seem to fit well just above the truncation point, do not adequately capture the tail of the underlying severity distribution. A curve that fits well on an excess basis might imply a loss ratio significantly, or implausibly different than expected. We examine the MLE for the ground up loss ratio. If the MLE is significantly different than the expected loss ratio, or ELR, the curve is not a good candidate for the underlying severity distribution. The expected loss ratio can be the result of a detailed analysis for the individual risk, or can be the result of an industry study. Given estimates of the ELR, uncertainty surrounding the ELR, and the process risk underlying the frequency and severity assumptions, we can calculate the likelihood of observing the MLE given these estimates. This is the second expression mentioned above.

For our case study, we assume ground up triangles are not available and we utilize a mean industry loss ratio of 100% with individual risks varying around this mean loss ratio. We will assume a lognormal distribution with $\sigma=40\%$. Our estimate of the dispersion around the expected loss ratio is considerably wider than the underlying 20%. Again, we will assume the MLE is lognormally distributed with mean 100% and $\sigma \approx (.4^2 + \text{process variance})^{.5}$, where the process variance is calculated assuming the underlying expected loss ratio of 100%. For the likelihood, we calculate the density of the MLE using the parameters above. For this calculation, I use Excel's built in function "LOGNORM.DIST". Exhibit 7 takes us through the calculation for curve #1. The calculation assumes that the premium charged for excess losses mirrors the expected loss ratio. A discussion of each term used to calculate the process variance is beyond the scope of this paper. However, analytical expressions are supplied as shown in Exhibit 7 below. The author uses simulation to validate the process variance.

Exhibit 8 – Posterior Weights for Severity Distributions

See Exhibit 4B													
Reported Claims	17	13	10	15	9	11	5	1					
% Distribution	21.0%	16.0%	12.3%	18.5%	11.1%	13.6%	6.2%	1.2%					
See Exh 2	See Exhibit 5												
(5)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)
Probability >\$100K	% Reported x [F(c _i ;θ)-F(c _{i-1} ;θ)]/Normalizing Constant								Group	Likelihood of	Ground Up	Excess	
	35,000 x	50,000 x	65,000 x	250,000 x	500,000 x	1,000,000 x	3,000,000 x	5,000,000 x	Likelihood	MLE given	Posterior Wgts	Posterior Wgts	MLE
	100,000	135,000	185,000	250,000	500,000	1,000,000	2,000,000	5,000,000	Function*	ELR & Tot Var	(24)x(25)/(5) #	(24)x(25) #	
0.117	0.231	0.197	0.149	0.227	0.116	0.050	0.023	0.007	9.57E-72	0.806	0.0081	0.0041	97.8%
0.126	0.220	0.190	0.146	0.229	0.123	0.056	0.027	0.010	4.56E-71	0.702	0.0312	0.0170	106.0%
0.137	0.208	0.182	0.142	0.230	0.129	0.062	0.033	0.013	1.57E-70	0.576	0.0813	0.0480	116.8%
0.148	0.195	0.174	0.138	0.230	0.136	0.069	0.039	0.018	3.68E-70	0.437	0.1338	0.0856	131.0%
0.161	0.183	0.165	0.133	0.229	0.142	0.076	0.047	0.025	5.54E-70	0.300	0.1274	0.0885	149.6%
0.175	0.170	0.156	0.128	0.226	0.147	0.084	0.055	0.035	4.89E-70	0.183	0.0629	0.0475	174.2%
0.190	0.157	0.146	0.122	0.221	0.151	0.091	0.064	0.049	2.25E-70	0.096	0.0140	0.0115	206.9%
0.208	0.143	0.135	0.115	0.214	0.153	0.098	0.074	0.067	4.58E-71	0.043	0.0012	0.0010	250.5%
0.213	0.214	0.189	0.148	0.236	0.125	0.055	0.025	0.007	2.31E-71	0.835	0.0112	0.0103	96.4%
0.222	0.204	0.182	0.145	0.237	0.132	0.061	0.030	0.010	8.65E-71	0.720	0.0345	0.0331	105.3%
0.232	0.193	0.175	0.141	0.237	0.138	0.067	0.035	0.013	2.36E-70	0.583	0.0731	0.0733	116.8%
0.243	0.182	0.167	0.136	0.235	0.144	0.074	0.042	0.019	4.49E-70	0.436	0.0992	0.1040	131.5%
0.254	0.171	0.159	0.132	0.233	0.149	0.081	0.050	0.026	5.54E-70	0.295	0.0792	0.0870	150.7%
0.267	0.160	0.150	0.126	0.229	0.153	0.089	0.058	0.035	4.09E-70	0.177	0.0335	0.0386	175.6%
0.280	0.148	0.141	0.120	0.223	0.156	0.096	0.068	0.049	1.61E-70	0.093	0.0065	0.0079	208.6%
0.295	0.136	0.131	0.113	0.216	0.158	0.102	0.078	0.067	2.87E-71	0.041	0.0005	0.0006	252.4%
0.378	0.188	0.177	0.146	0.249	0.142	0.063	0.028	0.007	5.45E-71	0.838	0.0149	0.0243	97.4%
0.383	0.179	0.170	0.142	0.248	0.147	0.070	0.033	0.010	1.45E-70	0.712	0.0333	0.0550	106.9%
0.389	0.171	0.164	0.138	0.246	0.152	0.076	0.040	0.013	2.90E-70	0.568	0.0523	0.0877	118.9%
0.394	0.162	0.156	0.133	0.243	0.157	0.083	0.047	0.019	4.13E-70	0.418	0.0541	0.0920	134.1%
0.400	0.153	0.149	0.128	0.239	0.161	0.090	0.055	0.026	3.95E-70	0.279	0.0340	0.0586	153.6%
0.405	0.144	0.141	0.123	0.233	0.163	0.097	0.064	0.035	2.32E-70	0.166	0.0117	0.0205	178.8%
0.412	0.135	0.132	0.116	0.226	0.165	0.103	0.073	0.049	7.57E-71	0.086	0.0020	0.0035	211.9%
0.418	0.125	0.123	0.109	0.217	0.165	0.109	0.083	0.068	1.17E-71	0.038	0.0001	0.0002	255.7%
Total:											1.0000	1.0000	

* Curve #1: 9.57E-72 = (.231^17)(.197^13)(.149^10)(.227^15)(.116^9)(.050^11)(.023^5)(.007^1)
#: Weights are normalized to sum to unity.

Exhibit 9 below displays the posterior severity distribution. To create Exhibit 9, we append columns (26) and (27) from Exhibit 8 to Exhibit 2. The last row of Columns (7) through (14) represents the posterior severity distribution conditioned on claims greater than the truncation. It is important to note that given the statistics in the last row of columns (3) through (6), and an estimate for the ground up loss ratio, claim count estimates for the eight bands can be produced. I will refer to these claim count estimates as Bayesian claim count estimates, or Bayesian estimates. It is also instructive to note that curve number #12 produces the greatest likelihood (posterior weight) of generating the observed data. If you look back at Exhibit 2, you can see that, relative to the other curves, the parameters and statistics for curve #12 are relatively close to the underlying, or ACTUAL.

Exhibit 9 – Posterior Severity Distribution

Curve	(3)	(4)	(5)	(6)	[F(c;θ)-F(c ₀ ;θ)]								(26)	(27)
	Excess Severity	Severity Ground Up	Probability >\$100K	% Loss>T [3]X[5]/[4]	35,000 x 100,000	50,000 x 135,000	65,000 x 185,000	250,000 x 250,000	500,000 x 500,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000	Ground Up Posterior Wgts	Excess Posterior Wgts
1	393,605	71,166	0.117	64.7%	0.219	0.189	0.146	0.228	0.121	0.058	0.029	0.010	0.0081	0.0041
2	453,173	83,162	0.126	68.9%	0.207	0.182	0.142	0.229	0.128	0.064	0.035	0.014	0.0312	0.0170
3	527,510	98,833	0.137	73.0%	0.195	0.173	0.138	0.229	0.134	0.071	0.041	0.019	0.0813	0.0480
4	621,171	119,601	0.148	77.0%	0.182	0.165	0.133	0.228	0.139	0.079	0.049	0.026	0.1338	0.0856
5	740,297	147,523	0.161	80.7%	0.169	0.155	0.128	0.225	0.144	0.086	0.058	0.035	0.1274	0.0885
6	893,192	185,605	0.175	84.1%	0.156	0.145	0.121	0.220	0.148	0.094	0.067	0.048	0.0629	0.0475
7	1,091,159	238,286	0.190	87.2%	0.142	0.135	0.114	0.213	0.151	0.101	0.078	0.066	0.0140	0.0115
8	1,349,634	312,195	0.208	89.8%	0.129	0.123	0.107	0.204	0.152	0.107	0.089	0.090	0.0012	0.0010
9	413,781	127,332	0.213	69.1%	0.202	0.182	0.145	0.237	0.131	0.063	0.031	0.010	0.0112	0.0103
10	473,573	144,931	0.222	72.6%	0.191	0.174	0.141	0.236	0.137	0.070	0.037	0.014	0.0345	0.0331
11	547,988	167,404	0.232	76.0%	0.180	0.166	0.136	0.235	0.142	0.077	0.044	0.019	0.0731	0.0733
12	641,529	196,470	0.243	79.3%	0.169	0.158	0.131	0.232	0.147	0.084	0.052	0.026	0.0992	0.1040
13	760,281	234,553	0.254	82.5%	0.158	0.149	0.126	0.228	0.151	0.092	0.061	0.035	0.0792	0.0870
14	912,485	285,091	0.267	85.4%	0.146	0.139	0.119	0.222	0.155	0.099	0.071	0.048	0.0335	0.0386
15	1,109,379	353,015	0.280	88.1%	0.134	0.129	0.112	0.215	0.156	0.106	0.082	0.066	0.0065	0.0079
16	1,366,336	445,459	0.295	90.4%	0.121	0.119	0.104	0.205	0.156	0.112	0.093	0.090	0.0005	0.0006
17	447,754	228,304	0.378	74.2%	0.176	0.169	0.142	0.249	0.147	0.073	0.035	0.010	0.0149	0.0243
18	507,550	253,208	0.383	76.8%	0.168	0.162	0.138	0.246	0.152	0.080	0.042	0.014	0.0333	0.0550
19	581,709	284,379	0.389	79.5%	0.159	0.155	0.133	0.243	0.156	0.087	0.049	0.019	0.0523	0.0877
20	674,658	323,842	0.394	82.1%	0.150	0.147	0.128	0.238	0.160	0.094	0.058	0.026	0.0541	0.0920
21	792,399	374,379	0.400	84.6%	0.141	0.139	0.122	0.233	0.163	0.101	0.067	0.035	0.0340	0.0586
22	943,086	439,834	0.405	86.9%	0.131	0.130	0.116	0.225	0.164	0.108	0.077	0.048	0.0117	0.0205
23	1,137,868	525,558	0.412	89.1%	0.121	0.121	0.109	0.217	0.164	0.114	0.088	0.066	0.0020	0.0035
24	1,392,048	639,044	0.418	91.1%	0.111	0.112	0.101	0.206	0.162	0.119	0.099	0.090	0.0001	0.0002
PRIOR:	651,163	261,632	0.269	66.9%	0.154	0.146	0.124	0.227	0.152	0.093	0.063	0.041		
ACTUAL:	709,952	262,713	0.303	81.9%	0.157	0.150	0.127	0.232	0.154	0.091	0.059	0.030		
POSTERIOR:	669,547	193,398	0.232	80.2%	0.166	0.156	0.130	0.232	0.149	0.086	0.054	0.028		

PRIOR (averages) for columns 7 thru 14 are weighted by probabilities in column (5).
 POSTERIOR (averages) for columns 7 thru 14 are weighted by column (27).

7. The Ground Up Loss Ratio (Mean of the Posterior Loss Ratio Distribution)

As described in section 5, a Maximum Likelihood Estimate for the ground up loss ratio is derived for each curve. We can view these MLEs as a set of priors each initially receiving equal weight. As with the severity distribution, the likelihood of each MLE depends on two expressions. The first expression can be found on line (21) of Exhibit 6. However, instead of transforming the likelihood to a negative loglikelihood, we will rely directly upon the likelihood function. To convert the 43.1159 negative loglikelihood found in Exhibit 6, we change sign and exponentiate: $e^{-43.1159}=1.884E-19$ The second expression is described below.

The eight frequency distributions, one for each layer, are a function of the MLE. As with a severity distribution, a frequency distribution that may be a good candidate for the excess frequency may not be good candidate for the underlying ground up frequency. Again, if the MLE is significantly different than the Expected Loss Ratio, the MLE is not a good candidate for the underlying frequency distribution. Whereas process risk should be recognized when identifying which severity curves could potentially produce the observed data, expected process risk underlying each curve should not be reflected when calculating the likelihood of the MLE given the expected loss ratio and the surrounding uncertainty (parameter uncertainty). Our second expression reflects only this parameter uncertainty.

As described above, we assume an industry loss ratio of 100% with individual risks lognormally distributed with $\sigma=40\%$.

Exhibit 10 displays the calculation of the ground up loss ratio which is a weighted average over all MLEs. In this example the underlying, ACTUAL, loss ratio is 133.4% while the posterior loss ratio is 146.1%.

Exhibit 10 – Mean of the Posterior Loss Ratio Distribution

	(1)	(2)	(3)	(4)
Curve	MLE (Prior)	Frequency Likelihood Function	Likelihood of MLE given ELR & Parameter Risk*	Posterior weights $(2) \times (3) / \sum (2) \times (3)$
1	97.8%	1.884E-19	1.009	0.001
2	106.0%	1.552E-18	0.886	0.006
3	116.8%	8.996E-18	0.718	0.027
4	131.0%	3.473E-17	0.519	0.074
5	149.6%	8.316E-17	0.321	0.110
6	174.2%	1.128E-16	0.162	0.075
7	206.9%	7.672E-17	0.063	0.020
8	250.5%	2.216E-17	0.018	0.002
9	96.4%	5.328E-19	1.029	0.002
10	105.3%	3.381E-18	0.897	0.012
11	116.8%	1.533E-17	0.718	0.045
12	131.5%	4.705E-17	0.512	0.099
13	150.7%	9.126E-17	0.313	0.117
14	175.6%	1.024E-16	0.156	0.066
15	208.6%	5.895E-17	0.060	0.015
16	252.4%	1.481E-17	0.017	0.001
17	97.4%	1.579E-18	1.015	0.007
18	106.9%	6.957E-18	0.872	0.025
19	118.9%	2.250E-17	0.687	0.063
20	134.1%	5.070E-17	0.481	0.100
21	153.6%	7.455E-17	0.289	0.089
22	178.8%	6.565E-17	0.142	0.038
23	211.9%	3.088E-17	0.054	0.007
24	255.7%	6.623E-18	0.015	0.000
POSTERIOR LR:	146.1%	=SUMPRODUCT{(1),(4)}		

*: Likelihood=LOGNORM.DIST(MLE, μ , σ ,0) where: $\mu = -.08$, $\sigma = .40\%$,

Exhibit 11 ties together the posterior severity distribution in Exhibit 9 with the Mean of the Posterior Loss Ratio Distribution in Exhibit 10 to arrive at Bayesian claim count estimates. You can see how the Bayesian count better mirrors “ACTUAL” (underlying frequency/parameter) than the Experience estimate. An experience approach is more responsive to the reported claims in the layer. Thus, it is more responsive to random variation. Exhibit 4.2 shows only one claim has been reported to the 5x5 layer, while the true underlying frequency is 2.78 claims. Experience rating is too responsive to the reported claim counts. Where reported claims are higher than the underlying expected reported,

experience rating tends to produce an estimate that is too high. Where reported claims are lower than the underlying expected reported, experience rating tends to produce an estimate that is too low. A Bayesian approach takes into account experience in all layers and does not over-react to random variation to the degree that experience rating does.

Exhibit 11 – Projected Claim Counts

See Exhibit 9						See Exhibit 9								See Exh 10
(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	
Excess Severity	Severity Ground Up	Probability >\$100K	% Loss>T [3]X[5]/[4]	# Claims>T \$60Mx[15]x[5]/[2]	35,000 x 100,000	50,000 x 135,000	65,000 x 185,000	250,000 x 250,000	500,000 x 500,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000	Loss Ratio	
ACTUAL:	709,952	262,713	0.303	81.9%	92.3	0.157	0.150	0.127	0.232	0.154	0.091	0.059	0.030	133.4%
POSTERIOR:	669,547	193,398	0.232	80.2%	105.0	0.166	0.156	0.130	0.232	0.149	0.086	0.054	0.028	146.1%
					Claim Counts									
					35,000 x 100,000	50,000 x 135,000	65,000 x 185,000	250,000 x 250,000	500,000 x 500,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000		
ACTUAL:					92.3	14.48	13.82	11.75	21.46	14.20	8.45	5.41	2.78	
POSTERIOR/Bayesian:					105.0	17.42	16.34	13.67	24.37	15.60	9.02	5.64	2.94	
"Experience" Estimate:						18.98	14.76	11.57	17.79	11.08	15.01	7.46	1.67	

8. Bayesian Estimates vs MLE and Experience Estimates

Exhibit 12 displays an extract of Bayesian, MLE and Experience fits (estimates) to 200 sets of simulated data. 200 simulations are enough simulations to be an effective illustration of expected results. The MLE approach utilizes parameters as estimated by the methods discussed above. Specifically, the adjusted likelihood function as in equation 3 and the maximum likelihood estimate for the ground up loss ratio as shown in section 5. The calculations for simulation 1 can be found in appendix A. For the experience estimates, estimated emergence patterns are utilized⁴. For ease of exposition, only the top three layers for each approach are shown. Columns 1 through 3 are Bayesian claim count estimates, columns 4 through 6 are MLE estimates, and 6 through 9 are experience estimates. Each set of claim count estimates rely on the same estimated emergence pattern. We also show the underlying, Bayesian and MLE loss ratios. The underlying loss ratio is the simulated loss ratio, while the Bayesian and MLE loss ratios are estimates. The last column, underlying \$5M x \$5M,

⁴ For layers without any claims reported over the entire experience period, the emergence pattern utilized is the pattern of the layer immediately below. However, for experience rating purposes, this would overstate the expected claim counts in higher layers. In these instances, the ratio of probability in the higher layer to the lower layer is applied to the expected claim counts in the lower layer. We use the probabilities associated with the posterior severity distribution.

is the underlying frequency for the \$5M x \$5M layer for each simulated set of data. This is the parameter of interest we are attempting to estimate.

The average of the Bayesian loss ratios shown in Exhibit 12 is 105.8%, while the underlying parameter for these 200 simulations is 98.6%. Some of this loss ratio bias is due to the choice of prior curves, however this bias may “force” expected counts to more closely mirror reported counts. It is difficult to come up with general rules for constructing prior severity distributions. The choice of curves may not only reduce or increase bias, but also affects variance. This trade-off should be tested in the development of prior curves. MLE loss ratios exhibit less bias than Bayesian loss ratios, but Mean Square Error (MSE) is significantly greater.

The “actual” number of claims for each layer is shown in column 13. “Actual” represents the underlying parameter, not the actual simulated counts. Simulation #1 represents estimates for the sample data. For each simulation, any deviation from “actual” is due to process risk. The goal of any model is to produce estimates that are both unbiased and minimum variance. In Exhibit 12 below, you can see the experience and MLE estimates tend to have slightly lower bias but much higher variance than the Bayesian estimates. It is worth noting that in some instances actuaries will price higher layers by applying ILFs to lower layers where there is more experience. In this case, exhibit 12 suggests that pricing would be improved by applying the ILF to the Bayesian estimate for the lower layer. However, deciding which ILF to utilize still requires judgment while a Bayesian framework is objective in nature.

Exhibit 12: Bayesian versus MLE & Experience Claim Count Estimates (over 10 years)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Avg of 200 simulations:	6.46	4.05	2.17	6.12	3.71	1.98	6.05	4.22	2.07	98.6%	105.8%	97.0%	2.05
Actual (@100%LR):	6.33	4.06	2.08	6.33	4.06	2.08	6.33	4.06	2.08	100.0%	100.0%	100.0%	2.08
Bias(@sim LR):	0.22	0.05	0.12	-0.12	-0.29	-0.07	-0.19	0.22	0.01	0.0%	7.2%	-1.6%	0.00
Variance:	0.80	0.62	0.57	1.91	1.55	1.91	8.07	4.11	2.81	0.0%	3.0%	8.3%	0.00
MSE (Var+Bias ²):	0.85	0.63	0.58	1.92	1.63	1.91	8.10	4.16	2.81	0.0%	3.5%	8.4%	0.00
Simulation	Bayesian Estimates			MLE Estimates			Experience Estimates			Underlying Loss Ratio	Bayesian Loss Ratio	MLE Loss Ratio	Underlying 5,000,000 x
	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000	1,000,000 x 1,000,000	3,000,000 x 2,000,000	5,000,000 x 5,000,000				
1	9.02	5.64	2.94	8.75	5.93	3.65	15.01	7.46	1.67	133.4%	146.1%	154.0%	2.78
2	5.05	3.35	2.00	4.13	2.93	2.15	3.79	4.35	1.67	81.9%	86.4%	81.7%	1.71
3	8.44	5.82	3.69	8.73	6.38	4.32	8.14	5.61	5.00	114.0%	146.2%	155.3%	2.37
4	7.54	4.46	2.02	7.94	4.02	1.09	10.33	5.52	0.69	103.9%	115.4%	94.2%	2.16
5	4.81	3.44	2.36	5.29	3.49	1.56	2.50	5.22	1.43	63.7%	86.6%	72.0%	1.33
6	6.01	3.94	2.24	6.47	3.51	1.03	8.93	1.25	1.67	88.8%	99.1%	75.0%	1.85
7	5.66	3.26	1.41	5.07	1.66	0.19	5.28	1.43	0.19	88.2%	86.1%	54.3%	1.84
8	4.08	2.56	1.37	3.01	1.63	0.59	1.43	4.76	0.95	67.0%	67.4%	43.4%	1.40
9	6.70	4.02	1.86	7.24	2.66	0.35	6.47	4.29	0.60	82.4%	102.0%	71.0%	1.72
10	8.40	5.36	2.89	8.40	5.65	3.23	9.47	5.87	2.78	126.9%	137.0%	138.7%	2.64
11	7.42	4.71	2.61	6.48	4.71	4.05	8.96	2.78	7.54	139.4%	125.4%	147.3%	2.90
12	7.25	4.95	3.10	6.71	5.05	4.19	9.43	2.86	4.68	99.6%	125.9%	139.9%	2.07
13	7.09	4.09	1.76	6.52	2.23	0.27	7.40	2.00	0.43	93.4%	107.1%	69.0%	1.94
14	5.34	3.03	1.30	4.56	2.13	0.55	7.04	1.43	0.18	90.6%	83.4%	60.1%	1.89
15	7.21	4.67	2.62	7.03	4.75	2.72	9.44	7.19	1.67	97.6%	119.3%	115.5%	2.03

Exhibits 13.1 and 13.2 are pricing exercises of the Bayesian framework contrasted against Experience and MLE approaches. We assume that an Insured will select the lower of prices offered in the market. For the 200 simulations, we assume a 0% expense and profit loads. For calculating pure premium, the Experience approach utilizes experience frequency and the actual underlying severity. For the Bayesian approach, severity is a posterior weighted average of the individual severities for each curve. For the MLE approach, severity can be calculated directly. For contrasting the Bayesian and experience approaches, using actual severity for the experience rate is conservative. The simulations have been sorted in descending order of loss ratio and separated into five quintiles. This will allow us to do some testing to determine if the Bayesian model is picking up differences in loss ratio.

Based on the 200 simulations found in Exhibits 13.1 and 13.2 below (I only show the first 10 simulations), if a competitor were to use rates indicated by the experience and MLE approaches they would write less business at a significantly higher loss ratio. We can also see that the Bayesian approach does a reasonably good job at reflecting loss ratio differences. Recall that we assume an expected loss ratio of 100% with a parameter CV of 40%. For risks with the highest loss ratios, top 20%, the model estimate is 130.2%, while the actual is 128.9%. This should alleviate some concern regarding adverse selection. There are 40 simulations in each quintile. The profit or loss is the sum of all simulations in the quintile. The other statistics are averages for the quintile. You can see the Bayesian loss ratio tracks the underlying loss ratio reasonably well, and that a Bayesian UW profit is produced for each quintile.

Exhibit 13.3 is a head to head comparison/competition of all three approaches.

8. Summary and Conclusions

XOL rating generally requires an exposure curve, loss ratio and a credibility procedure to weight experience and exposure estimates. In many instances, at least one and sometimes all three of these parameters are judgmentally selected. Just as experience or maximum likelihood estimates can overreact/overfit to random variation in the data, so can actuarial judgment. When we see higher than expected loss, we tend to pick conservatively. When we see lower than expected, we are optimistic. The author believes that relying on “judgment” may produce results no more accurate, and possibly even less so, than the experience and MLE approaches presented in this paper. Another drawback, approaches that utilize too much judgment are difficult to monitor or re-calibrate.

This paper provides a Bayesian framework for XOL rating. There are many possible variations to the approach provided here each providing possible advantages and/or disadvantages. These variations range from the family of prior distributions, how curves are calibrated, number of layers utilized and the introduction of credibility into the procedure. The approach taken here can be viewed as a starting point, or Barnett approach.

The author believes that a single set of prior severity distributions allow for the development of a model that could be monitored for performance and recalibrated over time (a technical price can be constructed and compared to actual results to potentially recalibrate the model). Based on the volume of losses simulated in the exhibits above, in my opinion it is preferable to develop a single set of prior severity distributions that fit well to a wide range of potential losses (severity distributions) and producing similar results to those above⁵.

Readers may be tempted to develop separate sets of curves for different risks, or LOBs, so I offer a word of caution: try to develop an objective set of standards that clearly define under what circumstances each risk will be assigned to a set of curves. Otherwise, a technical price is difficult to monitor and recalibration becomes problematic. Another issue with separate sets of curves is bias. Presumably, separate sets are developed to capture differences in severity potential. If a risk is not properly assigned, this could lead to potentially greater bias. However, if properly assigned, an estimate with lower MSE is generally found.

It should be noted that biased low estimates can still be profitable relative to traditional pricing techniques, suggesting that narrow priors with some bias is acceptable. However, this is a function of

⁵ In my private work, I have done so. Promulgating prior severity distributions here may be considered anti-competitive.

the relative variance. If a competitor can produce unbiased estimates with comparable variance, it is likely that the biased low portfolio will be unprofitable.

Even if the risk exposure is well known and an exposure curve is available, a Bayesian approach may be preferable if the exposure curve is not properly constructed. Improperly constructed curves may lead to bias in all business priced. A Bayesian approach could be preferable even when a properly constructed exposure curve is available. Recall that an exposure curve may be accurate for all risks in aggregate, but it is unlikely that all risks share the same underlying severity propensity.

The error statistics presented in this paper are a function of known underlying parameters. In practice, we do not know the underlying severity and loss ratio distributions. With the emergence of Bayesian Markov Chain Monte-Carlo (MCMC) models, actuaries can estimate the uncertainty associated with expected outcomes. It is the author's hope that readers can embed the Bayesian approach presented here within an MCMC framework.

On a final note, the technical price should be considered a lower bound. If a competitor price is significantly above the technical price, one need only to price below competitor price to write the business. In a head-to-head competition with experience rating, Bayesian profitability would even be more pronounced than shown in Exhibit 13 above.

Acknowledgments

Friends, family & colleagues.

Bibliography

Meyers, G., On Predictive Modeling for Claim Severity, Casualty Actuarial Society Forum, Spring 2005

Klugman, S., H. Panjer, and G. Willmot, Loss Models: From Data to Decisions, 3rd ed., Hoboken, NJ: Wiley, 2009.

Gelman, A., Prior distributions for variance parameters in hierarchical models. *{\em Bayesian Analysis}* **{\bf 1}**, 515--533

Gelman, A., Informative and noninformative priors: andrewgelman.com/2007/07/18/informative_and/

Boswell, D., Maximum Likelihood vs. Expected Value April 5, 2011: dustwell.com/maximum-likelihood-vs-expected-value.html

Patrik, G. and Mashitz, I, Credibility for Treaty Reinsurance Excess Pricing, 323—374, Casualty Actuarial Society Discussion Paper Program, Spring 1990

A. Appendix: MLE estimates

MLE: Lognormal Parameters

	Fit	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\hat{\mu}$	9.328									
$\hat{\sigma}$	2.484									
		Emergence	$F(c_{i-1};\hat{\theta})$	$F(c_i;\hat{\theta})$	layer probabilities		(5)/		Reported	negative-
		Pattern	lower	upper	$[F(c_i;\hat{\theta}) - F(c_{i-1};\hat{\theta})]$	(4)x(1)	constant	$\ln(6)$	Claims	log likelihood
c_{i-1}	c_i				$[1-F(100,000;\hat{\theta})]$				n_i	$-(7)*(8)$
100,000	135,000	0.896	0.810	0.841	0.1634	0.146	0.177	-1.731	17	29.421
135,000	185,000	0.881	0.841	0.870	0.1516	0.133	0.162	-1.823	13	23.696
185,000	250,000	0.864	0.870	0.894	0.1260	0.109	0.132	-2.026	10	20.260
250,000	500,000	0.843	0.894	0.937	0.2248	0.190	0.229	-1.472	15	22.082
500,000	1,000,000	0.812	0.937	0.965	0.1472	0.120	0.145	-1.933	9	17.398
1,000,000	2,000,000	0.733	0.965	0.981	0.0892	0.065	0.079	-2.537	11	27.902
2,000,000	5,000,000	0.671	0.981	0.993	0.0605	0.041	0.049	-3.015	5	15.073
5,000,000	10,000,000	0.600	0.993	1.000	0.0373	0.022	0.027	-3.610	1	3.610
				Totals:	1.0000	0.826	1.000		81	159.442
						constant				

MLE: Ground Up Loss Ratio

Estimated on-level premium ALL YEARS	60,000,000	(1)								
Ground up Loss Ratio (MLE):	154.0%	(2)								
Excess Loss Percentage:	81.9%	(3); See Exhibit 2 (2), Column (6)								
Excess Loss Rate :	126.2%	(4); (2)x(3)								
Excess Loss:	75,697,369	(5); (4)x(1)								
Excess Severity:	772,129	(6); See Exhibit 2 (2), Column (3)								
Expected Excess Claims:	98.04	(7); (5)/(6)								
Probability of attaching:	19.0%	(8); See Exhibit 2 (2), Column (5)								
Expected Ground Up Claims (over 10 year period):	517.1	(9); (7)/(8)								
Assumed Contagion:	1.0%	(10)								
Assumed VTM Ground Up:	1.517	(11); $1+(10)x(9)/10$								
B (scale parameter gamma distribution):	1.934	(12); $1/[(11)-1]$								
A (scale parameter gamma distribution):	100	(13); $(9)x(12)/10$								
			Layer1	Layer2	Layer3	Layer4	Layer5	Layer6	Layer7	Layer8
Curve #1_Probability of attaching [q]:			19.0%	15.9%	13.0%	10.6%	6.3%	3.5%	1.9%	0.7%
B' = B/q			10.20	12.19	14.89	18.25	30.52	54.56	104.39	273.78
VTM:			1.098	1.082	1.067	1.055	1.033	1.018	1.010	1.004
										(14); (12)/[(14) [by layer]
Layer Probability:			0.163	0.152	0.126	0.225	0.147	0.089	0.060	0.037
Emergence Pattern [% Reported]:			89.6%	88.1%	86.4%	84.3%	81.2%	73.3%	67.1%	60.0%
Expected Reported:			14.4	13.1	10.7	18.6	11.7	6.4	4.0	2.2
Actual Reported Claims (simulated):			17.0	13.0	10.0	15.0	9.0	11.0	5.0	1.0
										(17); (7)x(17)x(18)
Negative Binomial pdf:			0.0739	0.1057	0.1189	0.0705	0.0932	0.0314	0.1545	0.2450
Normalizing constants (adjusted layer probs):			0.177	0.162	0.132	0.229	0.145	0.079	0.049	0.027
log likelihood :			-4.34	-4.07	-4.16	-4.12	-4.31	-6.00	-4.88	-5.02
										(21); $\ln[(21)x(22)]$
Negative loglikelihood:			36.888							
										(24); $-\text{sum}[(23)]$