Approximating the Aggregate Loss Distribution

by Dmitry E. Papush, Aleksey S. Popelyukhin and Jasmine G. Zhang

Abstract:

Aggregate loss distributions have extensive applications in actuarial practice. Several approaches have been suggested to estimate the aggregate loss distribution, including the Heckman – Meyers method, Panjer algorithm, and Fast Fourier transformation, to name a few. All of these methods rely on separate assumptions about frequency and severity components of the aggregate losses. Quite often however, it is not practical to obtain frequency and severity expectations independently, and only aggregate information is available for analysis. In this case, the a priori assumption about the "shape" of the aggregate loss distribution becomes critical, especially for assessing the probability of very high aggregate loss values, in the "tail."

The goal of this work was to determine which statistical two-parameter distribution, out of several, would serve best to approximate aggregate loss distributions for property and casualty products. We focused on ground up losses limited by a per-occurrence limit. These results are relevant for quota share agreements. In addition, we considered layer losses, the results of which are important for umbrella quota share transactions.

We simulated samples of aggregate loss, fit statistical distributions to the samples, and then used goodness-of-fit tests to determine the best-fitting distribution. In all realistic scenarios with limited losses, we found that the Gamma distribution uniformly provided the most reasonable approximation to the aggregate loss.

Keywords: aggregate models, aggregate loss, collective risk model, compound distribution, simulation, Gamma distribution

1. Introduction.

Aggregate loss distributions have extensive applications in actuarial practice. The modelling of aggregate losses is one of the main fundamentals of actuarial work, as it is important to business decisions regarding many aspects of insurance and reinsurance contracts. The purpose of this study was to determine the best statistical distribution to approximate the aggregate loss distribution for property and casualty business with applications to quota share agreements.

When separate data on loss frequency and loss severity distributions is available, actuaries can approximate the aggregate loss distribution using methods such as the Heckman-Meyers method (Heckman and Meyers 1983), Panjer method (Panjer 1981), Fast Fourier transform (Robertson 1992), and stochastic simulations (Mohamed and Ismail 2010). However, sometimes only aggregate information is available for analysis. In this case, the choice of the shape of the aggregate loss distribution becomes very important, especially in the "tail" of the distribution. The "tail" is often the part of the distribution that is most affected by policy limits and a failure to model the tail correctly could lead to an overestimation of the discount given to loss ratios due to the application of policy limits.

Previously published papers debate the appropriateness of various aggregate loss distributions. Dropkin (1964) and Bickerstaff (1972) showed that the Lognormal distribution closely approximates certain types of homogenous loss data. Pentikainen (1987) suggested an improvement of the Normal approximation using so-called NP-method. He compared this method with the Gamma approximation and concluded that both methods yield reasonable approximations when the skewness of aggregate losses is less than 1, but that neither method is accurate when the skewness is greater than 1. Venter (1983) suggested the Transformed Gamma and Transformed Beta distributions for the approximation of aggregate loss, while Chaubey et al. (1998) suggested the Inverse Gaussian distribution.

Papush et al. (2001) analyzed several simulated samples of aggregate losses and compared the fit of the Normal, Lognormal, and Gamma distributions to the simulated data in the tails of the distributions. This was a deviation from previous research, which was based solely on theoretical considerations. In all seven scenarios that were tested by Papush et al., the Gamma distribution performed the best. Therefore, they recommended the Gamma as the most appropriate approximation of aggregate loss.

This research has expanded the 2001 study but retained the same general approach.

2. Method:

Overview of the Study

Initially, proceeding as in Papush et al. (2001), we limited our consideration to two-parameter probability distributions. Three-parameter distributions often provide a better fit, but observed data is often too sparse to reliably estimate a third parameter. We compared the fit of five candidate distributions, which are shown in Table 1 below.

Table 1: Distributions Used for the Approximation of Aggregate Loss

Distribution	Parameters	Probability Density Function	Mean	Variance
Normal	μ - location $\sigma > 0$ - scale	$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$	μ	σ^2
Logistic	μ - location $s > 0$ - scale	$\frac{e^{-\frac{x-\mu}{s}}}{s\left(1+e^{-\frac{x-\mu}{s}}\right)^2} = \frac{1}{4s}\operatorname{sech}^2\left(\frac{x-\mu}{2s}\right)$	μ	$\frac{s^2\pi^2}{3}$
Gamma	$\alpha > 0$ - shape $\beta > 0$ - rate	$\frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}$	$\frac{\alpha}{\beta}$	$\frac{\alpha}{\beta^2}$
Inverse Gauss	$\mu > 0$ - location $\lambda > 0$ - shape	$\left[\frac{\lambda}{2\pi x^3}\right]^{1/2} \exp\left\{\frac{-\lambda(x-\mu)^2}{2\mu^2 x}\right\}$	μ	$\frac{\mu^3}{\lambda}$
Lognormal	μ - scale $\sigma > 0$ - shape	$\frac{1}{x} \cdot \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right)$	$e^{(\mu+\sigma^2/2)}$	$e^{(2\mu+\sigma^2)}(e^{\sigma^2}-1)$

Our analysis procedure is summarized in the following formal steps:

1. Choose frequency distribution and obtain severity distributions from ISO's circulars or loss submission data
2. Simulate the number of claims (N) and the individual loss amounts $(X_1, \dots, X_N)$, put the individual loss amounts into per-occurrence layers $(X_1^l, \dots, X_N^l)$, and calculate the corresponding aggregate loss ($S^l = \sum_{i=1}^{N} X_i^l$) in each layer l
3. Repeat this analysis many times (50,000) to obtain a sample of aggregate loss
4. Fit the parameters of different candidate probability distributions
5. Test the goodness of fit of the distributions and compare results

Choice of Software

We used the open-source statistical software tool R (R Core Team 2017) to perform our analysis. R is widely used in different actuarial contexts and is also popular in the sciences and social sciences.

In addition to the base version, R allows users to develop packages of functions and compiled code and upload them to the Comprehensive R Archive Network (CRAN). These packages can be freely downloaded for use. In our project, we used packages **lhs** (Carnell 2016), **fitdistrplus** (Delignette-Muller and Dutang 2015), **NORMT3** (Nason 2012), **gsl** (Hankin 2006), **actuar** (Dutang et al, 2008) and **e1071** (Meyer et al, 1998).

Selection of Frequency Distribution

In actuarial science, the Poisson distribution is commonly used to represent the frequency of insurance claims. It has a memoryless property, i.e. the number of claims in any time interval should not affect the number of claims in any other interval. This is a good approximation to what we observe in real data on claim frequency. We selected different mean frequencies (λ's) to model small, large, and in some cases medium books of business. The selected mean frequencies can be seen in Tables 2a and 2b on the next page.

Selection of Severity Distributions

For casualty products, we used curves developed by actuaries of the Insurance Services Office (ISO). The majority of ISO's curves are Mixed Exponential distributions, which can be represented as sums of exponentials $f(x) = \sum_i w_i \frac{1}{\theta_i} e^{\frac{x}{\theta_i}}$. Each Exponential distribution in the mixture has a different mean θ_i and a weight w_i, with the weight corresponding to the probability of the individual Exponential distribution being chosen.

The table below shows the severity distributions we considered for casualty products, along with the means of the claim counts used, and the per-occurrence layers we divided the individual losses into. Additional detail may be found in the section "Layer Descriptions" below.

Table 2a: Distributions used for Casualty Products

Distribution	Type of Distribution	Mean frequency, λ	Per-Occurrence Layers
General Liability: Premises and Operations	Mixed Exponential (mean 35K)	100, 500, 1000	250K x 0, 500K x 0, 1M x 0 750K x 250K, 500K x 500K, 4M x 1M
General Liability: Products	Mixed Exponential (mean 135K)	Same as above	Same as above
Commercial Auto	Mixed Exponential (mean 45K)	Same as above	Same as above
E&O - Medium Lawyers	Lognormal (mean 250K)	50, 500	1M x 0, 5M x 0
D&O - Public - Non F500	Lognormal (mean 1.3M)	50, 500	10M x 25M

To model typical Small Commercial, Middle Market, and Large Commercial books of Non-Cat Property business, we chose distributions derived from ISO data for increasing amount of insurance (AOI) ranges. In addition, we simulated distributions based on representative samples of real-life losses.

Table 2b: Distributions used for Property, Non-Cat

Distribution	Type of Distribution	Mean frequency, λ	Layers
Small Commercial: AOI 5M to 6M	Mixed Exponential (mean 95K)	100, 500	1M x 0, Unlimited
Middle Market: AOI 25M to 30M	Mixed Exponential (mean 175K)	Same as above	Same as above
Large Commercial: AOI 100M to 125M	Mixed Exponential (mean 285K)	Same as above	Same as above
Non-Cat Property 1	Loss Sample (mean 100K)	Same as above	Same as above
Non-Cat Property 2	Loss Sample (mean 110K)	Same as above	Same as above

Simulation Method

We used Latin Hypercube Sampling (LHS) to sample frequencies from the Poisson distribution and severities from each exponential component of the mixed exponential distributions before applying corresponding weights. We chose LHS over Monte Carlo Simulation because it spreads sample points more evenly across all possible values so that samples drawn using LHS are more representative of the real variability in frequency or severity. In particular, since we were interested in studying the "tail" of the distribution for this study, LHS ensured that our simulation contained a reasonable sampling of high values. We implemented LHS using the *randomLHS()* function in the 'lhs' package in R.

We used bootstrapping, the base function *sample()* in R, to choose which exponential component of the mixed exponential distribution to use and to bootstrap losses from the property loss submissions. This function allows the bootstrapping procedure to run very efficiently.

We sampled the losses from the loss submissions without replacement because we believe this provided a better representation of a true "year" of data: the same loss to the same property should not occur in the same year. Sampling without replacement also ensures that we do not obtain too many small losses in each year's sample, and thus prevents the understatement of aggregate loss.

Layer Descriptions

We divided our simulated individual losses into the per-occurrence layers in the included in Tables 2a and 2b above. We determined the amount of penetration of each simulated loss within a layer by calculating:

$$Loss\ in\ Layer = Min(Max\big((Aggregate\ Loss - RETENTION), 0\big), LIMIT),$$

where $RETENTION$ is the lower bound of a layer, possibly equal to zero, and $LIMIT$ is the width of the layer. For instance, for the layer of \$750K excess of \$250K, $RETENTION$ would be \$250,000 and $LIMIT$ would be \$750,000. For the layer \$1M excess of \$0, $RETENTION$ would be \$0 and $LIMIT$ would be \$1,000,000.

If none of the claims in one simulation penetrated one of the excess layers, i.e. "\$4M xs \$1M", the aggregate loss for that simulation was zero. This created a mass at zero in our distribution. Thus, the aggregate loss distributions we fitted were of the form:

$$p_0 + p_1 * Candidate\ Distribution,$$

where $p_0 = \Pr\{Aggregate\ Loss = 0\}$, $p_1 \equiv 1 - p_0$ and

$$Candidate\ Distribution \in \{Normal, Logistic, Gamma, Inverse\ Gauss, Lognormal\}$$

Estimating the Number of Simulations to Run

Many studies (e.g. Papush 1997) use the Central Limit Theorem to estimate the number of simulations necessary to achieve a reasonable degree of accuracy in the estimate of the mean. However, this study concentrates on approximating the "tail" of the distribution. To the best of our knowledge, there is unfortunately no established method for finding the number of simulations to run when the values of interest relate to the "tail" of a distribution. To ensure that we achieved a reasonable degree of accuracy of our results, we followed the method described below.

We started by running 1,000,000 simulations and calculating the 99th percentile of the resulting aggregate distribution. We then re-ran the same 1,000,000 simulations this time monitoring two metrics as we went through the simulation process. After each increment of 2,500 simulations we measured two quantities: the change in the 99th percentile of the distribution and the difference between the 99th percentile of the distribution from the cumulative set of simulations, compared to the 99th percentile of the distribution obtained as the result of our first run.

Using this approach, we were able to verify the following two assumptions. First, as the mean frequency of the Poisson distribution used in our simulation increases, i.e. the size of our book, the accuracy of the 99th percentile will also increase. Second, as the per occurrence limit of the coverage decreases, the accuracy of the 99th percentile will increase.

Due to the these two assumptions, we only show results for small-book scenarios with high per-occurrence limits. Under 50,000 simulations, we have found that both metrics suggested an error of 1% or below for nearly all scenarios. For instance, we found that for premises and operations with a $1M per occurrence limit and mean frequency of 100, we had an error of 0.8% after 25,000 simulations, while 50,000 simulations resulted in an error of 0.6%. Graphs of the two metrics in this scenario (by number of simulations used) are shown below. It is noteworthy to mention that the error was over 3% in only one scenario.

Fig 1: 99th Percentile simulation error

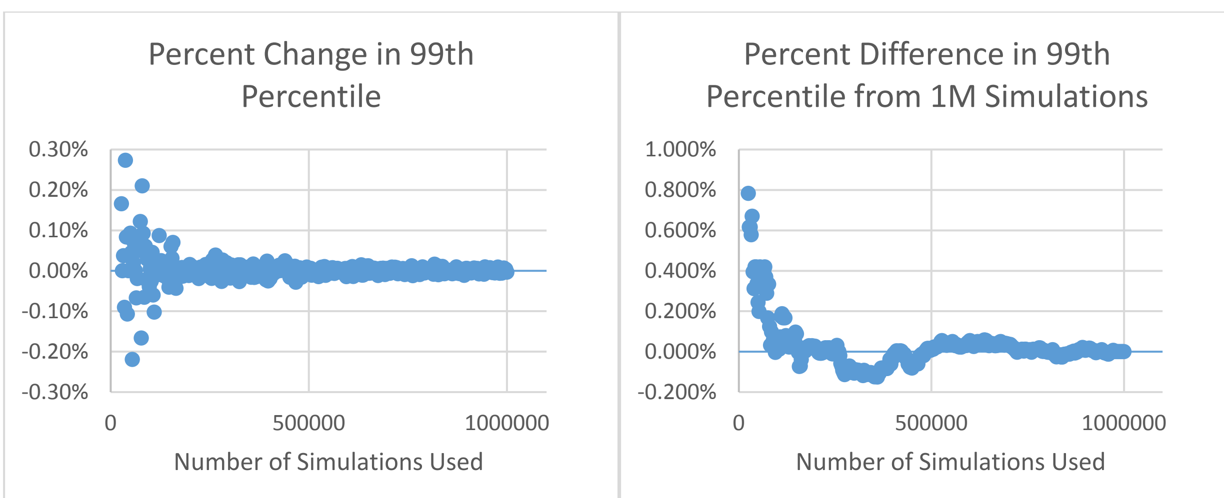

Parameter Estimation

Parameter estimation was implemented using the function *fitdist()* in the R package 'fitdistrplus' (Delignette-Muller and Dutang 2015). Initially, we used both the Maximum Likelihood Method and the Method of Moments to estimate parameters for approximating distributions. The parameter estimates for the two methods were similar to one another, but the parameters from the Method of Moments yielded a better-fitting distribution as measured by both the Percentile Matching and the Expected Excess Value tests. Therefore, we chose to use the Method of Moments.

Testing Goodness of Fit

Once we simulated the sample of aggregate losses and estimated the parameters for the distributions, we tested the goodness of fit. We created three tests to compare distributions in their "tails." These tests are also relevant to reinsurance pricing.

The aggregate features of proportional reinsurance treaties are usually expressed in terms of the loss ratios, which we substituted by calculating percentages of the mean of the simulated aggregate distribution (e.g. 100% of mean, 125% of mean, etc.). In other words, we expressed values x in the "tail" of the distributions in terms of the percentages p of means of the simulated distributions: $x = p * mean$.

The Percentile Matching Test compares the survival functions $Prob\{X > x\}$ of distributions at various values of the argument until the distributions effectively vanish. This test gives a transparent indication of where two distributions are different and by how much. For various percentages of the mean, we tested how much the survival functions $Prob\{X > x\}$ in the fitted distributions differed from those of the simulated sample.

The Excess Expected Loss Cost Test compares the conditional means of distributions in excess of different amounts. Specifically, it evaluates the conditional expectations $E[X - x|X > x] * Prob\{X > x\}$ for different values of x. These values are important for both the ceding company and the reinsurance carrier when considering aggregate loss ratio caps, stop loss coverage, annual aggregate deductible coverage (AAD), profit commission, sliding scale commission, and other types of aggregate reinsurance transactions with loss adjustable features. For instance, in the case of an aggregate loss ratio cap, it is important to accurately price the discount given to the cedent based on the cap. For various percentages of the mean, we test the percentage error in the Excess Expected Loss Cost in the fitted distributions as compared to the Excess Expected Loss Cost in the simulated sample. We call this "Error in Pricing of Aggregate Stop Loss" in the tables and charts presented in the Appendix.

Finally, we estimated the accuracy of the five candidate distributions in pricing loss corridors. In other words, we test the amounts $E[X^\wedge|x_2 > X > x_1] * Prob\{x_2 > X > x_1\}$, where X^ is loss in an aggregate layer, namely $\max(\min(X - x_1, 0), x_2 - x_1)$. In a loss corridor, the reinsurer returns the responsibility for losses between the two loss ratios x_1 and x_2 to the primary insurer.

3. Results and Conclusion:

To illustrate the results of our study we show the characteristics of the frequency, denoted "Size," and name, denoted "LOB," of severity distributions selected in each scenario, the mean aggregate loss in each layer, and the results of the three goodness-of-fit tests, see Appendix. The first column of each table shows the results of each test on the simulated data. In the top portion of each table, the other five columns show the difference in the percentile matching test for the various candidate distributions. The second and third portions in our tables show the error as a percent of mean aggregate loss for the different distributions in the Excess Expected Loss Cost Test and Error in Loss Corridor Pricing, respectively. The graphs show the results from the first two goodness-of-fit tests. One can clearly see that in every example Gamma (represented by a red line on charts and green background in the tables) shows the smallest error.

Surely, only 6 examples with measurements at just 6 points shown in Appendix should not be considered a very convincing argument. That is why we ran our tests for a large sample of reasonable severity distributions (253), for several reasonable layers (13), and for a few reasonable portfolio sizes (3). Among the severity distributions selected for the study there were several closed-form ones (ISO's PremOps Table 1, for example) as well as some empirical ones (losses from a large client's submission). In a latter case we use bootstrapping to generate an aggregate loss distribution. We measured results at every 5 percentage points from a minimum of 75% of Mean to a maximum of 250% of Mean.

We found that the Gamma distribution provides a fit that is almost always the best for both ground up and excess layers.

4. Additional Comments

Our choices of potential candidates for the aggregate distribution, see Table 1, were not, in fact, random. We chose the Normal distribution as a limiting distribution of sums of identically distributed independent losses. We considered the Gamma distribution as a distribution of sums of identical exponentials.[1] We included the Lognormal distribution as it is a popular distribution used to approximate the sum of unlimited losses. As we analyzed our choices through the differences in higher moments of the aggregate loss distribution[2] (since the first two moments were matched), we

---

[1] Indeed, the sum of identical exponentials has characteristic function $(1-\frac{it}{\lambda})^{-n}$ which is $Gamma(n,\lambda)$

[2] According to Pentikainen (1987), as long as the severity distribution is restricted to a limited interval, aggregate loss distributions with several matching moments approximate each other acceptably well.

decided to consider also the Logistic distribution, whose Skewness and Excess Kurtosis lie between that of the Normal and the Gamma distributions, as well as the Inverse Gaussian distribution, whose Skewness and Excess Kurtosis lie between that of the Gamma and the Lognormal distributions.

The following Table 3 demonstrates the behavior of higher moments of the listed distributions expressed in terms of their CVs, coefficients of variation. This table is important as it provides a comparison of the shape of the different theoretical curves with the same first two moments.

Table 3: Higher Moments (Skewness and Excess Kurtosis) in terms of CV

Distribution	**CV**	**Skewness**	**Ex. Kurtosis**
Normal	c	0	0
Logistic	c	0	1.2
Gamma	c	$2c$	$6c^2$
Inverse Gauss	c	$3c$	$15c^2$
Lognormal	c	$c + c^3$	$16c^2 + 15c^4 + 6c^6 + c^8$

Given that we narrowed our consideration to two-parameter distributions, we could only match the first two moments of the empirical distribution by varying parameters of the theoretical one. Consequently, the quality of the approximation depends mainly on how close the higher moments of the theoretical distribution are to the corresponding higher moments of the empirical distribution. Therefore, to decide which theoretical distribution is a better approximation, it is helpful to estimate the ratio of skewness, and, possibly, kurtosis, to the CV.

In the case of Poisson frequency and known severity assumptions it is possible to calculate these ratios exactly.[3] For example, the following Figure 2 illustrates the behavior of the Skewness-to-CV ratio using a GL Products severity curve along with a Poisson frequency distribution for different policy limits. The results of the empirical simulation are labeled as CV_agg, and the other three lines are derived from theoretical calculations. As one can see, the ratio of skewness to CV was much closer to 2 than to 3 in the simulation. We could expect the Gamma distribution to have a better fit than other distributions because its CV is 2.

[3] Indeed, for Poisson distributed number of claims with mean λ:

$$E[Agg] = \lambda E[Sev], \qquad Var[Agg] = \lambda E[Sev^2], \qquad CV[Agg] = \lambda^{-1/2}(E[Sev^2])^{1/2}/E[Sev],$$
$$Skew[Agg] = \lambda^{-1/2}E[Sev^3]/(E[Sev^2])^{3/2}$$

Fig 2: Skewness/CV as a Function of Primary Limit

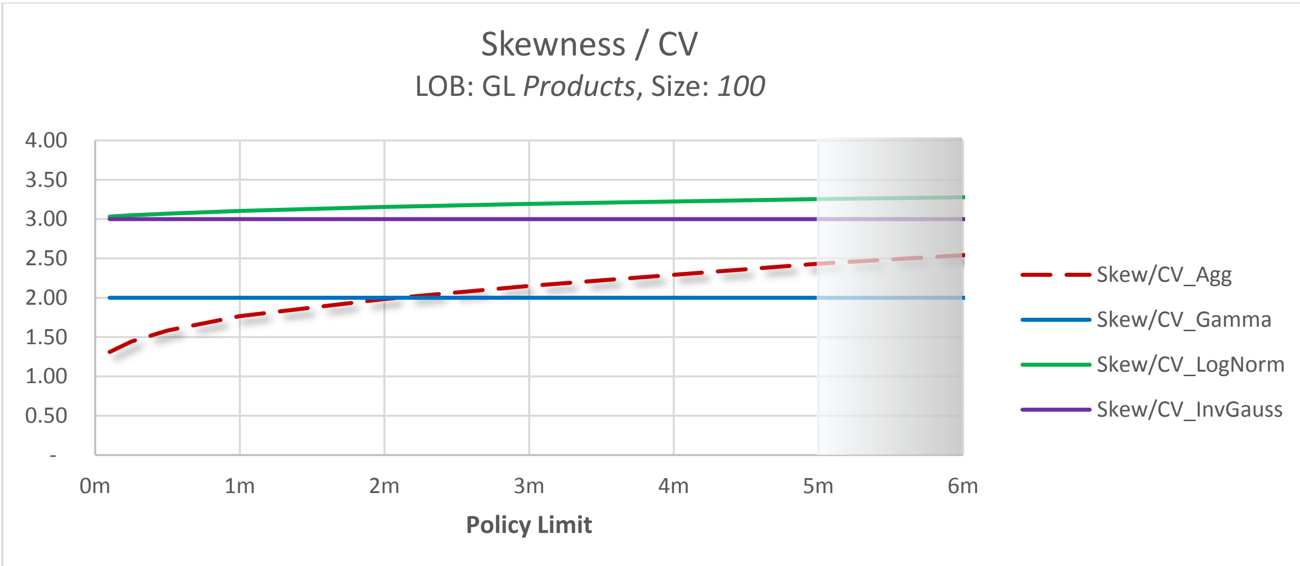

We ran a series of similar tests for a multitude of combinations of appropriate severity curves, layers and portfolio sizes, 253 x 13 x 3 scenarios overall. In deciding which theoretical distribution provides the best approximation to the simulated, empirical, one, we paid special attention to matching the "tail" of the distributions.

We were encouraged to see that closeness of higher moments translated to a good fit and that the overwhelming majority of lines of business can be, for all practical purposes, well approximated by a Gamma distribution. The same would hold true not only for a series of individual severity curves, but also for a mix of several of them. ***These results lead us to the general conclusion that the Gamma distribution provides a uniformly reasonable approximation to the aggregate loss on the interval from the mean to at least two means of the aggregate distribution.***

Appendix.

Examples of Ground Up Layers

LoB: 'Commercial Auto', Size: 100, Layer: 500x0

Mean Aggregate Loss: $2,038,218

Percentile matching LoB: 'Commercial Auto', Size: 100, Layer: 500x0

	P(X>x)					*difference*
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	70%	-0.18%	2.52%	4.74%	0.09%	-0.47%
90%	58%	-2.07%	3.40%	4.87%	-0.62%	-2.26%
100%	46%	-2.77%	3.79%	3.79%	-0.73%	-2.80%
110%	35%	-2.90%	3.02%	1.55%	-0.86%	-2.79%
120%	26%	-2.09%	2.04%	-0.18%	-0.46%	-1.89%

Error in Pricing of Aggregate Stop Loss LoB: 'Commercial Auto', Size: 100, Layer: 500x0

From	*To*	*E[X-x| X>x]*P(X>x)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	Inf	512,046	-0.76%	0.71%	0.41%	-0.19%	-0.73%
90%	Inf	381,743	-0.65%	0.40%	-0.09%	-0.17%	-0.59%
100%	Inf	275,393	-0.39%	0.03%	-0.54%	-0.10%	-0.33%
110%	Inf	192,640	-0.11%	-0.32%	-0.82%	-0.02%	-0.05%
120%	Inf	130,463	0.15%	-0.57%	-0.87%	0.05%	0.19%

Error in Corridor Pricing LoB: 'Commercial Auto', Size: 100, Layer: 500x0

From	*To*	*E[X^| x_2>X>x_1]*P(x_2>X>x_1)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
75%	80%	73,880	-0.01%	-0.11%	-0.22%	-0.01%	0.00%
80%	90%	130,303	0.11%	-0.31%	-0.50%	0.02%	0.14%
90%	100%	106,350	0.25%	-0.37%	-0.44%	0.07%	0.26%
100%	110%	82,753	0.28%	-0.36%	-0.28%	0.07%	0.28%
110%	120%	62,177	0.26%	-0.25%	-0.06%	0.07%	0.24%

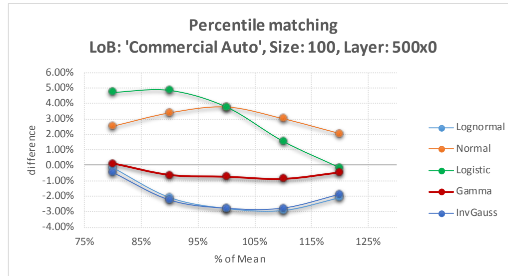

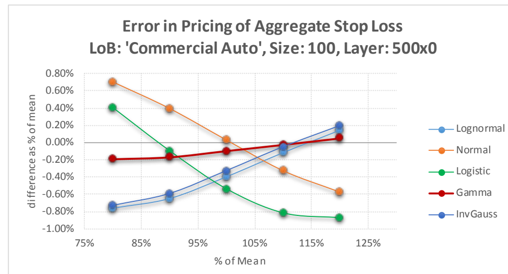

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

LoB: 'GL: Products', Size: 500, Layer: 1000x0

Mean Aggregate Loss: $38,737,178

Percentile matching LoB: 'GL: Products', Size: 500, Layer: 1000x0

	P(X>x)					*difference*
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	97%	0.56%	-0.55%	-0.60%	0.18%	0.57%
90%	81%	0.05%	0.20%	2.28%	0.03%	0.02%
100%	49%	-0.98%	1.26%	1.26%	-0.24%	-0.98%
110%	19%	-0.57%	-0.11%	-2.19%	-0.36%	-0.54%
120%	4%	0.47%	-0.39%	-0.34%	0.22%	0.47%

Error in Pricing of Aggregate Stop Loss LoB: 'GL: Products', Size: 500, Layer: 1000x0

From	*To*	*E[X-x| X>x]*P(X>x)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	Inf	7,795,920	-0.04%	0.05%	0.12%	-0.01%	-0.04%
90%	Inf	4,289,446	-0.08%	0.08%	0.06%	-0.02%	-0.08%
100%	Inf	1,744,726	-0.03%	-0.01%	-0.20%	-0.02%	-0.03%
110%	Inf	475,856	0.07%	-0.07%	-0.10%	0.02%	0.07%
120%	Inf	80,948	0.06%	-0.04%	0.03%	0.03%	0.06%

Error in Corridor Pricing LoB: 'GL: Products', Size: 500, Layer: 1000x0

From	*To*	*E[X^| x_2>X>x_1]*P(x_2>X>x_1)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
75%	80%	1,900,012	-0.02%	0.02%	0.04%	-0.01%	-0.02%
80%	90%	3,506,474	-0.04%	0.04%	-0.06%	-0.01%	-0.04%
90%	100%	2,544,720	0.05%	-0.09%	-0.26%	0.01%	0.05%
100%	110%	1,268,870	0.10%	-0.06%	0.10%	0.04%	0.09%
110%	120%	394,908	-0.01%	0.03%	0.13%	0.00%	-0.01%

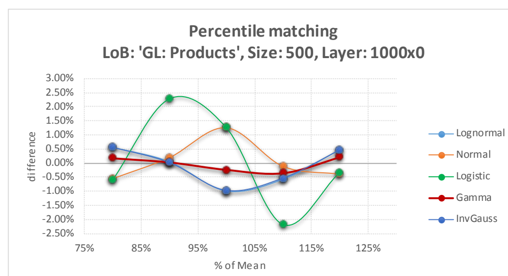

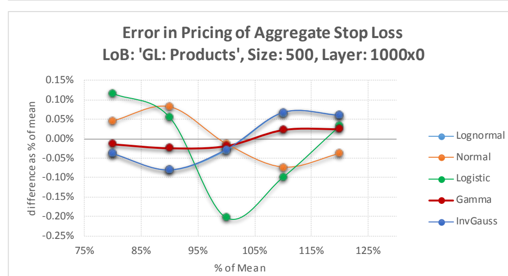

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

LoB: 'GL: PremOps ', Size: 1000, Layer: 250x0
Mean Aggregate Loss: $19,888,610

Percentile matching LoB: 'GL: PremOps ', Size: 1000, Layer: 250x0

	P(X>x)					*difference*
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	100%	0.13%	-0.13%	-0.55%	0.06%	0.13%
90%	91%	0.34%	-0.22%	0.83%	0.12%	0.33%
100%	49%	-0.74%	0.79%	0.79%	-0.23%	-0.74%
110%	10%	0.17%	-0.18%	-1.23%	0.08%	0.18%
120%	1%	0.13%	-0.20%	0.22%	0.01%	0.13%

Error in Pricing of Aggregate Stop Loss LoB: 'GL: PremOps ', Size: 1000, Layer: 250x0

From	*To*	*E[X-x| X>x]*P(X>x)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
80%	Inf	3,978,997	0.00%	0.01%	0.03%	0.00%	0.00%
90%	Inf	2,051,967	-0.03%	0.04%	0.07%	-0.01%	-0.03%
100%	Inf	610,646	0.00%	0.00%	-0.13%	0.00%	0.00%
110%	Inf	76,850	0.03%	-0.03%	0.00%	0.01%	0.03%
120%	Inf	3,840	0.01%	-0.01%	0.02%	0.00%	0.01%

Error in Corridor Pricing LoB: 'GL: PremOps ', Size: 1000, Layer: 250x0

From	*To*	*E[X^| x_2>X>x_1]*P(x_2>X>x_1)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
75%	80%	993,216	0.00%	0.00%	0.02%	0.00%	0.00%
80%	90%	1,927,030	-0.03%	0.03%	0.04%	-0.01%	-0.03%
90%	100%	1,441,321	0.03%	-0.03%	-0.19%	0.01%	0.03%
100%	110%	533,796	0.03%	-0.04%	0.12%	0.01%	0.03%
110%	120%	73,010	-0.03%	0.03%	0.02%	-0.01%	-0.03%

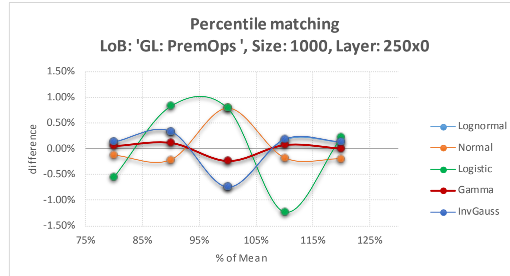

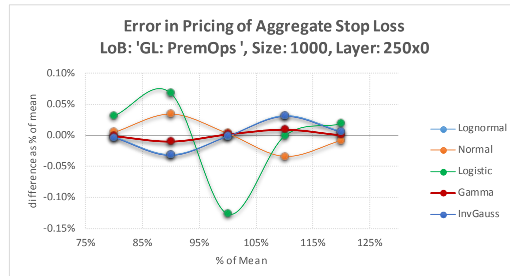

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

Examples of Excess Layers

LoB: 'Commercial Auto', Size: 100, Layer: 750x250
Mean Aggregate Loss: $1,015,045

Percentile matching LoB: 'Commercial Auto', Size: 100, Layer: 750x250

	P(X>x)					*difference*
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	42%	-4.16%	7.96%	7.96%	-1.09%	-4.41%
110%	37%	-5.14%	6.85%	6.07%	-1.72%	-5.12%
125%	31%	-6.04%	4.49%	2.76%	-2.58%	-5.71%
150%	20%	-3.65%	2.78%	0.52%	-0.86%	-3.08%
200%	9%	-1.33%	-1.84%	-2.43%	-0.33%	-0.85%

Error in Pricing of Aggregate Stop Loss LoB: 'Commercial Auto', Size: 100, Layer: 750x250

From	*To*	*E[X-x| X>x]*P(X>x)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	Inf	273,255	-2.68%	0.24%	-0.90%	-0.78%	-2.18%
110%	Inf	233,028	-2.22%	-0.50%	-1.61%	-0.64%	-1.70%
125%	Inf	181,017	-1.37%	-1.36%	-2.27%	-0.31%	-0.88%
150%	Inf	112,924	0.18%	-1.95%	-2.33%	0.44%	0.55%
200%	Inf	42,599	1.35%	-2.05%	-1.67%	0.68%	1.44%

Error in Corridor Pricing LoB: 'Commercial Auto', Size: 100, Layer: 750x250

From	*To*	*E[X^| x_2>X>x_1]*P(x_2>X>x_1)*				*difference as % of mean*	
% of Mean	*% of Mean*	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
90%	100%	45,231	0.34%	-0.84%	-0.88%	0.07%	0.39%
100%	110%	40,227	0.47%	-0.74%	-0.70%	0.14%	0.48%
110%	125%	52,011	0.85%	-0.86%	-0.66%	0.33%	0.82%
125%	150%	68,093	1.55%	-0.59%	-0.06%	0.76%	1.43%
150%	200%	70,325	1.16%	-0.10%	0.66%	0.24%	0.89%

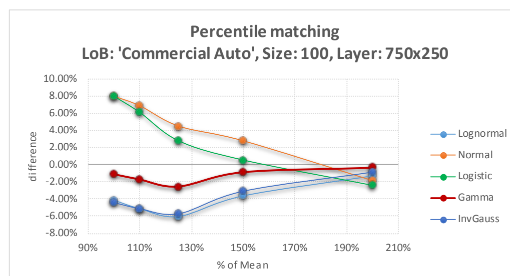

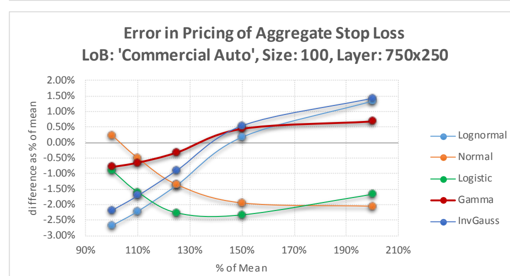

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

LoB: 'GL: Products', Size: 500, Layer: 4000x1000

Mean Aggregate Loss: $15,097,544

Percentile matching LoB: 'GL: Products', Size: 500, Layer: 4000x1000

	P(X>x)					difference
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	47%	-4.37%	3.47%	3.47%	-2.01%	-4.43%
110%	37%	-4.30%	3.06%	1.82%	-1.84%	-4.18%
125%	25%	-3.14%	1.73%	-0.51%	-1.20%	-2.88%
150%	12%	-0.65%	-0.49%	-1.77%	0.00%	-0.39%
200%	2%	0.88%	-0.92%	-0.47%	0.43%	0.90%

Error in Pricing of Aggregate Stop Loss LoB: 'GL: Products', Size: 500, Layer: 4000x1000

From	To	E[X-x| X>x]*P(X>x)				difference as % of mean	
% of Mean	% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	Inf	2,474,087	-0.71%	0.04%	-0.66%	-0.20%	-0.59%
110%	Inf	1,841,240	-0.26%	-0.29%	-0.92%	0.00%	-0.15%
125%	Inf	1,133,488	0.32%	-0.65%	-1.00%	0.25%	0.40%
150%	Inf	455,191	0.79%	-0.77%	-0.64%	0.39%	0.80%
200%	Inf	50,799	0.55%	-0.23%	-0.06%	0.20%	0.50%

Error in Corridor Pricing LoB: 'GL: Products', Size: 500, Layer: 4000x1000

From	To	E[X^| x_2>X>x_1]*P(x_2>X>x_1)				difference as % of mean	
% of Mean	% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
90%	100%	776,522	0.41%	-0.34%	-0.40%	0.20%	0.43%
100%	110%	632,847	0.44%	-0.33%	-0.26%	0.20%	0.44%
110%	125%	707,752	0.58%	-0.36%	-0.08%	0.24%	0.55%
125%	150%	678,297	0.47%	-0.13%	0.36%	0.15%	0.40%
150%	200%	404,392	-0.24%	0.54%	0.58%	-0.19%	-0.30%

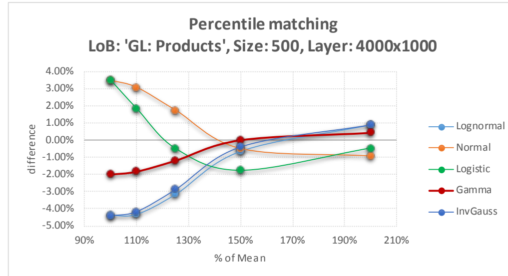

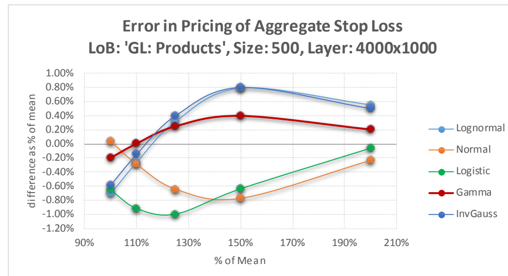

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

LoB: 'GL: PremOps ', Size: 1000, Layer: 500x500

Mean Aggregate Loss: $3,273,975

Percentile matching LoB: 'GL: PremOps ', Size: 1000, Layer: 500x500

	P(X>x)					difference
% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	47%	-4.39%	2.55%	2.55%	-2.25%	-4.43%
110%	37%	-4.15%	2.19%	0.80%	-1.99%	-4.04%
125%	23%	-2.53%	1.06%	-1.20%	-1.01%	-2.31%
150%	9%	0.14%	-0.69%	-1.48%	0.34%	0.31%
200%	1%	0.83%	-0.44%	-0.07%	0.39%	0.81%

Error in Pricing of Aggregate Stop Loss LoB: 'GL: PremOps ', Size: 1000, Layer: 500x500

From	To	E[X-x| X>x]*P(X>x)				difference as % of mean	
% of Mean	% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
100%	Inf	470,620	-0.50%	0.01%	-0.60%	-0.15%	-0.42%
110%	Inf	332,831	-0.07%	-0.23%	-0.77%	0.07%	0.01%
125%	Inf	186,092	0.45%	-0.48%	-0.71%	0.30%	0.50%
150%	Inf	60,980	0.72%	-0.50%	-0.32%	0.36%	0.71%
200%	Inf	3,629	0.32%	-0.08%	0.02%	0.12%	0.29%

Error in Corridor Pricing LoB: 'GL: PremOps ', Size: 1000, Layer: 500x500

From	To	E[X^| x_2>X>x_1]*P(x_2>X>x_1)				difference as % of mean	
% of Mean	% of Mean	Empirical	Lognormal	Normal	Logistic	Gamma	InvGauss
90%	100%	173,326	0.40%	-0.26%	-0.33%	0.21%	0.41%
100%	110%	137,789	0.43%	-0.24%	-0.17%	0.22%	0.43%
110%	125%	146,739	0.52%	-0.24%	0.06%	0.24%	0.49%
125%	150%	125,112	0.26%	-0.03%	0.39%	0.06%	0.21%
150%	200%	57,351	-0.39%	0.42%	0.34%	-0.24%	-0.42%

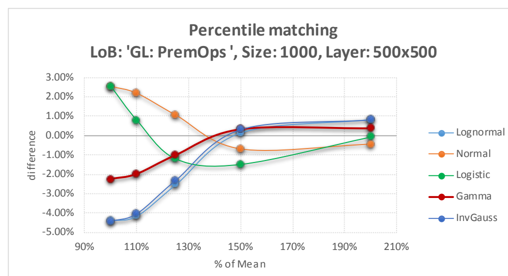

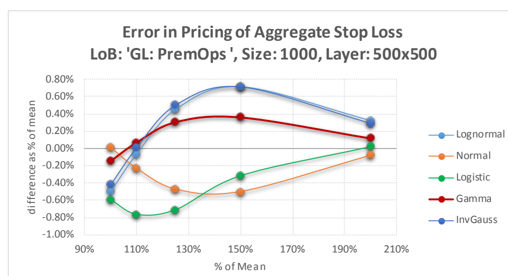

X is aggregate loss
X^ is loss in an aggregate layer, namely max(min(X-x_1, 0), x_2-x_1)

Bibliography

[1] Heckman, P. E., & Meyers, G. G. (1983). The calculation of aggregate loss distributions from claim severity and claim count distributions. In *Proceedings of the Casualty Actuarial Society* (Vol. 70, No. 133-134).

[2] Panjer, H. (1981). Recursive evaluation of a family of compound distributions. ASTIN Bull., 12, 22-26. *Mathematical Reviews (MathSciNet): MR632572*.

[3] Robertson, J. (1992). The computation of aggregate loss distributions. In *Proceedings of the Casualty Actuarial Society* (Vol. 79, No. 150, pp. 57-133).

[4] Mohamed, M. A., Razali, A. M., & Ismail, N. (2010). Approximation of aggregate losses using simulation. *Journal of Mathematics and Statistics*, *6*(3), 233-239.

[5] Dropkin, L. B. (1964). Size of loss distributions in Workmen's Compensation insurance. In *Proceedings of the Casualty Actuarial Society* (Vol. 51, pp. 198-223).

[6] Bickerstaff, D. R. (1972). Automobile Collision Deductibles and Repair Cost Groups: The Lognormal Model. *PCAS LIX*, 68.

[7] Pentikäinen, T. (1987). Approximative Evaluation of the Distribution Function of Aggregate Claims1. *ASTIN Bulletin: The Journal of the IAA*, *17*(1), 15-39.

[8] Venter, G. (1983). Transformed beta and gamma distributions and aggregate losses. In *Proceedings of the Casualty Actuarial Society* (Vol. 70, No. 133 & 134, pp. 289-308).

[9] Chaubey, Y. P., & Trudeau, S. (1998). On the computation of aggregate claims distributions: some new approximations. *Insurance: Mathematics and Economics*, 23(3), 215-230.

[10] Papush, D. E., Patrik, G. S., & Podgaits, F. (2001, March). Approximations of the aggregate loss distribution. *CAS Forum* (Winter, pp. 175-186).

[11] R Core Team (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

[12] Carnell, R. (2016). lhs: Latin Hypercube Samples. R package version 0.14. https://CRAN.R-project.org/package=lhs

[13] Delignette-Muller, M. L., & Dutang, C. (2015). fitdistrplus: An R package for fitting distributions. *Journal of Statistical Software*, 64(4), 1-34. URL http://www.jstatsoft.org/v64/i04/

[14] Nason, G. (2012). NORMT3: Evaluates complex erf, erfc, Faddeeva, and density of sum of Gaussian and Student's t. R package version 1.0-3. https://CRAN.R-project.org/package=NORMT3

[15] Hankin, R. K. (2006). Special functions in R: introducing the gsl package. *R News* 6(4)

[16] Dutang, C., Goulet, V., & Pigeon, M. (2008). actuar: An R package for actuarial science. *Journal of Statistical software*, 25(7), 1-37. URL http://www.actuar-projet.org.

[17] Meyer, D., Dimitriadou, E., Hornik, K., Weingessel, A., & Leisch, F. (2017). e1071: Misc Functions of the Department of Statistics, Probability Theory Group. R package version 1.6-8. https://CRAN.R-project.org/package=e1071

[18] Papush, D. E. (1997). A simulation approach in excess reinsurance pricing. Insurance Mathematics and Economics, 3(20), 266.

[19] Keatinge, C. L. (1999). Modeling losses with the mixed exponential distribution. In *Proceedings of the Casualty Actuarial Society* (Vol. 86, pp. 654-698).