

Dependencies in Stochastic Loss Reserve Models

by Glenn Meyers

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 two Bayesian model selection statistics to compare the fits of the two models.

KEYWORDS

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

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, $w = 1, \dots, K$ and development year $d = 1, \dots, K$. 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

$$\begin{aligned}\log(C_{wd}^X) &\sim \text{Normal}(\mu_{wd}^X, \sigma_d^X) \\ \log(C_{wd}^Y) &\sim \text{Normal}(\mu_{wd}^Y, \sigma_d^Y)\end{aligned}\quad (1.1)$$

¹As 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.

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

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 $\{\{\mu_{wd}^X\}, \{\sigma_d^X\}_{i=1}^{10000}\}$ 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).

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

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:

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

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 sub-optimal model might produce artificial dependencies. To illustrate, consider Figure 1.1, 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:

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

Figure 1.1. Illustration of artificial correlation



The upper plots of Figure 1.1 show the pairs $(x, y_1(x))$ and $(x, y_2(x))$ with the suboptimal constant models plotted in red, and the correct parabolic models plotted in green. The lower left plot shows a scatter plot of $y_1(x) - \mu_1$ and $y_2(x) - \mu_2$ for the suboptimal model. 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. Avanzi, Taylor and Wong (2016) discuss this phenomenon in greater detail.

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).

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 suit-

able for studying dependency models. Preferring to use loss triangles that had already been vetted, I decided to stick with these within-group pairs of triangles.

This paper will provide detailed analyses for two illustrative insurers (Group 620 - Lines CA and PA and Group 5185 Lines CA and OL). The complete loss triangles and outcomes are in Table 2.1. 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 zip archive titled "Appendix." The files in the Appendix contain:

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

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

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

Table 2.1. Data for illustrative insurers

Variance Advancing the Science of Risk

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 - Commercial 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 5185 - Other Liability											
AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	6,093	2,333	4,977	5,798	6,054	6,422	6,494	6,639	6,857	6,856	6,859
2	7,396	2,899	4,044	4,800	6,005	6,252	6,295	6,304	6,333	6,336	6,336
3	8,038	1,750	3,180	4,381	5,184	5,989	6,230	6,258	6,278	6,301	6,304
4	9,366	2,329	3,886	5,172	6,016	6,586	7,096	7,240	7,424	7,512	7,554
5	10,677	2,652	4,828	6,347	7,356	7,834	8,099	8,247	8,554	8,555	8,557
6	11,012	2,823	5,139	7,058	8,053	8,734	8,867	9,008	9,019	9,025	9,025
7	11,768	2,689	5,381	7,725	9,384	10,125	10,407	10,699	11,470	11,696	11,694
8	12,777	3,397	5,582	8,005	10,035	10,512	11,003	11,072	11,080	11,064	11,062
9	13,803	3,365	6,344	8,064	9,381	9,851	10,269	10,375	10,522	10,523	10,524
10	14,575	4,077	6,753	8,608	9,676	10,174	10,429	10,627	10,635	10,662	10,662
Group 5185 - Other Liability											
AY	Premium	DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
1	15,376	3,517	6,817	8,191	9,070	9,656	10,033	10,106	10,179	10,298	10,298
2	16,510	4,153	7,230	9,110	10,938	11,751	11,951	12,035	12,103	12,280	12,243
3	15,596	3,923	6,945	8,877	10,104	11,003	11,476	11,596	11,902	11,938	11,973
4	18,190	4,051	7,725	10,094	11,854	12,822	13,326	13,348	13,364	13,435	13,437
5	19,207	4,580	7,902	10,097	10,862	12,339	12,870	12,951	13,019	13,041	13,071
6	19,488	4,418	8,607	10,433	11,871	12,397	12,773	13,088	13,144	13,239	13,359
7	21,585	5,151	9,996	13,269	15,162	15,093	15,336	15,513	15,756	15,787	15,846
8	22,227	5,529	8,919	10,729	11,891	12,426	12,769	13,060	13,105	13,296	13,210
9	22,983	6,921	11,255	14,436	16,191	17,793	18,427	18,708	19,025	19,050	19,044
10	25,612	6,346	11,181	14,083	15,838	17,152	17,798	17,971	18,167	18,485	18,532

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 tends 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 ,

$$\left\{ \begin{array}{l} \{ \alpha_w^X \}_{w=2}^{10}, \{ \beta_d^X \}_{d=1}^9, \{ \sigma_d^X \}_{d=1}^{10}, \\ \logelr^X, \gamma^X, \delta^X \end{array} \right\}_{i=1}^{10000}.$$

Following Meyers (2015), the script then simulates 10,000 outcomes $\{C_{w,10}^X\}_{i=1}^{10000}$ 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 illustrative insurers in their two 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.

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 as 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, S , of (say) size 100, then calculate the standardized residuals for each w and d in the upper loss triangle, and s in the subsample

$$R_S^X = \left\{ \frac{\log(C_{wd}^X) - s\mu_{wd}^X}{s\sigma_d^X} \right\}_{s \in S} \quad (3.1)$$

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.

4. Fitting bivariate distributions

The previous 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 with the CSR model as marginal distributions.

Table 3.1. CSR Models on illustrative insurer data

CA Insurer Group 620			Outcome Percentile = 39.09		
w	Premium	Estimate	Std. Dev.	CV	Outcome
1	30,224	22,123	0	0	21,960
2	35,778	21,890	482	0.022	21,981
3	42,257	30,077	725	0.024	29,705
4	47,171	34,062	889	0.026	34,764
5	53,546	36,664	1,135	0.031	36,286
6	58,004	35,167	1,339	0.038	35,852
7	64,119	45,397	2,307	0.051	45,549
8	68,613	53,183	4,066	0.077	50,587
9	74,552	55,163	7,284	0.132	53,895
10	78,855	63,493	16,849	0.265	57,906
Total	553,119	397,219	26,901	0.068	388,485

PA Insurer Group 620			Outcome Percentile = 63.82		
w	Premium	Estimate	Std. Dev.	CV	Outcome
1	48,731	44,538	0	0	44,431
2	49,951	45,463	366	0.008	45,146
3	52,434	48,315	379	0.008	48,241
4	58,191	51,007	460	0.009	50,980
5	61,873	48,343	512	0.011	49,133
6	63,614	54,272	728	0.013	54,471
7	63,807	50,804	861	0.017	50,890
8	61,157	52,726	1,338	0.025	53,659
9	62,146	52,780	2,362	0.045	51,914
10	68,003	53,015	5,038	0.095	55,080
Total	589,907	501,264	8,529	0.017	503,945

CA Insurer Group 5185			Outcome Percentile = 21.82		
w	Premium	Estimate	Std. Dev.	CV	Outcome
1	6,093	6,937	0	0	6,859
2	7,396	6,517	214	0.033	6,336
3	8,038	6,228	227	0.036	6,304
4	9,366	7,102	292	0.041	7,554
5	10,677	8,491	401	0.047	8,557
6	11,012	9,100	543	0.060	9,025
7	11,768	10,718	813	0.076	11,694
8	12,777	11,181	1,245	0.111	11,062
9	13,803	12,732	2,469	0.194	10,524
10	14,575	17,192	6,606	0.384	10,662
Total	105,505	96,199	9,917	0.103	88,577

OL Insurer Group 5185			Outcome Percentile = 56.12		
w	Premium	Estimate	Std. Dev.	CV	Outcome
1	15,376	10,354	0	0	10,298
2	16,510	12,339	209	0.017	12,243
3	15,596	11,809	220	0.019	11,973
4	18,190	13,496	286	0.021	13,437
5	19,207	13,147	314	0.024	13,071
6	19,488	13,253	407	0.031	13,359
7	21,585	17,111	672	0.039	15,846
8	22,227	14,225	820	0.058	13,210
9	22,983	18,194	1,808	0.099	19,044
10	25,612	16,838	3,562	0.212	18,532
Total	196,774	140,766	6,055	0.043	141,013

Figure 3.1. Uniformity tests for CSR model






Figure 3.1. Uniformity tests for CSR model (continued)**Figure 3.2. Standardized residual plots for the CSR model**

To shorten the notation let

$${}_i\theta^X = \left\{ \begin{array}{l} \{ {}_i\alpha_w^X \}_{w=2}^{10}, \{ {}_i\beta_d^X \}_{d=1}^9, \{ {}_i\sigma_d^X \}_{d=1}^{10}, \\ {}_i\log\text{elr}^X, {}_i\gamma^X, {}_i\delta^X \end{array} \right\}$$

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

Let's first consider what happens when we try the Zhang-Dukic approach with the CSR model. Let's call this model the ZD-CSR model. Using Bayesian MCMC to get a sample from the posterior distribution for $X=CA$ and $Y=PA$ of Insurer #620 yields:

$$\{ {}_i\theta^{CA}, {}_i\theta^{PA}, \rho_i | \{ C_{wd}^{CA} \}, \{ C_{wd}^{PA} \} \}_{i=1}^{10,000}$$

The prior distributions for the θ^X parameters are the same as they are for the CSR model. The prior distribution for the ρ parameter is a $\beta(2,2)$ prior distribution translated from $(0,1)$ to $(-1,1)$.

Using the parameter sets

$$\{ {}_i\theta^{CA} \}_{i=1}^{10,000} \text{ and } \{ {}_i\theta^{PA} \}_{i=1}^{10,000}$$

I then simulated the marginal distributions, given in Table 4.1, for each line of insurance. Notice that the estimates and the standard deviations are significantly higher for marginal distributions of the ZD-CSR model than for the corresponding univariate distributions in Table 3.2. Moreover, the posterior distribution of ρ in the ZD-CSR model, given in Figure 4.1, appears to be clustered close to negative 1. This suggests that something is wrong with the ZD-CSR model.

One way I found to improve the ZD-CSR model is to allow the correlation parameter, ρ , to vary by settlement lag. Let's call this model the ZD2-CSR model.

Figure 4.2 shows that the ρ_d s vary by settlement lag. Table 4.2 shows that the marginal distributions implied by the ZD2-CSR model are much closer to the univariate distributions in Table 3.1. Outputs from the ZD2-CSR model for other insurers show similar results, but the pattern of the ρ_d s are quite different. For some insurers the ρ_d s are predominantly positive. Regardless of whether or not the ρ_d s are positive or

Table 4.1. Marginal distributions implied by the ZD-CSR model

CA Insurer Group 620				
w	Premium	Estimate	Std. Dev.	CV
1	30,224	22,487	0	0
2	35,778	22,193	1,964	0.089
3	42,257	33,353	4,831	0.145
4	47,171	36,883	4,712	0.128
5	53,546	36,886	3,996	0.108
6	58,004	36,776	6,131	0.167
7	64,119	43,263	7,585	0.175
8	68,613	55,233	14,135	0.256
9	74,552	48,469	19,779	0.408
10	78,855	75,545	109,914	1.455
Total	553,119	411,087	125,788	0.306

PA Insurer Group 620				
w	Premium	Estimate	Std. Dev.	CV
1	48,731	44,362	0	0
2	49,951	45,020	1,190	0.026
3	52,434	46,936	1,828	0.039
4	58,191	50,226	1,716	0.034
5	61,873	47,966	1,671	0.035
6	63,614	55,290	3,394	0.061
7	63,807	52,204	3,108	0.060
8	61,157	53,211	4,788	0.090
9	62,146	59,579	8,853	0.149
10	68,003	61,237	15,835	0.259
Total	589,907	516,030	26,359	0.051

Figure 4.1. The posterior distribution of ρ for the ZD-CSR model




Figure 4.2. Boxplot of posterior distribution of ρ_d s for ZD2-CSR model




Table 4.2. Marginal distributions implied by the ZD2-CSR model

CA Insurer Group 620				
w	Premium	Estimate	Std. Dev.	CV
1	30,224	22,123	0	0
2	35,778	21,693	455	0.021
3	42,257	30,382	725	0.024
4	47,171	34,608	886	0.026
5	53,546	36,304	1,047	0.029
6	58,004	36,060	1,357	0.038
7	64,119	45,315	2,280	0.050
8	68,613	52,114	3,953	0.076
9	74,552	55,389	7,482	0.135
10	78,855	65,215	17,864	0.274
Total	553,119	399,203	28,024	0.070
PA Insurer Group 620				
w	Premium	Estimate	Std. Dev.	CV
1	48,731	44,569	0	0
2	49,951	45,215	332	0.007
3	52,434	48,006	419	0.009
4	58,191	51,071	446	0.009
5	61,873	48,083	487	0.010
6	63,614	55,237	813	0.015
7	63,807	51,009	859	0.017
8	61,157	53,200	1,343	0.025
9	62,146	53,836	2,608	0.048
10	68,003	55,091	5,414	0.098
Total	589,907	505,317	9,217	0.018

negative, the “outlier” ρ_d s usually appear at different settlement lags.⁴

This erratic behavior of the ρ_d s suggests that a good practice would be to force those ρ_d s to be equal for all settlement lags. This is what the ZD-CSR model does! This model forces the other parameters to account for the differences in the data that indicate differing ρ_d s in the ZD2-CSR model. It turns out that the differences in the parameters are quite large. Figure 4.3 shows the differences between some of the parameters in the bivariate and the univariate models. If we are to produce a model that has a

⁴The outputs for the other insurers are not provided here, but the reader should be able to reproduce the outputs with the scripts for the ZD-CSR and ZD2-CSR models that are in the Appendix worksheet.

single ρ parameter, we need to find a way to preserve the marginal distributions produced by the univariate models.

Let’s now move on to a two-step approach that produces a bivariate model that preserves the univariate distributions as the marginal distributions.


The first step is to obtain the univariate samples, $\{\theta^x | \{C_{wd}^x\}_{i=1}^{10,000}\}$ and $\{\theta^y | \{C_{wd}^y\}_{i=1}^{10,000}\}$ 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 $\{\rho | \{C_{wd}^x\}, \{C_{wd}^y\}, \theta^x, \theta^y\}$ where ρ has a $\beta(2,2)$ prior distribution translated from $(0,1)$ to $(-1,1)$. Next we randomly select a single i ρ from that sample and use $\{\theta^x, \theta^y, \rho\}_{i=1}^{10,000}$ 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 $\{\theta^x, \theta^y, \rho | \{C_{wd}^x\}, \{C_{wd}^y\}\}_{i=1}^{10,000}$ 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, i $\rho = 0$. 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 fairly new high-end laptop usually turns out this bivariate

Figure 4.3. Selected parameter comparisons for ZD-CSR and univariate models



distribution in about 5 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 4.3 gives results analogous to Table 3.1 for the sum of the two lines for the two illustrative insurers. Also included are the sums of the two lines predicted under the assumption of independence. Figure 4.4 contains histograms of the two-step posterior distributions for ρ for the illustrative insurers.

Table 4.3 and Figure 4.4 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.

Typically, the posterior mean ρ over all the within-group pairs of lines is not all that different from zero. Figure 4.5 shows the frequency distribution of posterior mean ρ s from the insurer group sample.

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, so one would think that a valid bivariate model would also pass a uniformity test. Figures 4.6 and 4.7 show the results.

It turns out that the two-step bivariate model just barely fails the test and the independence assumption passes the uniformity test. 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.

5. Model selection

Let's start the discussion with a review of the Akaike information criteria (AIC).

Suppose that we have a model with a data vector, $\mathbf{x}$, and a parameter vector θ , with p parameters. Let $\hat{\theta}$ be the parameter value that maximizes the log-likelihood, L , of the data, $\mathbf{x} = \{x_j\}_{j=1}^n$. Then the AIC is defined as

$$AIC = 2 \cdot p - 2 \cdot L(\mathbf{x} | \hat{\theta}) \quad (5.1)$$



Given a choice of models, the model with the lowest AIC is to be preferred. This statistic rewards a model

Table 4.3. Bivariate CSR models on illustrative insurer data

	Insurer Group 620			Outcome Percentile = 40.92		
	w	Premium	Estimate	Std. Dev.	C.V.	Outcome
Two-Step Bivariate Model	1	78,955	66,656	0	0.0000	66,391
	2	85,729	66,928	583	0.0087	67,127
	3	94,691	78,353	872	0.0111	77,946
	4	105,362	85,558	1,043	0.0122	85,744
	5	115,419	84,336	1,199	0.0142	85,419
	6	121,618	90,952	1,521	0.0167	90,323
	7	127,926	96,004	2,445	0.0255	96,439
	8	129,770	105,826	4,240	0.0401	104,246
	9	136,698	107,563	7,596	0.0706	105,809
	10	146,858	118,353	17,932	0.1515	112,986
	Total	1,143,026	900,527	28,491	0.0316	892,430
	Insurer Group 620			Outcome Percentile = 43.79		
	w	Premium	Estimate	Std. Dev.	C.V.	Outcome
Independence Assumption	1	78,955	66,661	0	0.0000	66,391
	2	85,729	67,353	608	0.0090	67,127
	3	94,691	78,392	820	0.0105	77,946
	4	105,362	85,069	1,001	0.0118	85,744
	5	115,419	85,007	1,243	0.0146	85,419
	6	121,618	89,439	1,519	0.0170	90,323
	7	127,926	96,201	2,464	0.0256	96,439
	8	129,770	105,910	4,250	0.0401	104,246
	9	136,698	107,943	7,637	0.0708	105,809
	10	146,858	116,508	17,513	0.1503	112,986
	Total	1,143,026	898,483	28,141	0.0313	892,430
	Insurer Group 5185			Outcome Percentile = 24.35		
	w	Premium	Estimate	Std. Dev.	C.V.	Outcome
Two-Step Bivariate Model	1	21,469	17,317	0	0.0000	17,157
	2	23,906	18,628	286	0.0154	18,579
	3	23,634	17,932	346	0.0193	18,277
	4	27,556	20,766	414	0.0199	20,991
	5	29,884	21,313	506	0.0237	21,628
	6	30,500	23,088	688	0.0298	22,384
	7	33,353	27,777	1,048	0.0377	27,540
	8	35,004	25,453	1,496	0.0588	24,272
	9	36,786	30,761	3,036	0.0987	29,568
	10	40,187	34,870	7,765	0.2227	29,194
	Total	302,279	237,904	11,839	0.0498	229,590

Table 4.3. Bivariate CSR models on illustrative insurer data (continued)

	Insurer Group 5185			Outcome Percentile = 27.01		
	w	Premium	Estimate	Std. Dev.	C.V.	Outcome
Independence Assumption	1	21,469	17,292	0	0.0000	17,157
	2	23,906	18,856	300	0.0159	18,579
	3	23,634	18,037	318	0.0176	18,277
	4	27,556	20,599	412	0.0200	20,991
	5	29,884	21,638	512	0.0237	21,628
	6	30,500	22,353	686	0.0307	22,384
	7	33,353	27,828	1,051	0.0378	27,540
	8	35,004	25,406	1,493	0.0588	24,272
	9	36,786	30,927	3,056	0.0988	29,568
	10	40,187	34,030	7,531	0.2213	29,194
Total		302,279	236,966	11,647	0.0492	229,590

Figure 4.4. Posterior distribution of ρ for CSR 2-step bivariate model

Figure 4.5. Posterior mean ρ s from the 102 within-group pairs of lines


for having a high log-likelihood, but it penalizes the model for having more parameters.

There are problems with the AIC in a Bayesian 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.



Figure 4.6. Uniformity tests of outcome percentiles


Figure 4.7. Uniformity tests of outcome percentiles

To address these concerns, Gelman et al. (2014, Chapter 7) describe statistics that generalize the AIC in a way that is appropriate for Bayesian MCMC models.⁵ Here is a brief overview of two of these statistics.

First, given a stochastic model $p(\mathbf{x}|\vartheta)$, define the log of predictive point density as

$$elpd = \sum_{j=1}^J \log \left(\int p(x_j|\vartheta) \cdot f(\vartheta) d\vartheta \right). \quad (5.2)$$

where f is the unknown density function of ϑ .

If $\{\vartheta_i\}_{i=1}^I$ is a sample from the posterior distribution of ϑ , define the computed log predictive density as

$$\widehat{lpd} = \sum_{j=1}^J \log \left(\frac{1}{I} \sum_{i=1}^I p(x_j|\vartheta_i) \right). \quad (5.3)$$

Note that if we replace $\{\vartheta_i\}_{i=1}^I$ with the maximum likelihood estimate, $\widehat{lpd}$ is equal to $L(\mathbf{x}|\hat{\vartheta})$ in Equation (5.1).

⁵Another 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) 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.

If the data vector, $\mathbf{x}$, comes from a holdout sample, i.e., $\mathbf{x}$ was not used to generate the parameters $\{\vartheta_i\}_{i=1}^I$, then $\widehat{lpd}$ is an unbiased estimate of $elpd$.

But if the data vector, $\mathbf{x}$, comes from the training sample, i.e., $\mathbf{x}$ was used to generate the parameters $\{\vartheta_i\}_{i=1}^I$, then we expect that $\widehat{lpd}$ will be higher than $elpd$. The amount of the bias is called the “effective number of parameters.”

With the above as a high-level view, let’s first consider what is called “leave one out cross validation” or “loo” for short. For each data point x_k , one obtains a sample of parameters $\{\vartheta_i^{(-k)}\}_{i=1}^I$ by an MCMC simulation using all values of $\mathbf{x}$ except x_k . After doing this calculation for all observed data points $\{x_k\}_{k=1}^J$ one can then calculate an unbiased estimate of $elpd$

$$\widehat{elpd}_{loo} \equiv \widehat{lpd}_{loo} = \sum_{k=1}^J \log \left(\frac{1}{I} \sum_{i=1}^I p(x_k|\vartheta_i^{(-k)}) \right). \quad (5.4)$$

The effective number of parameters is then given by $\hat{p}_{loo} \equiv \widehat{lpd} - \widehat{elpd}_{loo}$.

As is done with the AIC in Equation (5.1), we can express $\widehat{elpd}_{loo}$ on the deviance scale by writing

$$\text{LOO-IC} = -2 \cdot \widehat{elpd}_{loo} = 2 \cdot \hat{p}_{loo} - 2 \cdot \widehat{lpd}_{loo}. \quad (5.5)$$

When comparing two models, the model with the higher $\widehat{elpd}_{loo} = \widehat{lpd}_{loo} - \hat{p}_{loo}$, or equivalently, the model with the lower LOO-IC should be preferred.

Our second model selection statistic, called the Watanabe-Akaike information criteria (WAIC), is calculated using only the training data, $\mathbf{x} = \{x_j\}_{j=1}^J$. Following Gelman et al. (2014) this statistic is given by

$$\widehat{elpd}_{waic} = \widehat{lpd} - \hat{p}_{waic} \quad (5.6)$$

where⁶

$$\hat{p}_{waic} = \sum_{j=1}^J \text{Var} \left[\left\{ \log(p(x_j|\vartheta_i)) \right\}_{i=1}^I \right]. \quad (5.7)$$

⁶Gelman et al. (2014) give an alternative mean-based formula for $\hat{p}_{waic}$. They recommend using Equation 5.5 “because its series expansion has a closer resemblance to the series expansion for $\hat{p}_{loo}$.”

On the deviance scale, Equation (5.6) becomes

$$\text{WAIC} = 2 \cdot \hat{p}_{\text{waic}} - 2 \cdot \widehat{\text{lpd}}. \quad (5.8)$$

Generally speaking, performance on holdout data is considered to be the gold standard for model evaluation and by that standard, the LOO-IC statistic should be favored over the WAIC statistic. An argument favoring the WAIC statistic is computing time. Applying Equation (5.4) for each of 55 data points can be time consuming. For example, in my current computing environment, it takes about five minutes to run the single two-step bivariate model described in Section 4 above. I really did not want to run this model for 55×5 minutes for each of 102 pairs of loss triangles. To address this problem, Vehtari, Gelman, and Gabry (2015) propose an approximation, $\widehat{\text{elpd}}_{\text{psis-loo}}$ to $\widehat{\text{elpd}}_{\text{loo}}$, that can be calculated with information available from single MCMC fit. PSIS is an abbreviation for Pareto smoothed importance sampling.

The calculations for this statistic, not described here, are captured in the R “loo” package. This package calculates estimates of both the PSIS-LOO and the WAIC statistics, along with their standard errors. This package also has a “compare” function that calculates the standard error of the difference between the respective (PSIS-LOO or WAIC) statistics for two models.

The goal of this section is to compare the two-step bivariate model with a model that assumes that the two lines of insurance are independent. To use the

“loo” package to calculate the PSIS-LOO and WAIC statistics for the two models we need to specify the corresponding $p(x_j | \vartheta)$ function for each model.

Let (C_{wd}^X, C_{wd}^Y) be an observation for accident year w and development year d in the pair of loss triangles for lines X and Y . Then for the two-step bivariate model the probability density function is given by

$$\begin{aligned} p((C_{wd}^X, C_{wd}^Y) | \theta^X, \theta^Y, \rho) \\ = \phi_M(\log(C_{wd}^X), \log(C_{wd}^Y) | \theta^X, \theta^Y, \rho) \end{aligned} \quad (5.9)$$

where ϕ_M is the multivariate normal distribution that is given in Equation (1.2) with parameters determined by Steps 5 and 6 of the model description in Section 3.

For the bivariate independent model

$$\begin{aligned} p((C_{wd}^X, C_{wd}^Y) | \theta^X, \theta^Y, \rho) \\ = \phi_U(\log(C_{wd}^X) | \theta^X) \cdot \phi_U(\log(C_{wd}^Y) | \theta^Y) \end{aligned} \quad (5.10)$$

where ϕ_U is the univariate normal distribution that is given in Equation (1.1) with parameters determined by Steps 5 and 6 of the model description in Section 3.

Table 5.1 gives the PSIS-LOO and WAIC statistics for the illustrative insurers using the CSR model. These tables show that the bivariate model assuming independence is the preferred model. The appendix spreadsheet shows these statistics for the each of the 102 pairs of loss triangles, with the result that the bivariate model assuming independence is the preferred model for *every* pair. When we compare the differences in the $\widehat{\text{elpd}}$ statistics for the two models

Table 5.1. Model selection statistics for the illustrative insurer groups for the CSR model

Group	Two-Step Bivariate		Independent Bivariate		Difference	Std Err.
	$\widehat{\text{elpd}}_{\text{psis-loo}}$	$\hat{p}_{\text{psis-loo}}$	$\widehat{\text{elpd}}_{\text{psis-loo}}$	$\hat{p}_{\text{psis-loo}}$		
620	218.17	37.63	222.23	31.19	4.07	1.36
5185	154.74	37.54	159.24	30.07	4.50	1.67
Group	Two-Step Bivariate		Independent Bivariate		Difference	Std Err.
	$\widehat{\text{elpd}}_{\text{waic}}$	$\hat{p}_{\text{waic}}$	$\widehat{\text{elpd}}_{\text{waic}}$	$\hat{p}_{\text{waic}}$		
620	224.51	31.29	225.98	27.44	1.47	0.88
5185	161.65	30.63	162.96	26.35	1.31	1.00

with their standard errors, the PSIS-LOO statistics are in general more significant.

6. Illustration of model sensitivity

In discussions with my actuarial colleagues over the years, I have sensed that 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, some insurers will find 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 and Hayes (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 $w = 1, \dots, 10$, $\gamma = 0$ and $\delta = 0$. Let's call this model the stochastic Cape Cod (SCC) 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. In Table 6.1 we see that both the PSIS-LOO and the WAIC statistics indicate that the independent bivariate model is the preferred

model for Group 620. The two-step bivariate model is the preferred model for Group 5185. The results for all 102 pairs of loss triangles are given in the appendix spreadsheet. For the SCC model, independence is preferred for 95 of the 102 pairs of loss triangles with the PSIS-LOO statistics and for 65 of the 102 pairs for the WAIC statistics. All seven pairs that were favored by the two-step bivariate model with the PSIS-LOO statistic were also favored by the WAIC statistics. Illustrative insurer group 5185 for CA and OL was one of those seven.

Figures 6.3 and 6.4 explain the underlying cause of the dependency for Group 5185. These figures consist of scatterplots of the standardized residuals for the first six accident years⁷ for the SCC and CSR models. The scatterplots are reproduced six times, with each accident year highlighted. In Figure 6.3 the individual accident year residuals occupy noticeably distinct regions of the scatterplots. The likely cause for these distinct regions are the (either random or systematic) deviations from the assumed constant expected loss ratio assumed in the SCC model. If the regions for the various accident years tend to fall in opposite quadrants (see, for example, accident years 1 and 3 in Figure 6.3) a significant correlation could be indicated. If, as happens with many pairs for the SCC model, the regions are more or less equally spread out in the various quadrants, “independence” will be preferred. By way of contrast, there are no distinct regions by accident year indicated by the CSR model in Figure 6.4.

It is also worth noting that, as shown in Figure 6.5, 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 CSR model leads to independence between lines for our sample of insurer loss triangles.

⁷The R scripts for the CSR and SCC models in the appendix spreadsheet create the scatterplots for all 10 accident years.

Figure 6.1. Standardized residual plots for the SSC model





Figure 6.2. Posterior distribution of ρ for SCC model



7. Summary and conclusions

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.

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

Table 6.1. Model selection statistics for the illustrative insurer groups for the SCC model

Group	Two-Step Bivariate		Independent Bivariate		Difference	Std Err.
	$\widehat{elpd}_{psis-loo}$	$\hat{p}_{psis-loo}$	$\widehat{elpd}_{psis-loo}$	$\hat{p}_{psis-loo}$		
620	102.83	19.68	104.2	17.79	1.37	0.47
5185	59.36	21.48	58.21	18.36	-1.15	2.59

Group	Two-Step Bivariate		Independent Bivariate		Difference	Std Err.
	$\widehat{elpd}_{waic}$	$\hat{p}_{waic}$	$\widehat{elpd}_{waic}$	$\hat{p}_{waic}$		
620	105.01	17.49	105.94	16.05	0.93	0.35
5185	62.03	18.8	60.26	16.31	-1.78	2.10



Figure 6.3. Bivariate standardized residual plots for the SCC model

Figure 6.4. Bivariate standardized residual plots for the CSR model

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 (2013).

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 did well in 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 for the independence assumption over the bivariate two-step model.

3. The prospective model selection statistics in Section 5 indicated a preference for the independence assumption for all 102 within-group pairs of lines.

The range of ρ s 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.

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.

As Bayesian MCMC models become used to set capital requirements, my advice would be test for

Figure 6.5. Uniformity test for the SCC model

dependence using the statistics described in this paper. If dependence is indicated, one should try to develop a more refined model.

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

- Avanzi, B., G. Taylor, and B. Wong, "Correlations between Insurance lines of Business: An Illusion or a Real Phenomenon? Some Methodological Considerations," *ASTIN Bulletin* 46:2, 2016, pp. 225–263. doi: 10.1017/asb.2015.31.
- Barnett, G., and B. Zehnwirth, "Best Estimates for Reserves," *Proceedings of the Casualty Actuarial Society* 87, 2000, pp. 254–321.
- Bornhuetter, R. L., and R. E. Ferguson, "The Actuary and IBNR," *Proceedings of the Casualty Actuarial Society* 59, 1972, p. 181.
- Clark, D., "Where Are We In the Market Cycle," Presentation to CAMAR, October 9, 2013. <http://www.casact.org/community/affiliates/camar/1013/Clark.pdf>.
- England, P. D., and R. J. Verrall, "Stochastic Claims Reserving in General Insurance," Meeting Paper. Presented to the Institute and Faculty of Actuaries, 28 January 2002.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed., CRC Press, 2014.
- Herzog, T. N., "Summary of CEIOPS Calibration Work on Standard Formula," 2011, http://www.naic.org/documents/index_smi_solvency_ii_calibration.pdf.

- Leong, J. W. K., and J. B. Hayes. "Property versus Casualty Risks," Presentation to CAMAR, October 9, 2013. <http://www.casact.org/community/affiliates/camar/1013/Leong-Hayes.pdf>
- Mack, T., "Distribution-Free Calculation of the Standard Error on Chain Ladder Reserve Estimates," *ASTIN Bulletin* 23(2), 1993, pp. 213–225.
- Mack, T., "Measuring the Variability of Chain Ladder Reserve Estimates." Casualty Actuarial Society *Forum* (Spring), 1994, pp. 101–182.
- Meyers, G. G., *Stochastic Loss Reserving Using Bayesian MCMC Models*, CAS Monograph Series, No. 1, 2015.
- Stanard, J. N., "A Simulation Test of Prediction Errors of Loss Reserve Estimation Techniques." *Proceedings of the Casualty Actuarial Society* 72, 1985, p. 124.
- Vehtari, A., A. Gelman, and J. Gabry, "Efficient Implementation of Leave-One-Out Cross-Validation and WAIC for Evaluating Fitted Bayesian Models," 2015, <http://arxiv.org/abs/1507.04544>.
- Zhang, Y., and V. Dukic, "Predicting Multivariate Insurance Loss Payments Under the Bayesian Copula Framework," *Journal of Risk and Insurance* 80: 4, 2013, pp. 891–919.