Dependencies in Stochastic Loss Reserve Models

Glenn Meyers, FCAS, MAAA, Ph.D.

Abstract

Given a Bayesian Markov Chain Monte Carlo (MCMC) stochastic loss reserve model for two separate lines of insurance, this paper describes how to fit a bivariate stochastic model that captures the dependencies between the two lines of insurance. A Bayesian MCMC model similar to the Changing Settlement Rate (CSR) model, as described in Meyers (2015), is initially fit to each line of insurance. Then taking a sample from the posterior distribution of parameters from each line, this paper shows how to produce a sample that represents a bivariate distribution that maintains the original univariate distributions as its marginal distributions. This paper goes on to compare the predicted distribution of outcomes by this model with the actual outcomes, and a bivariate model predicted under the assumption that the lines are independent. It then applies the Watanabe-Akaike Information Criterion to compare the fits of the two models.

Key Words: Bayesian MCMC, Stochastic Loss Reserving, Correlation, Dependencies.

1. INTRODUCTION

Recent attempts to apply enterprise risk management principles to insurance have placed a high degree of importance using stochastic models to quantify the uncertainty on the various estimates. For general insurers, the most important liability is the reserve for unpaid losses. Over the years, a number of stochastic models have been developed to address this problem. Some of the more prominent nonproprietary models are those of Mack (1993, 1994), England and Verrall (2002) and Meyers (2015).

As good as these models may be, they fall short of quantifying the uncertainty in the insurer's liability as they do not address the issue of correlation (or more generally – dependencies) between lines of insurance. The failure to resolve this problem analytically has resulted in judgmental adjustments to various risk-based capital formulas. Herzog (2011) provides a summary of some current practices.

Zhang and Dukic (2013) describe what I believe to be a very good attempt at solving this problem. As this paper uses their paper as a starting point, it would be good to provide an outline of their approach[1].

[1] As this paper deals with lognormal models of claim amounts, its description of the Zhang-Dukic ideas are not as general as they put forth in their paper. Their results apply for more general copulas, where this paper deals only with the more specialized multivariate lognormal distribution.

But first, we need to set our notation. Let C_{wd}^{X} be the cumulative paid claim amount in line of insurance X for accident year, w = 1, …. , K and development year d = 1, … , K. Since this paper works with Schedule P data taken from the CAS Loss Reserve Database,[2] we can set $K = 10$. In this paper, X will be CA for Commercial Auto, PA for Personal Auto, WC for Workers Compensation, or OL for Other Liability.

Now suppose that we have models for two different lines of insurance such as

$$\begin{aligned}\log\left(C_{wd}^{X}\right) &\sim Normal\left(\mu_{wd}^{X},\sigma_{d}^{X}\right)\\ \log\left(C_{wd}^{Y}\right) &\sim Normal\left(\mu_{wd}^{Y},\sigma_{d}^{Y}\right)\end{aligned} \tag{1.1}$$

As we shall see below, the parameters μ_{wd}^{X} will be functions of w and d and the parameters σ_{d}^{X} will be subject to constraints for each line X. That feature can be ignored for now as we are setting up the problem.

As shown in Meyers (2015), it is possible to use a Bayesian MCMC model to generate a large sample, say of size 10,000, from the posterior distributions of $\left\{\left\{{}_{i}\mu_{wd}^{X}\right\},\left\{{}_{i}\sigma_{d}^{X}\right\}\right\}_{i=1}^{10000}$ for each line of insurance X.

The idea put forth by Zhang and Dukic is to fit a bivariate Bayesian MCMC model of the following form given the Bayesian MCMC models described by Equation (1.1).

$$\begin{pmatrix}\log\left(C_{wd}^{X}\right)\\ \log\left(C_{wd}^{Y}\right)\end{pmatrix} \sim \text{Multivariate Normal}\left(\begin{pmatrix}\mu_{wd}^{X}\\ \mu_{wd}^{Y}\end{pmatrix},\begin{pmatrix}\left(\sigma_{d}^{X}\right)^{2} & \sigma_{d}^{X}\cdot\rho\cdot\sigma_{d}^{Y}\\ \sigma_{d}^{X}\cdot\rho\cdot\sigma_{d}^{Y} & \left(\sigma_{d}^{Y}\right)^{2}\end{pmatrix}\right) \tag{1.2}$$

The correlation parameter, ρ, describes the dependency between Line X and Line Y.

Zhang and Dukic then use a Bayesian MCMC model to obtain a large sample from the posterior distribution:

[2] The CAS Loss Reserve Database is on the CAS website at
http://www.casact.org/research/index.cfm?fa=loss_reserves_data

$$\left\{\left\{\begin{array}{c} {}_i\mu_{wd}^{X*} \\ {}_i\mu_{wd}^{Y*} \end{array}\right\},\left\{\begin{array}{cc} \left({}_i\sigma_d^{X*}\right)^2 & {}_i\sigma_d^{X*}\cdot{}_i\rho\cdot{}_i\sigma_d^{Y*} \\ {}_i\sigma_d^{X*}\cdot{}_i\rho\cdot{}_i\sigma_d^{Y*} & \left({}_i\sigma_d^{Y*}\right)^2 \end{array}\right\}\right\}_{i=1}^{10000}.$$

The asterisk (*) on the μ and σ parameters calls attention to the fact that the posterior distributions from the models in Equation (1.1) may, and often do, differ significantly from the corresponding marginal posterior distributions from the models in Equation (1.2). To the actuary who prepares loss reserve reports, this presents a problem. Typically actuaries analyze their reserves by individual line of insurance. With a Bayesian MCMC model, they can quantify the uncertainty of the outcomes for that line. Now suppose that there is a demand to quantify the uncertainty in the sum of losses for two or more lines of insurance using the Zhang-Dukic framework. They will need to explain, for example, why the univariate distribution for Commercial Auto produces different results than the marginal distribution for Commercial Auto when combined with Personal Auto. And the marginal distribution could be different still when combined with Workers Compensation.

Scalability is also a problem. For example, the univariate model used in this paper has 31 parameters. Using this model with the bivariate Zhang-Dukic framework yields a model with 31+31+1=63 parameters. In theory, Bayesian MCMC software can handle it, but in practice I have found that running times increase at a much faster rate than the number of parameters. I have coded models using the bivariate Zhang-Dukic framework that work well for some pairs of loss triangles, but others took several hours of running time to obtain convergence of the MCMC algorithm.

The purpose of this paper is to present a framework similar to that of Zhang and Dukic that preserves the univariate models as the marginal distributions.

Before we go there, we should note that a suboptimal model might produce artificial dependencies. To illustrate, consider Figure 1.1 below where y_1 and y_2 are independent random deviations off two parabolic functions of x. We want to fit a bivariate distribution to the ordered pair $(y_1(x), y_2(x))$ of the form:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \sim \text{Multivariate Normal}\left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix},\begin{pmatrix} \sigma_1^2 & \sigma_1\cdot\rho\cdot\sigma_2 \\ \sigma_1\cdot\rho\cdot\sigma_2 & \sigma_2^2 \end{pmatrix}\right) \qquad (1.3)$$

The lower left plot of Figure 1.1 shows a scatter plot of $y_1(x) - \mu_1$ and $y_2(x) - \mu_2$ for the (suboptimal) model $\mu_i(x)$ is a constant. The lower right plot is a scatter plot of $y_1(x) - \mu_1(x)$ and $y_2(x) - \mu_2(x)$ for the (correct) parabolic model. This example shows how suboptimal models for the marginal distribution can cause an artificial nonzero correlation the multivariate model.

The next section will describe the data used in this paper. Section 3 will describe the univariate (marginal) models and illustrate some diagnostics to test the appropriateness of the model. Section 4 will show how to obtain a random sample from the posterior distribution of parameters subject to the constraint that the marginal distribution is the same as those obtained by the corresponding univariate models. Section 5 will describe statistical tests to test the hypothesis that the correlation parameter, ρ, in the bivariate distribution is significantly different from zero. Section 6 will address the sensitivity of the results to the choice of models, and Section 7 will discuss the conclusions.

This paper assumes that the reader is familiar with Meyers (2015).

Figure 1.1 – Illustration of Artificial Correlation

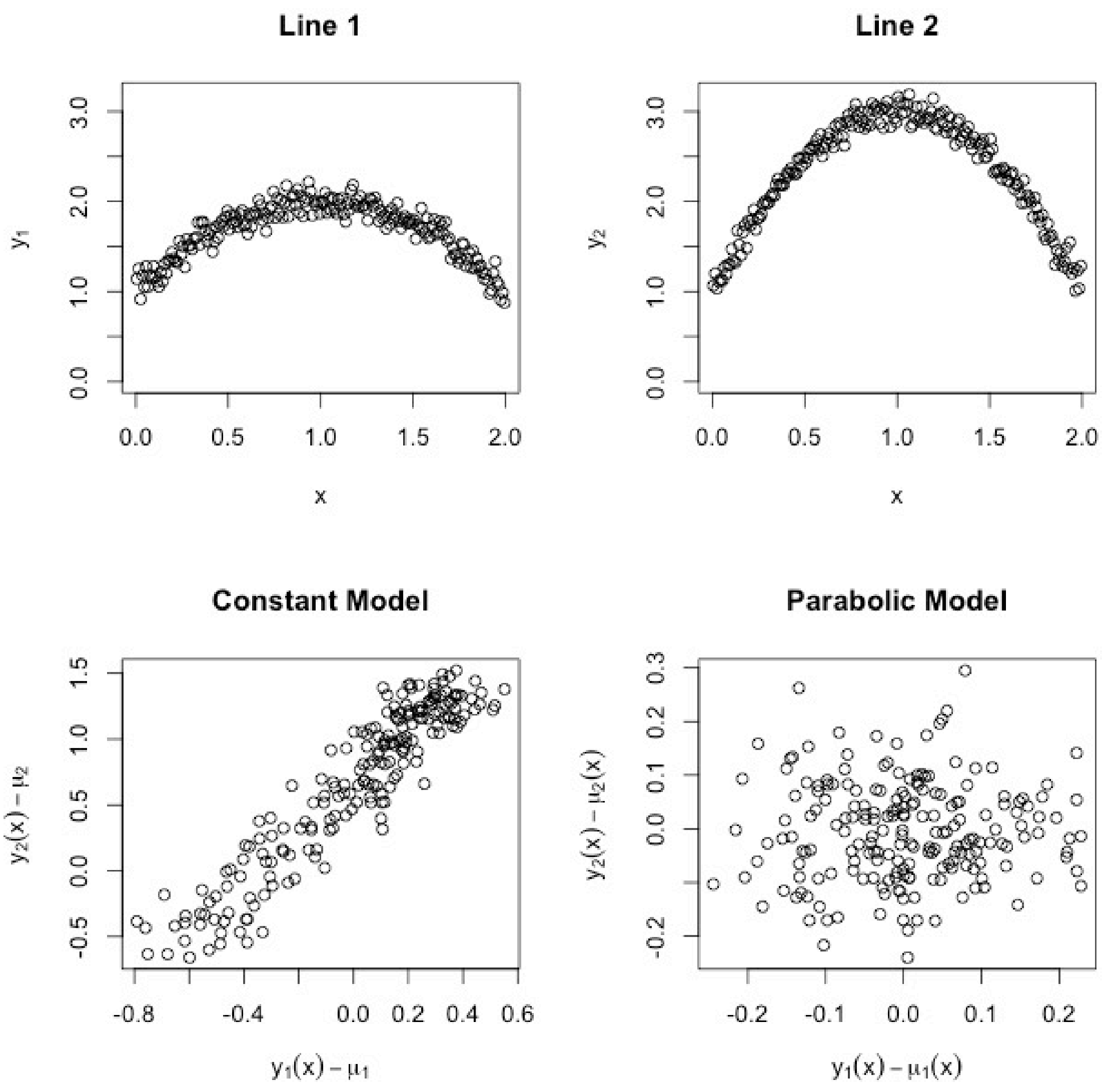

2. THE DATA

The data used in this paper comes from the CAS Loss Reserve Database[3]. The Schedule P loss triangles taken from this database are listed in Appendix A of Meyers (2015). There are 200 loss triangles, 50 each from the CA, PA, WC and OL lines of insurance. Univariate models from all 200 loss triangles will be analyzed in Section 3 and 6.

At the time of writing the monograph, Meyers (2015), I did not envision a dependency study. But it turned out that there were 102 within-group pairs of triangles (29 CA-PA, 17 CA-WC, 17 CA-OL, 14 PA-WC, 15 PA-OL and 10 WC-OL) that were suitable for studying dependency models. Preferring to use loss triangles that have already been vetted, I decided to stick with these within-group pairs of triangles.

This paper will provide detailed analyses for two illustrative insurers (Groups 620 and 1066) for the CA and PA lines of business. The complete loss triangles and outcomes are in Table 2.1 below. The upper data triangle used to fit each model is printed with the ordinary font. The lower data triangle used for retrospective testing is printed with bold and italicized font.

A complete list of the insurer groups used in this paper is included in a spreadsheet titled "Appendix." The sheets in the Appendix contain:

- The 200 groups along with the associated calculations in Section 3.
- The R scripts that produce the univariate model calculations described in Section 3 and 6.
- The 102 within-group pairs with associated calculations in Sections 4 and 5.
- The R scripts that produce the bivariate model calculations described in Sections 4, 5 and 6.

[3] http://www.casact.org/research/index.cfm?fa=loss_reserves_data

Table 2.1 – Data for Illustrative Insurers

Group 620 - Commercial Auto

AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	30,224	4,381	9,502	15,155	18,892	20,945	21,350	21,721	21,934	21,959	21,960
2	35,778	5,456	9,887	13,338	17,505	20,180	20,977	21,855	21,877	21,912	***21,981***
3	42,257	7,083	15,211	21,091	27,688	28,725	29,394	29,541	29,580	***29,595***	***29,705***
4	47,171	9,800	17,607	23,399	29,918	32,131	33,483	33,686	***34,702***	***34,749***	***34,764***
5	53,546	8,793	19,188	26,738	31,572	34,218	35,170	***36,154***	***36,201***	***36,256***	***36,286***
6	58,004	9,586	18,297	25,998	31,635	33,760	***34,785***	***35,653***	***35,779***	***35,837***	***35,852***
7	64,119	11,618	22,293	33,535	39,252	***42,614***	***44,385***	***44,643***	***44,771***	***45,241***	***45,549***
8	68,613	12,402	27,913	39,139	***45,057***	***47,650***	***50,274***	***50,505***	***50,554***	***50,587***	***50,587***
9	74,552	15,095	27,810	***35,521***	***44,066***	***48,308***	***50,061***	***51,337***	***51,904***	***52,016***	***53,895***
10	78,855	16,361	***28,545***	***40,940***	***50,449***	***54,212***	***56,722***	***57,658***	***57,734***	***57,883***	***57,906***

Group 620 - Personal Auto

AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	48,731	15,318	27,740	35,411	40,204	42,388	43,726	44,217	44,277	44,400	44,431
2	49,951	15,031	30,132	37,946	42,371	43,875	44,518	44,738	45,089	45,094	***45,146***
3	52,434	16,994	31,614	39,599	44,943	46,342	47,653	47,866	48,085	***48,097***	***48,241***
4	58,191	17,717	33,767	42,741	46,881	49,117	50,419	50,641	***50,787***	***50,942***	***50,980***
5	61,873	17,842	31,117	39,436	44,871	46,810	47,421	***48,209***	***48,724***	***48,815***	***49,133***
6	63,614	20,266	37,466	45,721	50,641	52,244	***53,241***	***53,794***	***54,093***	***54,468***	***54,471***
7	63,807	18,778	33,216	42,030	47,695	***49,252***	***50,002***	***50,546***	***50,799***	***50,887***	***50,890***
8	61,157	19,900	36,442	43,585	***49,177***	***52,052***	***53,150***	***53,420***	***53,488***	***53,649***	***53,659***
9	62,146	20,395	35,797	***43,816***	***47,687***	***50,468***	***51,085***	***51,598***	***51,754***	***51,756***	***51,914***
10	68,003	20,622	***36,466***	***44,589***	***50,539***	***52,860***	***53,886***	***54,610***	***54,796***	***55,048***	***55,080***

Group 1066 - Commercial Auto

AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	5,103	1,060	3,034	4,580	5,243	4,178	4,347	4,399	4,598	4,582	4,629
2	5,196	1,224	3,751	5,735	4,902	5,295	5,486	5,941	5,976	5,977	***5,977***
3	6,947	1,252	3,568	5,265	6,102	6,607	6,315	6,343	6,370	***6,445***	***6,419***
4	9,482	1,606	3,875	5,439	6,507	8,021	8,098	8,282	***8,300***	***8,328***	***8,378***
5	10,976	1,750	4,038	5,662	6,293	6,779	7,048	***7,048***	***7,047***	***7,047***	***7,047***
6	11,893	1,125	4,322	5,263	6,036	6,462	***6,617***	***6,647***	***6,649***	***6,654***	***6,654***
7	13,029	1,403	3,746	5,800	6,737	***7,078***	***7,110***	***7,225***	***7,346***	***7,366***	***7,366***
8	12,511	1,541	4,620	5,746	***6,171***	***6,462***	***6,680***	***6,714***	***6,713***	***6,728***	***6,729***
9	14,372	1,986	4,532	***4,817***	***5,653***	***5,932***	***5,988***	***6,036***	***6,038***	***6,051***	***6,043***
10	7,371	1,970	***2,730***	***3,214***	***3,376***	***3,502***	***3,605***	***3,744***	***3,750***	***3,777***	***3,780***

Group 1066 - Personal Auto

AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	24,988	5,135	11,980	16,368	18,163	20,189	20,462	20,715	20,749	20,720	20,813
2	26,082	5,655	15,108	19,498	23,097	23,819	24,296	24,622	24,735	24,736	***24,741***
3	29,606	6,648	17,982	23,078	25,334	26,596	26,983	27,096	27,150	***27,195***	***27,206***
4	33,802	5,722	14,677	19,356	21,906	22,497	22,732	23,149	***23,207***	***23,197***	***23,254***
5	37,261	5,906	14,864	18,305	20,075	21,779	22,277	***22,425***	***22,466***	***22,424***	***22,536***
6	35,849	6,439	15,146	19,187	21,576	22,539	***22,941***	***23,037***	***23,029***	***23,135***	***23,174***
7	35,053	6,934	15,703	19,748	21,300	***21,948***	***22,004***	***22,043***	***22,136***	***22,211***	***22,210***
8	33,254	6,194	12,183	15,282	***17,315***	***18,550***	***18,697***	***18,876***	***19,014***	***19,040***	***19,210***
9	29,101	5,314	10,915	***13,854***	***15,179***	***15,537***	***16,083***	***16,057***	***16,088***	***16,101***	***16,137***
10	29,149	4,301	***9,758***	***11,914***	***13,216***	***13,740***	***14,098***	***14,427***	***14,448***	***14,491***	***14,513***

3. THE CHANGING SETTLEMENT RATE (CSR) MODEL

The univariate model used in this paper will be a minor modification to the CSR model used in Meyers (2015). Here is the model. Let:

1. $\alpha_w \sim \text{Normal}\left(0, \sqrt{10}\right)$ for $w = 2, \ldots, 10$. $\alpha_1 = 0$.
2. $logelr \sim$ Uniform(-1, 0.5).
3. $\beta_d \sim$ Uniform(-5, 5) for $d = 1, \ldots, 9$. $\beta_{10} = 0$.
4. $S_1 = 1$, $S_w = S_{w-1} \cdot (1 - \gamma - (w-2) \cdot \delta)$ for $w = 2, \ldots, 10$. $\gamma \sim$ Normal (0, 0.05), $\delta \sim$ Normal(0, 0.01).
5. $\mu_{wd} = \log(\text{Premium}_w) + logelr + \alpha_w + \beta_d \cdot S_w$.
6. $\sigma_d^2 = \sum_{i=d}^{10} a_i$, $a_i \sim$ Uniform(0, 1).
7. $\log(C_{wd}) \sim \text{Normal}(\mu_{wd}, \sigma_d)$.

This model differs from the CSR model described in Meyers (2015) in three aspects.

1. The parameter γ, allows for a speedup (or slowdown when γ is negative) of the claim settlements. By including the δ parameter, this version of the CSR model allows the settlement rate to change over time.
2. Forcing $\alpha_1 = 0$ eliminates some overlap between the α_w parameters and the *logelr* parameter. In the Meyers (2015) version of the model, a constant addition to each α_w parameters could be offset by a subtraction in the *logelr* parameter. Correcting features of this sort tend to speed up convergence of the MCMC algorithm.
3. The MCMC software used for the calculation described in this paper is Stan. See http://mc-stan.org for installation instructions. I have found that, in general, the MCMC algorithm implemented by Stan converges faster than that of JAGS. Stan also allows one to compile a model (in C++) in advance of its use. Using a compiled model can greatly speed up the processing when one uses the same model repeatedly (as we will do below) with different inputs.

The R script that implements this version of the CSR model is available in the appendix spreadsheet. The script produces a sample from the posterior distribution of the parameters for line *X*,

$$\left\{ \left\{ {}_{i}\alpha_{w}^{X} \right\}_{w=2}^{10}, \left\{ {}_{i}\beta_{d}^{X} \right\}_{d=1}^{9}, \left\{ {}_{i}\sigma_{d}^{X} \right\}_{d=1}^{10}, {}_{i}logelr^{X}, {}_{i}\gamma^{X}, {}_{i}\delta^{X} \right\}_{i=1}^{10000}.$$

Following Meyers (2015), the script then simulates 10,000 outcomes $\left\{ {}_{i}C_{w,10}^{X} \right\}_{i=1}^{10000}$ from which we can calculate various summary statistics such as the predictive mean and standard deviation of the outcomes and the percentile of the actual outcome. Table 3.1 gives a summary of the result of these calculations for the Commercial Auto (*X*=*CA*) and the Personal Auto (*X*=*PA*) lines of business.

Figure 3.1 gives the test for uniformity of the predictive percentiles of this version of the CSR model. When compared with Meyers (2015) Figure 22, we see that allowing the claim settlement rate to change over time improves the model so that the percentiles are (within 95% statistical bounds) uniformly distributed for all four lines.

Table 3.1. CSR Models on Illustrative Insurer Data

CA Insurer Group 620			Outcome Percentile = 39.24		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	30,224	22,119	0	0.0000	21,960
2	35,778	21,896	453	0.0207	21,981
3	42,257	30,068	685	0.0228	29,705
4	47,171	34,052	852	0.0250	34,764
5	53,546	36,638	1,106	0.0302	36,286
6	58,004	35,192	1,342	0.0381	35,852
7	64,119	45,387	2,305	0.0508	45,549
8	68,613	53,215	4,061	0.0763	50,587
9	74,552	55,166	7,439	0.1348	53,895
10	78,855	63,922	17,493	0.2737	57,906
Total	553,119	397,656	27,378	0.0688	388,485

PA Insurer Group 620			Outcome Percentile = 65.14		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	48,731	44,535	0	0.0000	44,431
2	49,951	45,453	366	0.0081	45,146
3	52,134	48,304	386	0.0080	48,241
4	58,191	51,003	457	0.0090	50,980
5	61,873	48,335	511	0.0106	49,133
6	63,614	54,243	712	0.0131	54,471
7	63,807	50,779	877	0.0173	50,890
8	61,157	52,674	1,351	0.0256	53,659
9	62,146	52,704	2,437	0.0462	51,914
10	68,003	52,910	5,125	0.0969	55,080
Total	589,907	500,941	8,709	0.0174	503,945

CA Insurer Group 1066			Outcome Percentile = 12.59		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	5,103	4,727	0	0.0000	4,629
2	5,196	6,077	363	0.0597	5,977
3	6,947	6,439	415	0.0645	6,419
4	9,482	7,855	620	0.0789	8,378
5	10,976	7,300	606	0.0830	7,047
6	11,893	6,218	659	0.1060	6,654
7	13,029	7,117	867	0.1218	7,366
8	12,511	7,260	1,160	0.1598	6,729
9	14,372	8,305	2,013	0.2424	6,043
10	7,371	9,299	4,380	0.4710	3,780
Total	96,880	70,597	7,573	0.1073	63,022

PA Insurer Group 1066			Outcome Percentile = 81.50		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	24,988	20,888	0	0.0000	20,813
2	26,082	24,943	290	0.0116	24,741
3	29,606	27,471	367	0.0134	27,206
4	33,802	23,274	328	0.0141	23,254
5	37,261	22,564	367	0.0163	22,536
6	35,849	22,960	466	0.0203	23,174
7	35,053	23,370	605	0.0259	22,210
8	33,254	18,117	669	0.0369	19,210
9	29,101	15,515	985	0.0635	16,137
10	29,149	11,704	1,727	0.1476	14,513
Total	314,145	210,804	3,617	0.0172	213,794

Figure 3.1. Uniformity Tests for the CSR Model

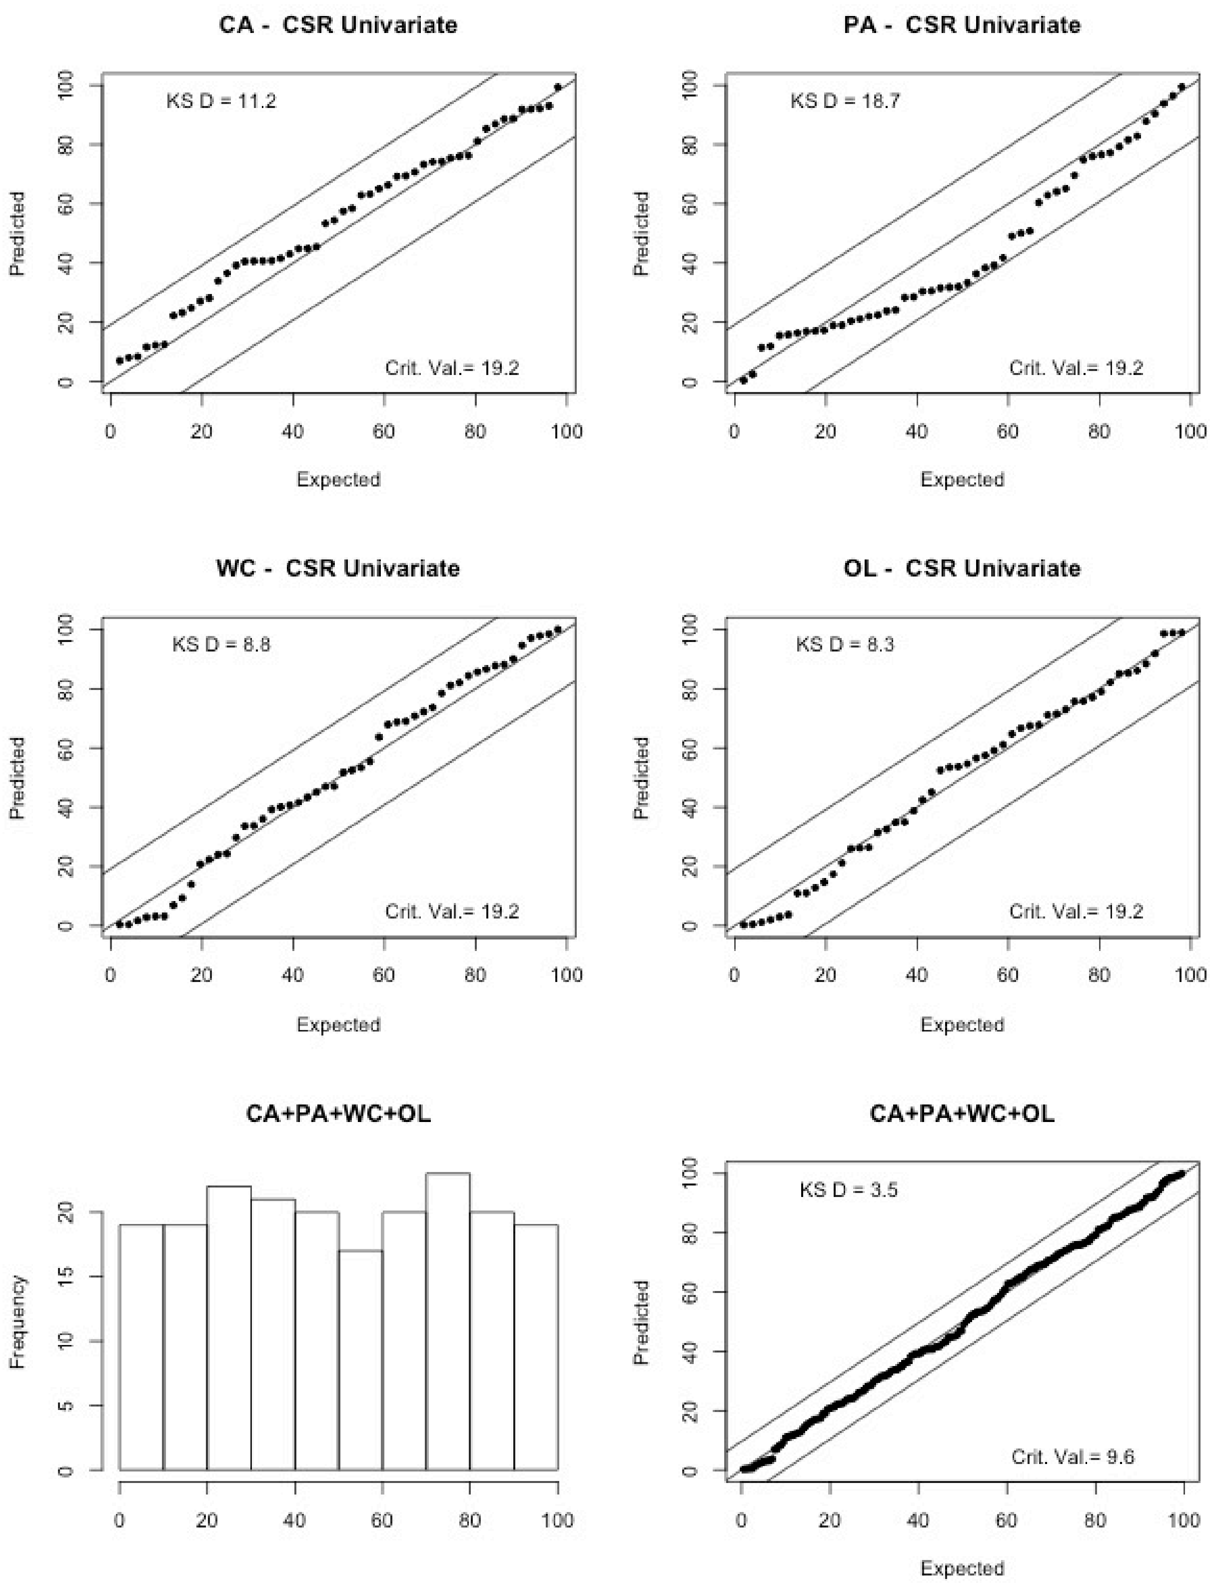

While the observation that the CSR model performs well on a large number of old triangles with outcome data is encouraging, it should not relieve the actuary from testing the assumptions underlying their model of their current data. Traditional tests, such has those provided by Barnett and Zehnwirth (2000) plot residuals (i.e. differences between observed and expected values) along accident year, development year and calendar year dimensions.

The Bayesian MCMC models in this paper provide a sample of size 10,000 from a posterior distribution of parameters. Given that we have this large sample, I consider it to be more informative if we take a subsample, *I*, of (say) size 100, then calculate the standardized residuals for each *w* and *d* in the upper loss triangle, and *i* in the subsample

$$\mathrm{R}_I^X = \left\{ \frac{\log\left(C_{wd}^X\right) - {}_i\mu_{wd}^X}{{}_i\sigma_d^X} \right\}_{i \in I} \qquad (3.1).$$

In general we should expect these residual plots to have a standard normal distribution with mean 0 and standard deviation 1. Figure 3.2 shows plots of these standardized residuals against the accident year, development year and calendar year for the illustrative insurers. I have made similar plots for other insurers as well. For accident years and development years, the plots have always behaved as expected. Deviations for the early calendar years as shown in two of the four plots are not uncommon. I have chosen to regard them as unimportant, and attach more importance to later calendar years.

If the standardized residual plots look like those of the illustrative insurers, we should not have to worry about artificial correlations.

Figure 3.2 – Standardized Residual Plots for the CSR Model

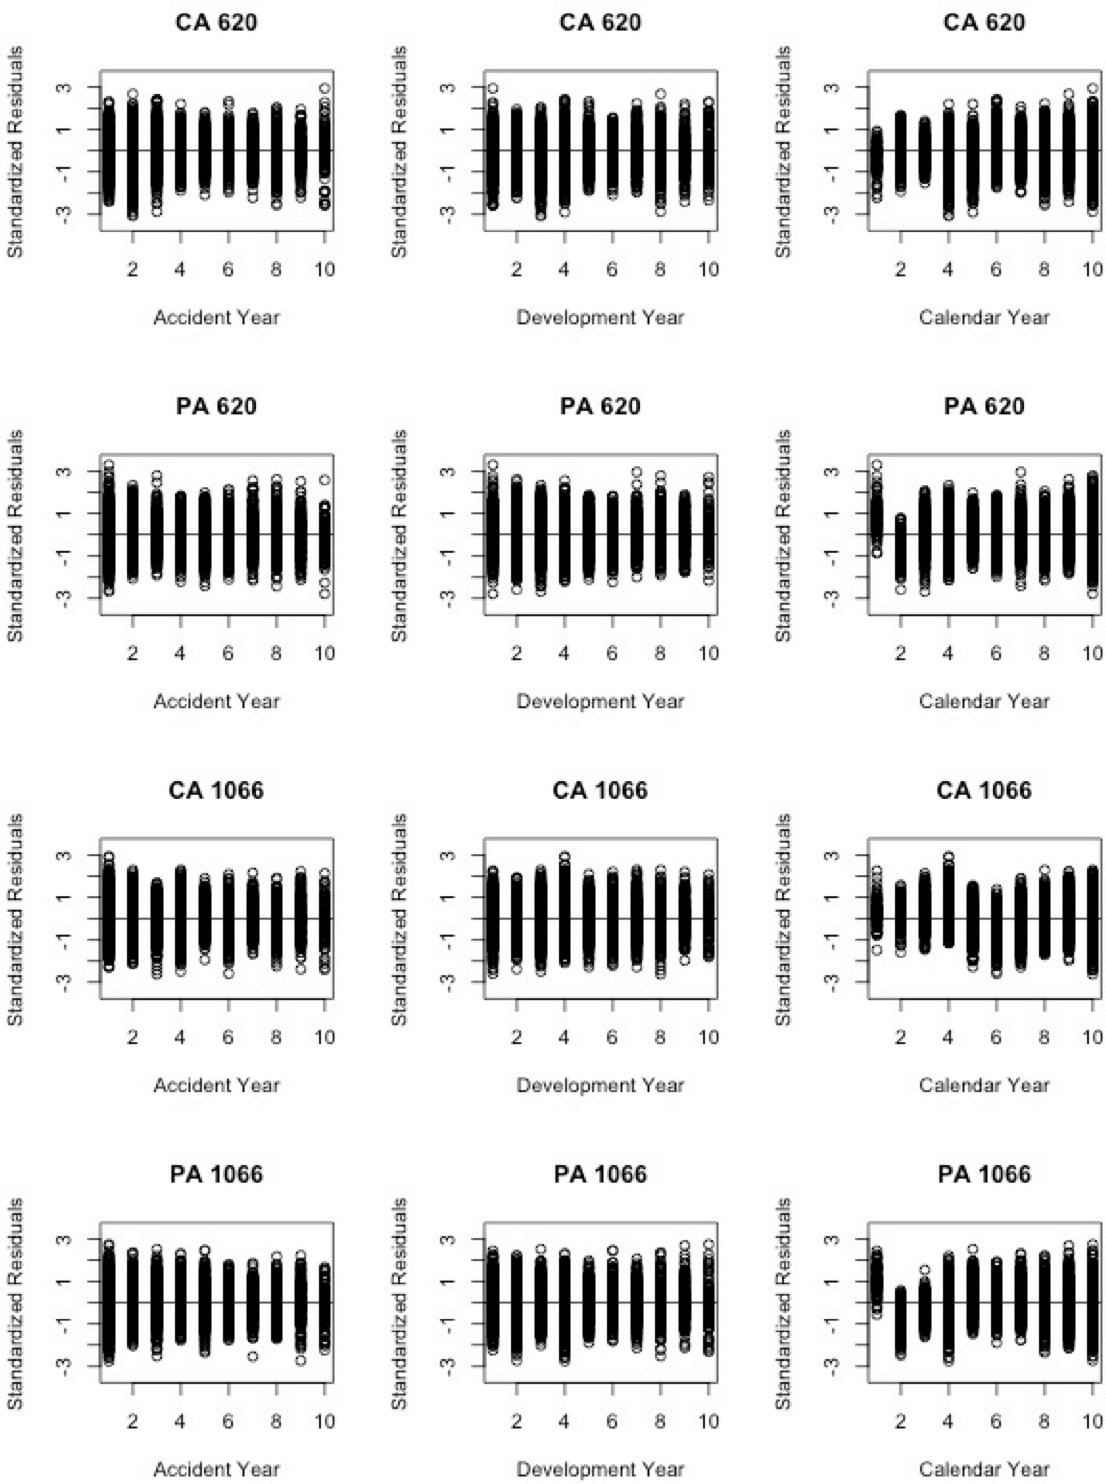

4. A TWO-STEP BIVARIATE MODEL

The last section presented a univariate model that performed well on data in the CAS Loss Reserve Database. This section shows how to construct a bivariate distribution that has the univariate distributions as marginal distributions.

To shorten the notation let

$${}_i\theta^X = \left\{ \left\{ {}_i\alpha_w^X \right\}_{w=2}^{10}, \left\{ {}_i\beta_d^X \right\}_{d=1}^{9}, \left\{ {}_i\sigma_d^X \right\}_{d=1}^{10}, {}_i\,\text{logelr}^X, {}_i\gamma^X, {}_i\delta^X \right\}$$

for line X and $i = 1, \ldots, 10000$.

The first step is to obtain the univariate samples, $\left\{ {}_i\theta^X \middle| \left\{ C_{wd}^X \right\} \right\}_{i=1}^{10,000}$ and $\left\{ {}_i\theta^Y \middle| \left\{ C_{wd}^Y \right\} \right\}_{i=1}^{10,000}$ where $C_{wd}^X \in$ Upper Triangle of line X and $C_{wd}^Y \in$ Upper Triangle of line Y. Then repeatedly for each i, use Bayesian MCMC to take a sample from the posterior distribution of $\left\{ \rho \middle| \left\{ C_{wd}^X \right\}, \left\{ C_{wd}^Y \right\}, {}_i\theta^X, {}_i\theta^Y \right\}$ where ρ has a $\beta(2,2)$ prior distribution translated from (0,1) to (-1,1). Next we randomly select a single ${}_i\rho$ from that sample and use $\left\{ {}_i\theta^X, {}_i\theta^Y, {}_i\rho \right\}_{i=1}^{10,000}$ to calculate the derived parameters in the bivariate distribution given by Equation (1.2). This amounts to using the two univariate distributions as the prior distribution for the second Bayesian step. From that two-step bivariate distribution, one can simulate outcomes from the "posterior" distribution of parameters and calculate any statistic of interest. Be reminded that this can be different from the usual Bayesian posterior distribution $\left\{ {}_i\theta^X, {}_i\theta^Y, {}_i\rho \middle| \left\{ C_{wd}^X \right\}, \left\{ C_{wd}^Y \right\} \right\}_{i=1}^{10,000}$ that comes out of the Zhang-Dukic approach.

At first glance, one might expect the run time for 10,000 Bayesian MCMC simulations to be unacceptably long. But there are a number of considerations that allow one to speed up the calculations.

1. The MCMC simulation is for a single parameter that runs much faster than a multi-parameter simulation that one normally runs with stochastic loss reserve models.
2. We have a good starting value, $\rho = 0$. The burn-in period is short and convergence is rapid.
3. Since we are repeatedly running the same model with different inputs, we need only compile the model once, which the Stan software permits.
4. Using the "parallel" package in R allows one to distribute the simulations to separate cores on a multi-core computer.

Taking these factors into account, my laptop[4] usually turns out this bivariate distribution in about 6 minutes. As I mentioned above, the R scripts that produce these calculations are made available to the reader in the Appendix.

The purpose of getting a bivariate distribution is to predict the distribution of the sum of the outcomes for the two lines of insurance. Table 4.1 gives results analogous to Table 3.1 for the sum of CA and PA lines for the two illustrative insurers. Also included are the sums of the two lines predicted under the assumption of independence. Figure 4.1 contains histograms of the two-step posterior distributions for ρ for the illustrative insurers.

[4] Apple MacBook Pro with quad-core processor – purchased in late 2013.

Table 4.1. Combined CSR Models on Illustrative Insurer Data[5]

Two-Step Bivariate Model

	Insurer Group 620		Outcome Percentile = 41.34		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	78,955	66,647	0	0.0000	66,391
2	85,729	66,928	567	0.0085	67,127
3	94,691	78,326	844	0.0108	77,946
4	105,362	85,532	1,006	0.0118	85,744
5	115,419	84,311	1,189	0.0141	85,419
6	121,618	90,935	1,516	0.0167	90,323
7	127,926	95,970	2,452	0.0255	96,439
8	129,770	105,814	4,259	0.0402	104,246
9	136,698	107,501	7,770	0.0723	105,809
10	146,858	118,650	18,581	0.1566	112,986
Total	1,143,026	900,614	28,999	0.0322	892,430

Independence Assumption

	Insurer Group 620		Outcome Percentile = 44.05		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	78,955	66,654	0	0.0000	66,391
2	85,729	67,349	586	0.0087	67,127
3	94,691	78,372	789	0.0101	77,946
4	105,362	85,054	972	0.0114	85,744
5	115,419	84,973	1,225	0.0144	85,419
6	121,618	89,435	1,528	0.0171	90,323
7	127,926	96,166	2,470	0.0257	96,439
8	129,770	105,890	4,269	0.0403	104,246
9	136,698	107,870	7,810	0.0724	105,809
10	146,858	116,833	18,172	0.1555	112,986
Total	1,143,026	898,597	28,667	0.0319	892,430

Two-Step Bivariate Model

	Insurer Group 1066		Outcome Percentile = 26.62		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	30,091	25,610	0	0.0000	25,442
2	31,278	30,680	444	0.0145	30,718
3	36,553	33,773	580	0.0172	33,625
4	43,284	31,429	710	0.0226	31,632
5	48,237	29,451	693	0.0235	29,583
6	47,742	30,056	813	0.0270	29,828
7	48,082	30,392	1,047	0.0344	29,576
8	45,765	25,397	1,329	0.0523	25,939
9	43,473	23,672	2,208	0.0933	22,180
10	36,520	21,726	5,018	0.2310	18,293
Total	411,025	282,187	8,617	0.0305	276,816

Independence Assumption

	Insurer Group 1066		Outcome Percentile = 29.78		
w	Premium	Estimate	Std. Dev.	C.V.	Outcome
1	30,091	25,615	0	0.0000	25,442
2	31,278	31,020	464	0.0150	30,718
3	36,553	33,910	552	0.0163	33,625
4	43,284	31,129	705	0.0226	31,632
5	48,237	29,864	709	0.0237	29,583
6	47,742	29,178	810	0.0278	29,828
7	48,082	30,487	1,056	0.0346	29,576
8	45,765	25,376	1,333	0.0525	25,939
9	43,473	23,819	2,236	0.0939	22,180
10	36,520	21,003	4,733	0.2253	18,293
Total	411,025	281,401	8,417	0.0299	276,816

[5] I attribute the differences in the "Estimate" column by insurer to simulation error. The expected values for the bivariate and independence assumptions are equal.

Figure 4.1 – Posterior Distribution of ρ for CSR Model

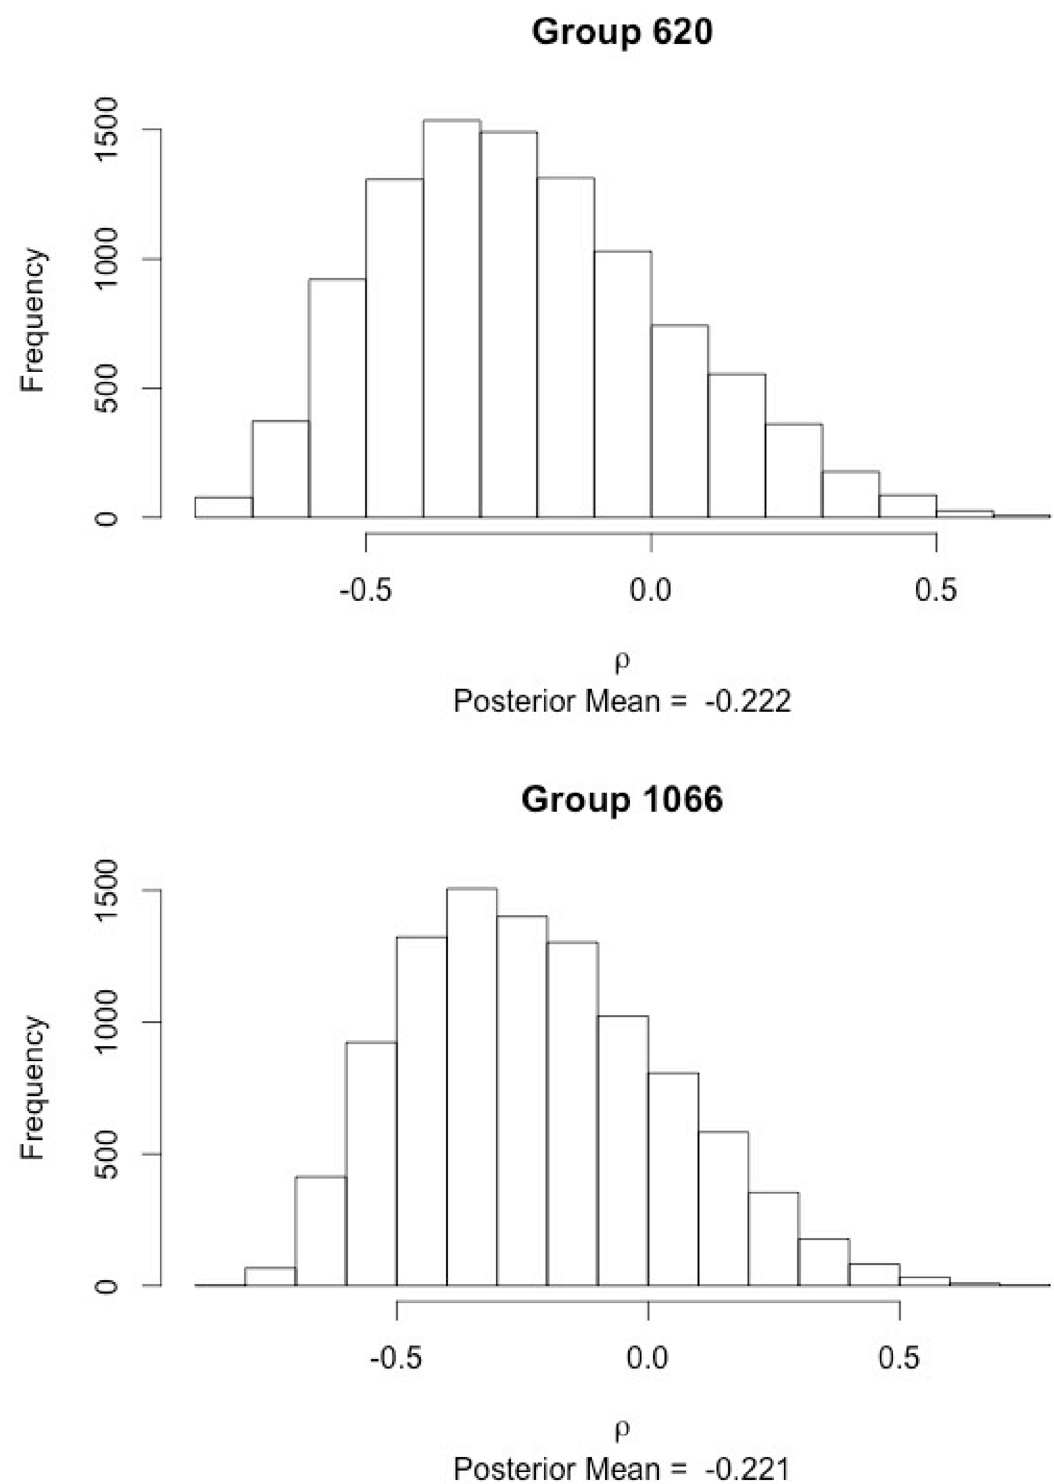

Table 4.1 and Figure 4.1 are notable in two aspects. First, the output from the bivariate model is not all that different from the output created by taking independent sums of losses from the univariate model. Second, the posterior distributions of ρ from the two-step bivariate model have a fairly wide range. The posterior distributions of ρ for both groups are predominantly negative.

Typically the posterior mean ρ over all the within-group pairs of lines is not all that different from zero. Figure 4.3 shows the frequency distribution of posterior mean ρs from the insurer group sample.

Figure 4.3 – Posterior Mean ρs from the Within-Group Pairs of Lines

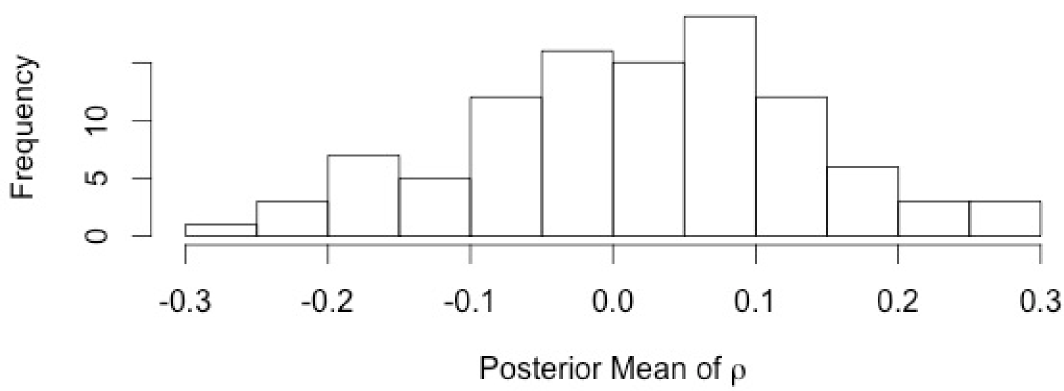

This section concludes with a test of uniformity of the outcome percentiles of the within-group pairs for the sum of two lines predicted by the two-step bivariate model and the independence assumption. As Figure 3.1 shows, the univariate models pass our uniformity test, one would think that a valid bivariate model would also pass a uniformity test. Figures 4.4 and 4.5 show the results.

It turns out that both the two-step bivariate model and the independence assumption pass the uniformity test, with the independence assumption performing slightly better. This suggests that the lines of insurance are independent for many, if not all, insurers. In the next section we will examine the independence assumption for individual pairs of loss triangles.

Figure 4.4 - Uniformity Tests of Outcome Percentiles

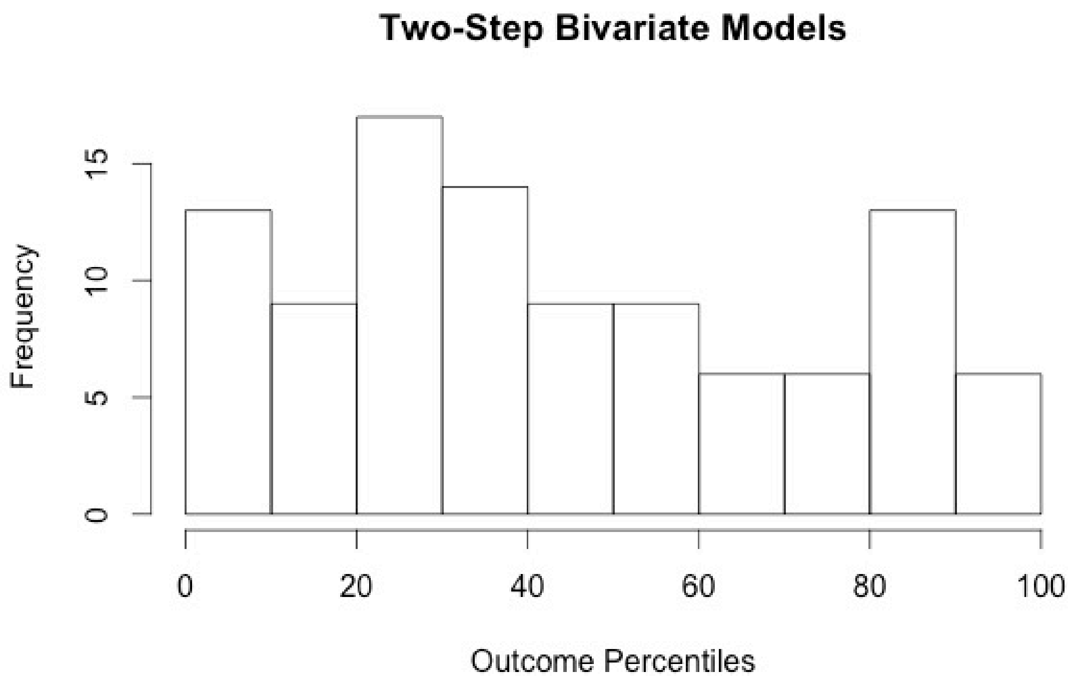

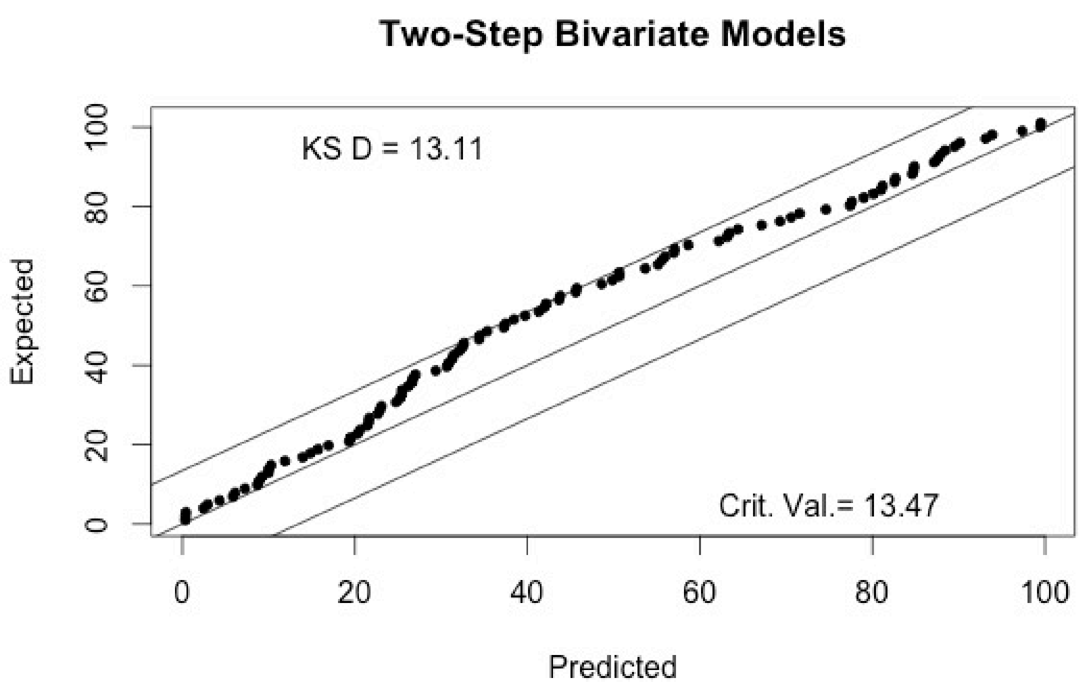

Figure 4.5 - Uniformity Tests of Outcome Percentiles

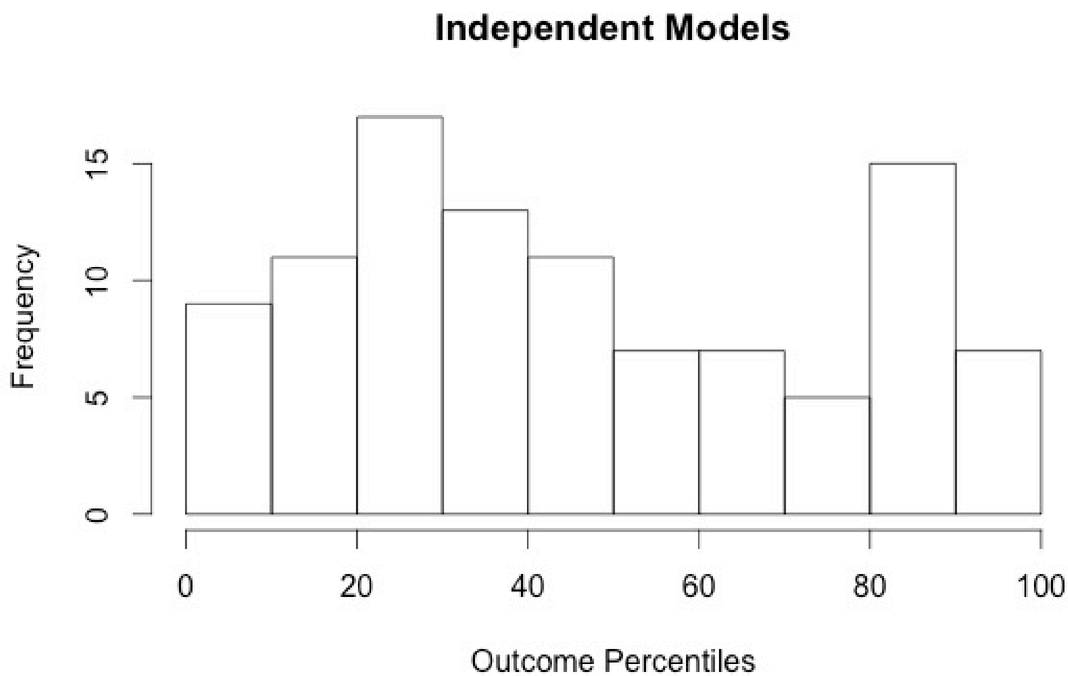

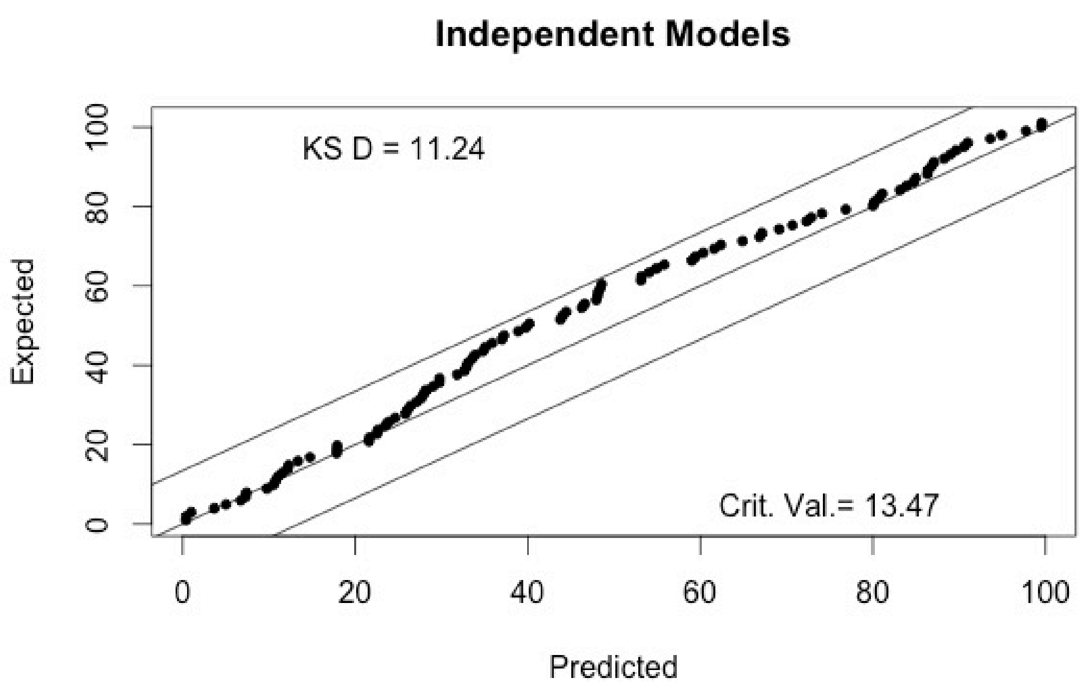

5. MODEL SELECTION[6]

Let's start the discussion with a review of the Akaike Information Criteria (AIC).

Suppose that we have a model with a data vector, $\boldsymbol{x}$, and a parameter vector θ, with p parameters. Let $\hat{\boldsymbol{\theta}}$ be the parameter value that maximizes the log-likelihood, L, of the data, $\boldsymbol{x}$. Then the AIC is defined as

$$\text{AIC} = 2 \cdot p - 2 \cdot L\left(\hat{\theta}\right). \tag{5.1}$$

Given a choice of models, the model with the lowest AIC is to be preferred. This statistic rewards a model for having a high log-likelihood, but it penalizes the model for having more parameters.

There are problems with the AIC in a Bayesian MCMC environment. Instead of a single maximum likelihood estimate of the parameter vector, there is an entire sample of parameter vectors taken from the model's posterior distribution. There is also the sense that the penalty for the number of parameters should not be as great in the presence of strong prior information.

To address these concerns, Gelman *et. al.* (2014) and Vehtari and Gelman (2014) describe a statistic, called the Watanabe-Akaike Information Criterion (WAIC) that generalizes the AIC in a way that is appropriate for Bayesian MCMC models[7].

First define the computed log pointwise predictive density (made specific for this paper) as

$$L_{\text{WAIC}} = \sum_{j=1}^{55} \log\left(\frac{1}{10000} \sum_{i=1}^{10000} \phi\left(\log\left(C^X_{w(j),d(j)}\right), \log\left(C^Y_{w(j),d(j)}\right) \Big|\, {}_i\theta^X, {}_i\theta^Y, {}_i\rho \right) \right). \tag{5.2}$$

where ϕ is a multivariate normal distribution such as that given in Equation(1.2). The L_{WAIC} statistic replaces the log-likelihood L in Equation(5.1) with an average log-likelihood taken over the sample from the posterior distribution.

Next, define the effective number of parameters p_{WAIC} as

[6] For more information about the model selection statistics in this section, see Section 7.2 of Gelman, *et. al.*.

[7] Another popular statistic designed for Bayesian MCMC models is the Deviance Information Criterion (DIC) that is available in the MCMC software WINBUGS and JAGS. Gelman *et. al.* (2014) and Vehtari and Gelman (2014) make the case that the WAIC is a better statistic as it is based on the entire sample from the posterior distribution as opposed to a point estimate.

$$p_{WAIC} = \sum_{j=1}^{55} Var_i \left[log\left(\phi\left(log\left(C^X_{w(j),d(j)}\right), log\left(C^Y_{w(j),d(j)}\right) \middle| {}_i\theta^X, {}_i\theta^Y, {}_i\rho\right)\right)\right]. \quad (5.3)$$

p_{WAIC} has the property that it decreases with the tightness of the prior distribution. Of possible general interest is that Vehtari and Gelman discuss situations, e.g. flat priors and a large number of data points, where p_{WAIC} is equal to the nominal number of parameters, p. But none of that applies to the examples in this paper where we have only 110 observations, some non-flat priors and, in addition, have some constraints between some of the parameters.

The final expression for the WAIC is analogous to Equation (5.1) and is given by

$$\mathrm{WAIC} = 2 \cdot p_{WAIC} - 2 \cdot L_{WAIC}. \quad (5.4)$$

The WAIC statistics for the bivariate two-step model were calculated with the posterior distribution $\left\{ {}_i\theta^X, {}_i\theta^Y, {}_i\rho \right\}_{i=1}^{10000}$. For the model that assumes independence of the univariate models, the WAIC statistics were calculated with the posterior distribution $\left\{ {}_i\theta^X, {}_i\theta^Y, 0 \right\}_{i=1}^{10000}$. Table 5.1 gives these statistics for the illustrative insurers. The lower WAIC statistic for the assumption of independence is the preferred model for both insurer groups.

Table 5.1 – WAIC Statistics for the Illustrative Insurer Groups For the CSR Model

Group	Model	p_{WAIC}	L_{WAIC}	WAIC
620	Bivariate	31.09	255.31	-448.44
620	Independent	27.23	252.92	-451.38
1066	Bivariate	30.89	180.41	-299.04
1066	Independent	27.12	178.03	-301.82

The WAIC statistics for the *all* the within-group pairs, given in the Appendix, indicate that the assumption of independence is the preferred model!

6. ILLUSTRATION OF MODEL SENSITIVITY

In discussions with my actuarial colleagues over the years, I have sensed a general consensus among most actuaries is that there is some degree of dependence between the various lines of insurance. But as pointed out in the introduction, using a suboptimal model can lead to artificial dependencies. This section takes a stochastic version of a currently popular model and demonstrates that it is suboptimal for our sample of insurers. It also shows that given this model, there are significant dependencies between the various lines of insurance suggesting that the "general consensus" is understandable given the state of the art that has existed over the years.

One of the most popular loss reserving methodologies is given by Bornhuetter-Ferguson (1972). A key input to the loss reserve formula given in that paper is the expected loss ratio, which must be judgmentally selected by the actuary. Presentations by Clark (2013) and Leong (2013) suggest that the Bornhuetter-Ferguson method that assumes a constant loss ratio provides a good fit to industry loss reserve data.

Actuaries who want to use data to select the expected loss ratio can use the "Cape Cod" model that is given by Stanard (1985). A stochastic version of the Cape Code model can be expressed as a special case of the CSR model by setting the parameters $\alpha_w = 0$ for $w = 1,\ldots, 10$, $\gamma = 0$ and $\delta = 0$. Let's call this model the Stochastic Cape Cod (SSC) model.

Figure 6.1 gives the standardized residual plots of the SCC model for the illustrative insurers that are analogous to those in Figure 3.2. Figure 6.2 gives the posterior distribution of the ρ parameters for the two-step bivariate SCC model. Noteworthy is that the mean ρ for Group 1066 is quite high compared to any of the results for the CSR model. In Table 6.1 we see that the WAIC statistic for Group 1066 is lower for the bivariate model than the independent model indicating that the bivariate model is favored. The reverse is true for Group 620.

Figure 6.1 – Standardized Residual Plots for the SSC Model

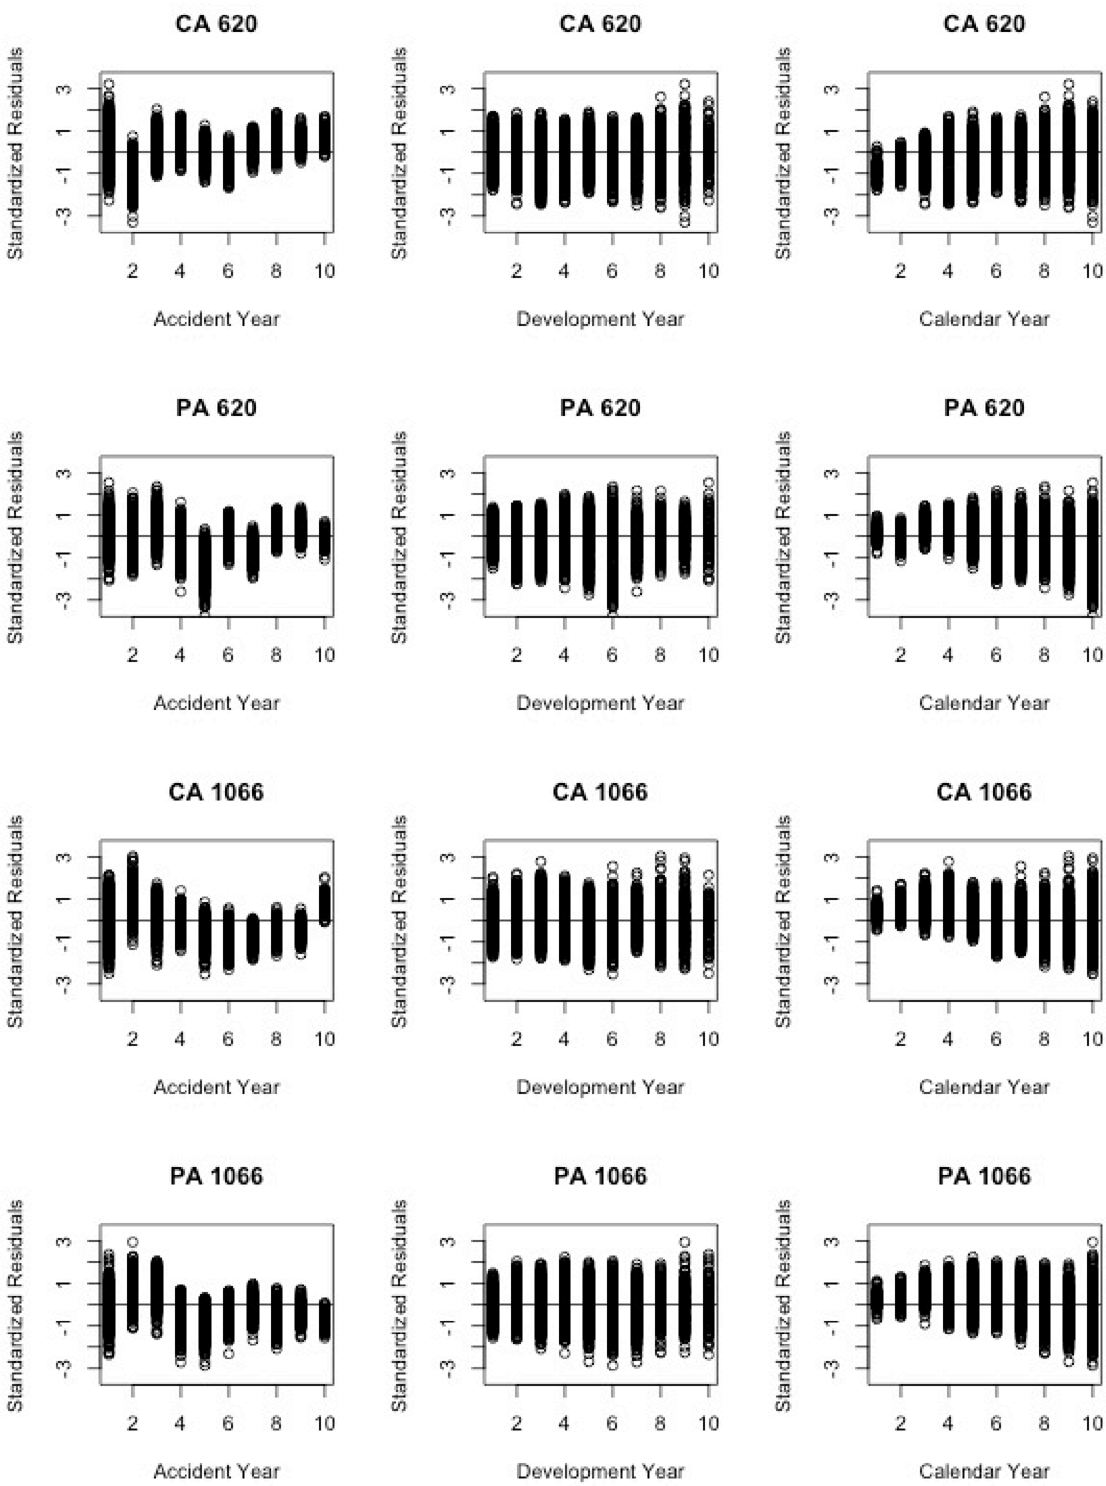

Figure 6.2 – Posterior Distribution of ρ for SCC Model

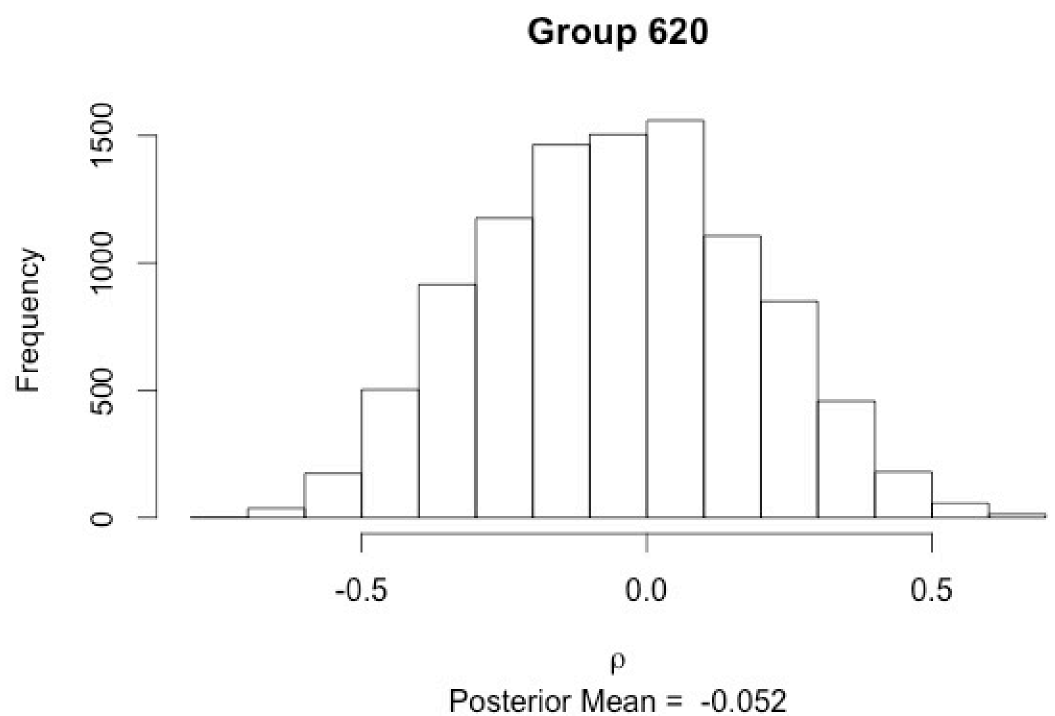

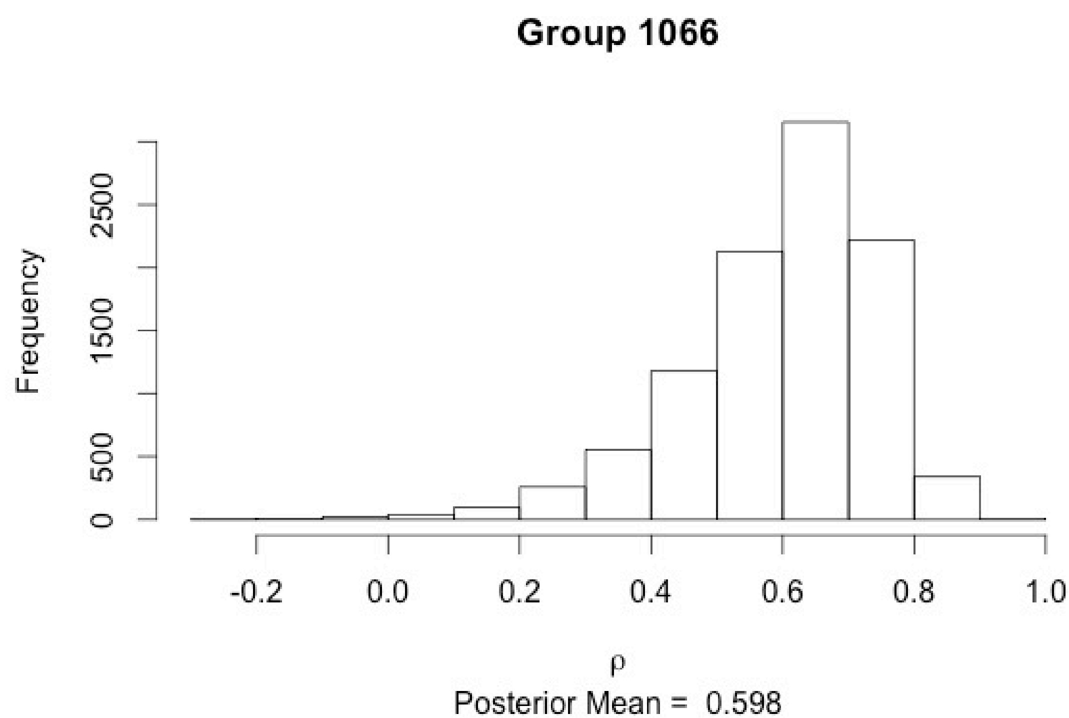

Table 6.1 – WAIC Statistics for the Illustrative Insurer Groups For the SCC Model

Group	Model	p_{WAIC}	L_{WAIC}	WAIC
620	Bivariate	17.79	122.10	-208.62
620	Independent	16.34	121.55	-210.42
1066	Bivariate	23.61	18.99	9.24
1066	Independent	14.53	6.68	15.7

By examining the standardized residual plots in Figure 6.1 across accident years we can see a possible explanation for these results. First note that the standardized residual plots for the SCC model are not as well behaved as the similar plots for the CSR model in Figure 3.2. But the pattern of the errors in the CA and PA plots are dissimilar for Group 620, but similar for Group 1066. The similarity of the plots for Group 1066 leads to the overwhelmingly positive posterior distribution of ρ, and the indicated preference of the bivariate model over the independent model.

Over the entire sample of insurer groups, the bivariate model was the preferred model for 39 of the 102 within-group pairs of triangles.

It is also worth noting that, as shown in Figure 6.3, the SCC model fails the uniformity test that the CSR model passed, as shown in Figure 3.1.

Here we see an example where the suboptimal SCC model leads to artificial dependencies between lines, whereas the less suboptimal CSR model leads to independence between lines for our sample of insurer loss triangles.

Figure 6.3. Uniformity Tests for the SCC Model

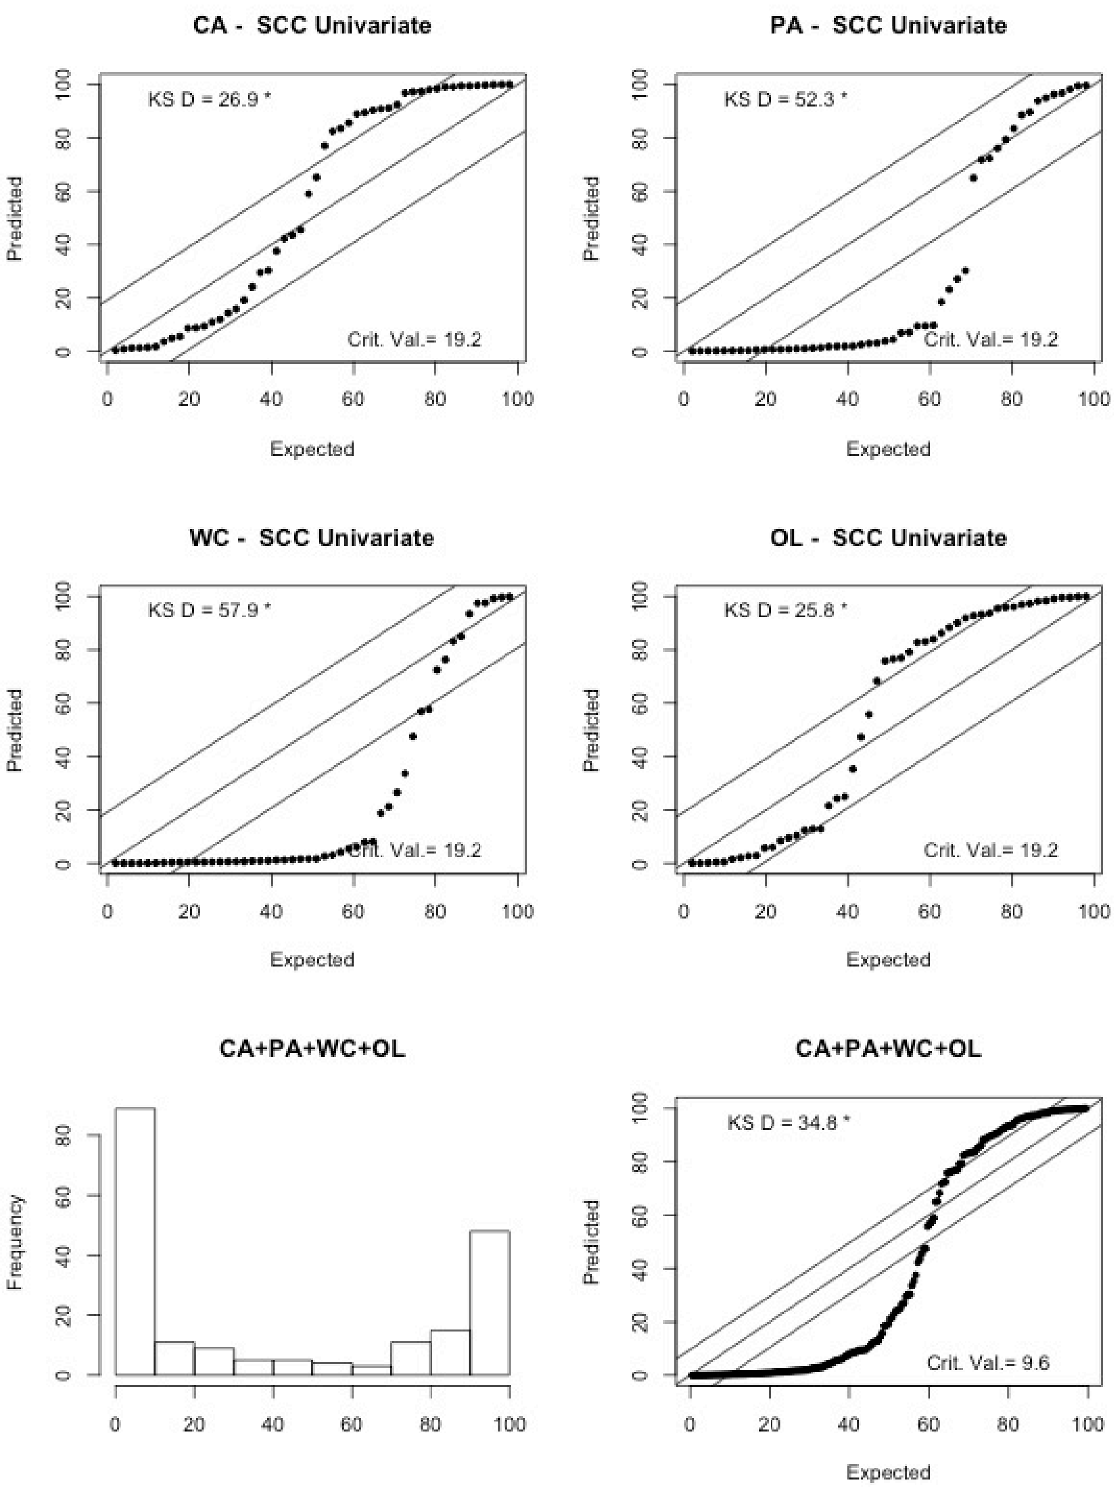

7. SUMMARY AND CONCLUSIONS

The purpose of this paper was to illustrate how to build a model that creates a bivariate distribution given two univariate Bayesian MCMC models that preserve the original univariate distributions. While this modeling technique was applied to lognormal stochastic loss reserve models, it should not be difficult to apply this two-step approach to other Bayesian MCMC models using bivariate copulas as was done by Zhang and Dukic.

While a statistical study such as that done in this paper can never carry the weight of a mathematical proof, its conclusion was derived from the analysis of a large number of within-group pairs of loss triangles. It should be noted that these loss triangles came from NAIC Schedule Ps reported in the same year.

The conclusion that the within-group pairs of loss triangles are independent for the CSR model may come as a surprise to some. But the evidence supporting this conclusion is as follows.

1. The univariate models pass two fairly restrictive tests (i.e. the retrospective test in Figure 3.1 and the standardized residual tests in Figure 3.2) that could disqualify many suboptimal models. Thus we should not expect to see an artificial appearance of dependency due to a bad model.
2. The retrospective results of Section 4 indicate support the independence assumption for the bivariate two-step model. The range of ρs for the 102 within-group pairs contained both positive and negative values, which appear to be random in light of the tests performed in this paper.

I feel fortunate that I was able to find a model that indicated independence between lines of insurance. Before taking on this line of research, there was no guarantee that I would be able to find such a model. In fact, initially I did not believe the independence results that I was getting. The lesson learned is that if one has a model with statistically significant dependences between lines of insurance, one should search for a more optimal model.

The reason that the dependency problem is so important is that risk-based insurer solvency standards are based on the total risk to the insurance company. Ignoring a true dependency could understate the total risk faced by an insurer. On the other hand, too stringent of a solvency standard could limit the supply of insurance. If this holds, then the current practice in some jurisdictions could limit the supply of insurance.

While retrospective results can be informative, there is a need for criteria testing the independence assumption that can be applied prospectively. That was the purpose of Section 5. The prospective test consists of; (1) fitting the two-step bivariate model; (2) fitting a bivariate model that assumes independence; and (3) calculating the WAIC statistic to see which model is favored. It turned out that the WAIC statistic favored the independence assumption in *every one* of the 102 within-group pairs of triangles.

So for now, the CSR model with the independence assumption is looking pretty good. But in light of the high stakes involved, assumptions of this sort need a stringent peer review and replication with new and different data. I look forward to seeing this happen.

REFERENCES

[1.] Barnett, Glen and Ben Zehnwirth. 2000. "Best Estimates for Reserves." *Proceedings of the Casualty Actuarial Society* 87:254-321.

[2.] Bornhuetter, Ronald L and Ronald E. Ferguson. 1972. "The Actuary and IBNR." *Proceedings of the Casualty Actuarial Society* LIX, p. 181.

[3.] Clark, David, 2013. "Where Are We In the Market Cycle." Presentation to CAMAR, October 9, 2013. http://www.casact.org/community/affiliates/camar/1013/Clark.pdf

[4.] England, P.D., and R. J. Verrall. 2002. "Stochastic Claims Reserving in General Insurance." Meeting Paper. Presented to the Institute and Faculty of Actuaries, 28 January.

[5.] Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2014. *Bayesian Data Analysis*. 3rd ed., CRC Press.

[6.] Herzog, Thomas N. 2011. "Summary of CEIOPS Calibration Work on Standard Formula." http://www.naic.org/documents/index_smi_solvency_ii_calibration.pdf.

[7.] Leong, Jessica (Wehn Kah). 2013. "Property versus Casualty Risks." Presentation to CAMAR, October 9, 2013. http://www.casact.org/community/affiliates/camar/1013/Leong-Hayes.pdf

[8.] Mack, Thomas. 1993. "Distribution-Free Calculation of the Standard Error on Chain Ladder Reserve Estimates." *ASTIN Bulletin* 23(2):213-225.

[9.] Mack, Thomas. 1994. "Measuring the Variability of Chain Ladder Reserve Estimates." Casualty Actuarial Society *Forum* (Spring):101:182.

[10.] Meyers, Glenn G. 2015. "Stochastic Loss Reserving Using Bayesian MCMC Models." CAS Monograph Series, No. 1.

[11.] Stanard, James N. 1985. "A Simulation Test of Prediction Errors of Loss Reserve Estimation Techniques." *Proceedings of the Casualty Actuarial Society* LXXII, p. 124.

[12.] Vehtari, Aki and Gelman, Andrew. 2014. "WAIC and Cross-Validation in Stan." http://www.stat.columbia.edu/~gelman/research/unpublished/waic_stan.pdf

[13.] Zhang, Yanwei and Vanja Dukic. 2013. "Predicting Mulitvariate Insurance Loss Payments Under the Bayesian Copula Framework." *The Journal of Risk and Insurance*, Vol. 80, No. 4, 891-919.