The Leveled Chain Ladder Model for Stochastic Loss Reserving

Glenn Meyers, FCAS, MAAA, CERA, Ph.D.

Abstract: The popular chain ladder model forms its estimate by applying age-to-age factors to the latest reported cumulative claims amount – fixed numbers. This paper proposes two models that replace these fixed claim amounts with estimated parameters, which are subject to parameter estimation error. This paper uses a Bayesian Markov-Chain Monte Carlo (MCMC) method to estimate the predictive distribution of the total reported claims amountsfor these models. Using the CAS Loss Reserve Database, it tests its performance in predicting the distribution of outcomes on holdout data, from several insurers, for both paid and incurred triangles on four different lines of insurance. Their performance is compared with the performance of the Mack model on these data.

Key Words – Chain Ladder Model, Bayesian MCMC estimation, JAGS, Mack Model, Retrospective Testing of Loss Reserve Estimates, The R ChainLadder Package

1. INTRODUCTION

This paper presents two more stochastic loss reserving models. Probably the most generally accepted stochastic models, as evidenced by their inclusion in the CAS Syllabus of Examinations, are those of Mack [3] and England and Verrall [1]. The former paper estimates the moments of the predictive distribution of ultimate claims based on cumulative triangles of claims data. While providing a nice overview of the research to date, the latter paper focuses on estimating the predictive distribution of ultimate claims based on incremental triangles using a Generalized Linear Model (GLM).

While each of the models has a reasonable rationale and when implemented produce a predictive distribution of outcomes, large scale testing of the predictive distributions on actual outcomes was almost nonexistent until recently. One of the first to address the problem was Jessica Leong in her 2010 CLRS presentation[1] where she concluded that the predictive distribution was too narrow for the homeowners' data she analyzed. Last year, Meyers and Shi [6] created the CAS Loss Reserve Database.[2] This database was constructed by linking Schedule P reported losses over a period of ten years to outcomes of predictions made based on data reported in the first year. Meyers and Shi then tested two different models based on paid incremental losses and found that the performance of

[1]Ms. Leong's presentation can be downloaded from the CAS website at http://www.casact.org/education/clrs/2010/handouts/VR6-Leong.pdf.

[2]The data and a complete description of its preparation can be found on the CAS Web site athttp://www.casact.org/research/index.cfm?fa=loss_reserves_data

these predictions left much to be desired. Moreover, they also compared the mean of their predictive distributions to the reserves actually posted by the insurers in their original statement and found that the reserves posted were closer to the reported outcomes than the means estimated by the two models. One has to wonder what the insurers saw that we did not see in the data.

I see two ways to try to remedy this situation. First, we can try to improve the model. Second, we can add information that we previously did not include. This paper attempts to do both. My proposals for improving the model will be described below. The new information is to use the reported losses that include both paid claims and the case reserves, which will be referred to as incurred claims. In Schedule P, this means the reported claims in Part 2 (Incurred Net Losses) minus the corresponding reported claims in Part 4 (Bulk and IBNR Reserves).

In my mind, using incurred claims should rule out the use of models based on incremental claims. Negative incremental claims cause a problem with these models and they are much more common in incurred claim data than they are in paid claim data. Thus this paper focuses on cumulative claims data and uses models that are appropriate for cumulative claims. A good place to start is with the popular chain ladder model.

This paper's proposed new models will make two departures from the standard chain ladder model as identified in Mack [3]. Its goal is to improve upon the performance of the predictive distribution given by Mack's formulas, as measured by the outcomes of 50 insurers in four separate lines of insurance in the CAS Loss Reserve Database.

As we proceed, the reader should keep in mind that this paper describes an attempt to solve a math problem – i.e., predict the distribution of the reported losses after ten years of development. This paper does not address the issue of setting a loss reserve liability. The loss reserve liability could be as simple as subtracting the claims already paid from the projected ultimate losses, but it could also involve discounting and a risk margin. These topics are beyond the scope of this paper.

2. THE HIDDEN PARAMETERS IN THE CHAIN LADDER MODEL.

First, let's describe the chain ladder model. Following Mack [3], let $C_{w,d}$ denote the accumulated claims amount, either paid or incurred, for accident year, w, and development period, d, for $1 \leq w \leq K$ and $1 \leq d \leq K$. $C_{w,d}$ is known for $w + d \leq K + 1$. The goal of the chain ladder model is to estimate $C_{w,K}$ for $w = 2,\ldots, K$. The chain ladder estimate of $C_{w,K}$ is given by

$$C_{w,K} = C_{w,K+1-w} \cdot f_{K+1-w} \cdot \ldots \cdot f_{K-1} \tag{2.1}$$

where the parameters $\{f_d\}$, generally called the age to age factors, are given by

$$f_d = \frac{\sum_{w=1}^{K-d} C_{w,d+1}}{\sum_{w=1}^{K-d} C_{w,d}}. \tag{2.2}$$

It will be helpful to view the chain ladder model in a regression context. In this view, the chain ladder model links $K-1$ separate, one for each d, weighted least-squares regressions through the origin with dependent variables $\{C_{w,d+1}\}$, independent variables $\{C_{w,d}\}$, and parameters f_d for $w=1,\ldots, K-1$. Since each parameter f_d is an estimate, it is possible to calculate the standard error of the estimate, and the standard error of various quantities that depend upon the set $\{f_d\}$. Mack [3] derives formulas for the standard error of each $C_{w,K}$ given by Equation (1) and of the sum of the $C_{w,K}$s for $w = 2,\ldots, K$.

Given a cumulative claims triangle $\{C_{w,d}\}$, the R "ChainLadder" package calculates the chain ladder estimates for each $C_{w,K}$ and the standard errors for each estimate of each $C_{w,K}$ and the sum of all the $C_{w,K}$s. This paper will use these calculations in the chain ladder examples that follow.

Now let's consider an alternative regression type formulation of the chain ladder model. This formulation treats each accident year, w, and each development year, d, as independent variables. The proposed models work in logarithmic space, and so the dependent variable will be the logarithm of the total cumulative (paid or incurred) claim amount for each w and d [3]. The first model takes the following form.

[3] If the reported claim amount is zero, we set the logarithm of the claim amount equal to zero. This should not be a serious problem as it is rare for reported claim amount to be zero, and in most cases, the claim amounts are much larger than zero.

$$C_{wd} \sim \text{lognormal}(\alpha_w + \beta_d, \sigma_d), \tag{2.3}$$

i.e., the mean of the logs of each claim amount is given by $\alpha_w + \beta_d$ and the standard deviation of the logs of each claim amount claim amount is given by σ_d.

Let's call the parameters $\{\alpha_w\}$ the level parameters and the parameters $\{\beta_d\}$ the development parameters. Also set $\beta_1 = 0$. As more claims are settled with increasing d, let's assume that σ_d decreases as d increases.

If we assume that the claim amounts have a lognormal distribution, we can see that this new model is a generalization of the chain ladder model in the sense that one can take the quantities on the right hand side of Equation (2.1) and algebraically translate them into the parameters in Equation 2.3 to get exactly the same estimate. One way to do this is to set

$$\beta_d = \sum_{i=1}^{d-1} \log\left(f_i\right) \text{ for } d=2, \ldots, K$$

$$\alpha_w = \log\left(C_{w,K+1-w}\right) - \sum_{i=1}^{K-w} \log\left(f_i\right) \tag{2.4}$$

$$\sigma_d = 0 .$$

Note that the chain ladder model treats the claims amounts $\{C_{w,K+1-w}\}$ as independent variables, that is to say, fixed values. In this model, the role of the claims amounts, $\{C_{w,K+1-w}\}$, is (indirectly) taken by the level parameters, $\{\alpha_w\}$, that are estimates and subject to error. From the point of view of this model, the chain ladder model "hides" the level parameters, and hence the title of this section. Due to its similarity with the chain ladder model and the fact that it explicitly recognizes the level parameters, let's now refer to the models in this paper as Leveled Chain Ladder (LCL), Versions 1 and 2, models.

Cross classified models such as the LCL models have been around for quite some time. For example, Taylor [8] discusses some of these models in his 1986 survey book. The cross classified model is often confused with the chain ladder model, but Mack [4] draws a clear distinction between the two types of models.

3. BAYESIAN ESTIMATION WITH MCMC SIMULATIONS

This paper uses a Bayesian Markov Chain Monte Carlo (MCMC) program, called JAGS (short for "Just Another Gibbs Sampler"), implemented from an R program to produce a simulated list of $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$ parameters from the posterior distribution. Meyers [7] illustrates how to use JAGS and R to produce such a list.

In an attempt to be unbiased, I chose the prior distributions for the $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$ parameters to be wide uniform distributions. Specifically,

$$\alpha_w \sim \text{uniform}(0, \log(2\cdot\max(C_{w,d}) \text{ for } w+d \leq K+1))$$

$$\beta_d \sim \text{uniform } (-5,5) \text{ for } d = 2,\ldots,10 \tag{3.1}$$

$$\sigma_d = \sum_{i=d}^{10} a_i \text{ , } a_i \sim \text{uniform } (0,1). \quad \text{(This forces } \sigma_d \text{ to decrease as } d \text{ increases.)}$$

The R/JAGS code distributed with this paper produces 10,000 parameters sets $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$ for 10 x 10 loss development triangles that are in the CAS Loss Reserve Database. For each set of parameters, it simulates 10 claim amounts, $C_{w,10}$ for $w = 1,\ldots,10$ from a lognormal distribution with log-mean = $\alpha_w + \beta_{10}$ and log-standard deviation σ_{10}. At a high-level, the code proceeds as follows.

1. The R code reads the CAS Loss Reserve Database, such as that given in Table 3.1, and arranges the data into a form suitable for exporting to the JAGS software.
2. The JAGS code contains the likelihood function (Equation 2.3) and the prior distributions of the parameters (Equation 3.1). JAGS produces 10,000 samples from the posterior distributions of $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$.
3. The R code takes the $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$ from the JAGS program and calculates 10,000 simulated losses from the lognormal distribution implied by these parameters.
4. With the 10,000 losses it calculates various statistics of interest such as the mean and standard deviation of the claims amounts, either by accident year or in total.

Let's consider a specific example. Table 3.1 has a triangle of incurred losses for the Commercial Auto line of insurance taken from the CAS Loss Reserve Database.

Table 3.1

$w \backslash d$	1	2	3	4	5	6	7	8	9	10
1	1,722	3,830	3,603	3,835	3,873	3,895	3,918	3,918	3,917	3,917
2	1,581	2,192	2,528	2,533	2,528	2,530	2,534	2,541	2,538	
3	1,834	3,009	3,488	4,000	4,105	4,087	4,112	4,170		
4	2,305	3,473	3,713	4,018	4,295	4,334	4,343			
5	1,832	2,625	3,086	3,493	3,521	3,563				
6	2,289	3,160	3,154	3,204	3,190					
7	2,881	4,254	4,841	5,176						
8	2,489	2,956	3,382							
9	2,541	3,307								
10	2,203									

Table 3.2 gives the first three (of 10,000) parameter sets $\{\alpha_w\}$, $\{\beta_d\}$ and $\{\sigma_d\}$that were calculated by the JAGS program. Table 3.3 shows the calculation of the mean of the lognormal distribution for the 10th development period. Table 3.4 shows the simulated claims amounts, $\{C_{w,10}\}$, given the log-means from Table 3.3 and the log-standard deviations, σ_d, in Table 3.2. This table also gives the mean and standard deviation of the claims amounts over all 10,000 simulations.

Table 3.2

Parameter	1st 3 of 10,000			
α_1	7.6199	7.6098	7.6223	…
α_2	7.1817	7.1806	7.1965	…
α_3	7.6588	7.6434	7.6720	…
α_4	7.7178	7.7072	7.7280	…
α_5	7.5112	7.5143	7.4643	…
α_6	7.4168	7.4145	7.4853	…
α_7	7.9104	7.8930	7.9435	…
α_8	7.6811	7.5237	7.6143	…
α_9	7.7174	7.6937	7.8590	…
α_{10}	7.8280	7.7604	7.8515	…
β_1	0	0	0	…
β_2	0.4836	0.4783	0.4069	…
β_3	0.5203	0.5545	0.5303	…
β_4	0.6348	0.6230	0.6285	…
β_5	0.6511	0.6593	0.6286	…
β_6	0.6518	0.6633	0.6731	…
β_7	0.6661	0.6689	0.6509	…
β_8	0.6615	0.6555	0.6460	…
β_9	0.6663	0.6607	0.6440	…
β_{10}	0.6580	0.6638	0.6534	…
σ_1	0.2270	0.3140	0.2790	…
σ_2	0.1736	0.1853	0.1198	…
σ_3	0.0956	0.0632	0.0597	…
σ_4	0.0373	0.0363	0.0520	…
σ_5	0.0186	0.0140	0.0455	…
σ_6	0.0180	0.0122	0.0430	…
σ_7	0.0169	0.0113	0.0210	…
σ_8	0.0157	0.0102	0.0188	…
σ_9	0.0155	0.0063	0.0142	…
σ_{10}	0.0055	0.0035	0.0121	…

Table 3.3

Calculation	1st 3 of 10,000			
$\alpha_1+\beta_{10}$	8.2779	8.2736	8.2757	…
$\alpha_2+\beta_{10}$	7.8398	7.8444	7.8499	…
$\alpha_3+\beta_{10}$	8.3168	8.3072	8.3254	…
$\alpha_4+\beta_{10}$	8.3759	8.3710	8.3814	…
$\alpha_5+\beta_{10}$	8.1692	8.1781	8.1177	…
$\alpha_6+\beta_{10}$	8.0749	8.0783	8.1387	…
$\alpha_7+\beta_{10}$	8.5685	8.5567	8.5969	…
$\alpha_8+\beta_{10}$	8.3391	8.1874	8.2677	…
$\alpha_9+\beta_{10}$	8.3754	8.3574	8.5124	…
$\alpha_{10}+\beta_{10}$	8.4861	8.4241	8.5049	…

Table 3.4

	1st 3 of 10,000				Mean	Std. Dev.
$C_{1,10}$	3,949	3,929	3,922	…	3,917	72
$C_{2,10}$	2,542	2,556	2,525	…	2,545	60
$C_{3,10}$	4,103	4,060	4,143	…	4,113	107
$C_{4,10}$	4,339	4,304	4,272	…	4,309	123
$C_{5,10}$	3,507	3,577	3,375	…	3,548	113
$C_{6,10}$	3,186	3,209	3,364	…	3,316	136
$C_{7,10}$	5,247	5,218	5,502	…	5,313	270
$C_{8,10}$	4,193	3,575	3,967	…	3,777	300
$C_{9,10}$	4,304	4,275	5,065	…	4,203	564
$C_{10,10}$	4,768	4,569	4,900	…	4,081	1,112

4. COMPARISIONS WITH THE MACK MODEL

This section compares results obtained on the example above from Version 1 of the LCL models with those obtained from the Mack [3] model as implemented in the R "ChainLadder" package. A summary of these results are in Table 4.1.

Table 4.1

	Leveled Chain Ladder – V1			Mack Chain Ladder			
w	Estimate	Std. Error	CV	Estimate	Std. Error	CV	Actual
1	3,917	72	0.0184	3,917	0	0.0000	3,917
2	2,545	60	0.0236	2,538	0	0.0000	2,532
3	4,113	107	0.0260	4,167	3	0.0007	4,279
4	4,309	123	0.0285	4,367	37	0.0085	4,341
5	3,548	113	0.0318	3,597	34	0.0095	3,587
6	3,316	136	0.0410	3,236	40	0.0124	3,268
7	5,313	270	0.0508	5,358	146	0.0272	5,684
8	3,777	300	0.0794	3,765	225	0.0598	4,128
9	4,203	564	0.1342	4,013	412	0.1027	4,144
10	4,081	1,112	0.2725	3,955	878	0.2220	4,181
Total w=2,…,10	35,206	1,524	0.0433	34,997	1,057	0.0302	36,144

What follows is a series of remarks describing the construction of Table 4.1

- The estimates in both models represent the expected claims amounts for $d = 10$.
- The LCL estimates and standard errors were calculated as described in Section 3 above.
- The Mack [3] standard errors represent, as described in the ChainLadder package user manual, "the total variability in the projection of future losses by the chain ladder method."
- The Mack [3] standard error for $w = 1$ will, by definition, always be zero. Since the α_1 and β_{10} parameters are estimates and hence have variability, the standard error for $C_{1,10}$ given by the LCL models will be positive. How to make use of this feature (e.g., uncertainty in further development) might make for an interesting discussion, but since our goal is to predict $\{C_{w,10}\}$, I chose to omit consideration of the variability of $C_{1,10}$ in any analyses of variability of the totals.
- The CAS Loss Reserve Database contains the completed triangles for the purpose of retrospective testing. The actual outcomes for $\{C_{w,10}\}$ are included here for those who might be curious.

Figure 4.1 is a graphical representation of the information in Table 4.1.

Figure 4.1

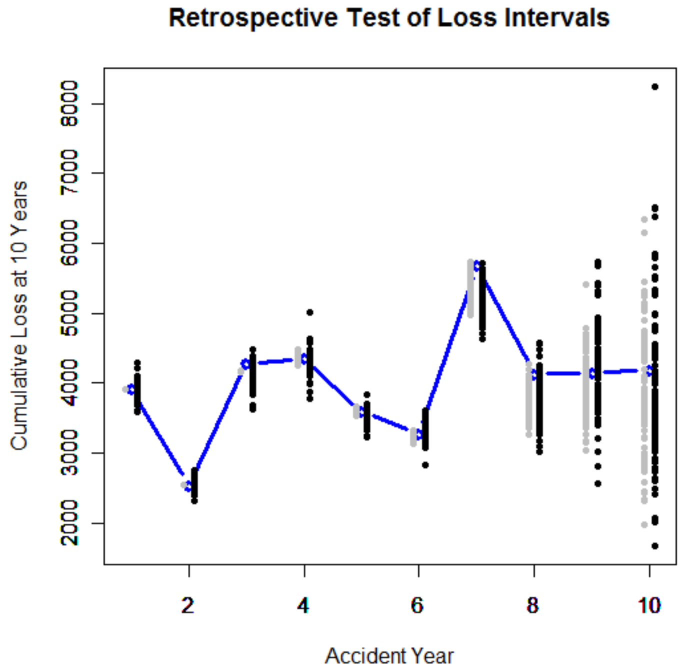

The actual claims amounts points are connected by the line. The darker colored points slightly to the right of the "actual" points are the result of a sample of 100 simulated claims amounts taken from the LCL model. The lighter colored points slightly to the left of the "actual" points are from 100 simulations from a lognormal distribution matching the first two moments given by the Mack [3] model.

The simulated points from the Mack [3] model have smaller standard errors than the standard errors of simulated points from the LCL model. This is to be expected, since the LCL model has more "estimated" parameters. In inspecting other triangles I have found that this is almost always the case, as illustrated in Figure 4.2, where most of the standard errors of the Mack [3] model lie below the diagonal line that represents equality of the standard errors.

At least for this triangle, the span of the simulated points from both models contains the actual outcomes. But for some accident years, this is barely the case.

For the total claims amount over w going from 2 to 10, the actual total, 36,144, lies at the 76th percentile as measured by the LCL predictive distribution. It lies at the 86th percentile as measured by the Mack predictive distribution. The Mack predictive distribution was determined by fitting a lognormal distribution to the first two moments of the total estimate and standard error. Taken by themselves, these observations do not favor one model over the other. To measure the relative performance of the models, we turn to fitting these models to a large number of triangles taken from the CAS Loss Reserve Database.

Figure 4.2

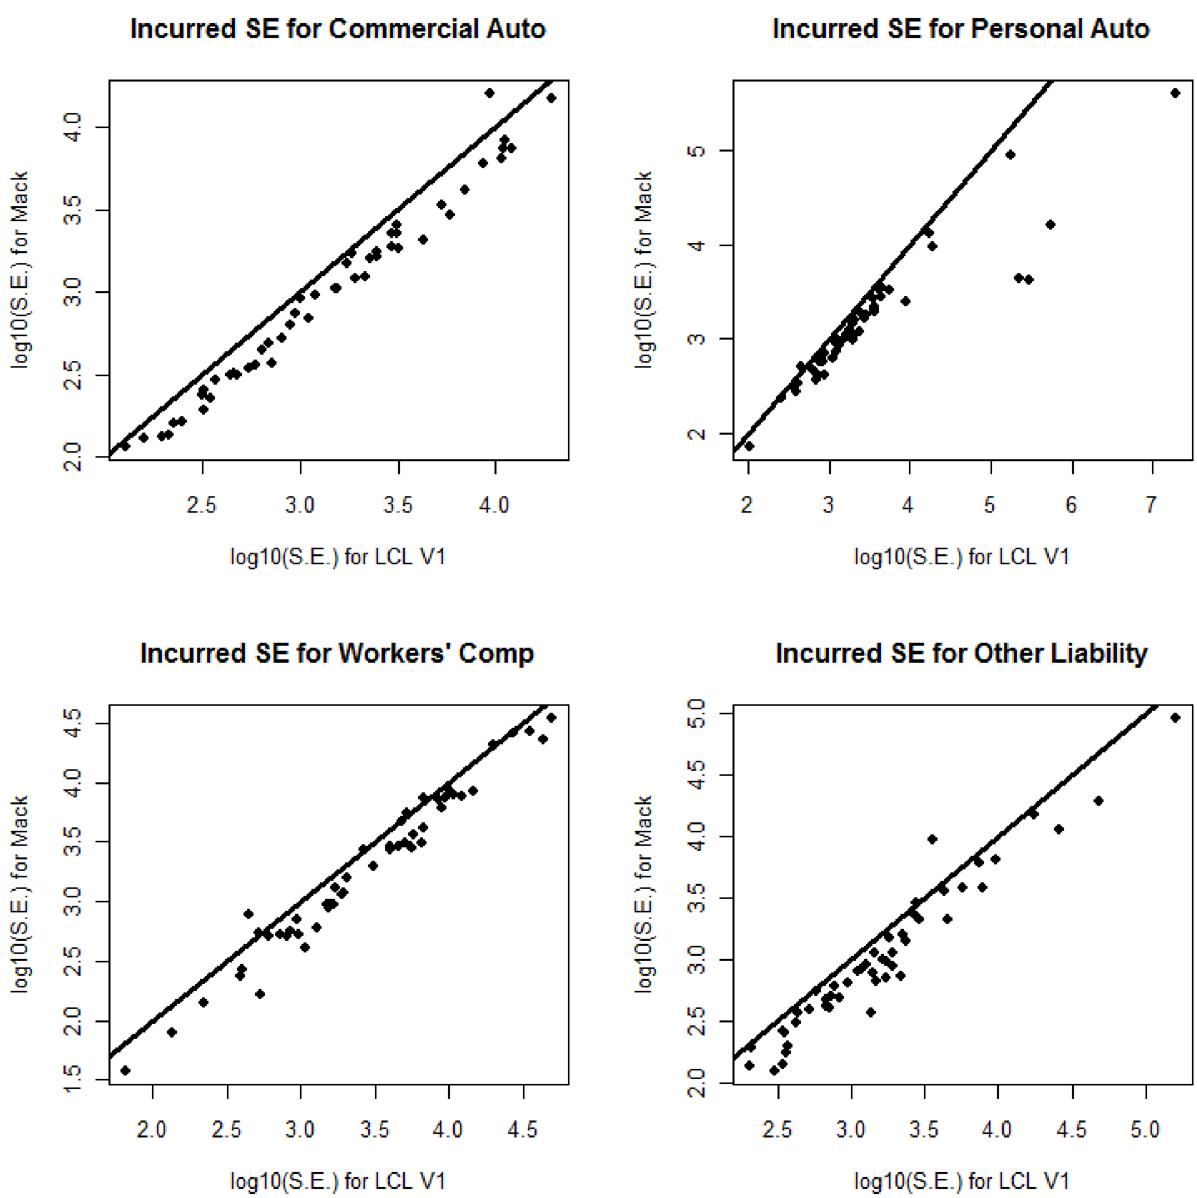

5. RETROSPECTIVE TESTS OF THE PREDICTIVE DISTRIBUTIONS

This section tests considers the LCL – Version 1 model that predict the distribution of unsettled claims using holdout data that is in the CAS Loss Reserve Database. As stated above, the model provides predictions for the sum of the losses $\{C_{w,10}\}$ for $w = 2, \ldots, 10$ using $\{C_{w,d}\}$ for $w + d \leq 11$ as observations. The database contains the actual outcomes available for testing.

This paper's goal is not to produce the smallest error. Instead it is to accurately predict the distribution of outcomes. For a given sum of claims amounts, $\sum_{w=2}^{10} C_{w,10}$, the model can calculate its percentile. If the model is appropriate, the set of percentiles that are calculated over a large sample of insurers should be uniformly distributed. And this is testable.

The most intuitive test for uniformity is to simply plot a histogram of the percentiles and see if the percentiles "look" uniform. If given a set of percentiles $\{p_i\}$ for $i = 1, \; , n$, a more rigorous test would be to use PP plots. To do a PP plot, one first sorts the calculated percentiles, $\{p_i\}$, in increasing order and plots them against the expected percentiles, i.e., the sequence $\{i/(n+1)\}$. If the model that produces the actual percentiles is appropriate, this plot should produce a straight line through the origin with slope one. In practice, the sorted percentiles will not lie exactly along the line due to random variation. But we can appeal to the Kolmogorov-Smirnov test. See, for example, Klugman [2] to account for the random variation. This test can be combined with the PP plot by adding lines with slope one and intercepts $\pm\, 1.36/\sqrt{n}$ to form a 95% confidence band within which the points in the PP plots must lie.

This section shows the results of the above uniformity tests for both paid and incurred losses reported in Schedule P for four lines of insurance, Commercial Auto, Personal Auto, Workers Compensation and Other Liability. After filtering out bad data, I selected 50 insurers for each line of insurance from the CAS Loss Reserve Database. Appendix A lists the insurers selected and describes the filtering criteria.

The results of the uniformity tests are in Figures 5.1-5.10.

Figure 5.1

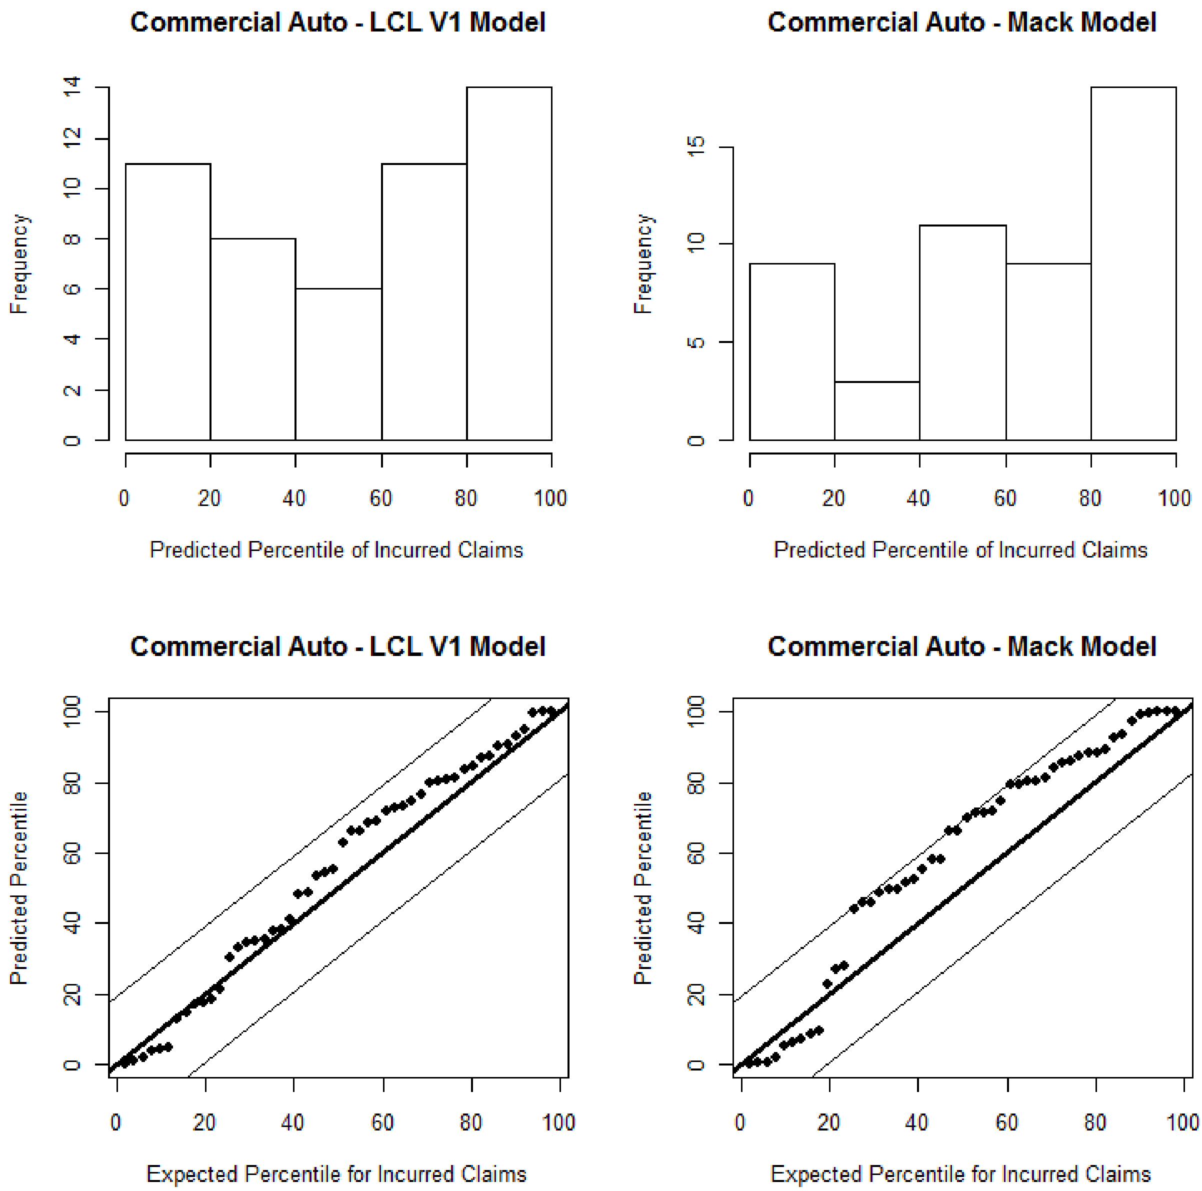
Commercial Auto - LCL V1 Model
Frequency
Predicted Percentile of Incurred Claims
Commercial Auto - Mack Model
Frequency
Predicted Percentile of Incurred Claims
Commercial Auto - LCL V1 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Commercial Auto - Mack Model
Predicted Percentile
Expected Percentile for Incurred Claims

Figure 5.2

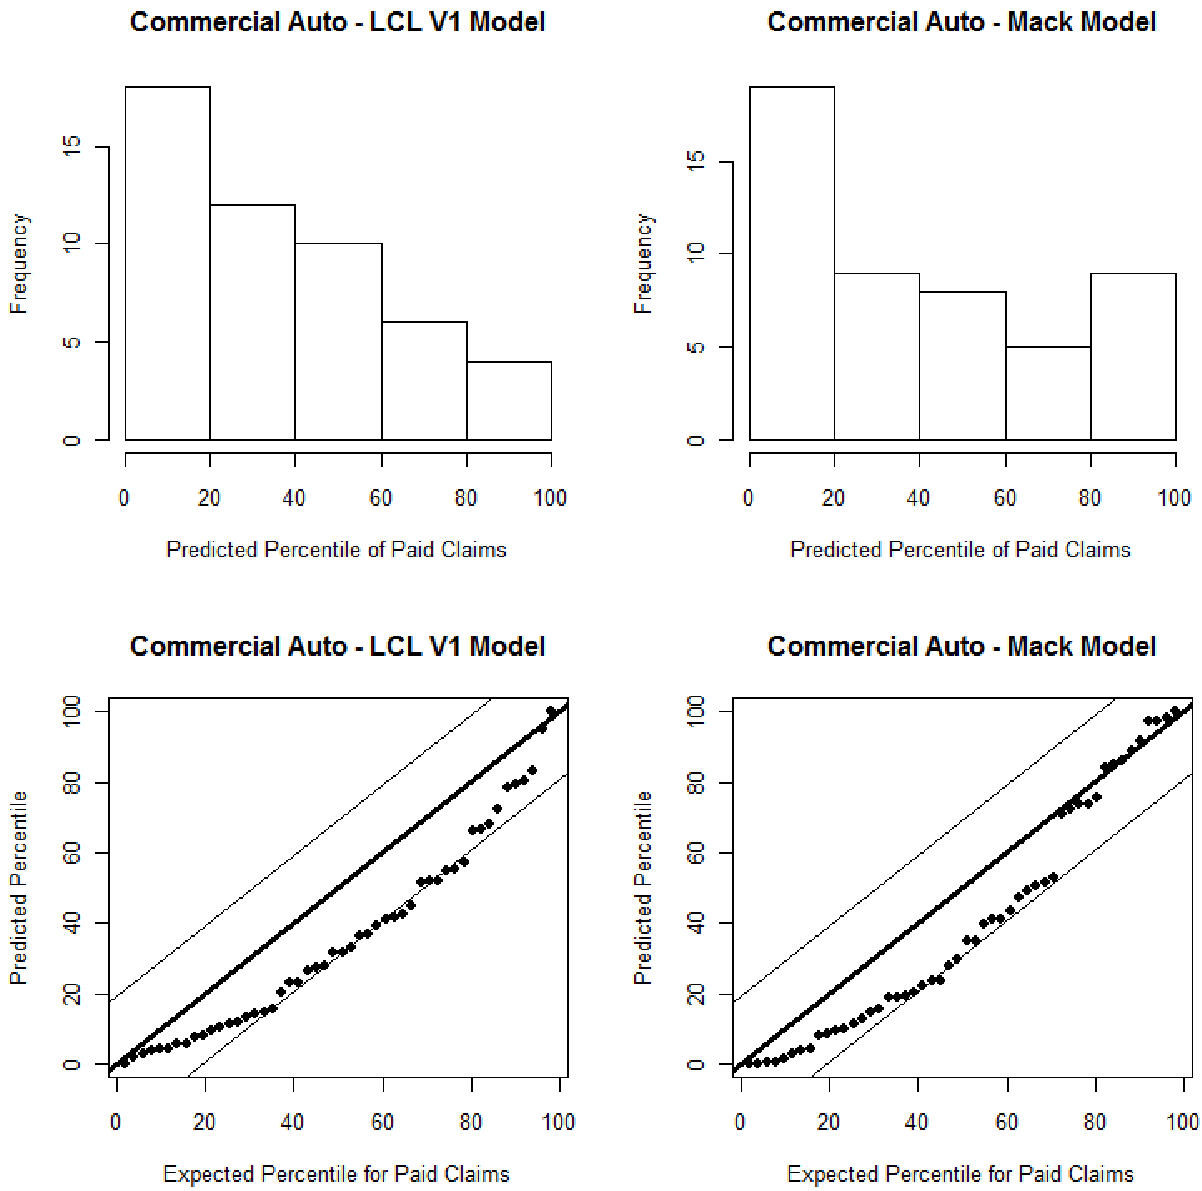
Commercial Auto - LCL V1 Model
Frequency
Predicted Percentile of Paid Claims
Commercial Auto - Mack Model
Frequency
Predicted Percentile of Paid Claims
Commercial Auto - LCL V1 Model
Predicted Percentile
Expected Percentile for Paid Claims
Commercial Auto - Mack Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 5.3

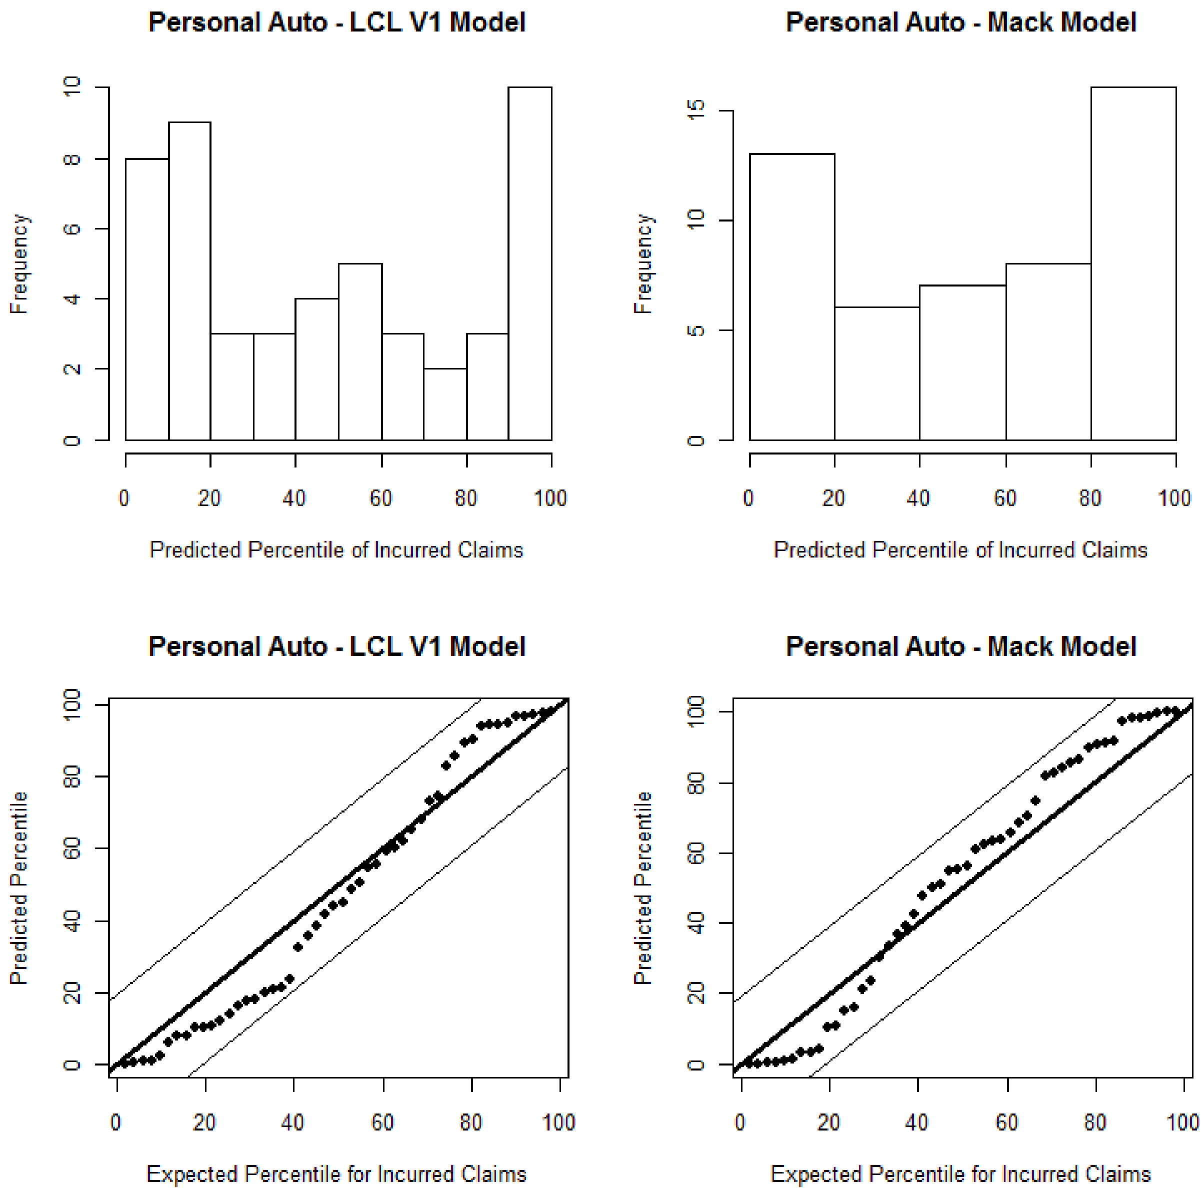
Personal Auto - LCL V1 Model
Frequency
Predicted Percentile of Incurred Claims
Personal Auto - Mack Model
Frequency
Predicted Percentile of Incurred Claims
Personal Auto - LCL V1 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Personal Auto - Mack Model
Predicted Percentile
Expected Percentile for Incurred Claims

Figure 5.4

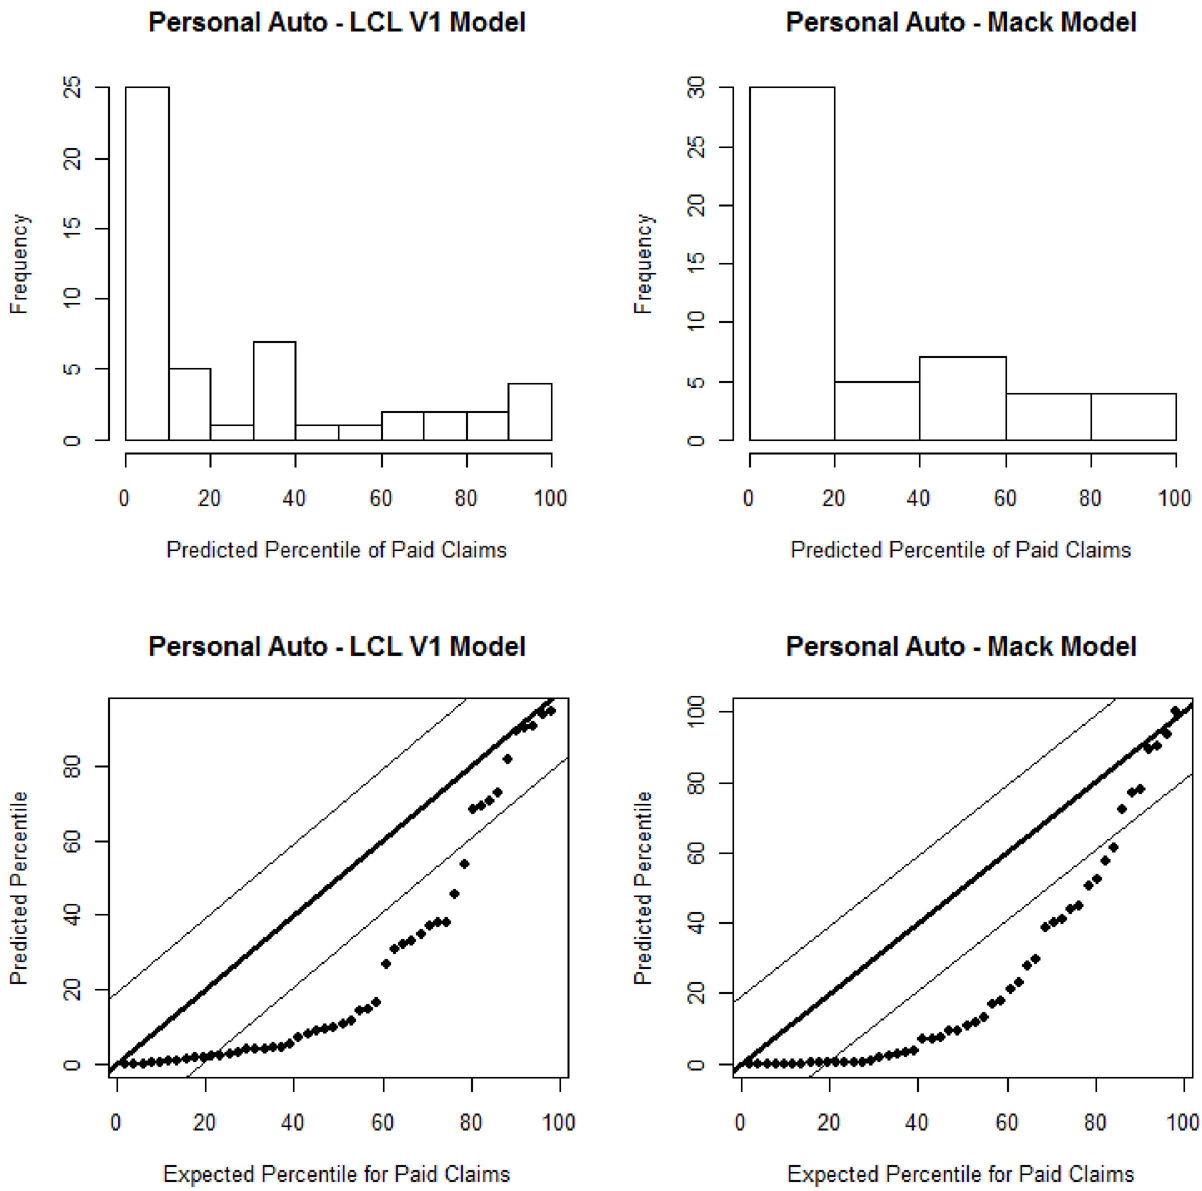
Personal Auto - LCL V1 Model
Frequency
Predicted Percentile of Paid Claims
Personal Auto - Mack Model
Frequency
Predicted Percentile of Paid Claims
Personal Auto - LCL V1 Model
Predicted Percentile
Expected Percentile for Paid Claims
Personal Auto - Mack Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 5.5

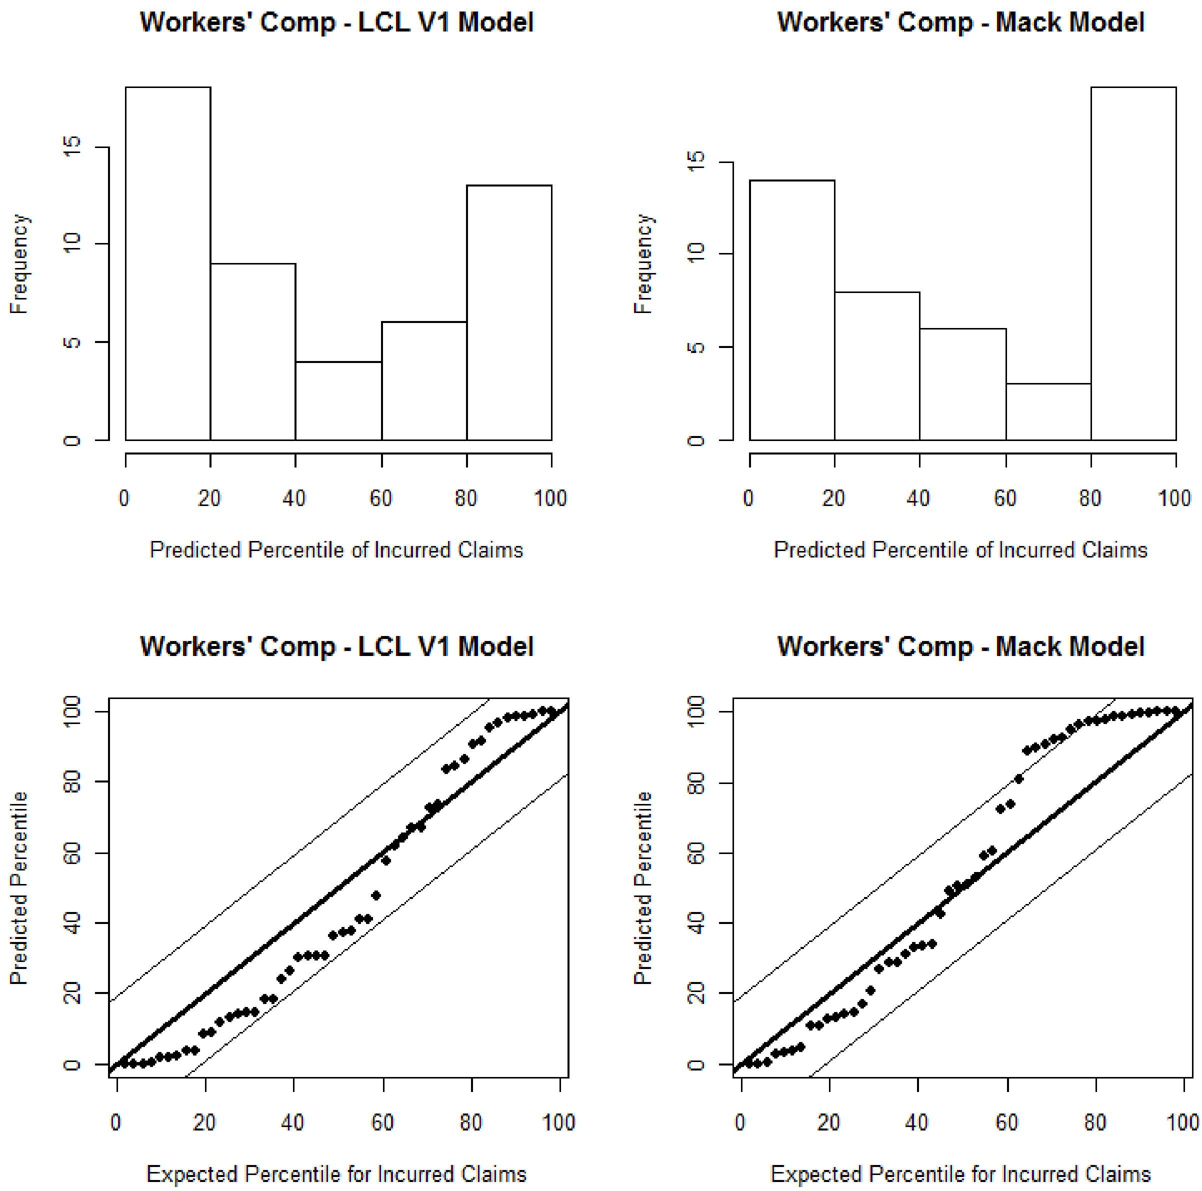
Workers' Comp - LCL V1 Model
Workers' Comp - Mack Model
Frequency
Predicted Percentile of Incurred Claims
Workers' Comp - LCL V1 Model
Workers' Comp - Mack Model
Predicted Percentile
Expected Percentile for Incurred Claims

Figure 5.6

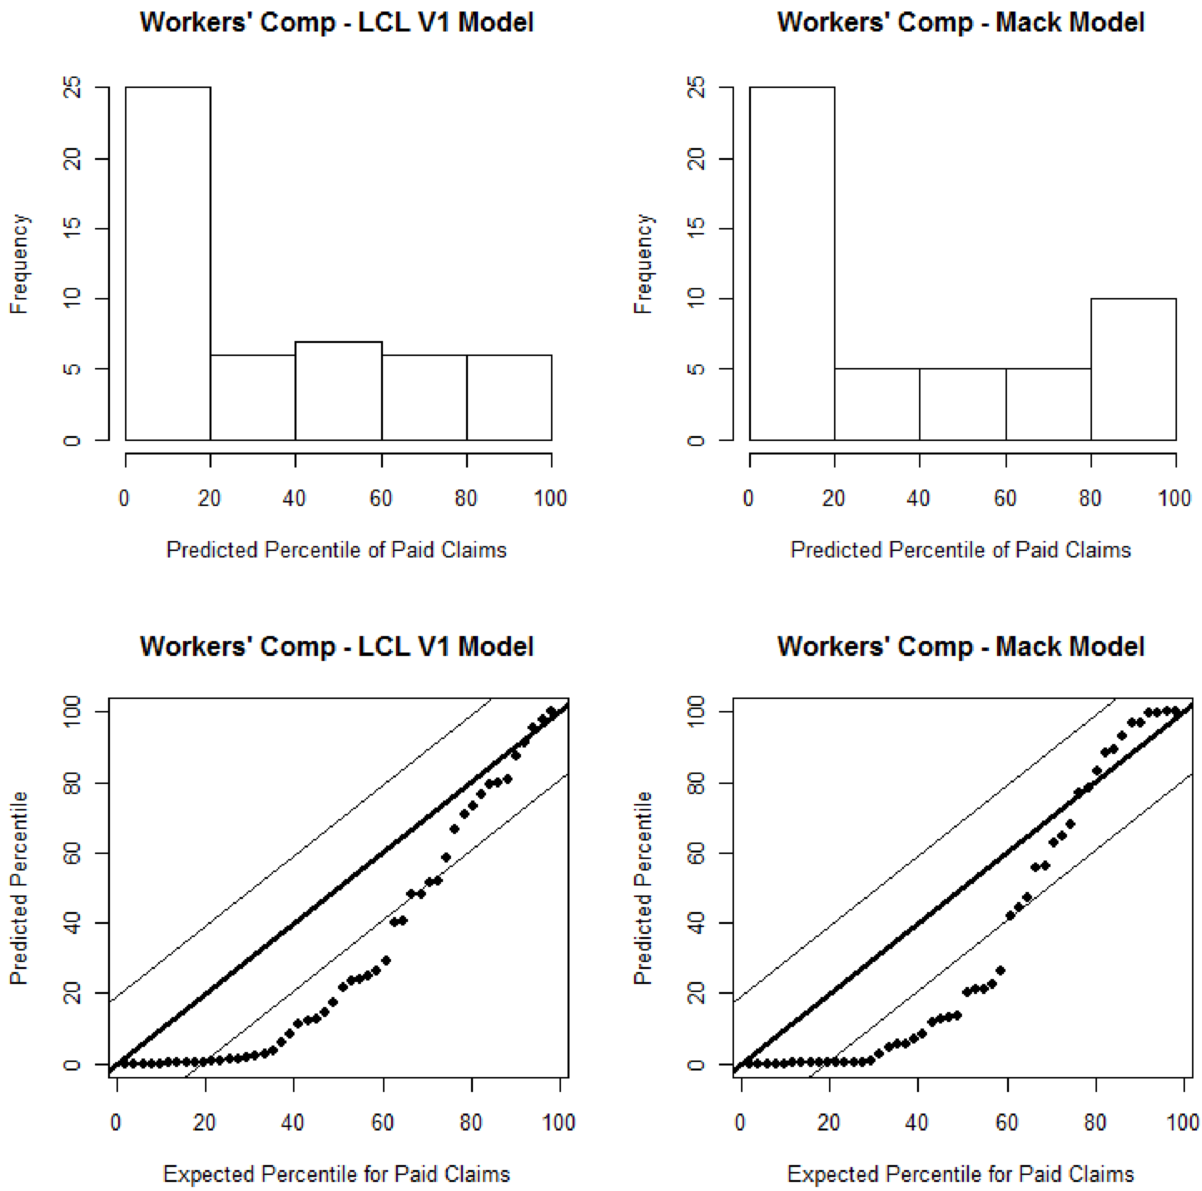
Workers' Comp - LCL V1 Model
Frequency
Predicted Percentile of Paid Claims
Workers' Comp - Mack Model
Frequency
Predicted Percentile of Paid Claims
Workers' Comp - LCL V1 Model
Predicted Percentile
Expected Percentile for Paid Claims
Workers' Comp - Mack Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 5.7

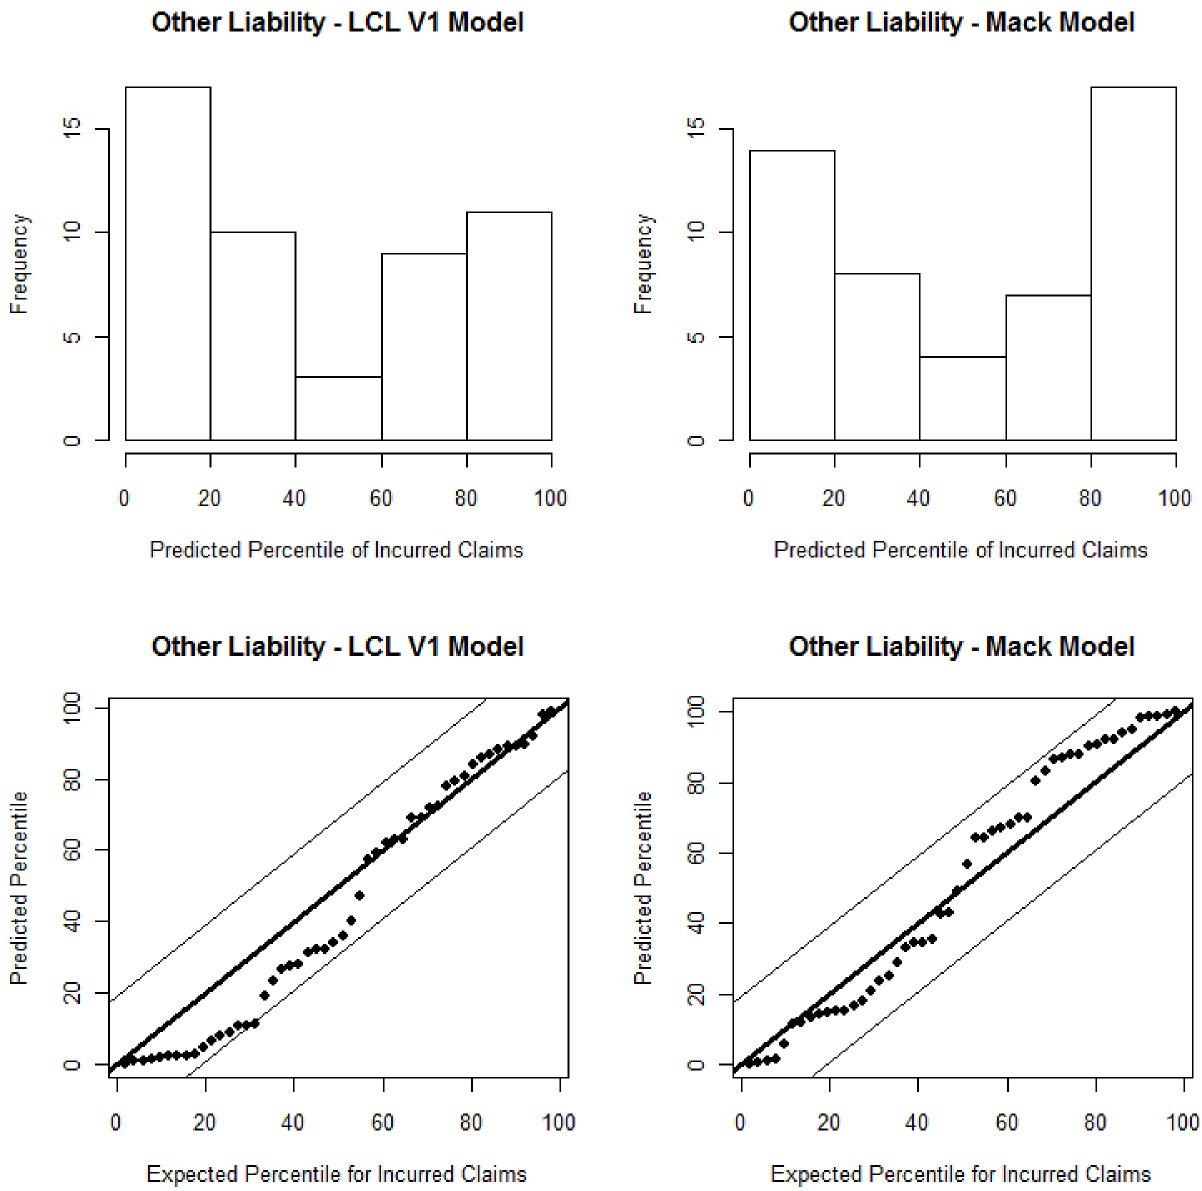
Other Liability - LCL V1 Model
Frequency
Predicted Percentile of Incurred Claims
Other Liability - Mack Model
Frequency
Predicted Percentile of Incurred Claims
Other Liability - LCL V1 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Other Liability - Mack Model
Predicted Percentile
Expected Percentile for Incurred Claims

Figure 5.8

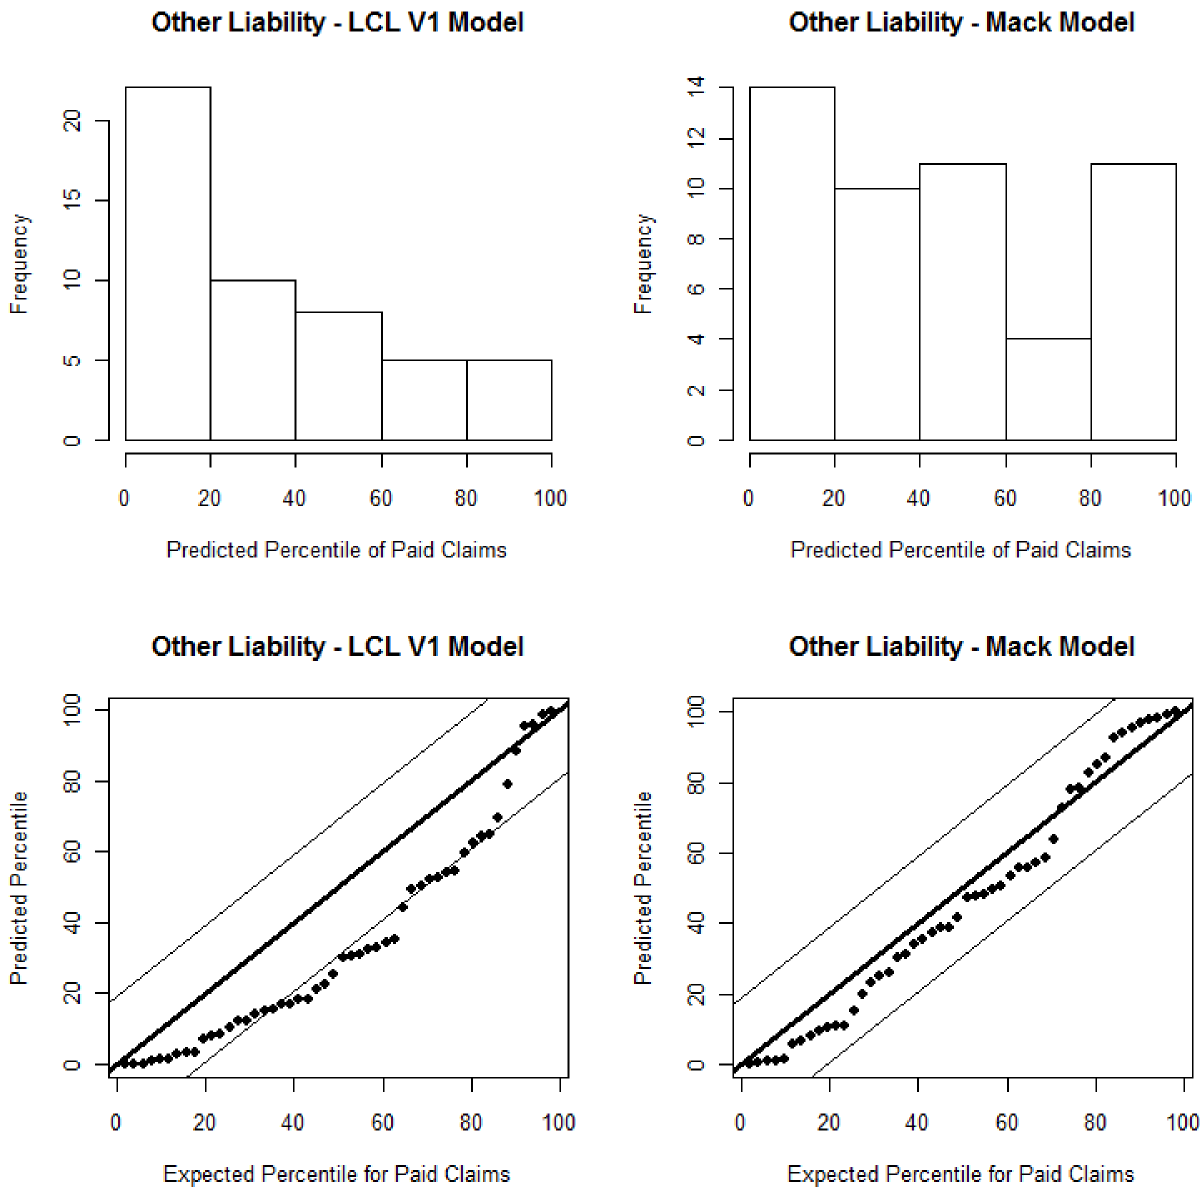
Other Liability - LCL V1 Model
Frequency
0
5
10
15
20
0
20
40
60
80
100
Predicted Percentile of Paid Claims
Other Liability - Mack Model
Frequency
0
2
4
6
8
10
12
14
0
20
40
60
80
100
Predicted Percentile of Paid Claims
Other Liability - LCL V1 Model
Predicted Percentile
0
20
40
60
80
100
0
20
40
60
80
100
Expected Percentile for Paid Claims
Other Liability - Mack Model
Predicted Percentile
0
20
40
60
80
100
0
20
40
60
80
100
Expected Percentile for Paid Claims

Figure 5.9

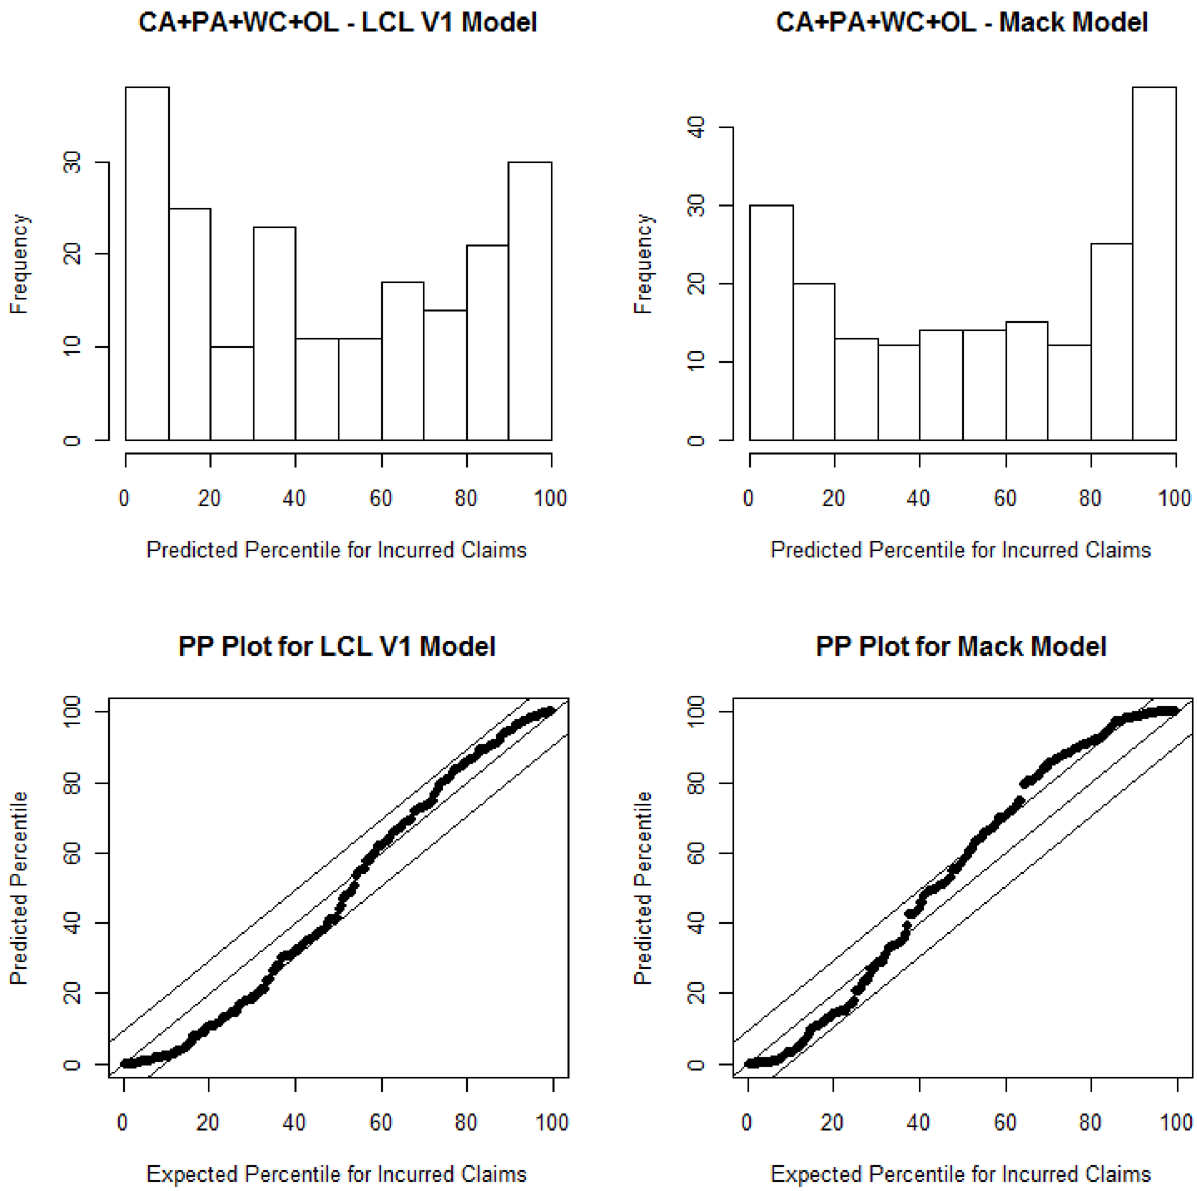
CA+PA+WC+OL - LCL V1 Model
Frequency
0
10
20
30
0
20
40
60
80
100
Predicted Percentile for Incurred Claims
CA+PA+WC+OL - Mack Model
Frequency
0
10
20
30
40
0
20
40
60
80
100
Predicted Percentile for Incurred Claims
PP Plot for LCL V1 Model
Predicted Percentile
0
20
40
60
80
100
0
20
40
60
80
100
Expected Percentile for Incurred Claims
PP Plot for Mack Model
Predicted Percentile
0
20
40
60
80
100
0
20
40
60
80
100
Expected Percentile for Incurred Claims

Figure 5.10

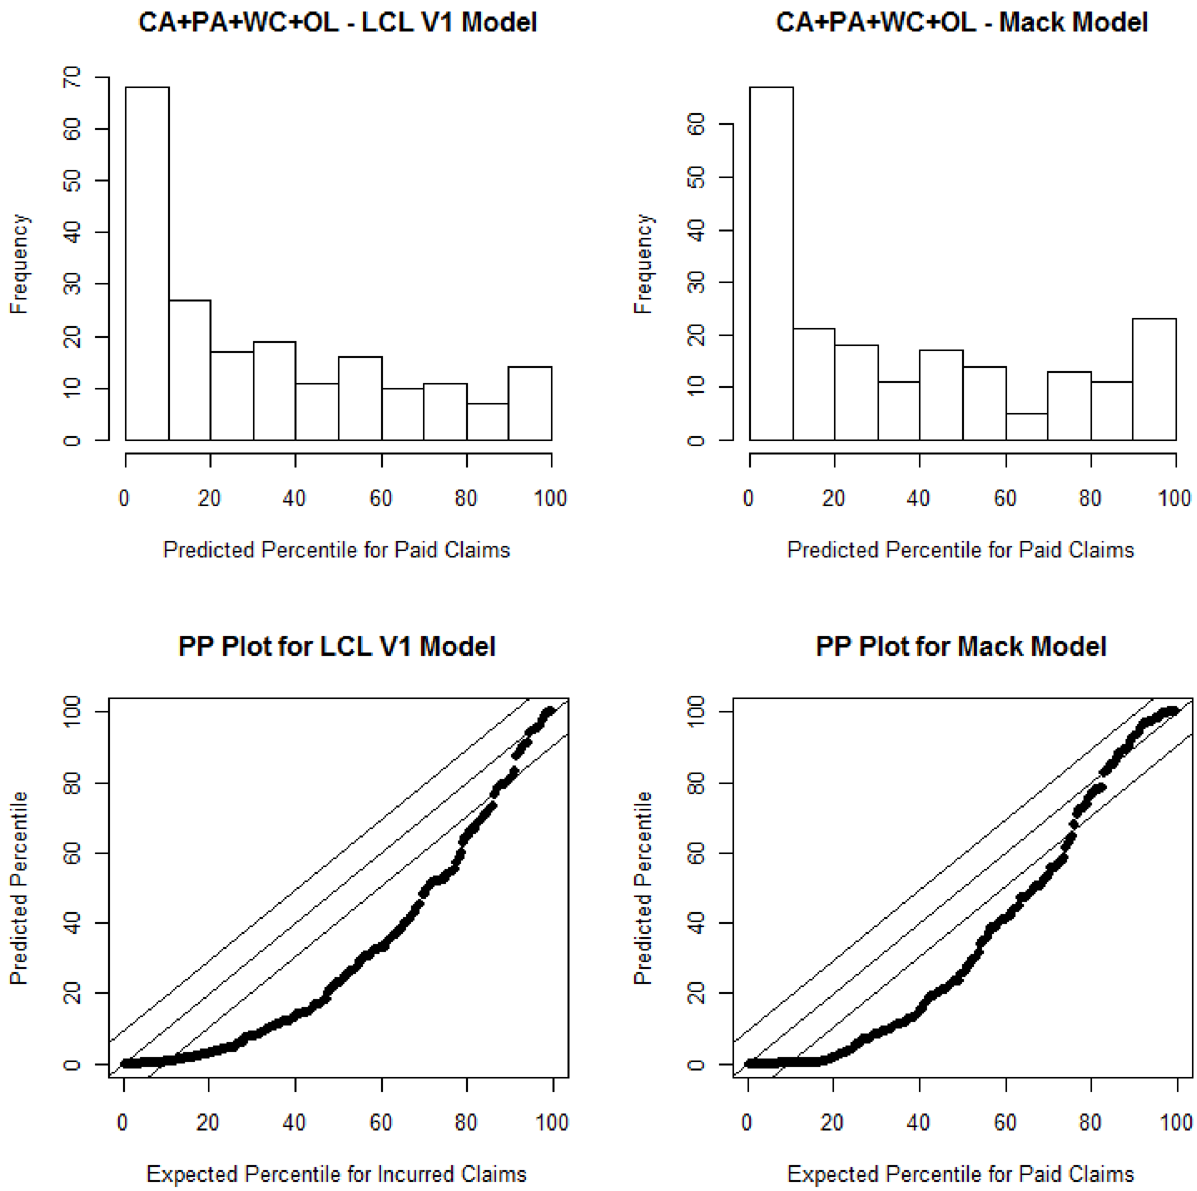
CA+PA+WC+OL - LCL V1 Model
Frequency
Predicted Percentile for Paid Claims
CA+PA+WC+OL - Mack Model
Frequency
Predicted Percentile for Paid Claims
PP Plot for LCL V1 Model
Predicted Percentile
Expected Percentile for Incurred Claims
PP Plot for Mack Model
Predicted Percentile
Expected Percentile for Paid Claims

The results are mixed when looking at the individual lines of insurance for these incurred claims data. The PP-plots lie within the 95% confidence bands for three of the four lines for the LCL Version 1 model. They lie within two of the 95% confidence bands for the four lines for the Mack model. The results are less mixed for these paid claims data. The PP-plots lie within the 95% confidence bands for only the line "Other Liability" for the Mack model. The remaining PP-plots for paid claims data lie well outside the 95% confidence bands.

The picture become clearer when we combine the percentiles in all four lines, as is done in Figures 5.9 and 5.10. While outside the 95% confidence bands, the PP-plots for the incurred claims are close to the band, with the Version 1 model performing somewhat better than the Mack model. The histograms of the percentiles indicate that there are more outcomes than expected in both the high and the low percentiles, i.e., the ranges indicated by both models are too narrow. As indicated by Figure 4.2, the Version 1 model estimates of the standard error are higher than the Mack model estimates, so it should come as no surprise that the Version 1 model performs better than the Mack model on these incurred claims data.

The plots for these paid claims data indicate that neither model is appropriate. I consider that the most likely explanation is that the paid data is missing some important information, some of which is included in the incurred data.

6. CORRELATION BETWEEN ACCIDENT YEARS

One possible reason that the LCL Version 1 model produces ranges that are too narrow is that it fails to recognize that there may be positive correlation between claims payments between accident years. In this section I will propose a model that allows for such correlations, and test the predictions of this model on the holdout data.

To motivate this model, let's suppose we are given random variables X and Y with means μ_X and μ_Y with common standard deviation σ. If we set $Y = \mu_Y + z\cdot(X - \mu_X)$ we can calculate the coefficient of correlation between X and Y as

$$\rho = \frac{E\left[\left(X-\mu_X\right)\cdot\left(Y-\mu_Y\right)\right]}{\sigma^2} = \frac{E\left[z\cdot\left(X-\mu_X\right)^2\right]}{\sigma^2} = z .$$

The proposed model will be one where the logarithms of the claims are correlated between successive accident years. We will refer this model as the LCLVersion 2 model.

$$C_{1,d} \sim \text{lognormal}\left(\alpha_1 + \beta_d, \sigma_d\right)$$
$$C_{w,d} \sim \text{lognormal}\left(\alpha_w + \beta_d + z \cdot \left(C_{w-1,d} - \alpha_{w-1} - \beta_d\right), \sigma_d\right) \text{ for } w = 2,\ldots,K \tag{6.1}$$

Equation 6.1 in Version 2 replaces Equation 2.3 in Version 1. The coefficient of correlation, z, is treated as a random variable with its prior distribution being uniformly distributed between -1 and +1. All other assumptions in Version 2 remain the same as in Version 1. The Bayesian MCMC simulation in Version 2 proceeds pretty much the same as described in Section 3, with the sole difference being the presence of the additional parameter z. Here is a more detailed description of the simulation.

1. Similar to Table 3.2, the JAGS program returns 10,000 vectors $\{\alpha_w\}$, $\{\beta_d\}$, $\{\sigma_d\}$ and z.
2. Similar to Table 3.3, the R program calculates the mean logs
   $\alpha_w + \beta_d + z \cdot \left(C_{w-1,d} - \alpha_{w-1} - \beta_d\right)$.
3. Similar to Table 3.4, the R program simulates claims (sequentially in order of increasing w) from a lognormal distribution with mean log $\alpha_w + \beta_d + z \cdot \left(C_{w-1,d} - \alpha_{w-1} - \beta_d\right)$ and standard deviation log σ_d.

While hypothesizing correlation between successive accident years, by choosing the prior distribution for z to be uniform between -1 and 1, this model does not force the correlation to be any particular value. If the correlation was spurious, the zs would cluster around zero. I ran the model on the data in Table 3.1. Figure 6.1 provides a histogram that strongly supports the presence of positive correlation. Table 6.1 shows that the predicted standard errors for Version 2 are significantly larger than those predicted by Version 1.

Tables 6.2 – 6.6 provide PP plots for Version 2 that are analogous to the Version 1 plots in Section 5. These plots show that the LCL Version 2 model percentile predictions lie within the bounds specified by the Kolmogorov-Smirnov test at the 95% level for incurred claims, but do not lie within the bounds for the paid claims.

Figure 6.1

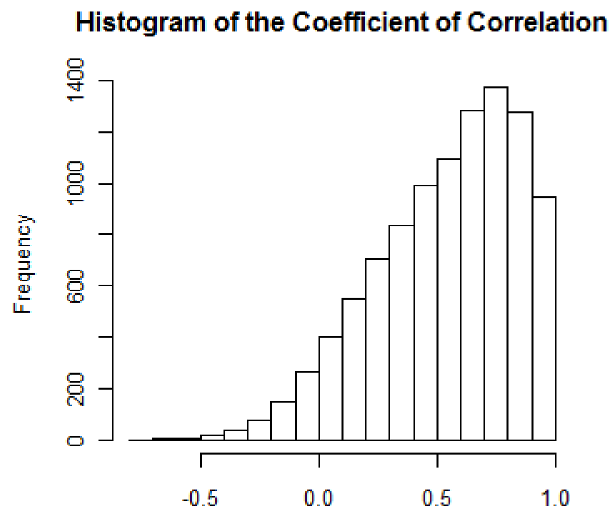

Table 6.1

	Leveled Chain Ladder V2			Leveled Chain Ladder V1			
w	Estimate	Std. Error	CV	Estimate	Std. Error	CV	Actual
1	3,918	86	0.0219	3,917	72	0.0184	3,917
2	2,546	74	0.0291	2,545	60	0.0236	2,532
3	4,113	135	0.0328	4,113	107	0.0260	4,279
4	4,324	162	0.0375	4,309	123	0.0285	4,341
5	3,565	154	0.0432	3,548	113	0.0318	3,587
6	3,338	179	0.0536	3,316	136	0.0410	3,268
7	5,237	356	0.0680	5,313	270	0.0508	5,684
8	3,736	377	0.1009	3,777	300	0.0794	4,128
9	4,122	699	0.1696	4,203	564	0.1342	4,144
10	3,937	1,367	0.3472	4,081	1,112	0.2725	4,181
Total w=2,…,10	34,918	2,192	0.0628	35,206	1,524	0.0433	36,144

Figure 6.2

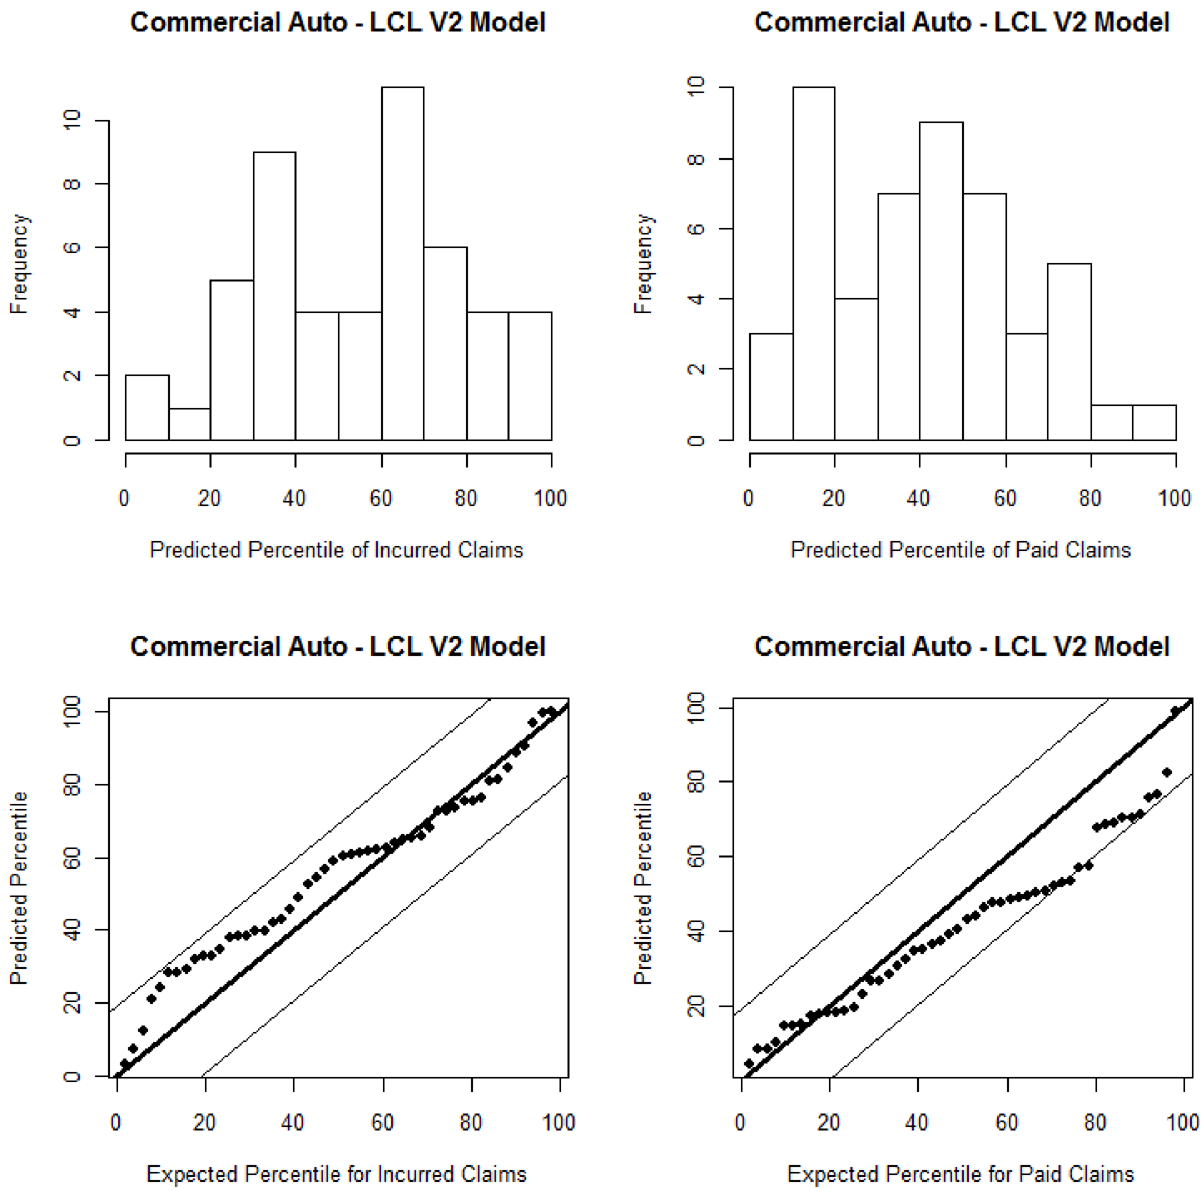
Commercial Auto - LCL V2 Model
Frequency
Predicted Percentile of Incurred Claims
Commercial Auto - LCL V2 Model
Frequency
Predicted Percentile of Paid Claims
Commercial Auto - LCL V2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Commercial Auto - LCL V2 Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 6.3

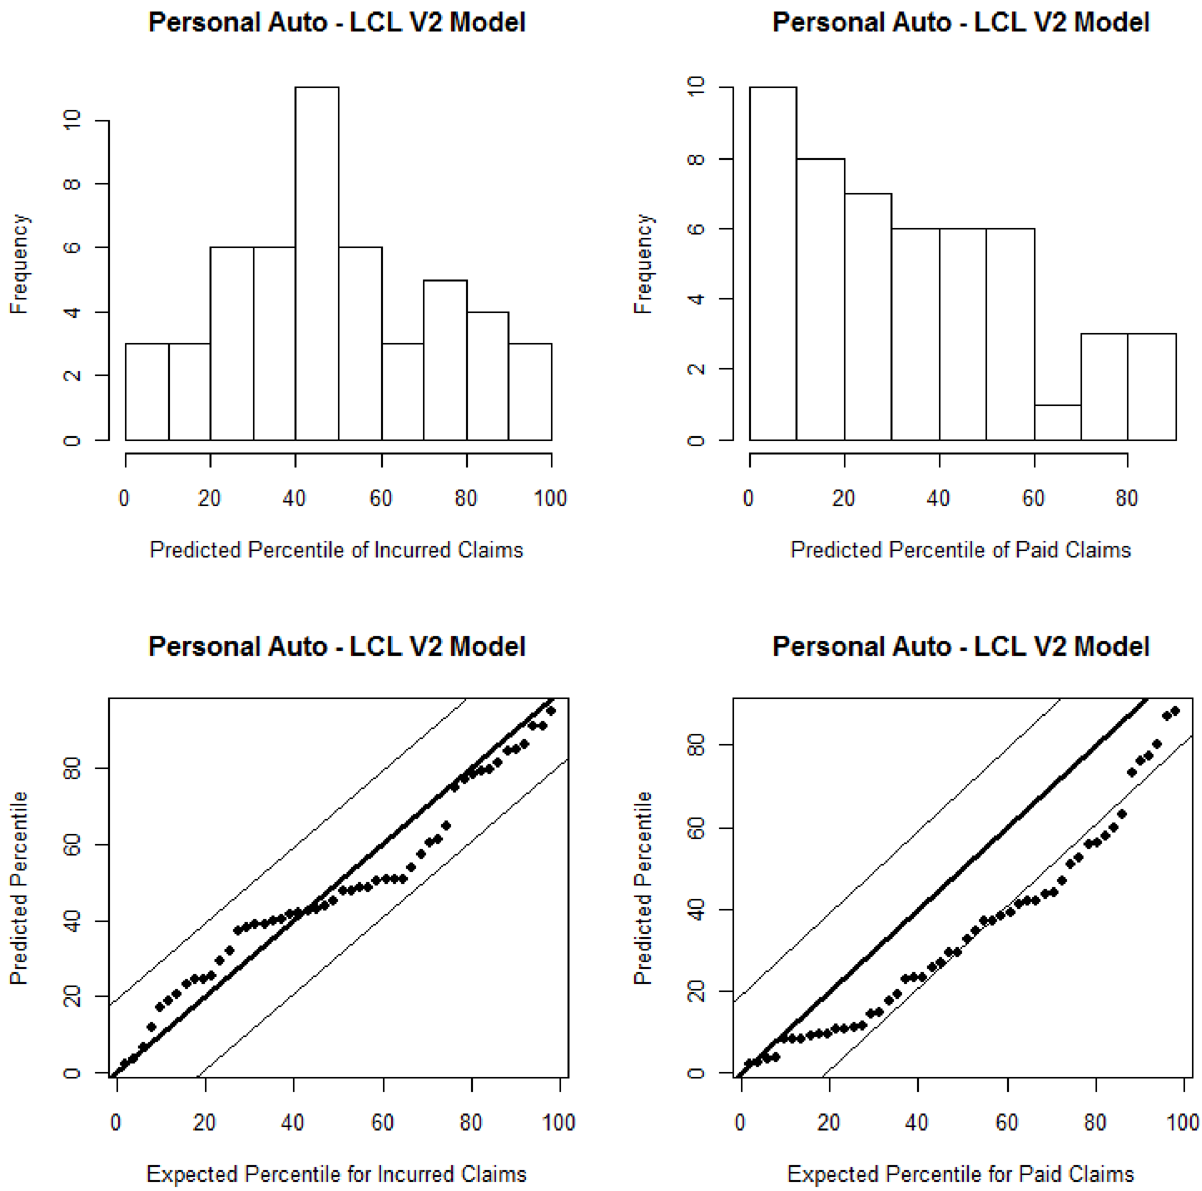
Personal Auto - LCL V2 Model
Frequency
Predicted Percentile of Incurred Claims
Personal Auto - LCL V2 Model
Frequency
Predicted Percentile of Paid Claims
Personal Auto - LCL V2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Personal Auto - LCL V2 Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 6.4

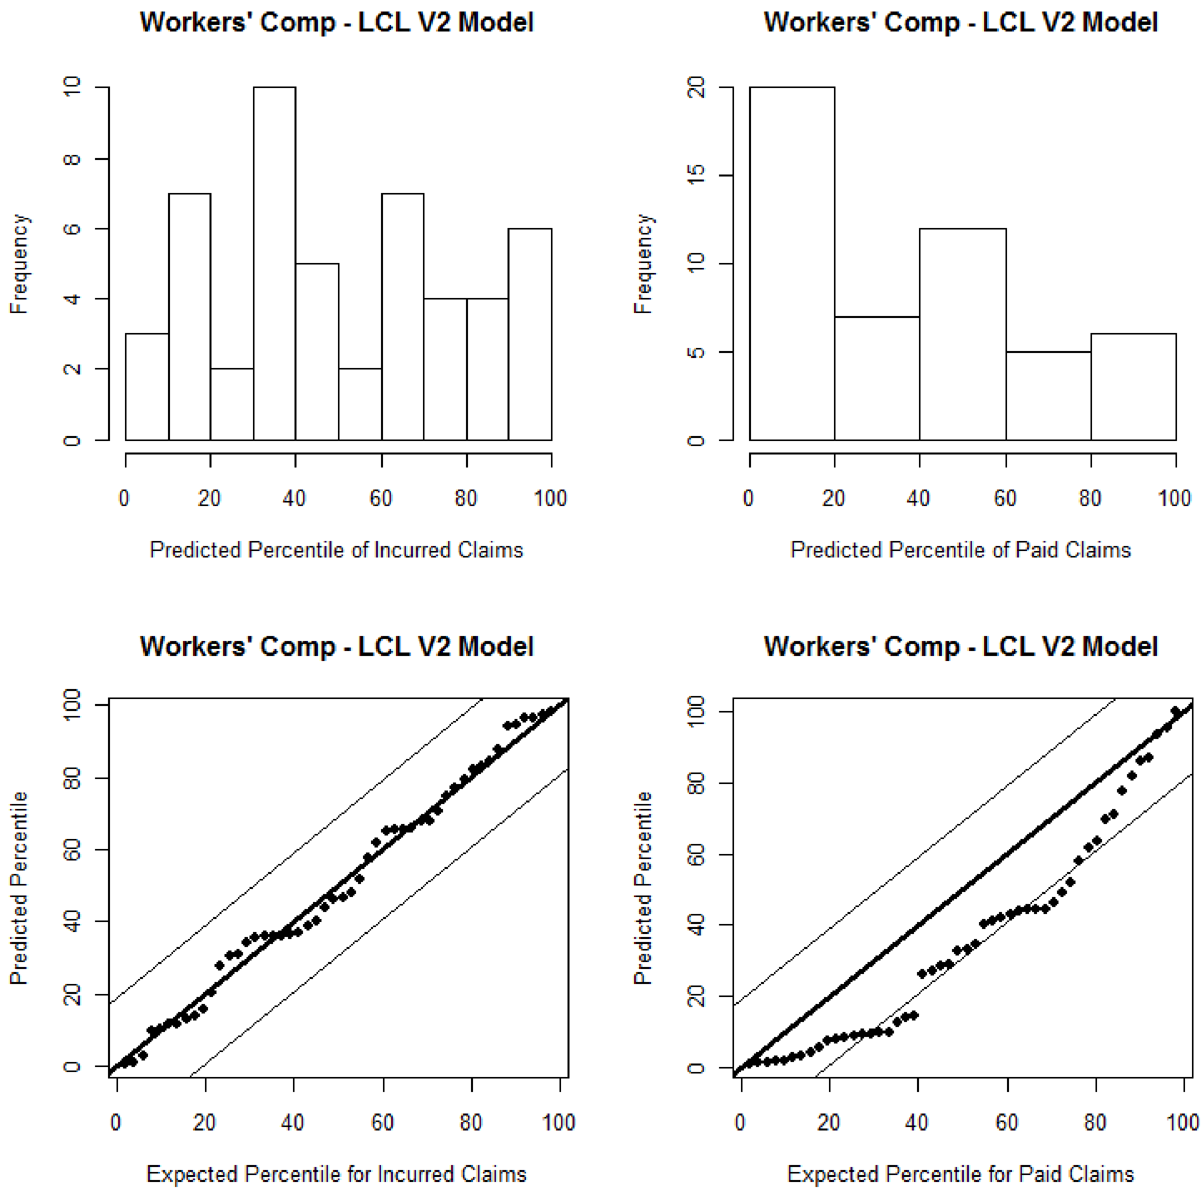
Workers' Comp - LCL V2 Model
Frequency
Predicted Percentile of Incurred Claims
Workers' Comp - LCL V2 Model
Frequency
Predicted Percentile of Paid Claims
Workers' Comp - LCL V2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Workers' Comp - LCL V2 Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 6.5

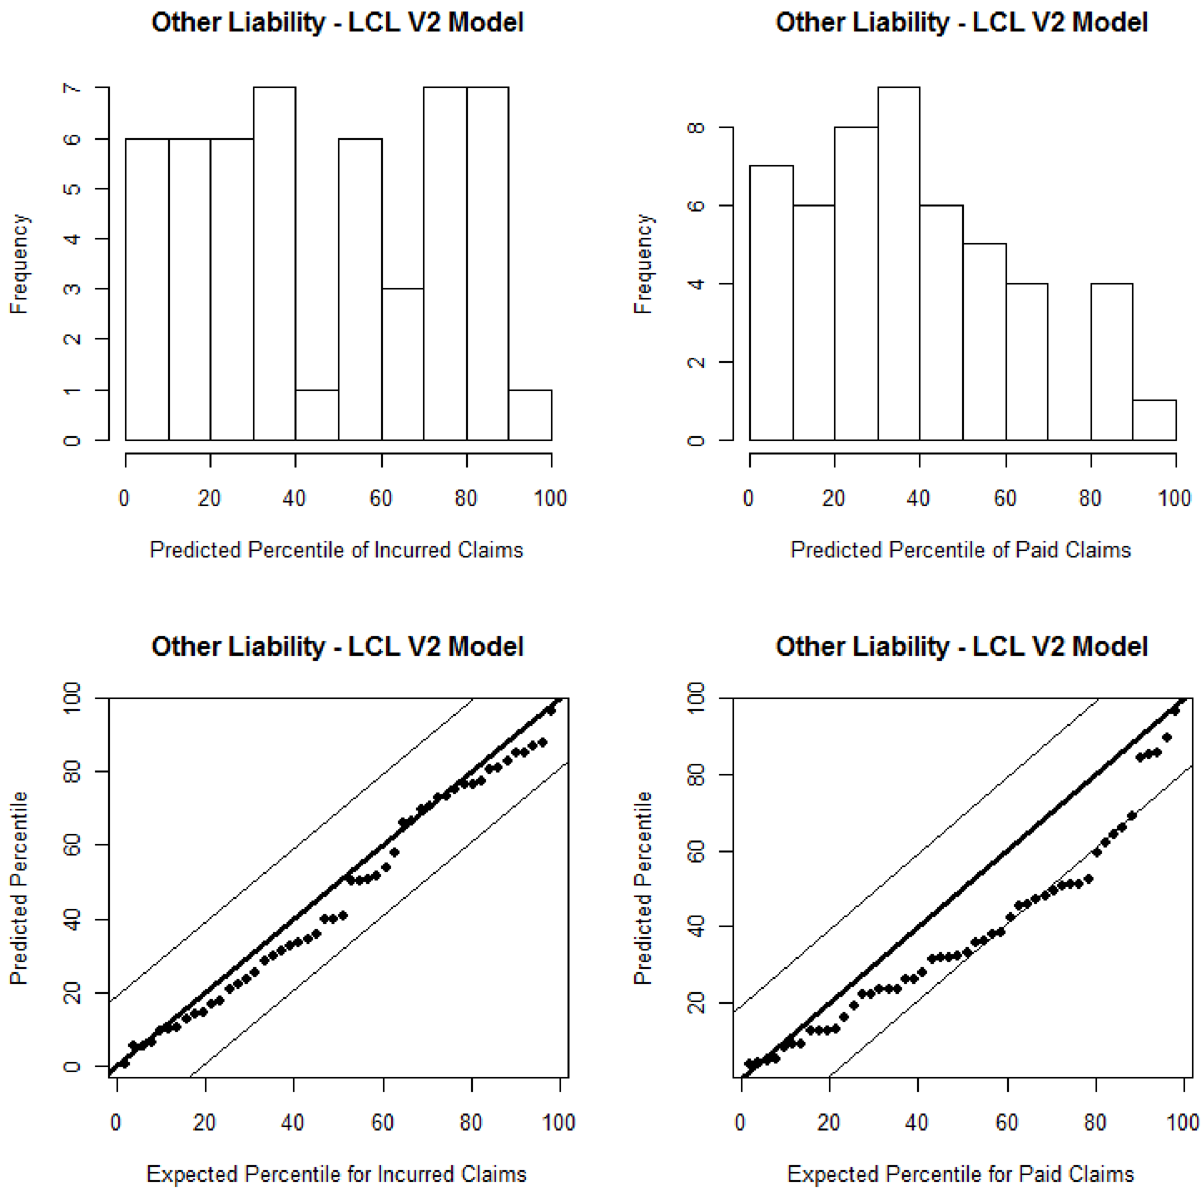
Other Liability - LCL V2 Model
Frequency
Predicted Percentile of Incurred Claims
Other Liability - LCL V2 Model
Frequency
Predicted Percentile of Paid Claims
Other Liability - LCL V2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Other Liability - LCL V2 Model
Predicted Percentile
Expected Percentile for Paid Claims

Figure 6.6

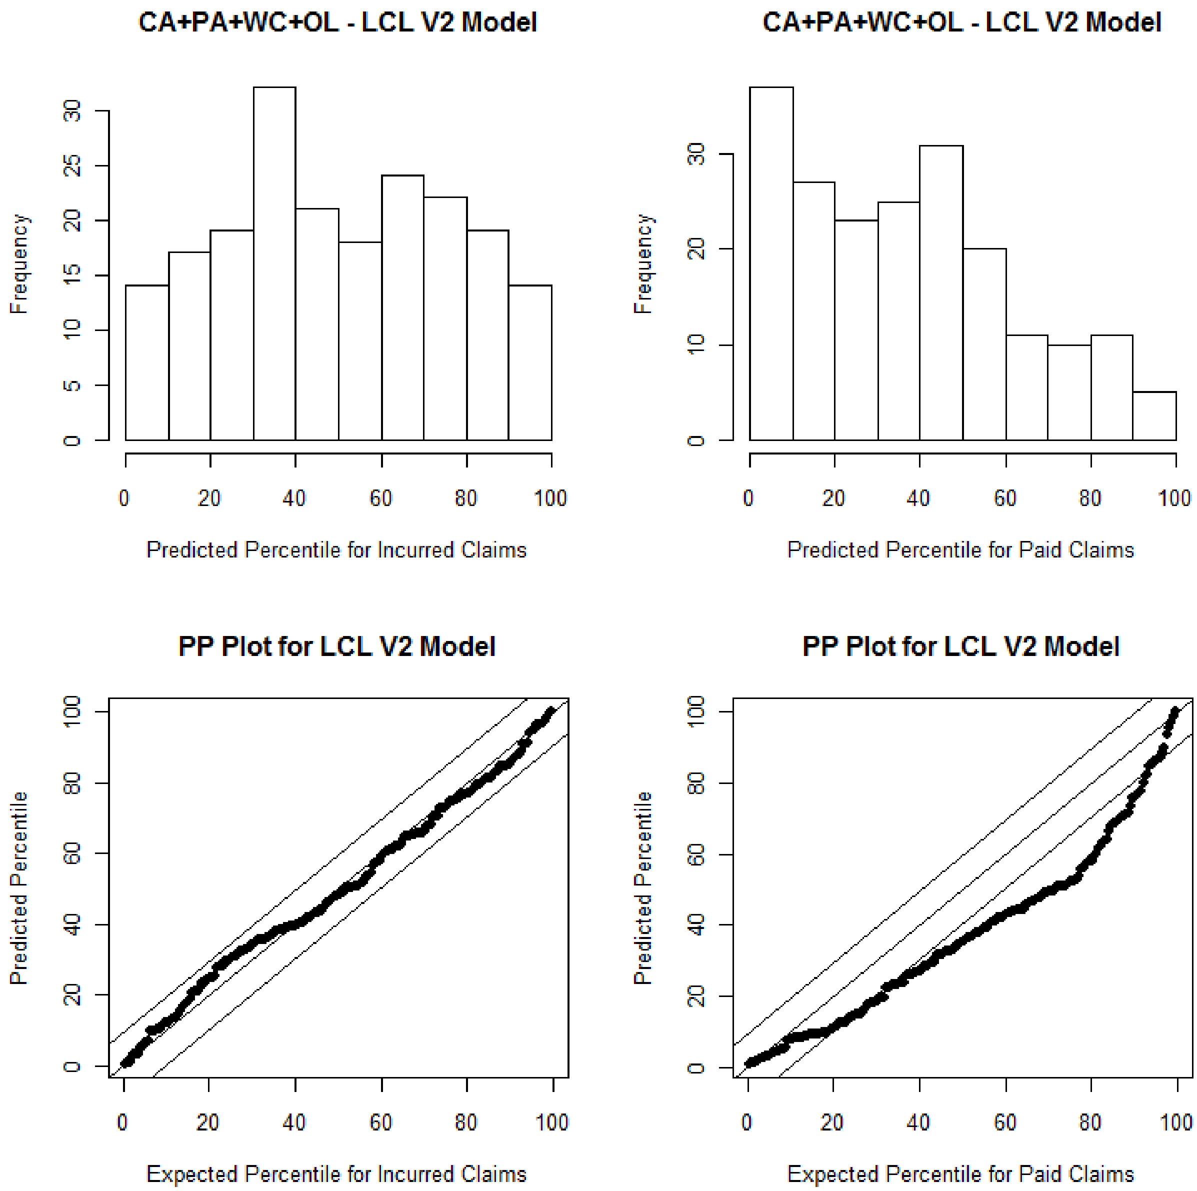
CA+PA+WC+OL - LCL V2 Model
Frequency
Predicted Percentile for Incurred Claims
CA+PA+WC+OL - LCL V2 Model
Frequency
Predicted Percentile for Paid Claims
PP Plot for LCL V2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
PP Plot for LCL V2 Model
Predicted Percentile
Expected Percentile for Paid Claims

7. CONCLUDING REMARKS

When a model fails to validate on holdout data one has two options. First, one can improve the model. Second, one can search for additional information to include in the model. This paper is the result of an iterative process where one proposes a model, watches it fail, identifies the weaknesses, and proposes another model. Successful modeling requires both intuition and failure.

The successful validation of the LCL Version 2 model on the incurred claims data was preceded by the failure of a quite elaborate model, Meyers-Shi [6], built with paid incremental data. This led to the decision to try a model based on cumulative incurred claims, and continued through Versions 1 and 2 of the LCL model.[4]

The simultaneous successful validation of Version 2 on incurred claims and the failure of any model (that I tried) to validate with paid claims suggest that there is real information in the case reserves that cannot be ignored in claims reserving.

A key element in the success of the LCL model is its Bayesian methodology. The simulations done in Meyers [5] suggest that models with a large number of parameters fit by maximum likelihood will understate the variability of outcomes, and that a Bayesian analysis can, at least in theory, fix the problem. The recent developments in the Bayesian MCMC methodology make the Bayesian solution practical.

The LCL models were designed to work with Schedule P claims data. Individual insurers often have access to information that is not published in their financial statements. We should all recall that stochastic models produce conditional probabilities that are not valid in the presence of additional information.That being said, I suspect that many insurers will find the LCL model useful, as it reveals what the outside world could see.

To the best of my knowledge, no stochastic loss reserve model has ever been validated on such a large scale. In any modeling endeavor, the first is always the hardest. Now that we have some idea of what it takes to build a successfully validated model, I would not be surprised to see better models follow.

[4]There were numerous other modeling attempts that will remain unreported.

8. The R/JAGS CODE

The code that produced Tables 4.1 and 6.1 and Figure 4.1 is included in the CAS eForum along with this paper. The code is written in R (freely downloadable from www.r-project.org) and JAGS (freely downloadable from www.mcmc-jags.sourceforge.net). The code requires that the CAS Loss Reserve Database (www.casact.org/research/index.cfm?fa=loss_reserves_data) be downloaded and placed on the user's computer. The code requires the use of the "rjags" and the "ChainLadder" packages in R.

The user should place the files "LCL1 Model.R," "LCL2 Model.R,""LCL1-JAGS.txt," and "LCL2-JAGS.txt" into a working directory. In the first four lines of the R code the user should specify: (1) the name of the working directory; (2) the name and location of the file in the CAS Loss Reserve Database; (3) the group code for the insurer of interest; and (4) the type of loss – either paid or incurred. Then run the code. The code takes about a minute to complete and two progress bars indicate how much of the processing has completed.

The code should work for any complete 10 x 10 triangle. Similar code has run for all the group ids listed in Appendix A.

APPENDIX A – GROUP CODES FOR SELECTED INSURERS

Commercial Auto	Personal Auto	Workers' Comp	Other Liab	Commercial Auto	Personal Auto	Workers' Comp	Other Liab
353	353	86	620	8559	13501	13501	11126
388	388	337	671	10022	13641	13528	11460
620	620	353	683	10308	13889	14176	12866
671	671	388	715	11037	14044	14257	13501
833	715	671	833	11118	14257	14320	13641
1066	965	715	1066	13439	14311	14370	13919
1090	1066	1066	1090	13641	14443	14508	14044
1538	1090	1252	1252	13889	15199	14974	14176
1767	1538	1538	1538	14044	15407	15148	14257
2003	1767	1767	1767	14176	15660	15199	14370
2135	2003	2135	2003	14257	16373	15334	14974
2208	2143	2712	2135	14320	16799	16446	15024
2623	3240	3034	2143	14974	18163	18309	15571
2712	4839	3240	2208	18163	18791	18767	16446
3240	5185	5185	2348	18767	23574	18791	18163
3492	5320	6408	3240	19020	25275	21172	18686
4839	5690	7080	5185	21270	25755	23108	18767
5185	6947	8559	5320	26077	27022	23140	26797
5320	8427	9466	6408	26433	27065	26433	27065
6408	8559	10385	6459	26905	29440	27529	28436
6459	10022	10699	6807	27065	31550	34576	35408
6777	11037	11126	6947	29440	34509	37370	37052
6947	11126	11347	8079	31550	34592	38687	38733
7080	13420	11703	10657	37036	35408	38733	41459
8427	13439	13439	11118	38733	42749	41300	41580

Selection Criteria

1. Removed all insurers with incomplete 10 x 10 triangles.
2. Sorted insurers in order of the coefficient of variation of the premium.
3. Visually inspected insurers and removed those (very few) with "funny behavior."
4. Kept the top 50.

REFERENCES

[1] England, P.D., and R.J. Verrall, "Stochastic Claims Reservingin General Insurance," Institute of Actuaries and Faculty of Actuaries, 28 January 2002.

[2] Klugman, Stuart A., Harry H. Panjer, and Gordon E. Willmot, *Loss Models, From Data to Decisions*, Second Edition, Wiley Series in Probability and Statistics, 2004.

[3] Mack, T., "Measuring the Variability of Chain Ladder Reserve Estimates," Casualty Actuarial Society *Forum*, Spring 1994.

[4] Mack, Thomas, "Which Stochastic Model is Underlying the Chain Ladder Method?",Casualty Actuarial Society *Forum*, Fall 1995.

[5] Meyers, Glenn G., "Thinking Outside the Triangle," ASTIN Colloquium, 2007

[6] Meyers, Glenn G. and Peng Shi, "The Retrospective Testing of Stochastic Loss Reserve Models,"Casualty Actuarial Society *E-Forum*, Summer 2011.

[7] Meyers, Glenn G., "Quantifying Uncertainty in Trend Estimates with JAGS," *The Actuarial Review*, Vol. 39, No. 2. May 2012.

[8] Taylor, G. C.,*Claims Reserving In Non-Life Insurance*,Elsevier Science Publishing Company, Inc., 1986.

Biography of the Author

Having worked as an actuary for over 37 years, Glenn Meyers retired at the end of 2011. His last 23 years of employment were spent working as a research actuary for ISO. In retirement, he still spends some of his time pursuing his continuing passion for actuarial research.