A Case Study
Using Bayesian MCMC Stochastic Loss Reserve Models

Glenn Meyers

Presentation to Casualty Loss Reserve Seminar
Austin Texas

September 18, 2019

Monographs

- Stochastic Loss Reserving Using Bayesian MCMC Models
  - First Edition — January 2015
  - Second Edition — September 2019
- MCMC — Markov Chain Monte Carlo
  - Given data and a model, MCMC extracts a sample (10,000) from the posterior distribution of parameters.
  - From the posterior sample of parameters, it produces a posterior distribution of loss reserves.
- Selected Topics Covered Here
  - Focus on the Changing Settlement Rate (CSR) model.
  - For 200 loss triangles, calculate the percentile of the holdout (lower triangle) outcome.
  - The percentiles are uniformly distributed — establishing the “reputation” of the CSR model.

The Wolf-Miller Triangles

- At the 2018 annual meeting of the Casualty Actuarial Society, Bob Wolf and Mary Frances Miller presented a loss reserve analysis[1] on real data (scaled to maintain anonymity). These data consisted of 16 x 16 paid and incurred loss triangles. Features of the data included.
  - Rapid premium growth
  - Change in claims philosophy?
  - Underestimates of outstanding liability in previous years

[1] Session C-24 - Learning Lounge Case Study: Material Adverse Reserve Development? When is it just that stuff happens?

The Wolf-Miller Triangles

- At the 2018 annual meeting of the Casualty Actuarial Society, Bob Wolf and Mary Frances Miller presented a loss reserve analysis[1] on real data (scaled to maintain anonymity). These data consisted of 16 x 16 paid and incurred loss triangles. Features of the data included.
  - Rapid premium growth
  - Change in claims philosophy?
  - Underestimates of outstanding liability in previous years
- The Meyers monograph triangles were selected from Schedule P (10x10) triangles.
  - Triangles that reflected "obvious" operational changes were eliminated from the set of 200 triangles.
- My question — How well do the models in Meyers (2019) work for these datasets?

[1] Session C-24 - Learning Lounge Case Study: Material Adverse Reserve Development? When is it just that stuff happens?

Stochastic Loss Reserve Models

- Start with the model framework in Meyers (2019).

$$C_{wd} \sim \text{lognormal}(\mu_{wd}, \sigma_d)$$

- where:
  - $w =$Accident Year (AY), $w = 1, \ldots, W$
  - $d =$Development Year (DY), $d = 1, \ldots, D$
  - Also, let $c =$Calendar Year (CY), $c = w + d - 1$
- This talk will initially examine models where:

$$\mu_{wd} = \log(Premium_w) + logelr + \alpha_w + \beta_d \cdot Sp(t)$$

- The model specification for the σ_d parameters and the prior distributions are the same as given in Meyers (2019).

Stochastic Loss Reserve Models - Continued

- Start with the model framework in Meyers (2019).

$$C_{wd} \sim \text{lognormal}(\mu_{wd}, \sigma_d)$$

- This talk will initially examine models where:

$$\mu_{wd} = \log(Premium_w) + logelr + \alpha_w + \beta_d \cdot Sp(t)$$

- The $Sp(t)$, i.e. the "Speedup", function specifies how the "development factors" change over the time, t, where t could be measured by accident year, or calendar year.
- This talk explores alternative $Sp(t)$ functions in an effort to find a model that makes better predictions of the ultimate losses.

Interpreting the Model Parameters

- To prevent overdetermining the model, set:

$$\alpha_1 \equiv 0 \text{ and } \beta_D \equiv 0$$

- Thus the expected ultimate loss, U_w for accident year w, is the mean of a lognormal distribution, i.e.

$$U_w \equiv Premium_w \cdot \exp(logelr + \alpha_w + \sigma_D^2/2) \qquad (1)$$

- If the reported losses are near ultimate, the parameter σ_D will be very small. Thus for $w = 1$ the ultimate loss is approximately equal to $Premium_1$ times the expected loss ratio, $\exp(logelr)$. The α_w parameters account for accident year differences in the loss ratio.

Interpreting the Speedup Function, $Sp(t)$

- What will distinguish the models in this talk is the choice of the $Sp(t)$ function. Let's discuss its meaning.
- Recall that $\beta_D = 0$. If $Sp(1) > Sp(2) > \cdots$, then the product $\beta_d \cdot Sp(t)$ is moving closer to 0 as t increases.
- For paid losses, this means losses are being settled more quickly over time.

Interpreting the Speedup Function, $Sp(t)$

- What will distinguish the models in this talk is the choice of the $Sp(t)$ function. Let's discuss its meaning.
- Recall that $\beta_D = 0$. If $Sp(1) > Sp(2) > \cdots$, then the product $\beta_d \cdot Sp(t)$ is moving closer to 0 as t increases.
- For paid losses, this means losses are being settled more quickly over time.
- For incurred losses, this means that losses are being recognized more quickly over time.

Interpreting the Speedup Function, $Sp(t)$

- What will distinguish the models in this talk is the choice of the $Sp(t)$ function. Let's discuss its meaning.
- Recall that $\beta_D = 0$. If $Sp(1) > Sp(2) > \cdots$, then the product $\beta_d \cdot Sp(t)$ is moving closer to 0 as t increases.
- For paid losses, this means losses are being settled more quickly over time.
- For incurred losses, this means that losses are being recognized more quickly over time.
- The reverse is true if $Sp(1) < Sp(2) < \cdots$. That is, paid losses are being settled more slowly over time, and incurred losses are being recognized more slowly over time.

Models Considered in This Talk

- The CRC Model — $Sp(w) \equiv 1$
  - This model most closely resembles the standard actuarial models that do not allow the development patterns to change over time.
- The CSR-w Model — $Sp(w) = (1-\gamma)^{w-1}$
  - $\gamma > 0$ gives us a decreasing $Sp(w)$ as the accident year, w increases from 1 to W. $\gamma < 0$ gives us an increasing $Sp(w)$.
- The CSR-c Model — $Sp(c) = (1+\gamma)^{C-c}$
  - $\gamma < 0$ gives us a increasing $Sp(c)$ as the calendar year, c, increases from 1 to $C-1$. $\gamma > 0$ gives us a decreasing $Sp(c)$
- We refer to the γ parameter as the speedup rate. We call a negative speedup rate a slowdown.

Models Considered in This Talk - Continued

- The CSR-vc Model —

$$
\begin{aligned}
Sp(C) &= 1 \\
Sp(C-i) &= Sp(C-i+1)\cdot(1+\gamma_{C-i}) \\
&\quad \text{for } i = 1,\cdots,C-1
\end{aligned}
$$

  - This model allows the speedup rate to vary by calendar year.
- The first two models are described in Meyers (2019). The next two were developed during the research that led to this talk. As we shall see, analyses of the shortcomings of these models point to another model, the POS-vc model that I will describe below.
- The "Case_Study_CLRS_Dist" zip folder has the R scripts for these models.

The Run ID

- The various model runs in this talk will be fit on a given set of calendar years of either the paid or incurred loss triangle.
- Each model run will have an identifier with three components.
  1. The model name
  2. The loss triangle used — either "P" or "I"
  3. The calendar year range.
- For example, the run id "CSR-vc P-7:16" means that the CSR-vc model was fit to the paid loss triangle using data from the calendar years from 7 to 16.

Invoking Bayesian MCMC

- As described in Meyers (2019), the Bayesian MCMC fitting algorithm produces 10,000 equally likely parameter sets[2] $\{logelr\}$, $\{\alpha_w\}_{w=1}^{W}$, $\{\beta_d\}_{d=1}^{D}$, $\{\gamma\}$ and $\{\sigma_d\}_{d=1}^{D}$.
- With a sample of 10,000 parameter sets, one can use Equation 1 to to obtain a sample of 10,000 expected ultimate losses, $\{U_w\}$
- Define $\{U_{Tot}\} = \sum_{w=1}^{16}\{U_w\}$.
- Also of interest is a sample of 10,000 possible unpaid losses (ultimate loss less current paid loss), $\{R_c\}$, at calendar year c where:

$$R_c = \sum_{w=1}^{c} U_w - \sum_{d=1}^{c} C_{c+1-d,d} \tag{2}$$

[2]The presence of brackets $\{\cdot\}$ around a parameter will indicate that it is a sample of 10,000 values from the posterior distribution.

Statistics of Interest

- From the samples $\{R_c\}$ and $\{U_{Tot}\}$, we can calculate statistics of interest, such as:
  - Ultimate Loss = mean$\{U_{Tot}\}$
  - Ultimate Standard Error = standard deviation$\{U_{Tot}\}$
  - Reserve Low = 2.5th percentile of $\{R_{16}\}$
  - Reserve = mean$\{R_{16}\}$
  - Reserve High = 97.5th percentile of $\{R_{16}\}$

The Expected Log Predictive Density — $\widehat{elpd}_{loo}$[3]

- For each observation, C_{wd} in the loss triangle:
  1. Remove that observation from the data.
  2. Fit the selected model to the data in the triangle that remains and obtain the parameter sets $\{\theta(-wd)\}$ (consisting of all the $\{\alpha_w\}$s, $\{\beta_d\}$s, etc.)
  3. Calculate the average likelihood, $p(C_{wd}|\{\theta(-wd)\})$ over all 10,000 parameter sets.
- Then

$$\widehat{elpd}_{loo} = \sum_{w,d} log(p(C_{wd}|\{\theta(-wd)\}))$$

[3] More details about this statistic are in Section 6 of Meyers (2019).

$\widehat{elpd}_{loo}$ — Continued

- The "loo" term refers to the "leave one out" feature in bullet #1 above.
- Since the likelihoods are calculated on holdout data, there is no penalty for fitting models with a large number of parameters.

Standardized Residual Boxplots

- The models in this talk all assume a lognormal distribution with the parameters μ_{wd} and σ_d. Thus we expect that

$$\frac{log(C_{wd}) - \{\mu_{wd}\}}{\{\sigma_d\}}$$

will have a normal(0,1) distribution.

- To test this graphically we split the residuals, in turn by accident year, development year and calendar year and plot a sample of size 200 in each "year" with the R "boxplot" function.

Running the MCMC Models on Paid Data

Run ID	Ult Loss	Ult SE	Res Low	Reserve	Res High	$\widehat{elpd}_{loo}$
CRC P-7:16	1,186,501	26,857	175,749	226,573	283,168	234.94
CSR-w P-7:16	1,233,188	46,782	185,870	273,260	369,290	235.05
CSR-c P-7:16	1,280,649	53,024	227,348	320,721	439,200	236.06
CSR-vc P-7:16	1,256,436	73,076	179,454	296,510	460,811	240.85

Some observations

- The CSR-vc model had the highest $\widehat{elpd}_{loo}$ statistic.
- The mean reserve estimates vary significantly by model.
- The mean speedup rate is -0.0156 for the CSR-w model, -0.0291 for the CSR-c model. For the CSR-vc model it starts as 0.0131 and moves down to fluctuate between the -0.012 to -0.035 range for the later calendar years.
- A negative speedup rate means a slowdown in claim settlements, and hence a higher predicted ultimate loss.

Expected Results with the R “boxplot” Function

- The following four pages contain the standardized residual Boxplots for the four models on the P-7:16 data.
- The gray bars correspond to the interquartile range. Ideally the bars should be centered on 0. The endpoints of those bars should be touching the black lines representing the interquartile range of the standard normal distribution.
- Most of the remaining residuals should be between $\pm$ 2. A few could be in the (-3,-2) or the (2,3) ranges. Very few should be outside the $\pm$ 3 range.
- Now flip through the next four pages to see how close the Boxplots are to the “ideal” Boxplot. I will give my take on the other side.

Boxplots P-7:16

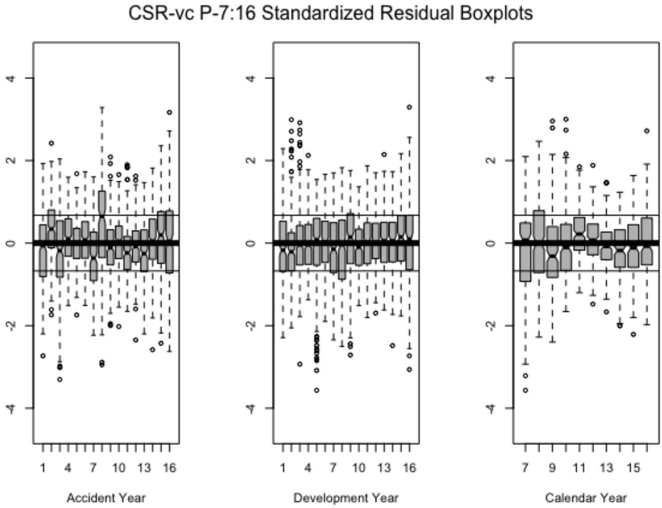

Boxplots P-7:16

A Case Study
Glenn Meyers
Introduction
Models
MCMC
Fit Evaluation
$\widehat{elpd}_{loo}$
Boxplots
P-7:16 Data
Boxplots
I-7:16 Data
Boxplots
Speedup Rates
POS Model
Boxplots
Speedup Rates
Bad Stuff?
Commentary

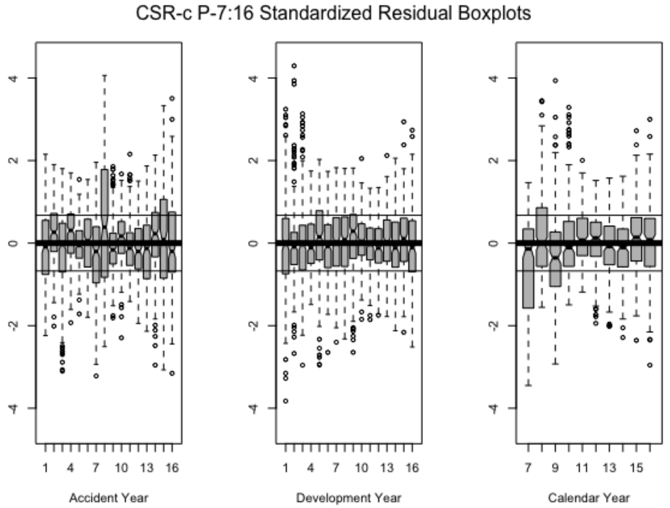

Boxplots P-7:16

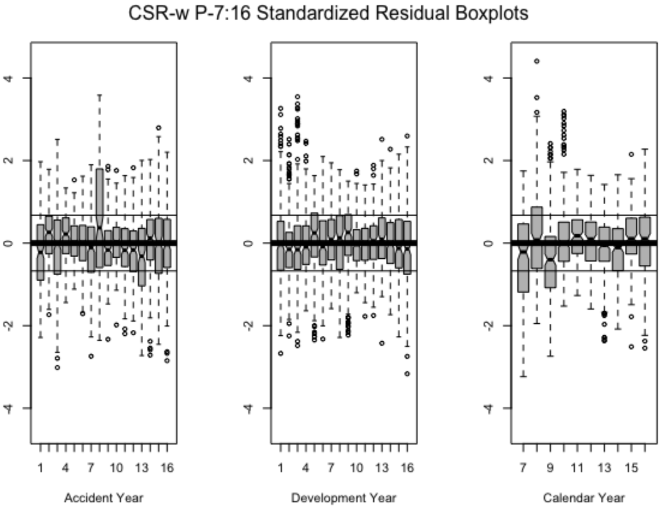

Boxplots P-7:16

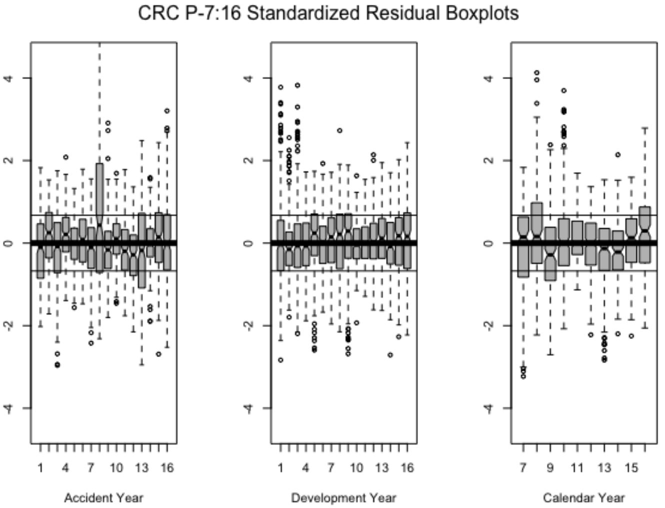

P-7:16 Discussion

- I judge the CSR-vc model to have the best Boxplots.
  - The interquartile ranges are about the same and all pretty good.
  - The CSR-vc model has noticeably fewer outliers in the Boxplots, i.e. outside the ± 2 range.
- This combined with its having the highest $\widehat{elpd}_{loo}$ statistic make it the model of choice for the paid data.

Running the MCMC Models on I-7:16 Data

Run ID	Ult Loss	Ult SE	Res Low	Reserve	Res High	$\widehat{elpd}_{loo}$
CRC I-7:16	1,230,151	29,800	214,940	270,223	333,281	235.64
CSR-w I-7:16	1,193,518	37,085	167,331	233,590	313,557	232.50
CSR-c I-7:16	1,317,128	75,412	243,610	357,200	531,381	235.19
CSR-vc I-7:16	1,262,187	58,618	201,157	302,261	430,947	241.27

Some observations

- The CSR-vc model had the highest $\widehat{elpd}_{loo}$ statistic.
- The mean reserve varies significantly by model.
- The mean speedup rate is a *positive* 0.0375 for the CSR-w model, a *negative* 0.0675 for the CSR-c model. For the CSR-vc model it starts close to zero and moves up around the -0.02 to -0.04 range for the later calendar years.
- A negative speedup rate for incurred losses can also indicate a decreasing recognition of outstanding losses, and hence a higher predicted ultimate loss.
- Now scroll through the Boxplots for these models.

Plots I-7:16

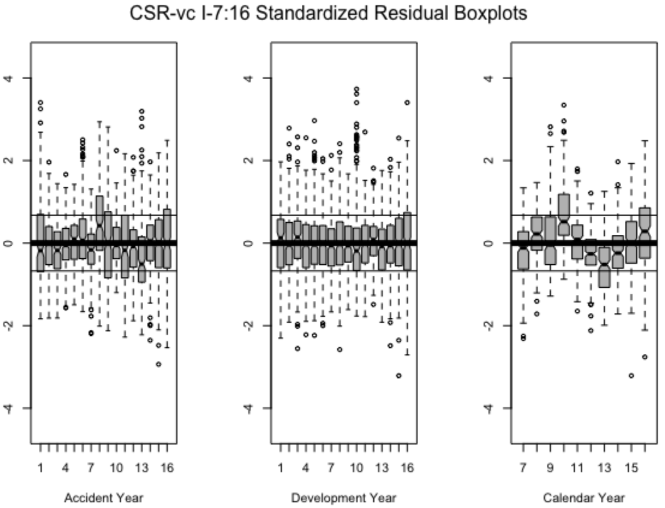
CSR-vc I-7:16 Standardized Residual Boxplots
Accident Year
Development Year
Calendar Year

Plots I-7:16

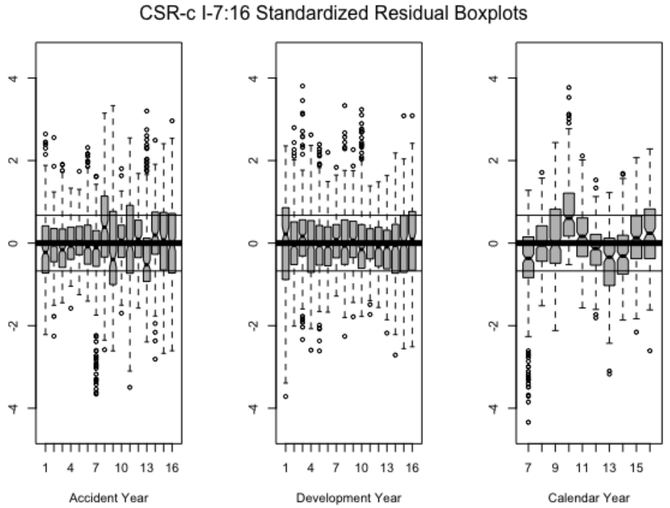
CSR-c I-7:16 Standardized Residual Boxplots
4
2
0
-2
-4
1
4
7
10
13
16
Accident Year
Development Year
7
9
11
13
15
Calendar Year

Plots I-7:16

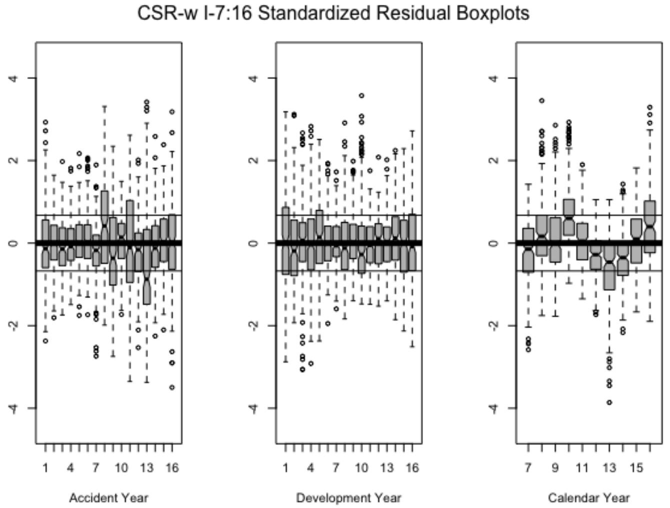
CSR-w I-7:16 Standardized Residual Boxplots
4
2
0
-2
-4
1
4
7
10
13
16
Accident Year
Development Year
7
9
11
13
15
Calendar Year

Plots I-7:16

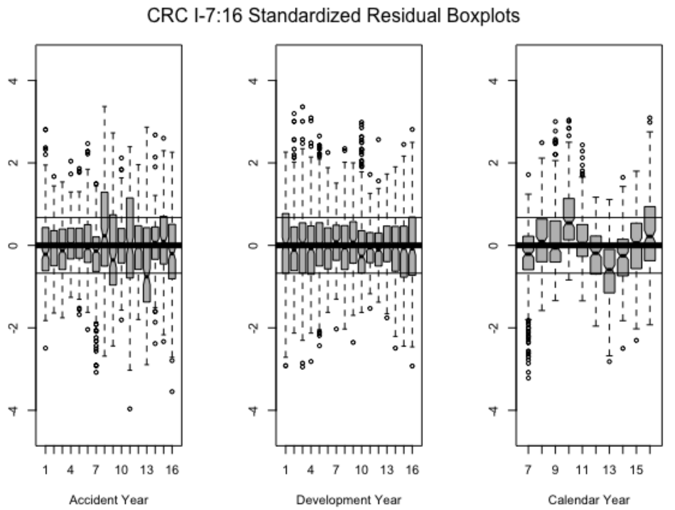
CRC I-7:16 Standardized Residual Boxplots
4
2
0
-2
-4
1 4 7 10 13 16
Accident Year
1 4 7 10 13 16
Development Year
7 9 11 13 15
Calendar Year

I-7:16 Discussion

- I judge that the CSR-vc model has the best Boxplots.
  - Slightly better by accident year and development year.
- The Boxplots by calendar year suggests that there as been a change in case reserving practices.
- The next page shows plots of the mean speedup rates for the paid and the incurred models. One would expect to see the plots track closely with each other as a substantial portion of the incurred losses are already paid.
- But — As we can see from these plots, there is a noticeable difference between the plots. And moreover, they cross.
- This suggests that there should be separate $\{\gamma\}$ parameters for paid and outstanding losses.

Mean Speedup Rates for the CSR-vc P-7:16 and the CSR-vc I-7:16 Models

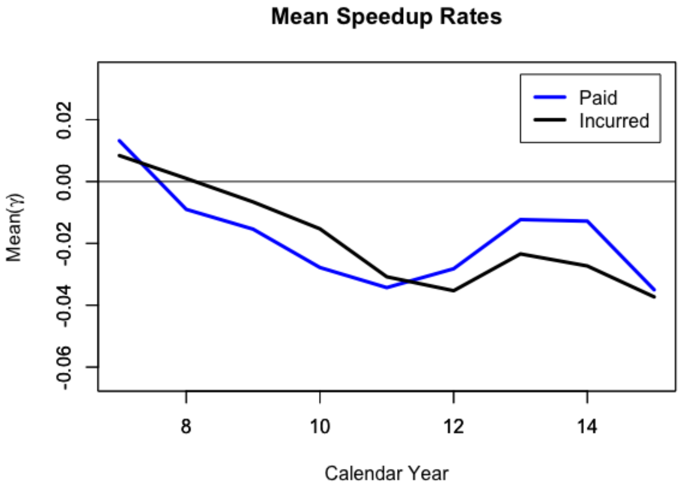

Integrated Paid and Outstanding (POS) Models

- This section proposes a model that simultaneously fits both paid and incurred losses.[4]
- This model has lognormal distributions for each of the paid and incurred losses.
  - The μ_{wd} parameter of the distribution for paid losses is the same as above.
  - The μ_{wd} of the incurred losses are equal to the sum of the μ_{wd} for the paid losses, plus a separate factor representing outstanding losses.
- More details on the next page.

[4] A more detailed discussion of fitting models simultaneously to paid and incurred is discussed in Section 9 of Meyers (2019)

The POS-vc Model

The prefixes P, I and OS denote "Paid", "Incurred" and "Outstanding" respectively.

$$
\begin{aligned}
{}_P\mu_{wd} &= \log(Premium_w) + logelr + \alpha_w + {}_P\beta_d \cdot {}_PSp(c) \\
{}_I\mu_{wd} &= {}_P\mu_{wd} + {}_{OS}\beta_d \cdot {}_{OS}Sp(c)
\end{aligned}
$$

$$ {}_P\beta_D \equiv 0, \text{ and } {}_{OS}\beta_D \not\equiv 0 $$

Where

$$
\begin{aligned}
{}_XSp(C) &= 1 \\
{}_XSp(C-i) &= Sp(C-i+1) \cdot (1 + {}_X\gamma_{C-i}) \\
&\quad \text{for } i = 1, \cdots, C-1 \text{ and } X = P \text{ or } OS
\end{aligned}
$$

Then

$$
\begin{aligned}
{}_PC_{wd} &\sim \text{lognormal}({}_P\mu_{wd}, {}_P\sigma_d) \\
{}_IC_{wd} &\sim \text{lognormal}({}_I\mu_{wd}, {}_I\sigma_d)
\end{aligned}
$$

Running the POS Model on PI-7:16 Data

- Results for the comparable CSR model runs are also given.

Run ID	Ult Loss	Ult SE	Res Low	Reserve	Res High	$\widehat{elpd}_{loo}$
CSR-vc P-7:16	1,256,436	73,076	179,454	296,510	460,811	240.85
CSR-vc I-7:16	1,262,187	58,618	201,157	302,261	430,947	241.27
POS-vcp PI-7:16	1,262,897	58,400	205,683	302,969	432,516	251.85
POS-vci PI-7:16	1,262,897	58,400	205,691	302,969	432,529	255.56

- The $\widehat{elpd}_{loo}$ statistics are calculated separately on the paid and incurred data in the POS model. These statistics are significantly better for the POS-vc model than they are for the corresponding CSR-vc models.
- The standardized residual Boxplots are on the following three pages. Compared with the corresponding CSR-vc Boxplots:
  - The POS-vc plots look a bit worse for the paid losses.
  - They look a bit better for the incurred losses.
  - They look pretty good for the combined losses.

Boxplots for POS-vc Model

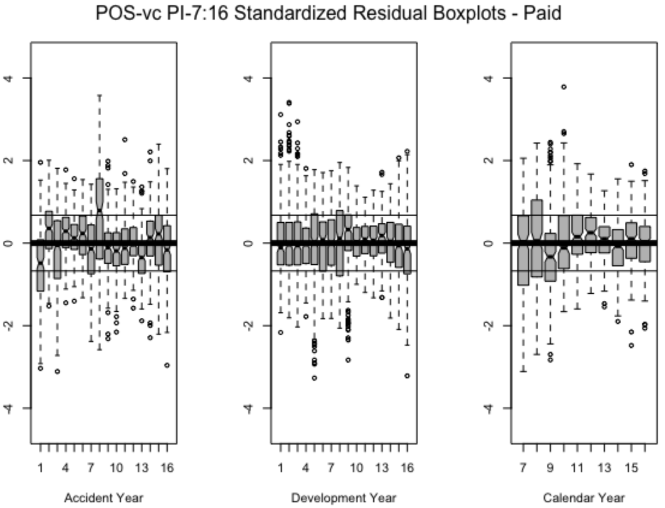

Boxplots for POS-vc Model

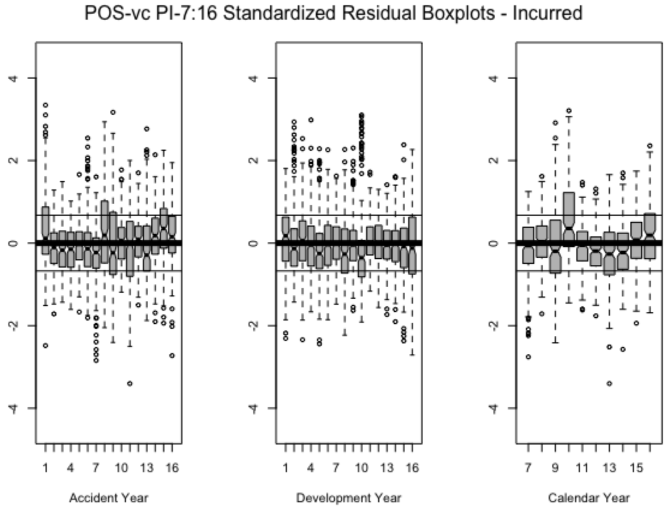

Boxplots for POS-vc Model

A Case Study
Glenn Meyers
Introduction
Models
MCMC
Fit Evaluation
$\widehat{elpd}_{loo}$
Boxplots
P-7:16 Data
Boxplots
I-7:16 Data
Boxplots
Speedup Rates
POS Model
Boxplots
Speedup Rates
Bad Stuff?
Commentary

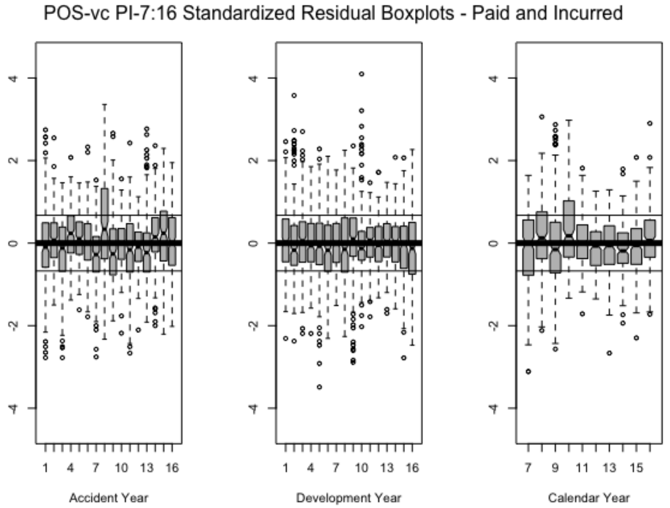

Claims Department Practices

- The following page has plots of the the mean paid claim speedup rate, mean$\{_P\gamma\}$, and the mean outstanding claim speedup rate, mean$\{_{OS}\gamma\}$.
- The claims department appears to be slowing down the paid claim settlements, while speeding up the recognition of outstanding claims, and vice versa.
- This observation should be discussed with the claims department.

Mean Speedup Rates for the POS-vc Model

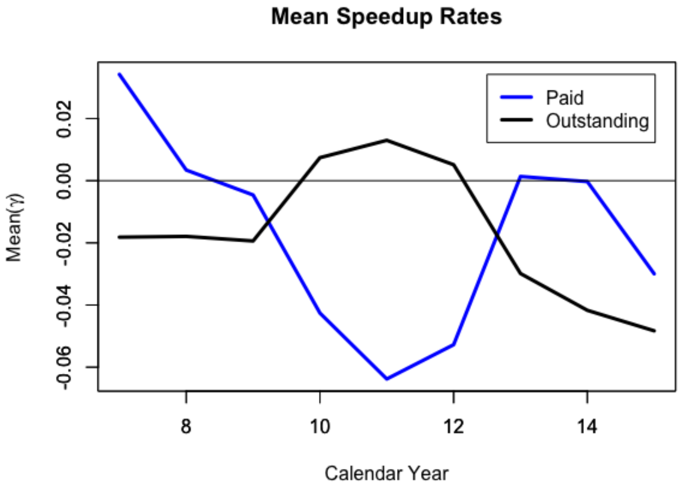

Estimating Ultimate Losses

- Recall from Equation 1 that the ultimate expected loss for accident year w is equal to the expected value

$$Premium_w \cdot \mathrm{E}\{\exp(logelr + \alpha_w + {}_X\sigma_D^2/2)\}$$

where X can refer to either paid, P, or incurred, I, losses.

- For the POS-vc model the expected ultimate incurred loss is slightly more complicated.

$$Premium_w \cdot \mathrm{E}\{\exp(logelr + \alpha_w + {}_{OS}\beta_D +_I \sigma_D^2/2)\}$$

- After 16 years of development, the values of ${}_{OS}\beta_D$ and ${}_X\sigma_D$ are close to zero. So the paid and incurred loss estimates are very close to each other.
- The following three pages give the ultimate loss estimates by accident year for the CSR-vc and POS-vc models.

Accident Year Exhibit for CSR-vc P-7:16

AY	Premium	Estimate	SE	CV
2002	13,750	7,035	36	0.0051
2003	28,052	11,172	89	0.0080
2004	44,853	27,882	237	0.0085
2005	70,507	42,229	397	0.0094
2006	80,285	45,451	459	0.0101
2007	96,286	58,149	659	0.0113
2008	130,481	66,126	817	0.0124
2009	142,059	49,960	715	0.0143
2010	131,024	70,952	1,150	0.0162
2011	131,870	89,695	1,702	0.0190
2012	122,125	83,745	2,025	0.0242
2013	125,456	88,474	2,794	0.0316
2014	201,129	105,300	4,505	0.0428
2015	271,351	148,458	10,143	0.0683
2016	297,237	180,482	22,320	0.1237
2017	292,035	181,328	51,519	0.2841
Total	2,178,500	1,256,436	73,076	0.0582

Accident Year Exhibit for CSR-vc I-7:16

AY	Premium	Estimate	SE	CV
2002	13,750	7,037	35	0.0050
2003	28,052	11,056	84	0.0076
2004	44,853	27,643	231	0.0084
2005	70,507	41,556	376	0.0090
2006	80,285	44,764	439	0.0098
2007	96,286	57,033	619	0.0109
2008	130,481	65,057	765	0.0118
2009	142,059	49,100	658	0.0134
2010	131,024	70,228	1,078	0.0154
2011	131,870	89,487	1,542	0.0172
2012	122,125	82,593	1,863	0.0226
2013	125,456	87,515	2,526	0.0289
2014	201,129	105,847	4,581	0.0433
2015	271,351	148,245	9,905	0.0668
2016	297,237	184,843	21,850	0.1182
2017	292,035	190,185	46,733	0.2457
Total	2,178,500	1,262,187	58,618	0.0464

Accident Year Exhibit for POS-vc PI-7:16

AY	Premium	Estimate(p)	Estimate(i)	SE(p)	SE(i)	CV(p)	CV(i)
2002	13,750	7,038	7,038	34	34	0.0048	0.0048
2003	28,052	11,117	11,117	68	68	0.0061	0.0061
2004	44,853	27,779	27,779	184	184	0.0066	0.0066
2005	70,507	41,907	41,907	299	299	0.0071	0.0071
2006	80,285	45,135	45,135	347	347	0.0077	0.0077
2007	96,286	57,636	57,636	489	489	0.0085	0.0085
2008	130,481	65,630	65,630	603	603	0.0092	0.0092
2009	142,059	49,658	49,658	513	513	0.0103	0.0103
2010	131,024	70,692	70,692	792	792	0.0112	0.0112
2011	131,870	89,930	89,930	1,175	1,175	0.0131	0.0131
2012	122,125	83,456	83,456	1,373	1,373	0.0165	0.0165
2013	125,456	88,619	88,619	1,909	1,909	0.0215	0.0215
2014	201,129	106,701	106,701	3,372	3,372	0.0316	0.0316
2015	271,351	149,816	149,816	7,690	7,690	0.0513	0.0513
2016	297,237	183,299	183,299	17,254	17,254	0.0941	0.0941
2017	292,035	184,484	184,484	37,469	37,469	0.2031	0.2031
Total	2,178,500	1,262,897	1,262,897	58,400	58,400	0.0462	0.0462

Predictive Distribution of Loss Reserve Liability

- Following Equation 2, a sample of the predictive distribution of the outstanding losses is given by:

$$\{_X R_C\} = \sum_{w=1}^{C} \{_X U_w\} - \sum_{d=1}^{C} C_{c+1-d,d}$$

where X =CSR-vc P-7:16, CSR-vc I-7:16 or POS-vc 7:16.

- Histograms of the predictive distributions for these models are given in the next page.
- Note that the POS-vc model reduces the range of ultimate estimates, by a lot for paid losses, and by a little for incurred losses.

Predictive Distribution of Loss Reserve Liability

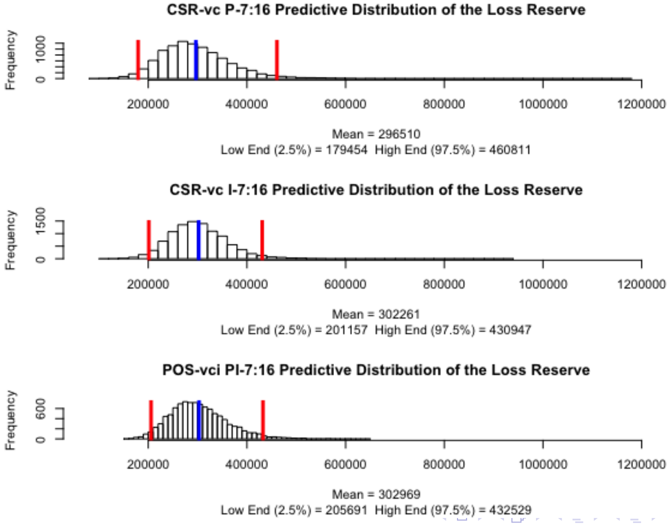

The Question Addressed by This Talk

- In prior years the original opining actuary underestimated the loss reserve liability.
- Was this a case of "bad stuff" that sometimes happens?
- Or was it the case that there is a loss reserve model that does a better job of predicting the "bad stuff?"
- The Learning Lounge presentation mentioned a number of red flags, e.g. declining paid to current ultimate and declining incurred to current ultimate ratios, and slowdown in claim settlement due to rapid premium growth.
- It appears that the opining actuary and the Learning Lounge presenters recognized by 2016 and 2017, that earlier reserve estimates were understated because of the slowdown in the claim settlements.

The Question Addressed by This Talk

- In prior years the original opining actuary underestimated the loss reserve liability.
- Was this a case of "bad stuff" that sometimes happens?
- Or was it the case that there is a loss reserve model that does a better job of predicting the "bad stuff?"
- The Learning Lounge presentation mentioned a number of red flags, e.g. declining paid to current ultimate and declining incurred to current ultimate ratios, and slowdown in claim settlement due to rapid premium growth.
- It appears that the opining actuary and the Learning Lounge presenters recognized by 2016 and 2017, that earlier reserve estimates were understated because of the slowdown in the claim settlements.
- To my way of thinking, this means that they needed a better model.

A Proposal for the “Bad Stuff”

- This talk proposes the “vc” models that explicitly recognize changes in the claim speedup rate by calendar year.
- The CRS-vc model works well with paid losses, but not very well with incurred losses.
- POS-vc model obtains a better fit with the incurred losses.

A Proposal for the "Bad Stuff"

- This talk proposes the "vc" models that explicitly recognize changes in the claim speedup rate by calendar year.
- The CRS-vc model works well with paid losses, but not very well with incurred losses.
- POS-vc model obtains a better fit with the incurred losses.
- My opinion — Declaring victory would be premature. We need further testing.

My Loss Reserving Philosophy

- I view loss reserving as a dialogue between an actuarial department and its corresponding claims department. One way this dialogue might work is as follows.
  1. In talking with the claims department, the actuaries try to find out how the claims adjustment process works.
  2. They then formulate a model that describe the claims adjustment process. Then test the model thoroughly.
  3. If testing reveals unexpected differences between the model and the data, repeat Steps 1-2 above as necessary.
- Advantages of using Bayesian MCMC for model building
  1. Flexibility in model building — If you can code the likelihood function, you can run the model.
  2. Bayesian models are transparent and reproducible. Your judgments are made explicit in your choice of models and prior distributions.
  3. Bayesian models provide output that can be used for calculating risk margins. See Section 11 of Meyers (2019).

Comments

- The flexibility of Bayesian MCMC models was very helpful in this exercise. It enabled me to to easily explore beyond my existing collection of models.
- Over time, I expect that I, and others, will add to our collection of such models in the future.
- I want to thank Bob and Mary Frances for making these data available to the public. It was interesting to see how well the estimates derived from a Bayesian MCMC model tracked with the estimates from experienced reserving actuaries. I was glad for my model and for the actuarial profession, to see that the estimates were reasonably close.

Comments

- The flexibility of Bayesian MCMC models was very helpful in this exercise. It enabled me to to easily explore beyond my existing collection of models.
- Over time, I expect that I, and others, will add to our collection of such models in the future.
- I want to thank Bob and Mary Frances for making these data available to the public. It was interesting to see how well the estimates derived from a Bayesian MCMC model tracked with the estimates from experienced reserving actuaries. I was glad for my model and for the actuarial profession, to see that the estimates were reasonably close.
- I want to make a call out to Ben Zehnwirth, who for years has been insisting on a calendar year model for loss reserving. See, for example, Barnett and Zehnwirth (2000).

For Discussion

Actuaries have long recognized the effect of changing settlement/recognition rates.

- Actuaries have traditionally adjusted the data to reflect changing claim settlement practices.
- Bayesian think of the data as being reality, and the model is random.
- Don't mess with the data!

For Discussion

Actuaries have long recognized the effect of changing settlement/recognition rates.

- Actuaries have traditionally adjusted the data to reflect changing claim settlement practices.
- Bayesian think of the data as being reality, and the model is random.
- Don't mess with the data!
- In this talk I have attempted to recognize the changing claim settlement practices explicitly in the model.