The Leveled Chain Ladder Model for Stochastic Loss Reserving

Appears in CAS E-Forum with R/JAGS Code

http://casact.org/pubs/forum/12sumforum/

Glenn Meyers

Presented at the ASTIN Colloquim

October 2, 2012

S&P Report, November 2003

Insurance Actuaries – A Crisis in Credibility

"Actuaries are signing off on reserves that turn out to be wildly inaccurate."

How to Approach the Problem

- There are a LOT of different approaches to stochastic loss reserving derived from any number of “reasonable” assumptions.
- First step – Decide on how to evaluate stochastic loss reserve models.
  - See “Estimating the Predictive Distribution for Loss Reserve Models” *Variance* 2007
- Appeal to the “Gold Standard” for predictive modeling – Performance on holdout data.

Criteria for a "Good" Stochastic Loss Reserve Model

- Using the upper triangle "training" data, predict the distribution of the outcomes in the lower triangle
- Using the predictive distributions, find the percentiles of the outcome data.
- The percentiles should be uniformly distributed.
  - Histograms
  - Test with PP Plots/KS tests
    - Plot Expected vs Predicted Percentiles

Illustrative Tests of Uniformity

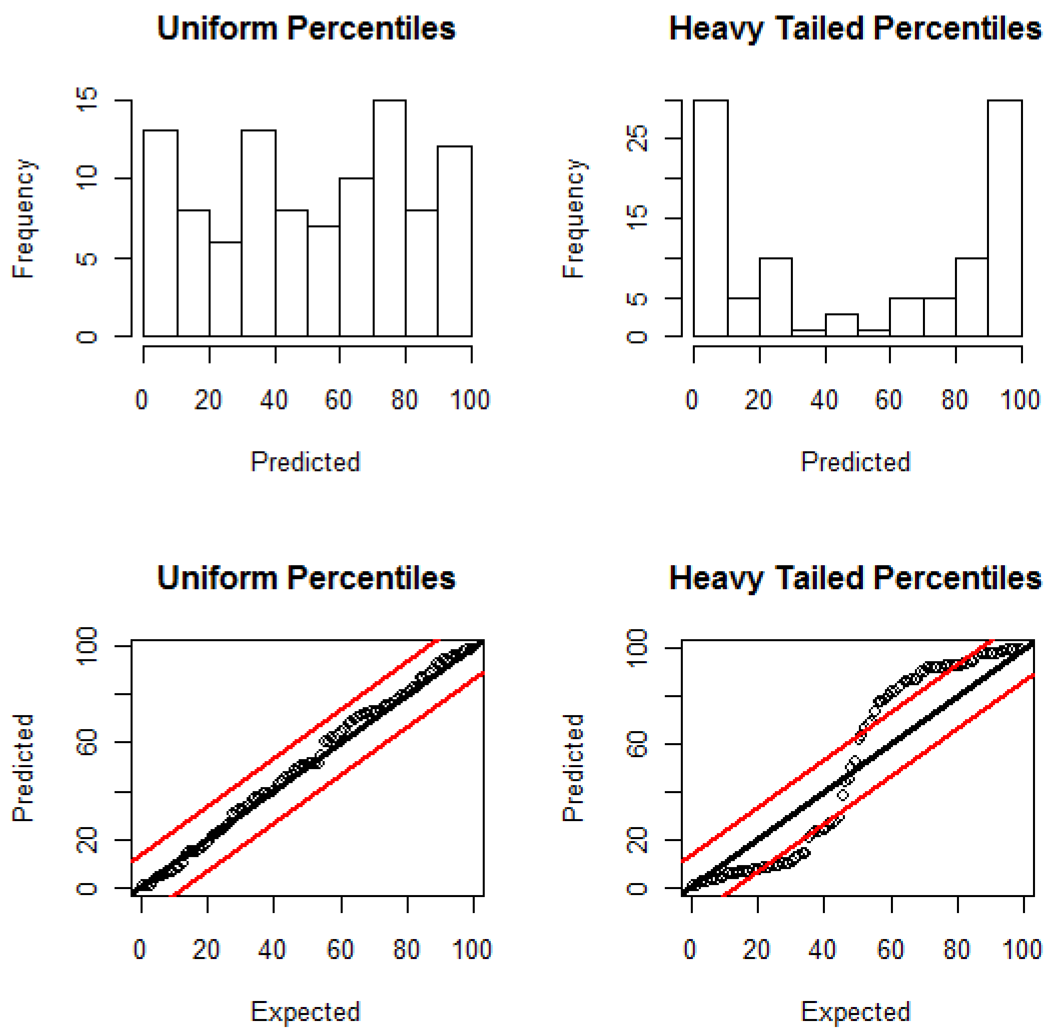

The CAS Loss Reserve Database

- Schedule P (Data from Parts 1-4) for several US Insurers
  - Private Passenger Auto
  - Commercial Auto
  - Workers' Compensation
  - General Liability
  - Product Liability
  - Medical Malpractice (Claims Made)
- Available on CAS Website

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

The CAS Loss Reserve Database

		Settlement Lag										
Accident Year	Premium	1	2	3	4	5	6	7	8	9	10	
1988	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	
1989	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx		← 1998
1990	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx			← 1999
1991	xxx	xxx	xxx	xxx	xxx	xxx	xxx	xxx				← 2000
1992	xxx	xxx	xxx	xxx	xxx	xxx	xxx					← 2001
1993	xxx	xxx	xxx	xxx	xxx	xxx						← 2002
1994	xxx	xxx	xxx	xxx	xxx							← 2003
1995	xxx	xxx	xxx	xxx								← 2004
1996	xxx	xxx	xxx									← 2005
1997	xxx	xxx										← 2006

Training Data from 1997 Schedule P

Outcome Data from Later Schedule Ps

- Can we predict the distribution of outcomes? Or sums of outcomes?

History

- "The Retrospective Testing of Stochastic Loss Reserve Models." CAS 2011 Summer e-Forum
  - Coauthored with Peng Shi

    http://casact.org/pubs/forum/11sumforum/Meyers-Shi.pdf
  - Described the CAS Loss Reserve Database and (1) demonstrated two models that don't validate on incremental paid data, and (2) Found that reported loss reserve estimates were more accurate than the models above.

Meyers – Shi 2011

Tested Two Models on Incremental Paid Data for Commercial Auto

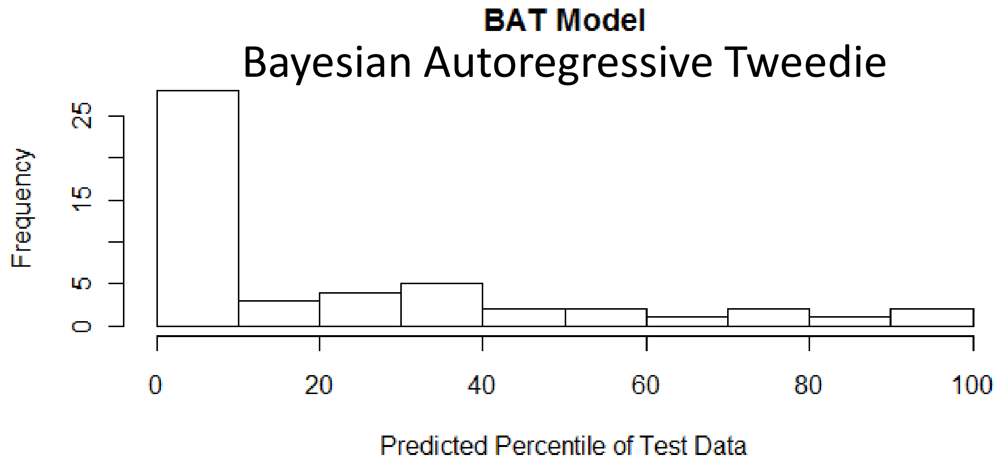

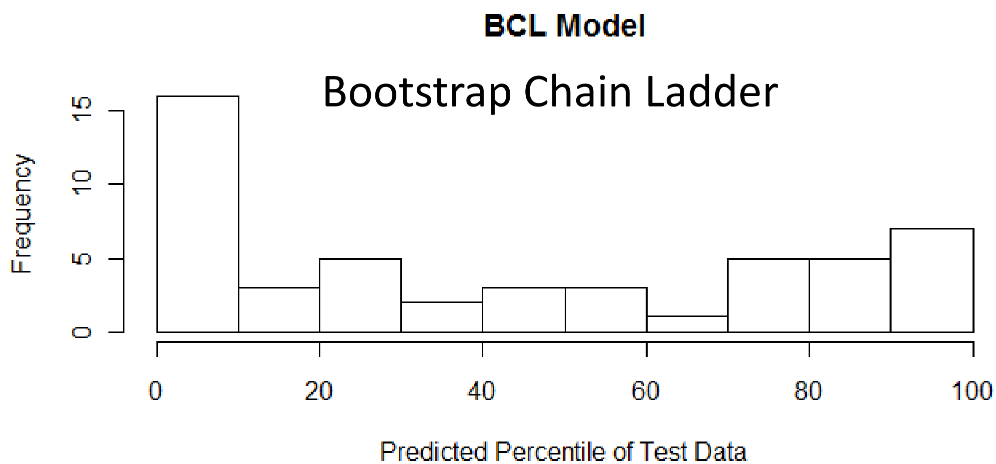

Lessons from History

- Subsequent to Meyers-Shi, I investigated what might have gone wrong.
  - Concluded that the incurred information available to company actuaries could explain the increased accuracy.
  - Decided to start over armed with the lessons learned
- Features of the start over
  - Work with both paid and incurred data and compare results
    - Use models based on cumulative rather than incremental data.
  - Use Bayesian models implemented with the latest MCMC algorithms.
  - Aggressive retrospective testing of the predictive distribution on actual outcomes
    - Use the CAS Loss Reserve database

Design of Retrospective Test For 50 Insurers in CA, PA, WC and OL

- Estimate the predictive distribution of the cumulative reported claims, $C_{w,10}$, at accident year w and development year 10, for each insurer using both models.

$$\sum_{w=2}^{10} C_{w,10}$$

- Calculate the percentile of the reported sum for each insurer using a model.
- Test the uniformity of the calculated percentiles for both models
- Goal is to find a model that passes this uniformity test
- It may not be possible. Consider an environment where losses are repeatedly influenced by different "black swan" events that ***we, the modelers, and the "data" do not see coming.***

Analysis of the Mack Model

- Example on illustrative triangle
  - On real data from CAS Loss Reserve Database
  - Use the R ChainLadder Package
- Fit with both incurred and paid triangles for all 50 insurers in each line of insurance
- Calculate percentiles of outcomes
- Test for the uniformity of percentiles
  - For each line
  - For all lines combined

The Illustrative Triangle

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

Results for Mack on the Illustrative Triangle

	Chain Ladder/Mack			
w	Estimate	Std. Error	CV	Actual
1	3,917	0	0.0000	3,917
2	2,538	0	0.0000	2,532
3	4,167	3	0.0007	4,279
4	4,367	37	0.0085	4,341
5	3,597	34	0.0095	3,587
6	3,236	40	0.0124	3,268
7	5,358	146	0.0272	5,684
8	3,765	225	0.0598	4,128
9	4,013	412	0.1027	4,144
10	3,955	878	0.2220	4,181
Total *w*=2-10	34,997	1,057	0.0302	36,144

Uniformity Tests with 50 Insurers

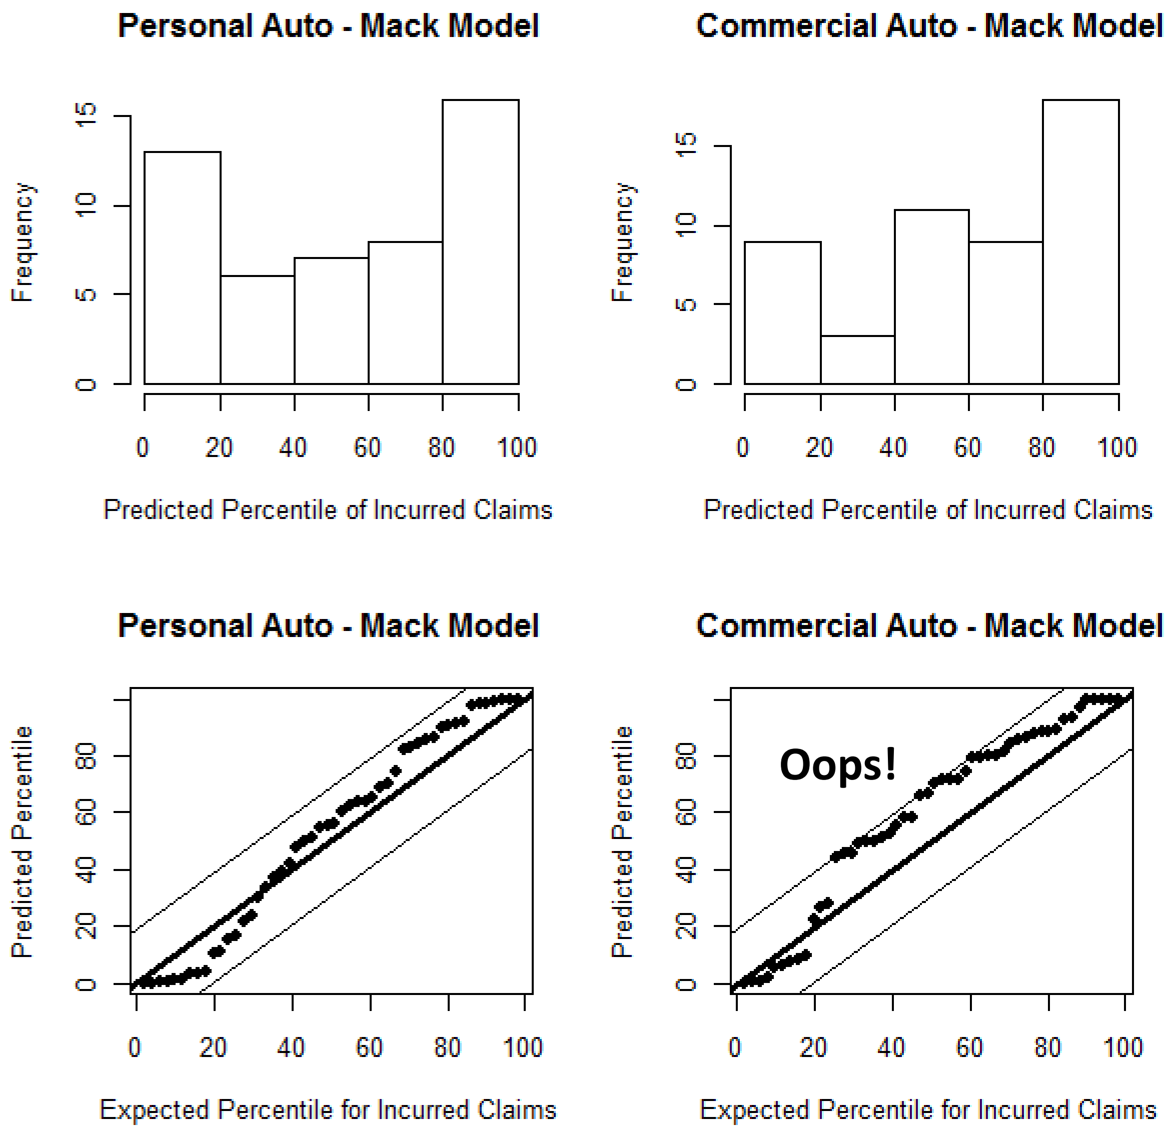

Uniformity Tests with 50 Insurers

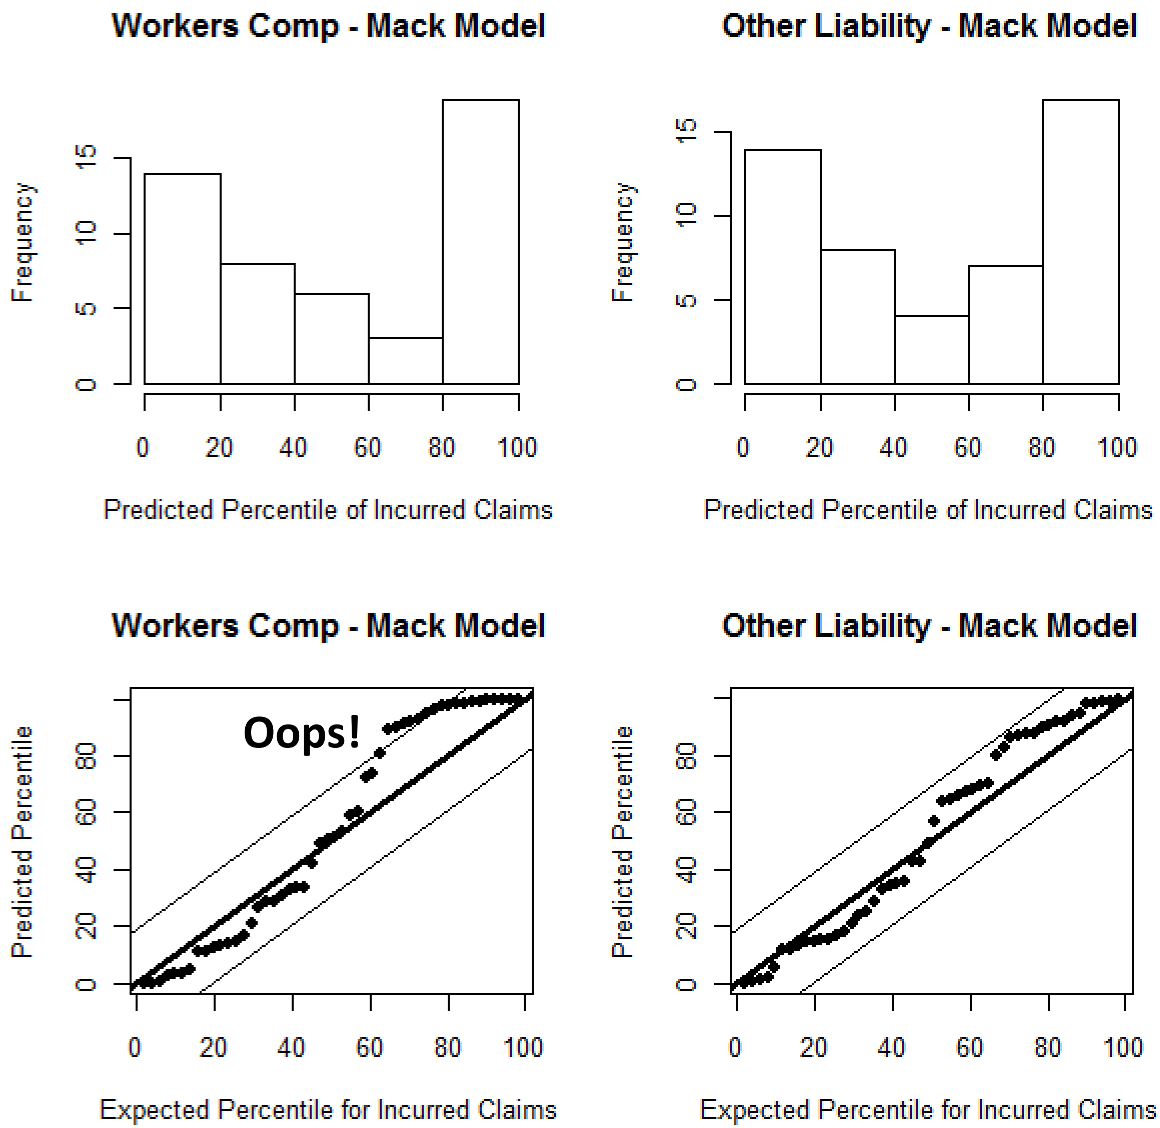

Uniformity Tests with 4 x 50 Insurers

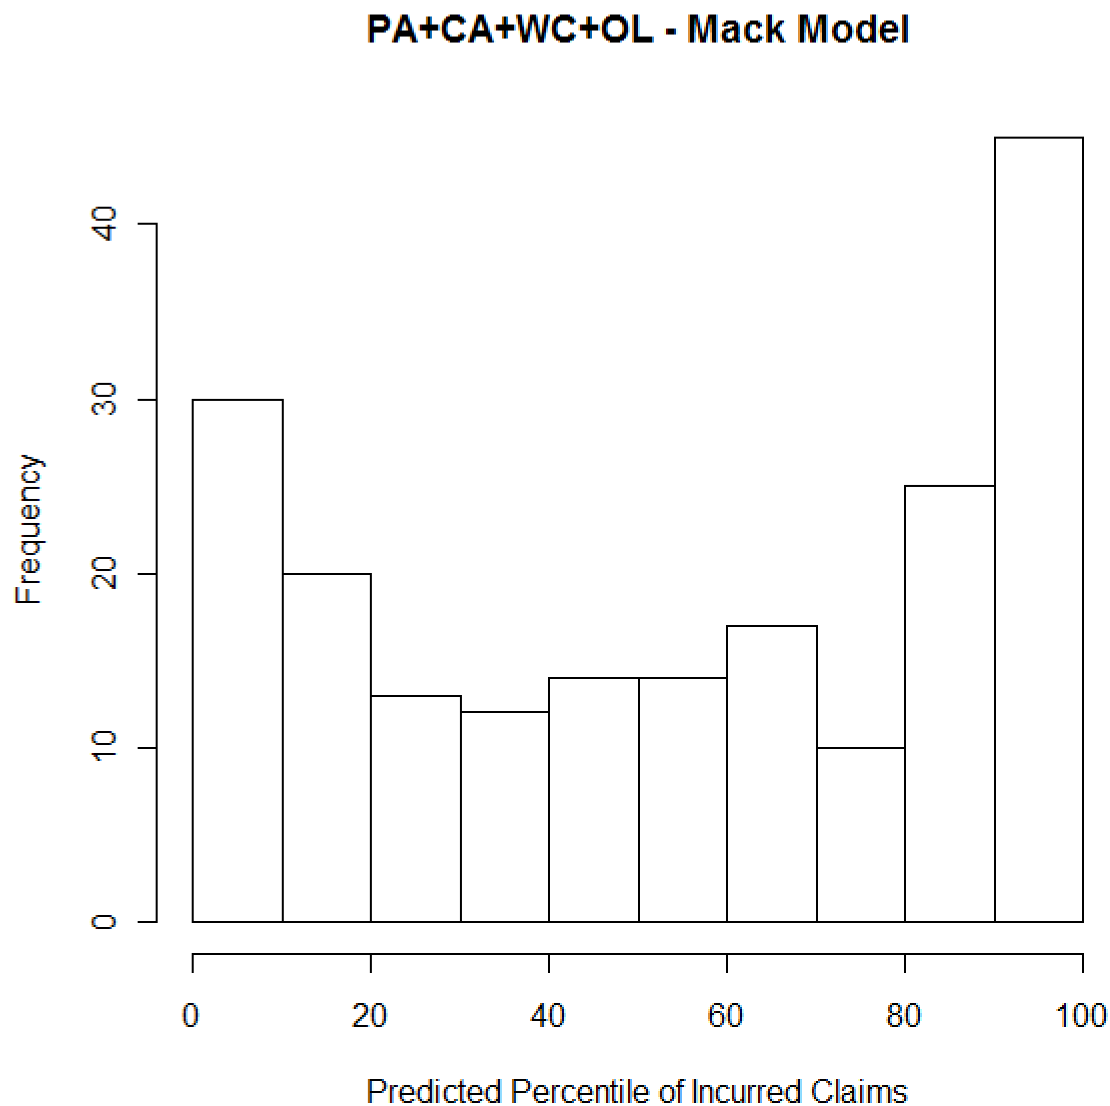

Uniformity Tests with 4 x 50 Insurers

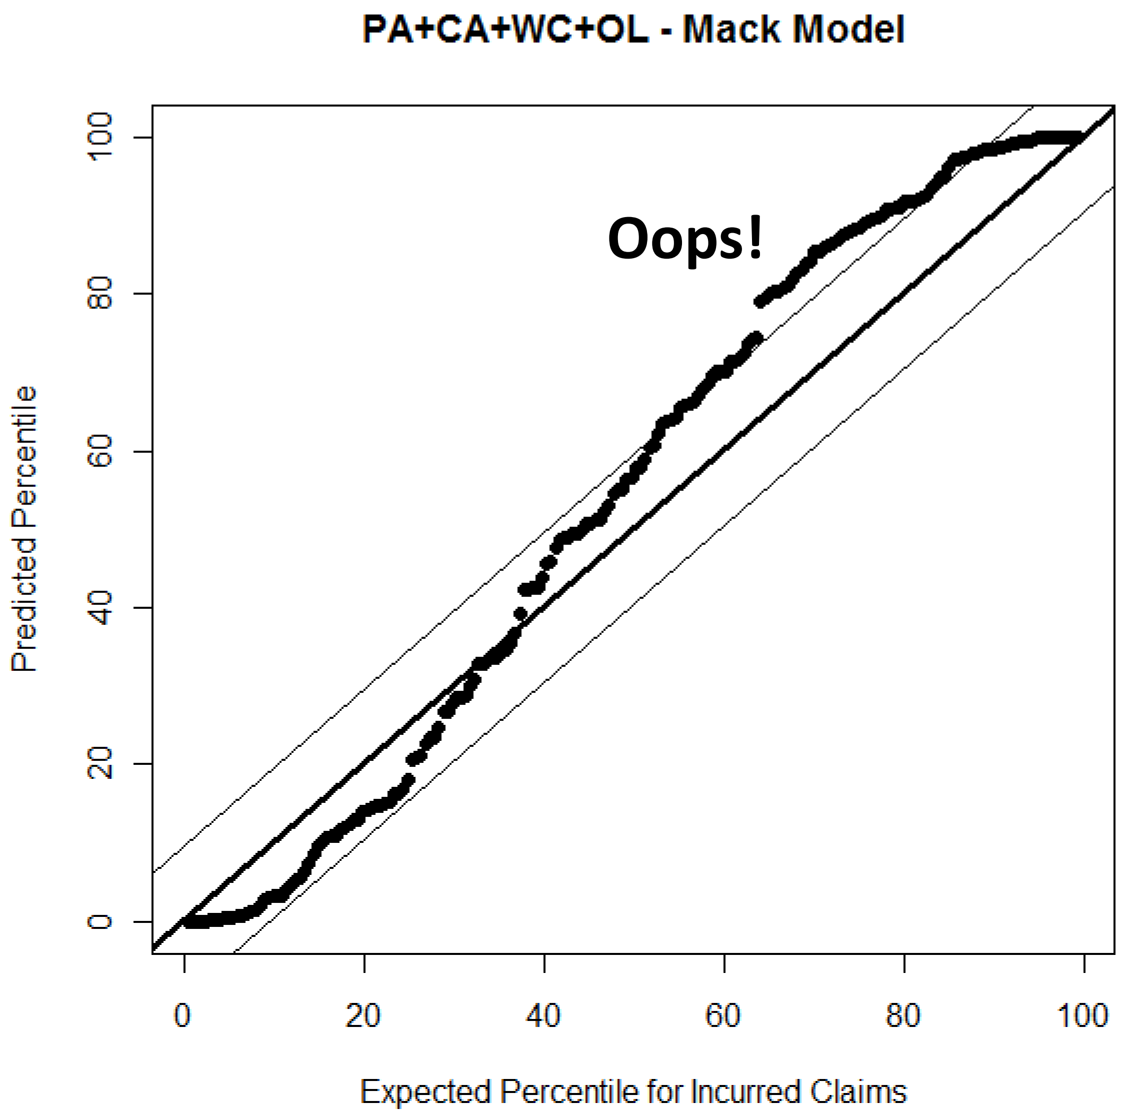

Conclusion - Mack Model

- Incurred data
  - Mack model underestimates the thickness of the tails
- A better model will have greater variability in the tails
  - Ways to increase variability
    - Bayesian Estimation
    - Correlations

The Case for Bayesian Estimation

- Mack gives a point estimation of variance
- Bayesian estimate of variance with posterior *P*:

$$Var\left[C_{w,d}\right] = E_P\left[Var\left[C_{w,d} \mid P\right]\right] + Var_P\left[E\left[C_{w,d} \mid P\right]\right]$$

- The last term: $Var_P\left[E\left[C_{w,d} \mid P\right]\right]$
  - Is equal to ZERO for a point estimate
  - Increases as we increase the number of parameters
  - Decreases as we increase the number of data points
- Loss Reserve Models – Many parameters, few data points
- See the “Thinking Outside the Triangle” paper.

http://www.actuaries.org/ASTIN/Colloquia/Orlando/Papers/Meyers.pdf

Calculating the Bayesian Posterior Using MCMC Sampling

- Generates samples from the posterior distribution
- MCMC sampling can be used to quantify the variability for many Bayesian models with few restrictions on the model specification.
  - There are freely available software packages to help with this. I use JAGS.
- JAGS generates a list (or distribution) of parameters, R then takes over to simulate a distribution of quantities of interest, such as loss reserve outcomes.
- See my AR column

http://casact.org/newsletter/index.cfm?fa=viewart&id=6352

The Leveled Chain Ladder Model Version 2

- $C_{1,d} \sim \text{lognormal}(\alpha_w + \beta_d, \sigma_d)$
- $C_{w,d} \sim \text{lognormal}(\alpha_w + \beta_d + z \cdot (\log(C_{w-1,d}) - \alpha_{w-1} - \beta_d), \sigma_d)$

σ

- $z \sim U(-1,1)$
- α_w and β_d are uniformly distributed ($\beta_1 = 0$).
- $\sigma_d = \sum_{i=d}^{10} a_i \quad a_i \sim U(0,1)$

Guarantees σ_d decreases as d increases

- Estimate distribution of $\sum_{w=2}^{10} C_{w,10}$
- If we set z ≡ 0 we get LCL Version 1

Using MCMC Models

1. Get $\{\alpha_w\}$, $\{\beta_d\}$, $\{\sigma_d\}$ and $\{z\}$ from JAGS output
2. Randomly select $C_{w,10}$ from a lognormal distribution with
   - log mean = $\alpha_w + \beta_{10} + z \cdot (\log(C_{w-1,10}) - \alpha_{w-1} - \beta_{10})$,
   - log standard deviation = σ_{10}
3. Form the sum $\sum_{w=2}^{10} C_{w,10}$
4. Repeat the above 10,000 times.

LCL v2 on Illustrative Triangle

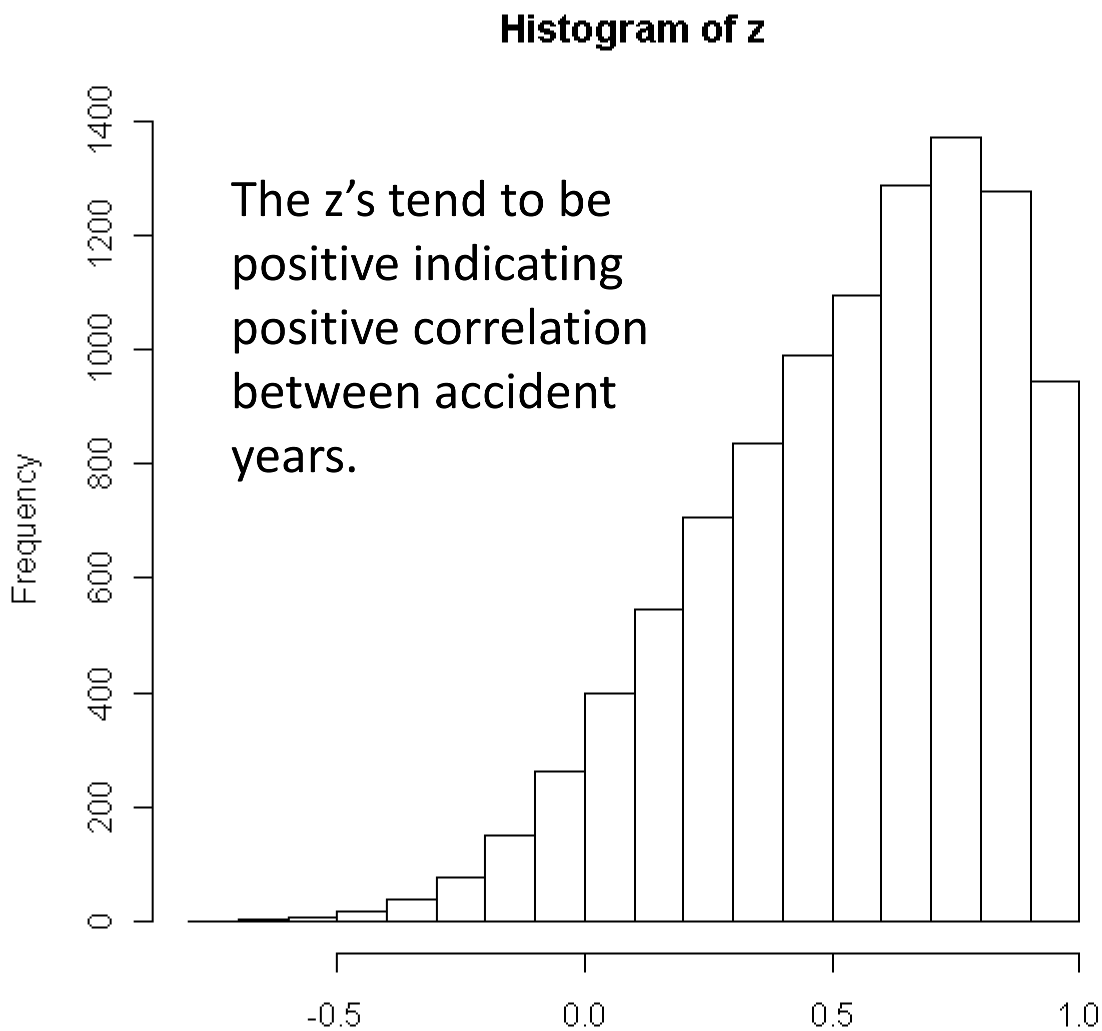

The z's tend to be positive indicating positive correlation between accident years.

Results for LCL v2 and Mack on the Illustrative Triangle

	Leveled Chain Ladder V2			Chain Ladder/Mack			
w	Estimate	Std. Error	CV	Estimate	Std. Error	CV	Actual
1	3,918	0	0.0219	3,917	0	0.0000	3,917
2	2,546	74	0.0291	2,538	0	0.0000	2,532
3	4,113	135	0.0328	4,167	3	0.0007	4,279
4	4,324	162	0.0375	4,367	37	0.0085	4,341
5	3,565	154	0.0432	3,597	34	0.0095	3,587
6	3,338	179	0.0536	3,236	40	0.0124	3,268
7	5,237	356	0.0680	5,358	146	0.0272	5,684
8	3,736	377	0.1009	3,765	225	0.0598	4,128
9	4,122	699	0.1696	4,013	412	0.1027	4,144
10	3,937	1,367	0.3472	3,955	878	0.2220	4,181
Total *w*=2-10	34,918	2,192	0.0628	34,997	1,057	0.0302	36,144

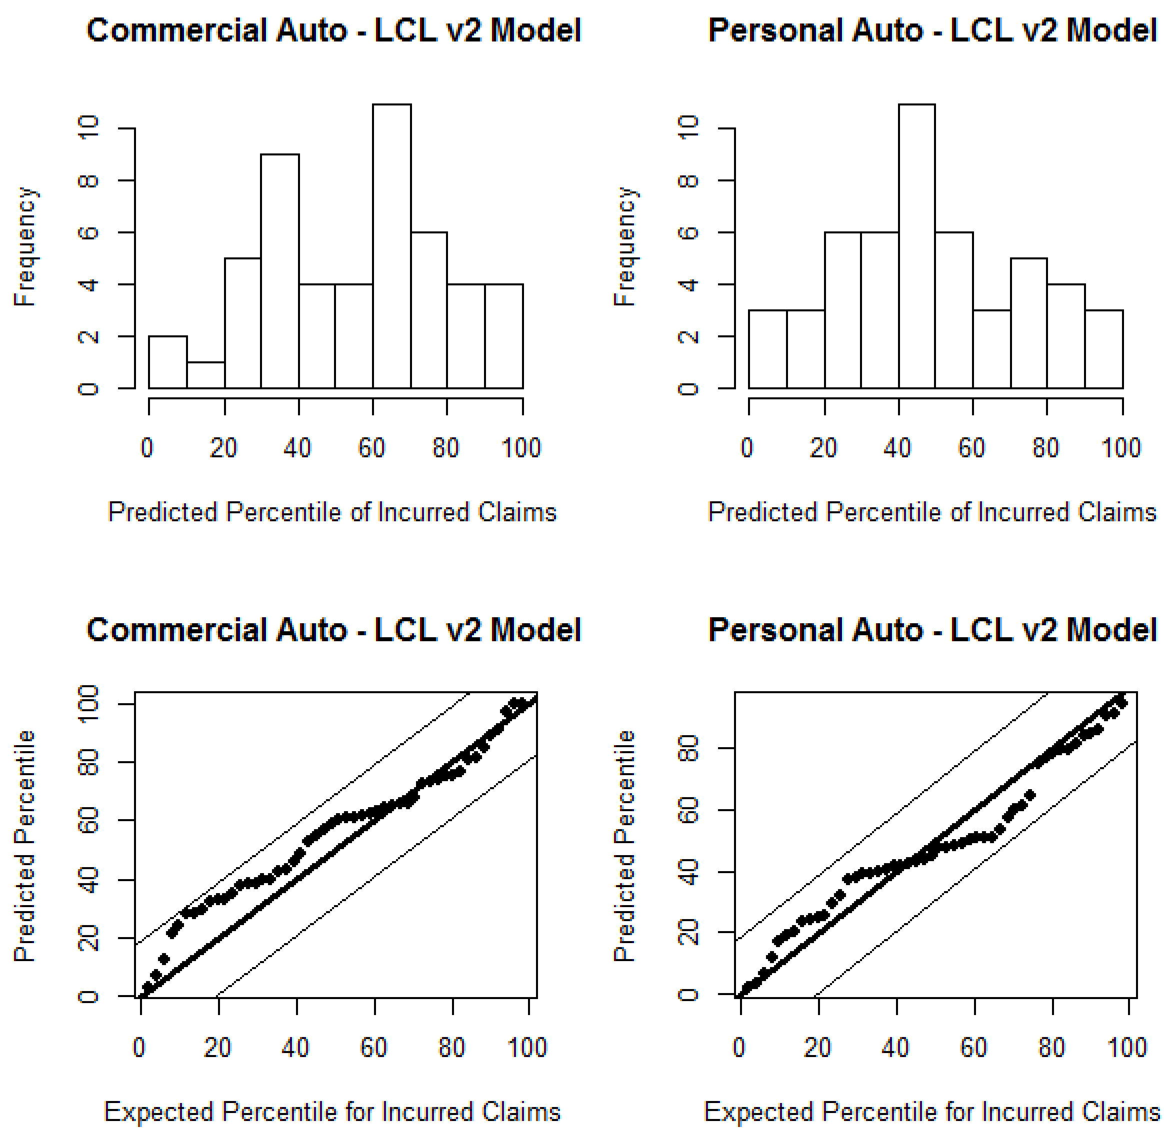

Commercial Auto - LCL v2 Model
Personal Auto - LCL v2 Model
Frequency
Predicted Percentile of Incurred Claims
Commercial Auto - LCL v2 Model
Personal Auto - LCL v2 Model
Predicted Percentile
Expected Percentile for Incurred Claims

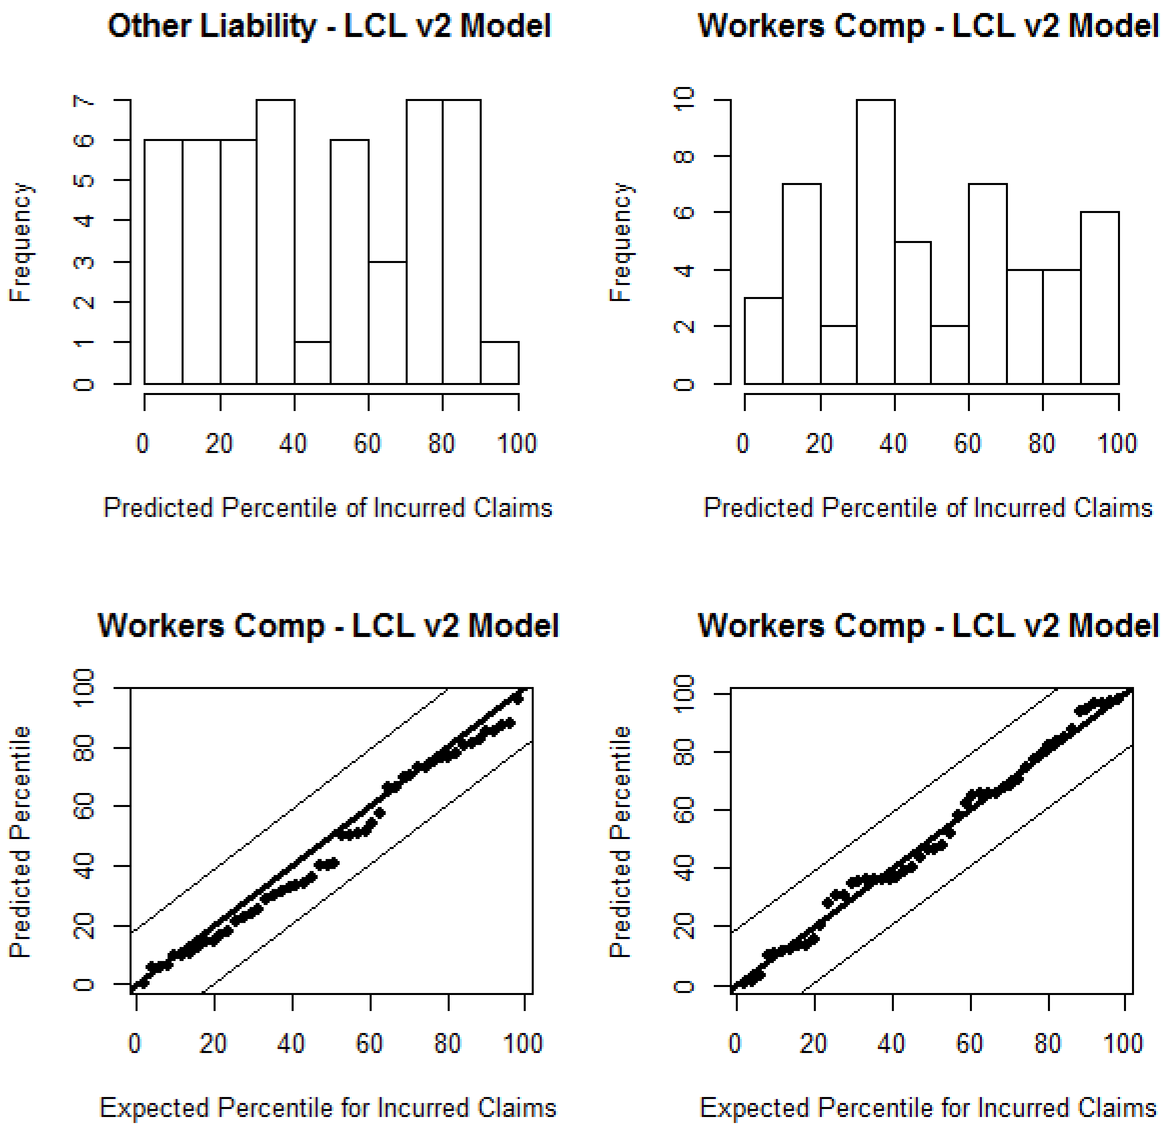
Other Liability - LCL v2 Model
Frequency
Predicted Percentile of Incurred Claims
Workers Comp - LCL v2 Model
Frequency
Predicted Percentile of Incurred Claims
Workers Comp - LCL v2 Model
Predicted Percentile
Expected Percentile for Incurred Claims
Workers Comp - LCL v2 Model
Predicted Percentile
Expected Percentile for Incurred Claims

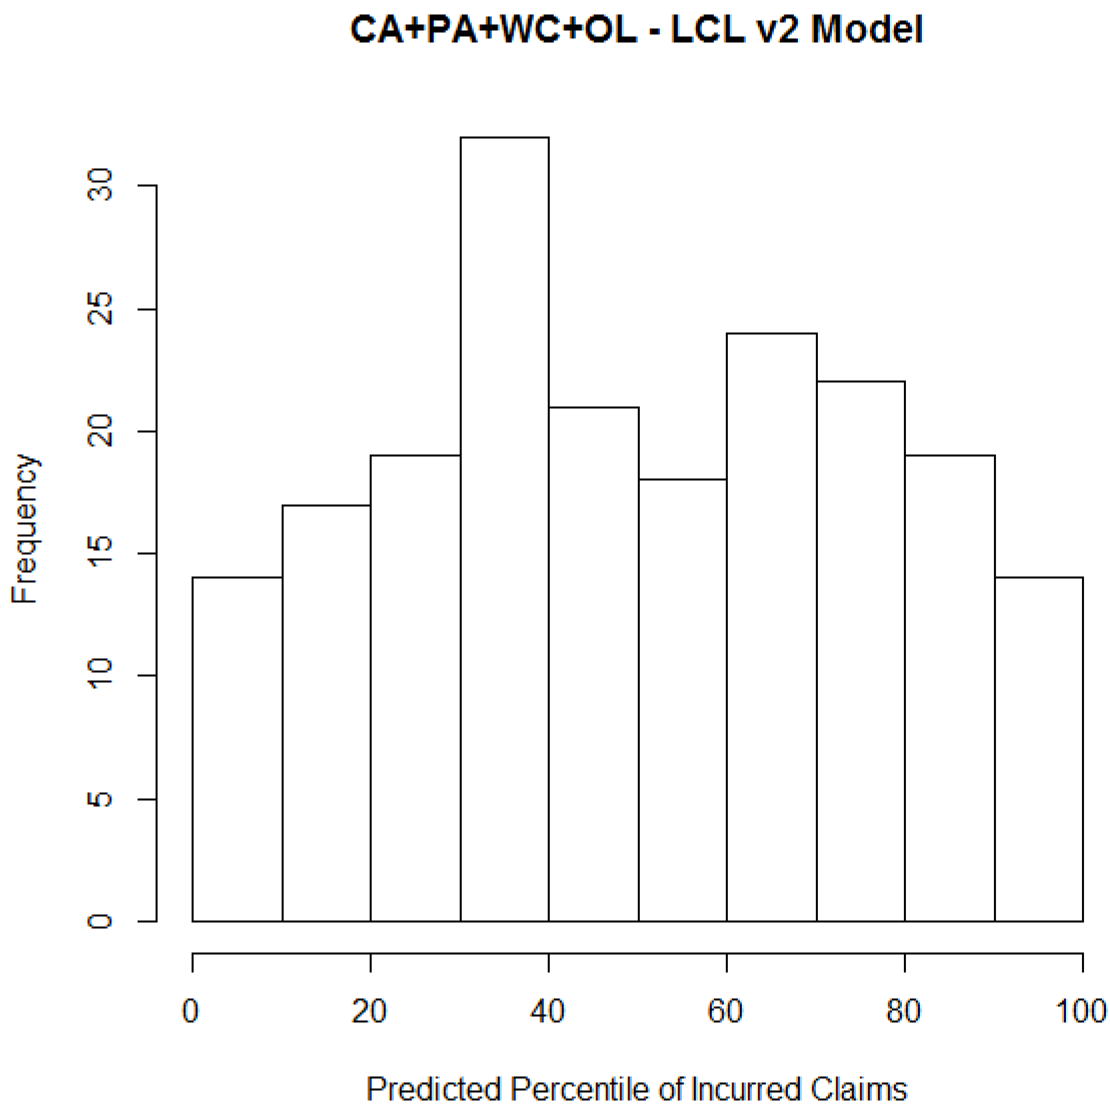
CA+PA+WC+OL - LCL v2 Model
Frequency
0
5
10
15
20
25
30
0
20
40
60
80
100
Predicted Percentile of Incurred Claims

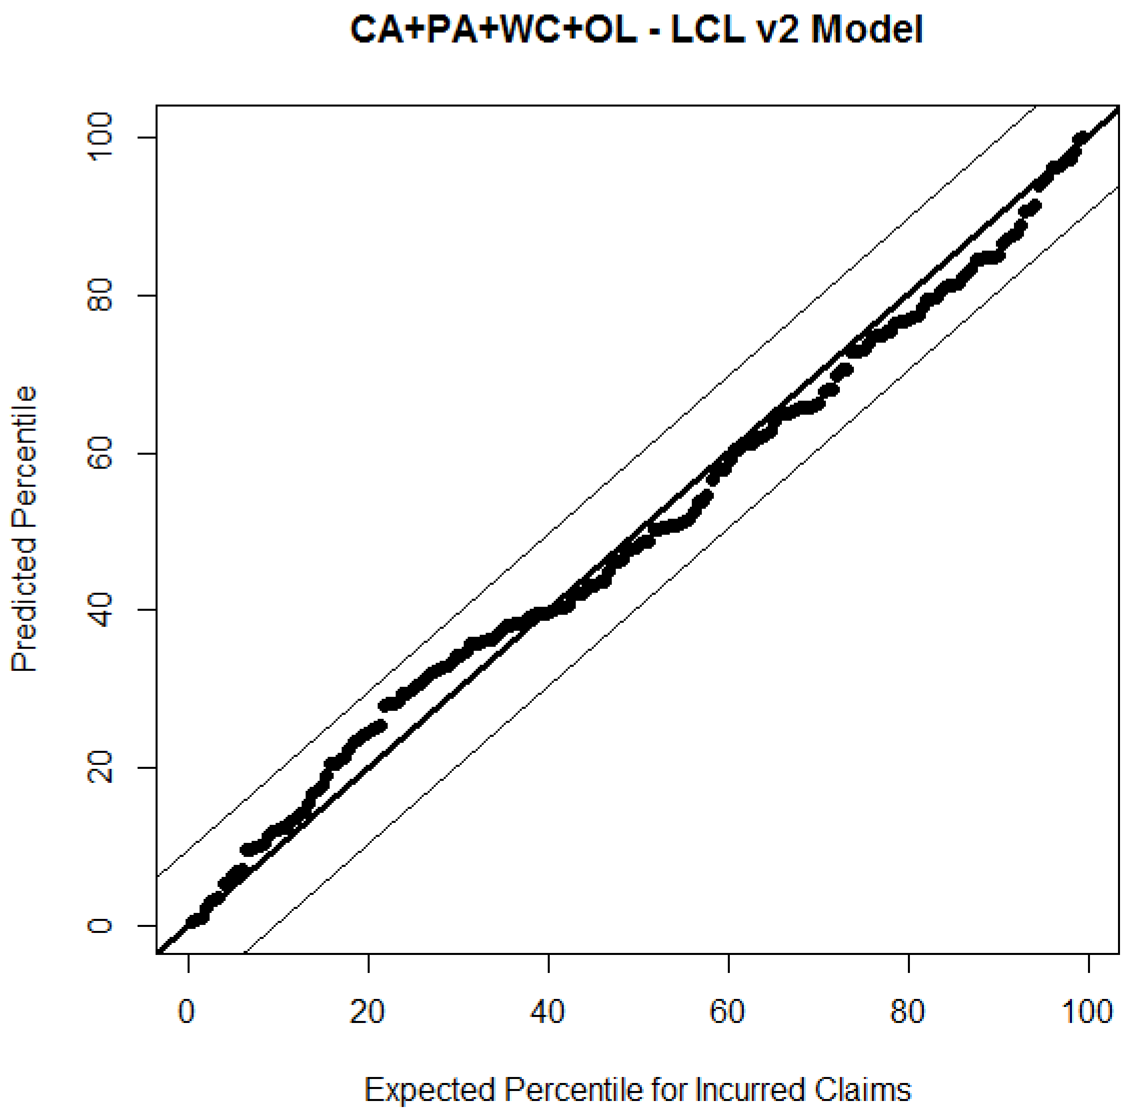

CA+PA+WC+OL - LCL v2 Model
Predicted Percentile
0
20
40
60
80
100
Expected Percentile for Incurred Claims

What Have We Accomplished?

- Found a model that validates on holdout data consisting of 200 separate triangles!
  - Bayesian
  - Correlations across development year
    - Consequence of random level parameters
  - Correlations across accident years
    - Autoregressive model linking prior accident years
  - Works with incurred claim data

Paid Claims Data

- Differences between Paid and Incurred Data in Schedule P
- Incurred = Paid + Case Reserves
- Case Reserves take into account information known to claims adjusters.
- Frequent decreases in Case Reserves resulting in negative incremental claims amounts.

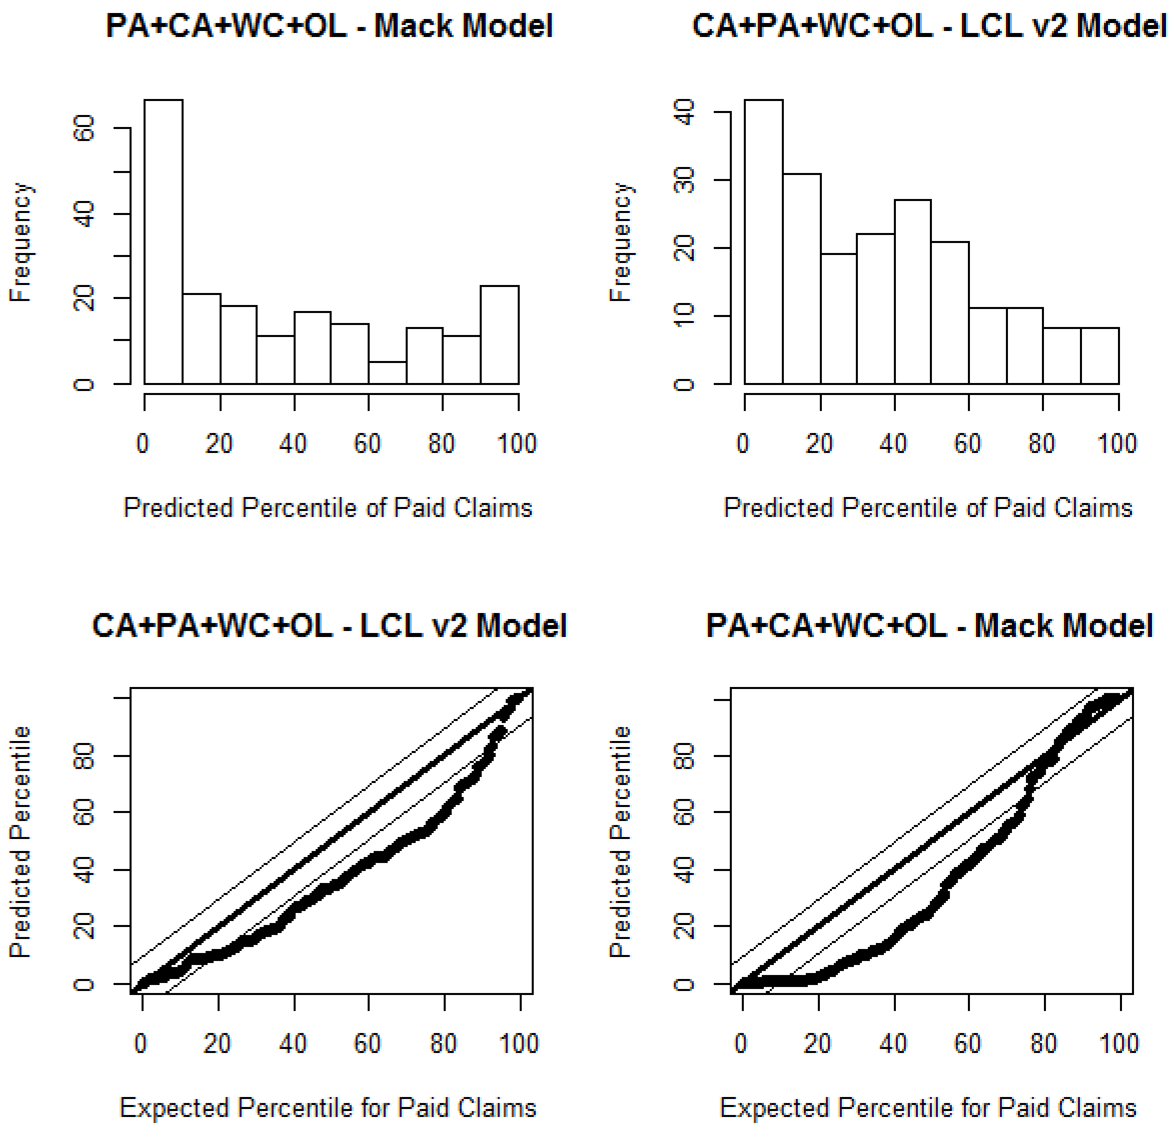
PA+CA+WC+OL - Mack Model
Frequency
0
20
40
60
0
20
40
60
80
100
Predicted Percentile of Paid Claims
CA+PA+WC+OL - LCL v2 Model
Frequency
0
10
20
30
40
0
20
40
60
80
100
Predicted Percentile of Paid Claims
CA+PA+WC+OL - LCL v2 Model
Predicted Percentile
0
20
40
60
80
0
20
40
60
80
100
Expected Percentile for Paid Claims
PA+CA+WC+OL - Mack Model
Predicted Percentile
0
20
40
60
80
0
20
40
60
80
100
Expected Percentile for Paid Claims

Meyers – Shi 2011

Tested Two Models on Incremental Paid Data for Commercial Auto

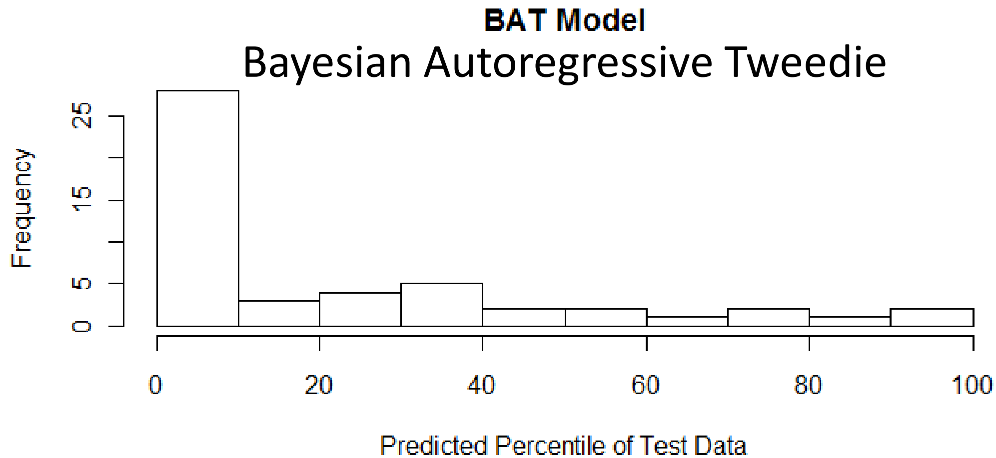

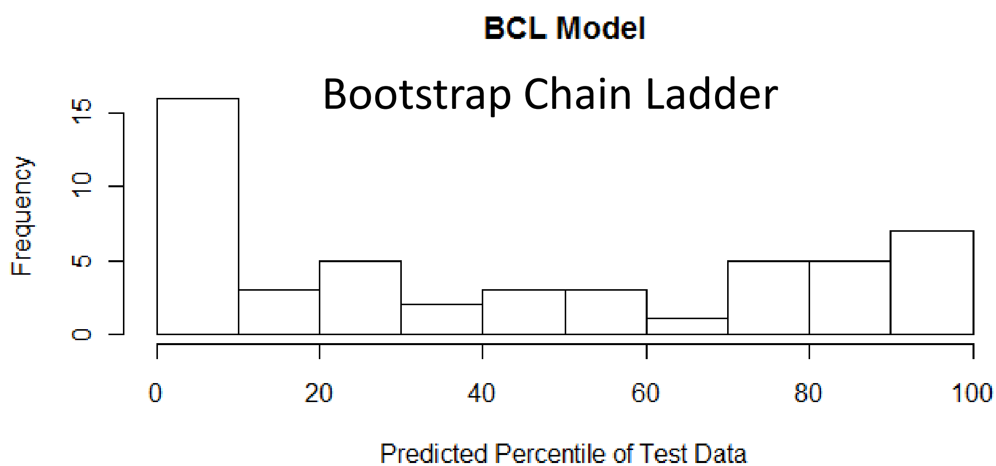

Suggested Conclusions Subject to Debate

- Incurred data contain information that is relevant for stochastic loss reserving that paid data do not.
- Mack model understates variability.
- Models (e.g. ODP/Bootstrap) that depend on incremental paid data will fail because they look at the wrong data.

“Aggressive” Retrospective Testing

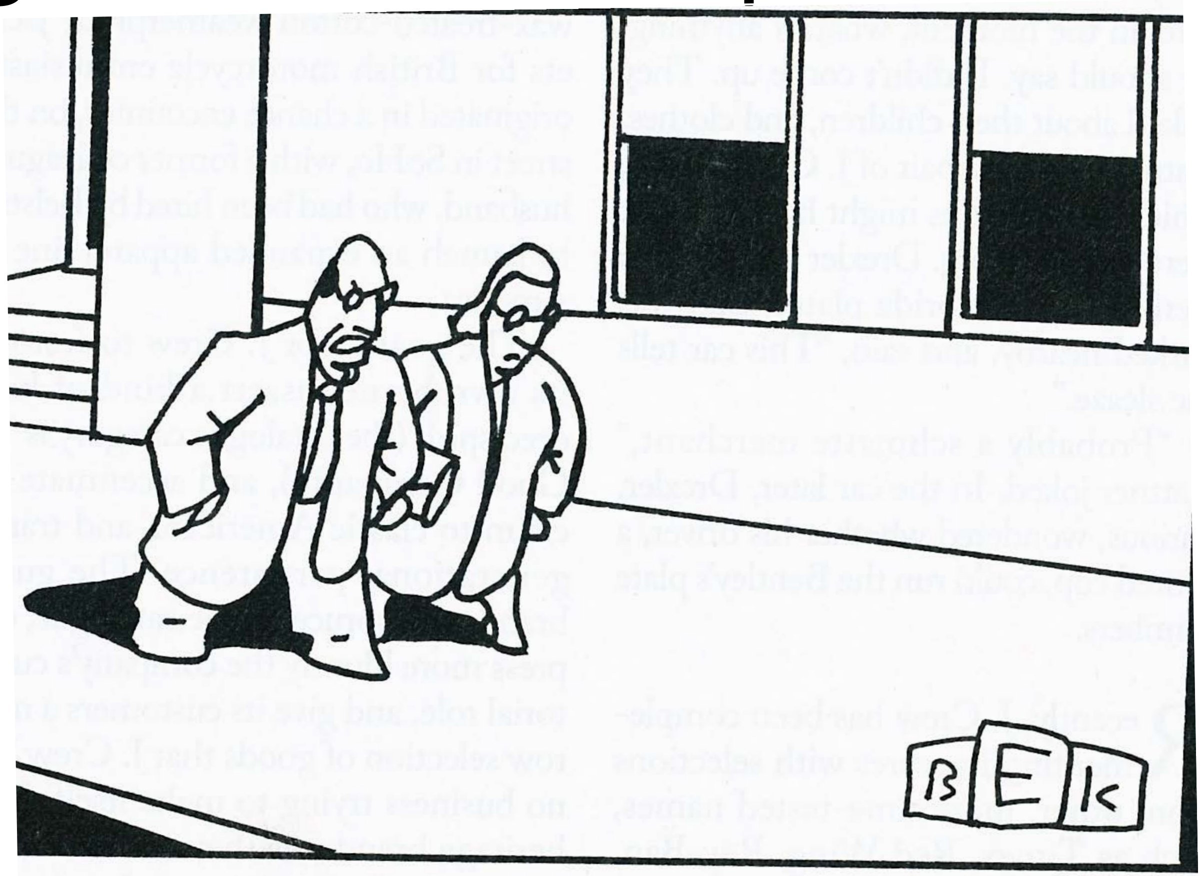

“If you want something done right, you have to live in the past.”

But will it work in the future??