Two Studies on Stochastic Loss Reserving

Between Line Dependencies
Cost of Capital Risk Margin for Loss Reserve Liabilities

Glenn Meyers

Presentation to Casualty Loss Reserve Seminar

September 18-20, 2016

Recent History

- Set up CAS Loss Reserve Database in 2011
  - Both upper and lower Schedule P triangles for hundreds of insurers.
  - Purpose was to enable "Aggressive Retrospective Testing" of stochastic loss reserve models.

Recent History

- Set up CAS Loss Reserve Database in 2011
  - Both upper and lower Schedule P triangles for hundreds of insurers.
  - Purpose was to enable "Aggressive Retrospective Testing" of stochastic loss reserve models.
- Published the monograph "Stochastic Loss Reserving Using Bayesian MCMC Models" in 2015
  - Used retrospective testing to identify shortcomings in two currently popular stochastic loss reserve models.
  - Proposed new models to address these shortcomings.

Unfinished Business

- Original purpose for creating the database - quantify dependencies
  - Joint project with the Australian Institute of Actuaries.
  - Project did not succeed!
  - Lesson learned - Pointless to quantify dependencies until we had a good univariate (i.e. single-line) model.
  - That was the purpose of the monograph!

Unfinished Business

- Original purpose for creating the database - quantify dependencies
  - Joint project with the Australian Institute of Actuaries.
  - Project did not succeed!
  - Lesson learned - Pointless to quantify dependencies until we had a good univariate (i.e. single-line) model.
  - That was the purpose of the monograph!
- Australian Objective - Risk margin for total loss reserve liability
  - Dependencies matter!

Unfinished Business

- Original purpose for creating the database - quantify dependencies
  - Joint project with the Australian Institute of Actuaries.
  - Project did not succeed!
  - Lesson learned - Pointless to quantify dependencies until we had a good univariate (i.e. single-line) model.
  - That was the purpose of the monograph!
- Australian Objective - Risk margin for total loss reserve liability
  - Dependencies matter!
- Risk Margins are also a concern for Solvency II and the Swiss Solvency Test (SST).
  - Add risk margins by line - assuming perfect dependencies.

Outline of Presentation

- Update Changing Settlement Rate (CSR) model for paid loss triangles
    - Reason - Risk margins deal with discounted loss reserves.

Outline of Presentation

- Update Changing Settlement Rate (CSR) model for paid loss triangles
  - Reason - Risk margins deal with discounted loss reserves.
- Propose a model to deal with dependencies between CSR models by line.
- Compare this model with one that assumes independence between CSR models by line.

Outline of Presentation

- Update Changing Settlement Rate (CSR) model for paid loss triangles
  - Reason - Risk margins deal with discounted loss reserves.
- Propose a model to deal with dependencies between CSR models by line.
- Compare this model with one that assumes independence between CSR models by line.
- Apply the above results to produce a cost of capital risk margin formula.
  - Similarities to Solvency II and SST approaches should be noted.
  - Bayesian MCMC provides a significant enhancement to these approaches.

The Changing Settlement Rate (CSR) Model

- The monograph makes the case that the Mack (chain ladder) model and the bootstrap ODP are biased upward on the CAS Loss Reserve Database data.
- CSR is an attempt to correct this bias.
- Modification of CSR model in the monograph
  - Monograph version assumes constant change in settlement rate.
  - New version allows settlement rate to change.
- Notation
  - w = Accident Year
  - d = Development Year
  - X_{wd} = Cumulative Paid Loss

The Changing Settlement Rate (CSR) Model

- $logelr$ ~uniform(-1.5,0.5)
- $\alpha_1 = 0, \alpha_w \sim$ normal$(0,\sqrt{10})$ for $w = 2,\ldots,10$.
- $\beta_{10} = 0, \beta_d \sim$ uniform(-5,5) for $d = 1,\ldots,9$.
- $\sigma_d^2 \sim \sum_{i=d}^{10} a_i$ for $d = 1,\ldots,10$, where $a_i \sim$ uniform(0,1)
- $\mu_{w,d} =\log(Premium_w)+logelr + \alpha_w + \beta_d \cdot speedup_w$
- $X_{w,d} \sim$ lognormal$(\mu_{w,d}, \sigma_d)$

Features of the CSR Model

- $\mu_{w,d} = \log(Premium_w) + logelr + \alpha_w + \beta_d \cdot speedup_w$
  - The α_w parameter allows the expected loss ratio to change by accident year.
  - The $\beta_d \cdot speedup_w$ product (or interaction) allows the loss development factors to change by accident year.
- $speedup_1 = 1$
- $speedup_w = speedup_{w-1} \cdot (1 - \gamma - (w-2)) \cdot \delta)$
- Speedup Rate $= \gamma - (w-2) \cdot \delta$.
  - $\gamma \sim$ normal(0,0.05), $\delta \sim$ normal(0,0.01)
  - If positive, claim settlement speeds up.
  - If negative, claim settlement slows down
  - The δ parameter allows the speedup rate to change over time.

Posterior Sample of Size 10,000 with Bayesian MCMC

- For each parameter set in the sample get
  - $\{\alpha_w\}_{w=2}^{10}, \{\beta_d\}_{d=1}^{9}, \{\sigma_d\}_{d=1}^{10}, logelr, \gamma, \delta$
- Calculate $\mu_{w,10}$
- Simulate $X_{w,10} \sim \text{lognormal}(\mu_{w,10}, \sigma_{10})$
- Calculate $\sum_{w=1}^{10} X_{w,10}$

Result is a sample of 10,000 outcomes from the predictive distribution of total losses.

Posterior Means of γ Over All Insurers

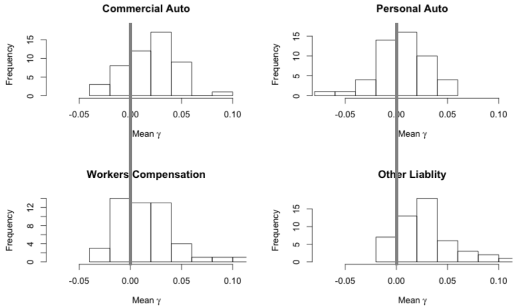

Generally, claim settlement is speeding up.

Criteria for Testing Stochastic Models

- Using the predictive distributions, find the percentiles of the outcome data for several loss triangles.
- The percentiles should be uniformly distributed.
  - Histograms
  - PP Plots and the Kolmogorov-Smirnov Test
    - Plot Expected vs Predicted Percentiles
    - KS Critical Values - 19.2 for $N = 50$ or 9.6 for $N = 200$

Illustrative Tests of Uniformity

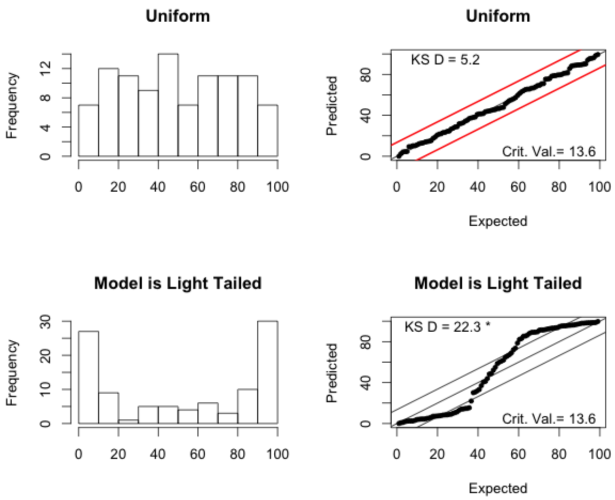

Illustrative Tests of Uniformity

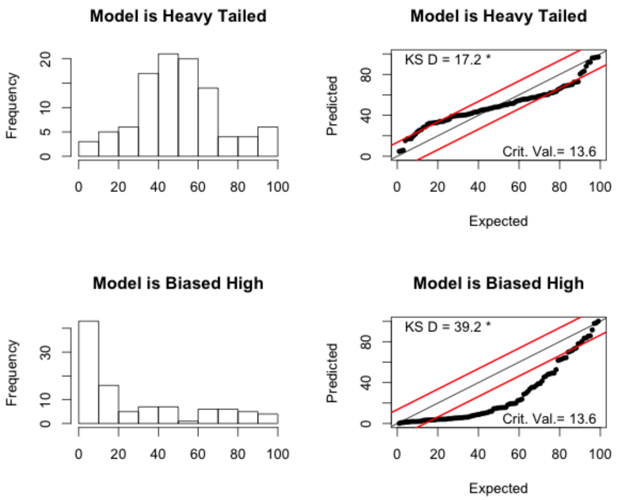

Mack Model on Paid Data

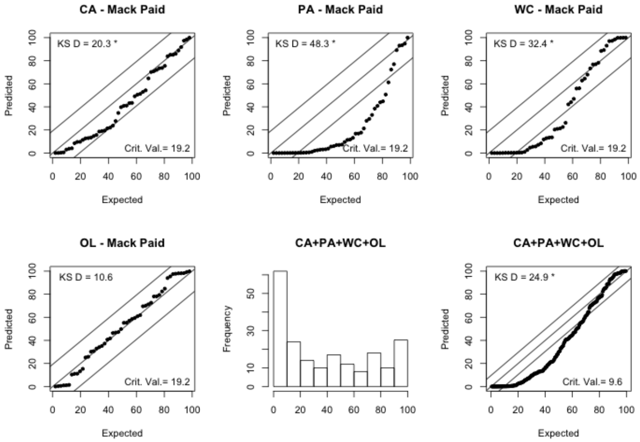

Conclusion - Mack model is biased upward.

CSR on Paid Data

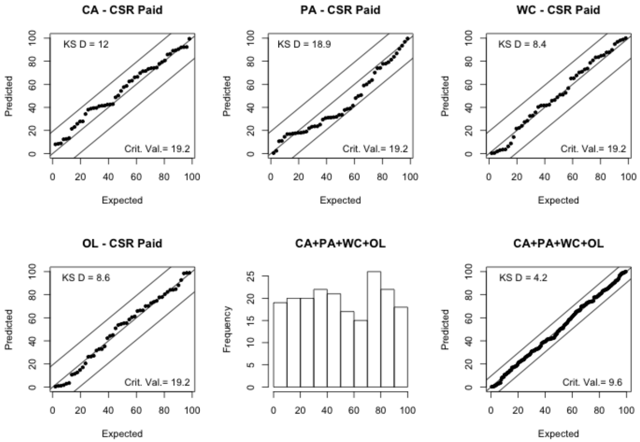

Conclusion - Validates within KS Boundaries.

Meaning of the Successful Validation

- Recall the lesson learned from the dependency project with the Australian Institute of Actuaries.
- Pointless to quantify dependencies until we had a good univariate (i.e. single-line) model.
- We have a univariate model that is suitable for the study of dependencies.

First - Beware of Artificial Correlations

- Red model - Constant — Green model - Parabolic

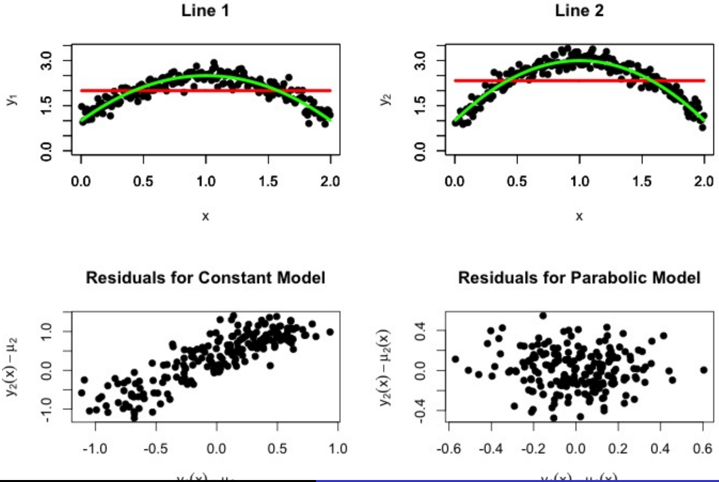

Additional Reference

- "Correlations between Insurance Lines of Business: An Illusion or a Real Phenomenon? Some Methodological Considerations"
- Authors - Benjamin Avanzi, Greg Taylor, Bernard Wong
- Appears in May 2016 *ASTIN Bulletin*

Downloadable from UNSW website

Dependencies - A Recent Development

- "Predicting Multivariate Insurance Loss Payments Under a Bayesian Copula Framework"
  - by Yanwei (Wayne) Zhang - FCAS and Vanja Dukic
  - Awarded the 2014 ARIA Prize by CAS

The General Idea Behind Zhang/Dukic

Given Bayesian MCMC models:

- $X_1 \sim$ Bayesian MCMC Model 1
- $X_2 \sim$ Bayesian MCMC Model 2, then:
- Fit the joint (X_1, X_2) with a joint Bayesian MCMC model.

The General Idea Behind Zhang/Dukic

Given Bayesian MCMC models:

- $X_1 \sim$ Bayesian MCMC Model 1
- $X_2 \sim$ Bayesian MCMC Model 2, then:
- Fit the joint (X_1, X_2) with a joint Bayesian MCMC model.
  - The marginal distributional model is of the same parametric form as the original models.
  - However the parameters of the univariate and marginal models may differ.

The General Idea Behind Zhang/Dukic

Given Bayesian MCMC models:

- $X_1 \sim$ Bayesian MCMC Model 1
- $X_2 \sim$ Bayesian MCMC Model 2, then:
- Fit the joint (X_1, X_2) with a joint Bayesian MCMC model.
  - The marginal distributional model is of the same parametric form as the original models.
  - However the parameters of the univariate and marginal models may differ.
- Marginal and univariate parameters were significantly different when I applied their approach with the CSR model.

The General Idea Behind Zhang/Dukic

Given Bayesian MCMC models:

- $X_1 \sim$ Bayesian MCMC Model 1
- $X_2 \sim$ Bayesian MCMC Model 2, then:
- Fit the joint (X_1, X_2) with a joint Bayesian MCMC model.
  - The marginal distributional model is of the same parametric form as the original models.
  - However the parameters of the univariate and marginal models may differ.
- Marginal and univariate parameters were significantly different when I applied their approach with the CSR model.
- I obtained better agreement between the marginal and univariate parameters with the model that Zhang/Dukic used in their paper.

Two Steps to Fitting a Bivariate Model That Preserves Univariate Fits

Joint Lognormal Distribution

$$\begin{pmatrix} \log\left(X_{wd}^{1}\right) \\ \log\left(X_{wd}^{2}\right) \end{pmatrix} \sim \text{Normal}\left(\begin{pmatrix} \mu_{wd}^{1} \\ \mu_{wd}^{2} \end{pmatrix}, \begin{pmatrix} \left(\sigma_d^1\right)^2 & \rho \cdot \sigma_d^1 \cdot \sigma_d^2 \\ \rho \cdot \sigma_d^1 \cdot \sigma_d^2 & \left(\sigma_d^2\right)^2 \end{pmatrix} \right)$$

Two Steps to Fitting a Bivariate Model That Preserves Univariate Fits

Joint Lognormal Distribution

$$\begin{pmatrix} \log\left(X_{wd}^1\right) \\ \log\left(X_{wd}^2\right) \end{pmatrix} \sim \text{Normal}\left(\begin{pmatrix} \mu_{wd}^1 \\ \mu_{wd}^2 \end{pmatrix}, \begin{pmatrix} \left(\sigma_d^1\right)^2 & \rho \cdot \sigma_d^1 \cdot \sigma_d^2 \\ \rho \cdot \sigma_d^1 \cdot \sigma_d^2 & \left(\sigma_d^2\right)^2 \end{pmatrix} \right)$$

1. Use Bayesian MCMC to get a sample of 10,000 μ_{wd}s and σ_ds for each line 1 and 2 (= CA, PA, WC and OL).

Two Steps to Fitting a Bivariate Model That Preserves Univariate Fits

Joint Lognormal Distribution

$$\begin{pmatrix} \log\left(X_{wd}^{1}\right) \\ \log\left(X_{wd}^{2}\right) \end{pmatrix} \sim \text{Normal}\left(\begin{pmatrix} \mu_{wd}^{1} \\ \mu_{wd}^{2} \end{pmatrix}, \begin{pmatrix} \left(\sigma_d^1\right)^2 & \rho \cdot \sigma_d^1 \cdot \sigma_d^2 \\ \rho \cdot \sigma_d^1 \cdot \sigma_d^2 & \left(\sigma_d^2\right)^2 \end{pmatrix}\right)$$

1. Use Bayesian MCMC to get a sample of 10,000 μ_{wd}s and σ_ds for each line 1 and 2 (= CA, PA, WC and OL).
2. For each **parameter set** in the univariate sample for each line, use Bayesian MCMC to get a single ρ from the bivariate distribution of $\left(\log(X_{wd}^1), \log(X_{wd}^2)\right)$.

Posterior Mean of ρ for 102 Pairs of Triangles

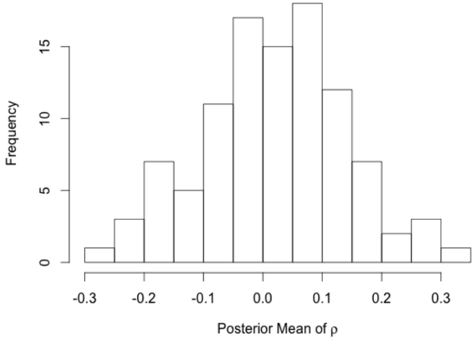

Note - $\overline{\rho}$ is fairly symmetric around 0.

Of Particular Interest - The Distribution of the Sum of Losses for Two Lines of Insurance

$$\sum_{w=1}^{10} X_{w,10}^{\{1\}} + \sum_{w=1}^{10} X_{w,10}^{\{2\}}$$

- From the 2-step bivariate model
- From the independent model formed as a random sum of losses from the univariate models

Retro Test of the Sum from the Two-Step Bivariate Model on 102 Pairs of Lines

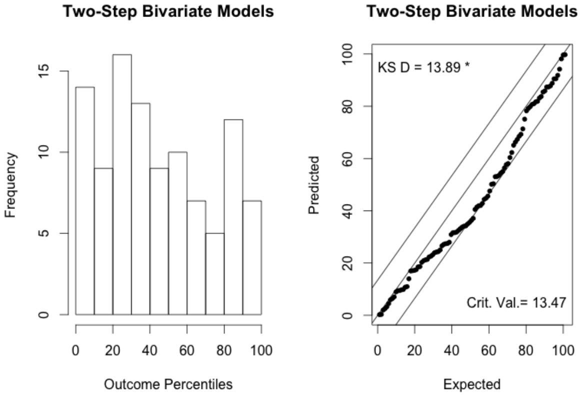

Just outside the 95% confidence band.

Retro Test of the Sum from the Independent Model on 102 Pairs of Lines

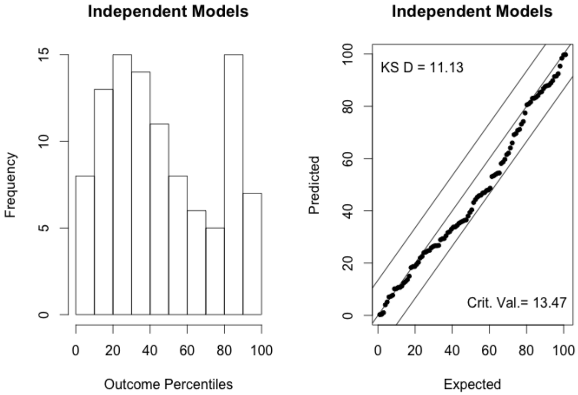

Just inside the 95% confidence band.

Model Selection on Training Data

- If we fit model, f, by maximum likelihood define

$$AIC = 2 \cdot p - 2 \cdot L\left(x|\hat{\theta}\right)$$

- Where
  - p is the number of parameters.
  - $L\left(x|\hat{\theta}\right)$ is the maximum log-likelihood of the model specified by f.
- Lower AIC indicates a better fit.
  - Encourages larger log-likelihood
  - Penalizes for increasing the number of parameters

Bayesian Model Selection the WAIC Statistic

- Given an MCMC model with parameters $\{\theta_i\}_{i=1}^{10{,}000}$

$$WAIC = 2 \cdot \hat{p} - 2 \cdot \overline{\{L(x|\theta_i)\}}_{i=1}^{10{,}000}$$

- Where
  - $\hat{p}$ is the **effective** number of parameters.
  - $\hat{p}$ decreases as the prior distribution becomes more "informative" i.e. less influenced by the data.
  - $\overline{\{L(x|\theta_i)\}}_{i=1}^{10{,}000}$ = average log-likelihood.
- WAIC is calculated with the "loo" package in R.

The Leave One Out Information Criteria (LOOIC)

- Given an MCMC model with data vector, x, and parameter vectors $\{\theta_i\}_{i=1}^{10,000}$, define:

$$\text{LOOIC} = 2 \cdot \hat{p}_{LOOIC} - 2 \cdot \overline{\left\{L(x|\theta_i)\right\}}_{i=1}^{10,000}$$

- L denotes the log-likelihood of x.
- $\hat{p}_{LOOIC}$ is the **effective number of parameters**.

The Leave One Out Information Criteria (LOOIC)

- Given an MCMC model with data vector, x, and parameter vectors $\{\theta_i\}_{i=1}^{10,000}$, define:

$$\text{LOOIC} = 2 \cdot \hat{p}_{LOOIC} - 2 \cdot \overline{\left\{L(x|\theta_i)\right\}_{i=1}^{10,000}}$$

- L denotes the log-likelihood of x.
- $\hat{p}_{LOOIC}$ is the **<u>effective</u> number of parameters**.

$$\hat{p}_{LOOIC} = \overline{\{\left\{L(x|\theta_i)\right\}_{i=1}^{10,000}} - \overline{\left\{\sum_{j=1}^{J}\left\{L(x_j|x_{-j},\theta_i)\right\}\right\}_{i=1}^{10,000}}$$

- $x_{-j} = (x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_J)$.
- LOOIC is approximated with the "loo" package in R.

Choosing Between 2-Step and Independent Models

- The WAIC and LOOIC statistics indicate that the independent model is preferred

Choosing Between 2-Step and Independent Models

- The WAIC and LOOIC statistics indicate that the independent model is preferred

For all 102 pairs of lines!

Choosing Between 2-Step and Independent Models

- The WAIC and LOOIC statistics indicate that the independent model is preferred

For all 102 pairs of lines!

- Counterintuitive to many actuaries.
  - Inflation affects all claims simultaneously.
  - Underwriting cycle effects

Choosing Between 2-Step and Independent Models

- The WAIC and LOOIC statistics indicate that the independent model is preferred

For all 102 pairs of lines!

- Counterintuitive to many actuaries.
  - Inflation affects all claims simultaneously.
  - Underwriting cycle effects
- I think I owe an explanation.

The Changing Settlement Rate (CSR) Model

- $logelr \sim$uniform(-5,0)
- $\alpha_1 = 0, \alpha_w \sim$ normal$(0, \sqrt{10})$ for $w = 2, \ldots, 10$.
- $\beta_{10} = 0, \beta_d \sim$ uniform(-5,5) for $d = 1, \ldots, 9$.
- $\sigma_d^2 \sim \sum_{i=d}^{10} a_i$ for $d = 1, \ldots, 10$, where $a_i \sim$ uniform(0,1)
- $\mu_{w,d} = \log(Premium_w) + logelr + \alpha_w + \beta_d \cdot speedup_w$
- $X_{w,d} \sim$ lognormal$(\mu_{w,d}, \sigma_d)$

The Stochastic Cape Cod (SCC) Model

- *logelr* ~uniform(-5,0)
- ~~$\alpha_1 = 0, \alpha_w \sim$ normal$(0,\sqrt{10})$ for $w = 2,\ldots,10$.~~
- $\beta_{10} = 0, \beta_d \sim$ uniform(-5,5) for $d = 1,\ldots,9$.
- $\sigma_d^2 \sim \sum_{i=d}^{10} a_i$ for $d = 1,\ldots,10$, where $a_i \sim$ uniform(0,1)
- $\mu_{w,d} = \log(Premium_w) + logelr$ ~~$+ \alpha_w + \beta_d \cdot speedup_w$~~
- $X_{w,d} \sim$ lognormal$(\mu_{w,d}, \sigma_d)$

The Stochastic Cape Cod (SCC) Model

- Simpler than the CSR model
- Resembles an industry standard
  - Bornhuetter Ferguson with a constant ELR
  - Source Dave Clark and Jessica Leong in the references
- 2-Step SCC model is preferred for some insurers
- Look at a sample of standardized residual plots
- Insurer 5185 for CA and OL favors 2-Step
  - Picked as an illustration

Posterior Distribution of ρ for Insurer 5185

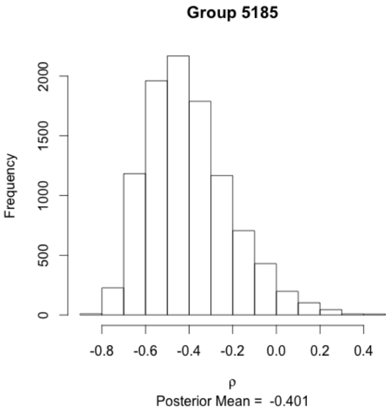

Note the negative posterior mean ρ.

Standardized Residual Plots for Insurer 5185

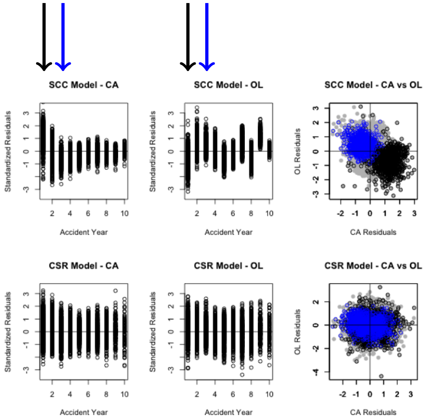

AY 1 borders are black

AY 3 borders are blue

In general, SCC residuals tend to find their own corner. If many are in the NW-SE corner, we see a negative mean ρ.

Implications of Independence

- Cost of capital risk margins should have a "diversification" credit. As an example, the EU Solvency II adds risk margins by line of business, implicitly denying a diversification credit.
- With a properly valid MCMC stochastic loss reserve model, one can get 10,000 stochastic scenarios of the future and calculate a cost of capital risk margin, and reflect diversification.

A Proposed "Law" for Dependency Modeling

- Using the 2-Step procedure, we can fit bivariate distributions.
- We can compare the 2-Step model to a model that assumes independence.

A Proposed "Law" for Dependency Modeling

- Using the 2-Step procedure, we can fit bivariate distributions.
- We can compare the 2-Step model to a model that assumes independence.

The Law

- If your dependent bivariate model is "better" than the independent model, you should look for something that is missing from your model.

What do we do with a predictive distribution of outcomes?

- Actuary presents a range to management.
- Management selects the posted liability.

Or

What do we do with a predictive distribution of outcomes?

- Actuary presents a range to management.
- Management selects the posted liability.

Or

- The “best estimate” is defined as the probability weighted average of the discounted future loss payouts.
- The loss reserve liability (a.k.a. “technical provision”) is equal to the best estimate plus a risk margin.

What do we do with a predictive distribution of outcomes?

- Actuary presents a range to management.
- Management selects the posted liability.

Or

- The "best estimate" is defined as the probability weighted average of the discounted future loss payouts.
- The loss reserve liability (a.k.a. "technical provision") is equal to the best estimate plus a risk margin.
  - Wider range increases the risk margin.
  - Longer payout increases the risk margin.
- A cost of capital risk margin has these properties.

The Insurer Capital Cash Flow

- At the end of the current calendar year, $t = 0$:
  - The insurer posts the amount of capital, C_0, needed to contain the uncertainty in its loss estimate. It invests C_0 in a fund that earns income at the risk-free interest rate i.

The Insurer Capital Cash Flow

- At the end of the current calendar year, $t = 0$:
  - The insurer posts the amount of capital, C_0, needed to contain the uncertainty in its loss estimate. It invests C_0 in a fund that earns income at the risk-free interest rate i.
- At $t = 1$, the insurer uses its next year of loss experience to reevaluate its liability.
  - The insurer posts the updated estimate of its capital, C_1, needed to contain the uncertainty its revised loss estimate.
  - C_0 is invested at the risk-free interest rate i.
  - The difference, $C_0 \cdot (1 + i) - C_1$, is returned to the investor. If that difference is negative, the investor is expected to contribute an amount to make up that difference.

The Insurer Capital Cash Flow - Continued

- The process continues for future calendar years, t, with the amount, $C_{t-1} \cdot (1+i) - C_t$, being returned to (or being contributed by) the investor.

The Insurer Capital Cash Flow - Continued

- The process continues for future calendar years, t, with the amount, $C_{t-1} \cdot (1+i) - C_t$, being returned to (or being contributed by) the investor.
- At some time $t = u$, the loss is deemed to at ultimate, i.e. no significant changes in the loss is anticipated. Then we set $C_t = 0$ for $t > u$. For the examples in this presentation, $u = 9$.

The Cost of Capital Risk Margin

- The present value of the capital returned, discounted at risky rate, r is equal to:

$$\sum_{t=1}^{u} \frac{C_{t-1} \cdot (1+i) - C_t}{(1+r)^t}$$

- The cost of capital risk margin is defined as the difference between the initial investment, C_0, and the present value of the capital returned.

$$R_{COC} \equiv C_0 - \sum_{t=1}^{u} \frac{C_{t-1} \cdot (1+i) - C_t}{(1+r)^t} = \frac{1}{r-i} \sum_{t=0}^{u} \frac{C_t}{(1+r)^t}$$

Compare R_{COC} with Solvency II Risk Margin

$$R_{COC} = \frac{1}{r-i}\sum_{t=0}^{u}\frac{C_t}{(1+r)^t}$$

$$R_{SII} = \frac{1}{r-i}\sum_{t=0}^{u}\frac{C_t}{(1+i)^t}$$

Similar, but not identical.

Bayesian MCMC with Risk Margins

- Bayesian MCMC provides a list of scenarios that represent the future.
- The set of parameters of possible lognormal distributions of loss, $X_{w,d}^{j}$, for each w and d is represented by

$$\{\mu_{w,d}^{j}, \sigma_{d}^{j}\}_{j=1}^{10,000}$$

Bayesian MCMC with Risk Margins

- Bayesian MCMC provides a list of scenarios that represent the future.
- The set of parameters of possible lognormal distributions of loss, $X_{w,d}^{j}$, for each w and d is represented by

$$\{\mu_{w,d}^{j}, \sigma_{d}^{j}\}_{j=1}^{10,000}$$

- Use the parameters to calculate statistics of interest.
- In particular - those statistics that are needed for risk margins.

The Best Estimate

- The "best estimate" is defined as the probability weighted average of the discounted future loss payouts.
- Recall - the mean of a lognormal distribution $= e^{\mu+\sigma^2/2}$.
- The mathematical expression for best estimate:

$$E_{Best} = \sum_{j=1}^{10,000} \sum_{w=2}^{10} \sum_{12-w}^{10} \frac{e^{\mu^j_{w,d}+(\sigma^j_d)^2/2} - e^{\mu^j_{w,d-1}+(\sigma^j_{d-1})^2/2}}{10,000 \cdot (1+i)^{w+d-11.5}}$$

- Sum is over the expected payout in the lower triangle.
- Assumes that loss is paid out one half year before the end of calendar year $t = w + d - 11$.

Ultimate Losses

- Capital requirements are based on the variability of the *estimates of the ultimate loss*.
- If we know the scenario, j, a pretty good estimate of the ultimate loss is given by:

$$U_j \equiv \sum_{w=1}^{10} e^{\mu^j_{w,10} + (\sigma^j_{10})^2/2}$$

Ultimate Losses

- Capital requirements are based on the variability of the *estimates of the ultimate loss*.
- If we know the scenario, j, a pretty good estimate of the ultimate loss is given by:

$$U_j \equiv \sum_{w=1}^{10} e^{\mu_{w,10}^{j}+(\sigma_{10}^{j})^2/2}$$

- By "pretty good" I mean that after discounting, uncertainty in further loss development should have a small effect on the risk margin.

Updating the Scenario Probabilities

- We don't know the scenario, j, but we can update the scenario probabilities as more losses come in at the end of each calendar year.
- Initially, all scenarios are equally likely (=1/10,000).
- For future calendar year $t = 1, \ldots, 9$, define the loss trapezoid, T_t as the set of the top t diagonals in the lower triangle.
- Then using Bayes' theorem:

$$Pr[J = j | T_t] = \frac{\prod_{X_{wd} \in T_t} \phi(\log(X_{wd}) | \mu_{wd}^j, \sigma_d^j)}{\sum_{k=1}^{10{,}000} \prod_{X_{wd} \in T_t} \phi(\log(X_{wd}) | \mu_{wd}^k, \sigma_d^k)}$$

Calculating Statistics of Interest for Risk Margins

1. Calculate $P_t \equiv \left\{ Pr[J = j | T_t] \right\}_{j=1}^{10{,}000}$
2. Take a sample, S_t, of size 10,000 from $\left\{ U_j \right\}_{j=1}^{10{,}000}$ with sampling probabilities P_t. Easy to do in R.

Calculating Statistics of Interest for Risk Margins

1. Calculate $P_t \equiv \left\{Pr[J = j|\, T_t]\right\}_{j=1}^{10{,}000}$
2. Take a sample, S_t, of size 10,000 from $\left\{U_j\right\}_{j=1}^{10{,}000}$ with sampling probabilities P_t. Easy to do in R.

- $E_t = \text{mean}(S_t) =$ Expected ultimate loss at the end of future calendar year T_t.
- $A_t = TVaR@99\% = \text{mean}\big(sort(S_t)[9901 : 10000]\big) =$ Required assets at the end of future calendar year T_t.
- $C_t = A_t - E_t =$ Capital required at the end of future calendar year T_t.
- $R_t = C_t - C_{t-1} =$ Capital released at the end of future calendar year t.

Steps for Calculating the Risk Margin

Do the following 10,000 times.

1. Select a parameter scenario at random.
2. Calculate the required capital, C_t, for future calendar years $t = 0, \ldots, 9$.
3. Calculate the risk margin using the following formula:

$$R_{COC} \equiv C_0 - \sum_{t=1}^{9} \frac{C_{t-1} \cdot (1+i) - C_t}{(1+r)^t}$$

- Posted Risk Margin is equal the the average of the 10,000 risk margins calculated above.

Example - Insurer 353 in Commercial Auto
$i = 4\%$ and $r = 10\%$

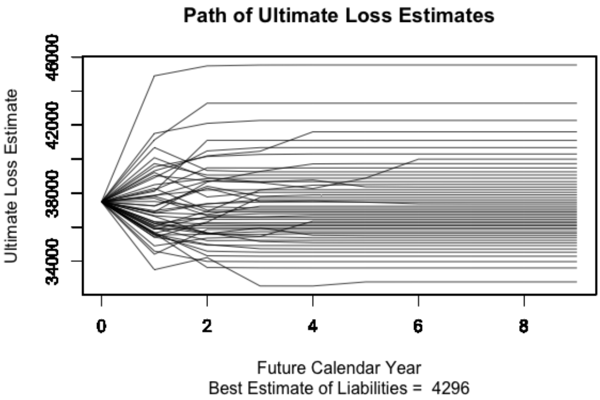

Note the leveling off of the paths.

Example - Insurer 353 in Commercial Auto $i = 4\%$ and $r = 10\%$

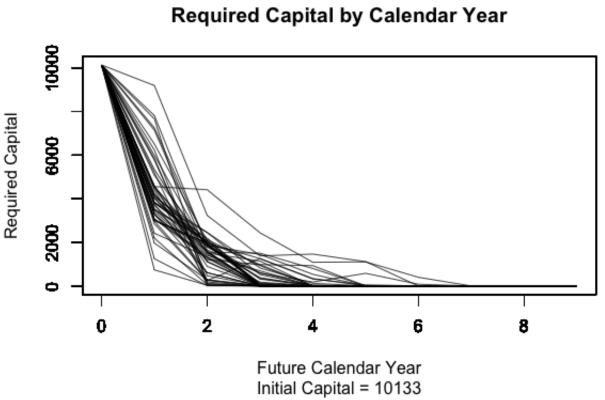

Note that capital requirement decrease to very near zero.

Example - Insurer 353 in Commercial Auto $i = 4\%$ and $r = 10\%$

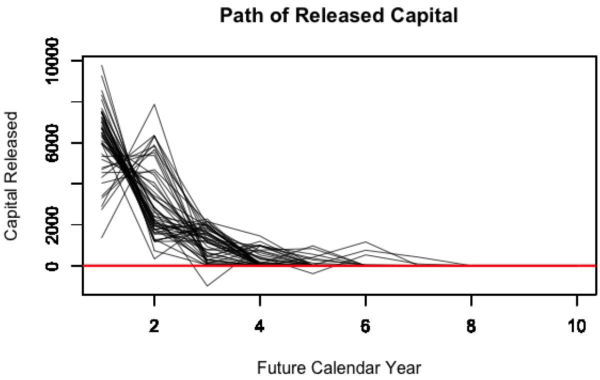

Occasionally the insurer must add capital.

Example - Insurer 353 in Commercial Auto $i = 4\%$ and $r = 10\%$

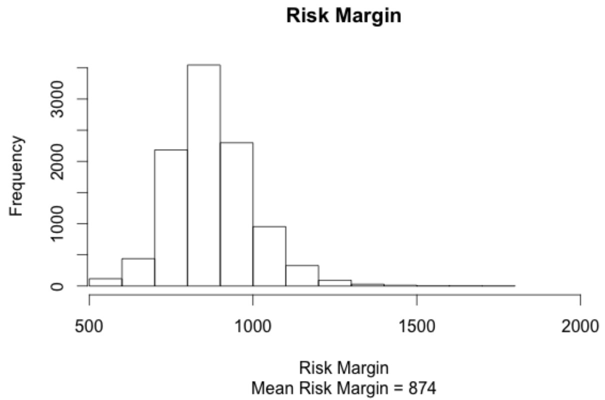

Example - Insurer 353 in Commercial Auto $i = 4\%$ and $r = 10\%$

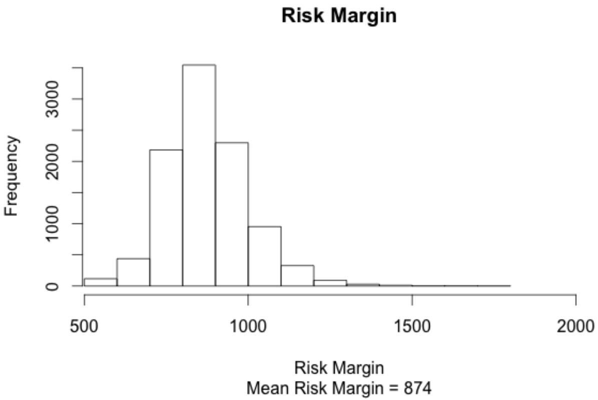

Should the variability of the risk margin influence r?

From an Ultimate to a One-Year Time Horizon

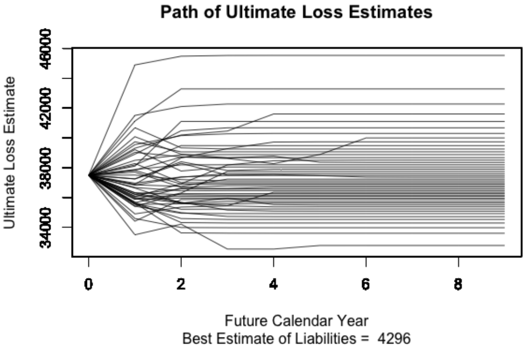

- A difference in the spread as t goes from 1 to 9.
- C_t calculated assuming the "ultimate" time horizon.
- Many regulators want to use a one-year time horizon.

Adopting the MCMC Strategy to a One Year Horizon

- Assigning the ultimate, U_j, to a scenario, j, is a good approximation for the "ultimate" time horizon.
- Not a good approximation for a one-year time horizon.
- We need to assign the expected value of the estimate given the scenario.
- Do the following 12 times for each scenario, j.
    1. Simulate the lower triangle with the scenario parameters.
    2. Calculate E_ts as described above.
- Set O_j equal to the average of the E_ts.
- Replace each U_j with O_j and proceed as above.

Our Example with a One-Year Time Horizon

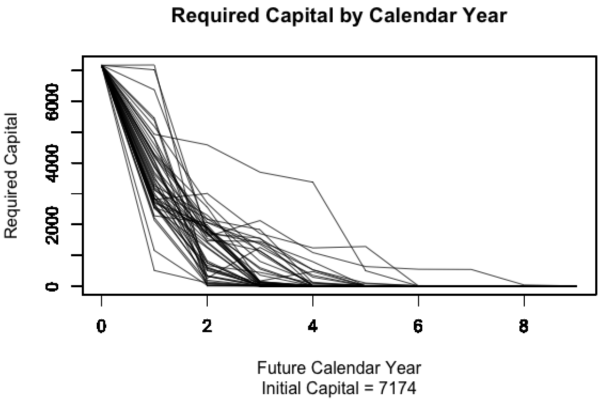

Initial capital for an "ultimate" time horizon = 10,133.

Our Example with a One-Year Time Horizon

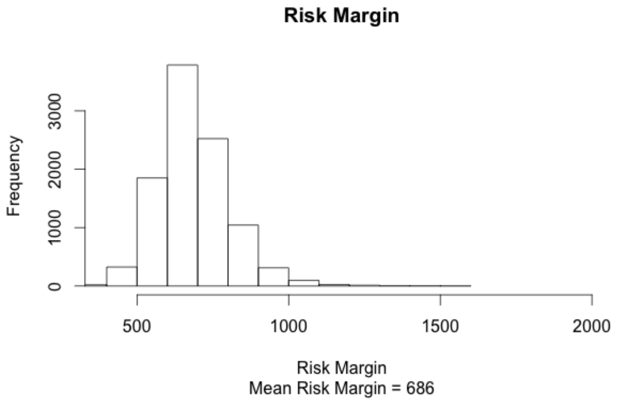

Risk margin for an "ultimate" time horizon $= 874$.

Which Time Horizon?

- Significantly longer run time for one-year - 2.5 minutes vs 20 minutes.
- Why not lower TVaR level on ultimate - say 97%?
- Solvency II Experience
  - Capital requirement based on one-year time horizon.

Which Time Horizon?

- Significantly longer run time for one-year - 2.5 minutes vs 20 minutes.
- Why not lower TVaR level on ultimate - say 97%?
- Solvency II Experience
  - Capital requirement based on one-year time horizon.
  - Need to validate internal model.
  - It takes a long time to validate estimates of ultimate.

Recall Independence Between Lines Result Above

- Example - Insurer 353 - Commercial and Personal Auto

Risk Margin for Commercial Auto	874
Risk Margin for Personal Auto	924
Solvency II sums risk margins by line	1,798
Independence sums the samples $S_t^{CA} + S_t^{PA}$	1,233

- Significant Diversification!

Concluding Remarks

- Bayesian MCMC models provide unprecedented flexibility in developing stochastic loss reserve models that validate of lower triangle data.
  - Correlated Chain Ladder (CCL) model on incurred triangles.
  - Changing Settlement Rate (CSR) model on paid triangles.

Concluding Remarks

- Bayesian MCMC models provide unprecedented flexibility in developing stochastic loss reserve models that validate of lower triangle data.
  - Correlated Chain Ladder (CCL) model on incurred triangles.
  - Changing Settlement Rate (CSR) model on paid triangles.
- Developed a model to fit bivariate model on paid loss triangles from two lines of insurance.
- Showed that a bivariate model that assumed independence provided a better fit to the bivariate loss triangles.
  - You need a good univariate (CSR) model to get independence.

Concluding Remarks

- A validated Bayesian MCMC model provides a sample of parameters that represents the future loss development.
  - Can model future statistics of interest.
  - Did it for the statistics relevant for cost of capital risk margins.
- Applying the dependency results to risk margins indicate that there should be a diversification credit for risk margins.