



# Antitrust Notice

- **The Casualty Actuarial Society is committed to adhering strictly to the letter and spirit of the antitrust laws. Seminars conducted under the auspices of the CAS are designed solely to provide a forum for the expression of various points of view on topics described in the programs or agendas for such meetings.**
- **Under no circumstances shall CAS seminars be used as a means for competing companies or firms to reach any understanding – expressed or implied – that restricts competition or in any way impairs the ability of members to exercise independent business judgment regarding matters affecting competition.**
- **It is the responsibility of all seminar participants to be aware of antitrust regulations, to prevent any written or verbal discussions that appear to violate these laws, and to adhere in every respect to the CAS antitrust compliance policy.**



# Total Credibility

Frank Schmid, Dr. habil.

CLRS

September 5-7, 2012

Denver, CO

# Overview

- The Research Problem
- Total Credibility
- The Data
- The Model
- Findings
- Strengths and Weaknesses
- Conclusion
- References



# The Research Problem

- Compute credibility-adjusted link ratios(\*) for pool reserving triangles
  - For research purposes, annual data is used
- There are two credibility aspects to pool reserving
  - First, the degree of variability of the link ratios may vary across states, possibly (but not necessarily entirely) due to variation in size across states of the Assigned Risk market
  - Second, the number of link ratios available varies greatly across states (from fully populated diagonals to no observations at all)

(\*) The link ratio from time  $t$  to time  $t+1$  is defined as the ratio of cumulative payments up to (and inclusive of) time  $t+1$  to the cumulative payments up to (and inclusive of) time  $t$



# Total Credibility

- Total Credibility is about credibility-adjusting the data-generating process in a comprehensive manner (as opposed to credibility-adjusting the outcome)
- The workhorse of Total Credibility is multilevel (hierarchical) modeling
- Although there are frequentist approaches to multilevel modeling, Bayesian statistics is particularly well-suited to building comprehensive (yet parsimonious) credibility-based models



# Multilevel Modeling

- Guszczka (2008), Zhang, Dukic, and Guszczka (2010), and Meyers (2011) discuss the use of multilevel modeling in reserving
  - Guszczka (2008) fits growth curves to cumulative losses using frequentist approaches with random effects in parameters
  - Zhang, Dukic, and Guszczka (2010) take a Bayesian approach to fitting growth curves to cumulative losses—the authors estimate multiple triangles simultaneously, thus accounting for correlation across loss triangles within an industry
  - Meyers, in reference to Guszczka (2008), fits an autoregressive process to loss ratios in a Bayesian model—yet, the model is not of multilevel nature
- Neither the models discussed by Guszczka (2008) nor the one suggested by Meyers (2011) serve our purpose—whereas the former cannot handle missing values, the latter is not multilevel
  - The approach closest to the Total Credibility model presented below is the Bayesian framework developed by Zhang, Dukic, and Guszczka (2010)

# Credibility and Multilevel Modeling

- In multilevel modeling, credibility is implemented by means of partial pooling (or, synonymously, shrinkage)
  - Let alpha be one of the parameters that govern the growth process of the cumulative loss in pool reserving
  - In partial pooling, the parameter alpha is allowed to vary across states, but all state-specific alphas must be drawn from the same, common distribution—the parameters that define this common distribution are called hyperparameters
    - Shrinkage is an adjustment toward the expected value of the alphas (that is, the expected value of the common distribution)
    - For any given state, the fewer observations there are and the noisier these observations are, the more shrinkage is applied to the alpha of that state



# Partial Pooling and Bühlmann Credibility

- In the normal linear model, partial pooling is equivalent to Bühlmann Credibility
- Following Gelman and Hill (2007), let  $y$  be a normally distributed variable:

$$y_i \sim \mathbf{N}(\alpha_{j[i]}, \sigma_y^2) \quad (j \text{ indicates the state; } i \text{ indicates the observation})$$

- Multilevel modeling assumes that the parameter  $\alpha_j$  that governs the process in state  $j$  is a draw from a distribution that is common to all states:

$$\alpha_j \sim \mathbf{N}(\mu_\alpha, \sigma_\alpha^2)$$

- It can be shown that the multilevel estimator for  $\alpha_j$  reads:

$$\hat{\alpha}_j = \omega_j \cdot \mu_\alpha + (1 - \omega_j) \bar{y}_j, \quad \omega_j = 1 - \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \frac{\sigma_y^2}{n_j}}$$

where  $\bar{y}_j$  is the sample mean for state  $j$  based on  $n_j$  observations



# The Data

- The data set consists of (paid plus case and, alternatively, paid) link ratios of the latest five diagonals (2005-2009) of the pool triangles of 45 states:
  - AK, AL, AR, AZ, CA, CO, CT, DC, DE, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MA, MD, ME, MI, MO, MS, MT, NC, NE, NH, NJ, NM, NV, OK, OR, PA, RI, SC, SD, TN, UT, VT, VA, WA, WV, WY
- For many states, the data are sparse—below the properties of the 2009 paid plus case diagonals:
  - A total of 18 states (or 40 percent) have a complete history of link ratios (AK, AL, AR, CT, DC, DE, GA, IA, IL, KS, MI, NC, NH, NJ, NM, SD, VA, VT)
  - There are four states for which all available link ratios are unity (between 4 and 13 unity link ratios per diagonal; CO, OK, WA, and WY)
  - There are two states with one observation (CA and WV)
  - There are two states with no data (MT and UT)

# The Growth Curve

- Log link ratios represent log rates of growth of cumulative losses, thus resembling a biological growth process
- A growth curve is fit to the natural logarithms of (paid plus case and, alternatively, paid) link ratios:

$$y_{i,j} = \beta_i \cdot \gamma_i^{q_i \cdot \log(j) + (1-q_i) \cdot (j-1)}, \quad j = 1, \dots, N, \quad \beta_i > 0, \quad 0 < \gamma_i < 1, \quad 0 \leq q_i \leq 1$$

where  $i$  indicates the state and  $j$  indicates the maturity (year);  $y_{i,j}$  is the log link ratio and  $N$  stands for the number of observations (per state) in the data set

- The parameter  $\beta_i$  delivers an estimate of the first-to-second link ratio
- The parameter  $q_i$  is a weighting factor between log-linear and linear influences
- The growth curve has an asymptote at zero
- The growth curve displayed above has been introduced to Bayesian modeling by Gelfand and Carlin (1991)

# Challenges Arising from the Data

- Missing link ratio observations are treated as parameters that need estimating
- The likelihood of the growth curve can handle observations of link ratios that are equal to or less than unity, even though the growth curve for the log link ratio has an asymptote at zero
- The biggest modeling challenge is the abundance of unity link ratios—for instance, for four states, all link ratios are equal to 1
  - Due to the lack of variation in the dependent variable, Gibbs sampling breaks down
  - To overcome this problem, the (entire) data set is jittered (by means of adding a small random error to the link ratios)—this way, three data sets are generated, and the growth curve is fit simultaneously to these three data sets
  - The noise introduced by jittering the data set is miniscule: there is a 67 percent probability that any given link ratio changes by less than one-hundredth of one percent



# The Likelihood

- The likelihood consists of a double exponential (or, equivalently, Laplace) distribution
  - The double exponential distribution is heavy-tailed and minimizes the sum of absolute errors (as opposed to the sum of squared errors), which makes this distribution robust to outliers
    - Minimizing the sum of absolute errors implies estimating the conditional median (as opposed to the traditional approach of modeling the conditional mean)
    - The double exponential likelihood does not account for skewness
  - The precision (which is the reciprocal of the variance) of the double exponential is credibility-adjusted
    - The variance of the log link ratios is allowed to vary across states
    - Credibility-adjusting these variances is critical for states that have no observations
  - The likelihood allows for heteroskedasticity
    - The variances of the first and second log link ratios are allowed to differ from each other and from the variance that applies to all subsequent link ratios
  - The three parameters of the growth curve are credibility-adjusted



# Model Validation


- The model is validated using a one-year holdout period
  - The model is fit to the diagonals of calendar years 2004 through 2008
  - Based on the estimated model parameters, link ratios for the calendar year 2009 are simulated
  - By comparing the simulated link ratios to the 2009 observed values, the mean absolute forecast error is calculated
  - The process of model validation is repeated using the diagonals of calendar years 2003 through 2007—the forecast errors are calculated based on the observed 2008 diagonal



# Distribution of the Raw Data

## Paid + Case Boxplots, All 45 States


2009 Diagonal



# Distribution of the Raw Data

## Paid + Case Boxplots, All 45 States


2008 Diagonal



# Distribution of the Raw Data

## Paid + Case Boxplots, All 45 States

2007 Diagonal






# Distribution of the Raw Data

## Paid + Case Boxplots, All 45 States


2006 Diagonal



# Distribution of the Raw Data

## Paid + Case Boxplots, All 45 States


2005 Diagonal



# Distribution of the Raw Data

## Paid Boxplots, All 45 States


2009 Diagonal



# Distribution of the Raw Data

## Paid Boxplots, All 45 States


2008 Diagonal



# Distribution of the Raw Data

## Paid Boxplots, All 45 States


2007 Diagonal



# Distribution of the Raw Data

## Paid Boxplots, All 45 States


2006 Diagonal



# Distribution of the Raw Data

## Paid Boxplots, All 45 States

2005 Diagonal



# Charts and Diagnostics for Selected States


- New Jersey
- Massachusetts
- Michigan
- New Mexico
- Tennessee
- West Virginia






# New Jersey

## Paid + Case




# Massachusetts

## Paid + Case




# Michigan Paid + Case




# New Mexico

## Paid + Case




# Tennessee

## Paid + Case



# West Virginia


## Paid + Case



# In-Sample Diagnostics

## Paid + Case


2009 Diagonal



# In-Sample Diagnostics

## Paid + Case

2008 Diagonal






# In-Sample Diagnostics

## Paid + Case


2007 Diagonal



# In-Sample Diagnostics

## Paid + Case


2006 Diagonal



# In-Sample Diagnostics

## Paid + Case


2005 Diagonal



# Forecast Diagnostics

## Paid + Case


2009 Diagonal



# Forecast Diagnostics

## Paid + Case

2008 Diagonal



# Strength and Weaknesses

- Total Credibility is an innovative concept—only a handful of papers are available on multilevel modeling
- The model is not built to replicate the skewness of logarithmic paid link ratios—the model is suitable for paid plus case link ratios only
- The results of the model are replicable, as no human judgment is involved
  - The model is transparent—the computer code of the core model is brief and simple in structure (see Technical Appendix)
- In model engineering terms, credibility is built-in, not bolted-on
- The model is robust to outliers



# Conclusion

- The model is part of a family of multilevel reserving models that have recently been discussed in actuarial literature
- The model is capable of estimating not only link ratios but also tail factors
- The model offers credible intervals(\*) for the estimated link ratios and tail factors
- The model can be extended to process quarterly (instead of annual) data

(\*) Credible intervals in Bayesian statistics are the equivalent to confidence intervals in frequentist statistics. Yet, there are important conceptual differences; see Bradley P. Carlin, and Thomas A. Louis, pp. 6-7



# References

- Carlin, Bradley P., and Thomas A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis*, 2<sup>nd</sup> ed., Boca Raton (FL): Chapman & Hall/CRC, 2000
- Gelfand, Alan E. and Bradley P. Carlin, "An Iterative Monte Carlo Method for Nonconjugate Bayesian Analysis," *Statistics and Computing* Vol. 1, pp. 119-128 , 1991
- Gelman, Andrew, and Jennifer, Hill, *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Cambridge (MA): Cambridge University Press, 2007
- Guszcza, James, "Hierarchical Growth Curve Models for Loss Reserving," *Casualty Actuarial Society Forum*, Fall 2008, pp. 146-173,  
<http://www.casact.org/pubs/forum/08fforum/7Guszcza.pdf>
- Meyers, Glenn, "Predicting Loss Ratios with a Hierarchical Bayesian Model," *Actuarial Review*, Vol. 38, No. 1, February 2011, pp. 8-9,  
<http://www.casact.org/newsletter/index.cfm?fa=viewart&id=6109>
- Zhang, Yanwei, Vanja Dukic, and James Guszcza, "A Bayesian Nonlinear Model for Forecasting Insurance Loss Payments," August 2010,  
<http://www.actuaryzhang.com/publication/bayesianNonlinear.pdf>



# Technical Appendix

## Markov Chain Monte Carlo Simulation Settings

- The model was estimated by means of Markov Chain Monte Carlo simulation (MCMC)—three Markov chains were run
- The model was estimated using R 2.13.1 64 bit (2011-07-08 ) and JAGS 3.0.0 (2011-07-21)
- The random number generators of both R and JAGS were seeded, thus making the results replicable
- An adaption phase of 10,000 draws and burnin of 10,000 draws were followed by a sample of 500,000 draws (of which every 100th draw was selected)

# Technical Appendix

## JAGS Code (Core Model)

```

model
{
  #shrinkage
  beta.mu ~ dnorm(0,1,E-2)T(0,)
  beta.tau <- pow(beta.sigma,-2)
  beta.sigma ~ dunif(0,2)

  gamma.mu ~ dbeta(1,1)
  gamma.sigma ~ dunif(0,1)
  gamma.tau <- pow(gamma.sigma,-2)
  gamma.alpha <- gamma.mu * gamma.tau
  gamma.beta <- (1-gamma.mu) * gamma.tau

  q.mu ~ dbeta(1,1)
  q.sigma ~ dunif(0,1)
  q.tau <- pow(q.sigma,-2)

  for(m in 1:3){ #different variances for first two development years
    tau.alpha[m] ~ dexp(1.0)
    tau.beta[m] ~ dgamma(0.1,0.1)
  }

  #likelihood
  for(i in 1:L){ #rows (states)
    beta[i] ~ dnorm(beta.mu,beta.tau)T(0,)
    gamma[i] ~ dbeta(gamma.alpha,gamma.beta)
    q[i] ~ dnorm(q.mu,q.tau)T(0,1) #using normal instead of beta eases convergence

    for(m in 1:3){
      tau[i,m] ~ dgamma(tau.alpha[m],tau.beta[m])
      sigma[i,m] <- sqrt(2)/tau[i,m] #double exponential errors
    }


    for(j in 1:T){ #columns (development years)
      y.pred[i,j] ~ ddexp(mu[i,j],tau[i,tau.index[j]]) #double-indexing for tau
      cdf[i,j] <- sum(mu[i,j:T])
      mu[i,j] <- beta[i]*pow(gamma[i],q[i]*log(j)+(1-q[i])*(j-1))
    }

    for(j in 1:N){ #columns (development years)
      for(k in 1:K){ #jitter
        y.2009[i,j,k] ~ ddexp(mu[i,j],tau[i,tau.index[j]]) #double-indexing for tau
        y.2008[i,j,k] ~ ddexp(mu[i,j],tau[i,tau.index[j]])
        y.2007[i,j,k] ~ ddexp(mu[i,j],tau[i,tau.index[j]])
        y.2006[i,j,k] ~ ddexp(mu[i,j],tau[i,tau.index[j]])
        y.2005[i,j,k] ~ ddexp(mu[i,j],tau[i,tau.index[j]])
      }
    }
  }
}

```

# Technical Appendix

## Boxplots



- The box comprises 50 percent of the data—its upper and lower bounds indicate the innerquartile range (IQR)
- The horizontal bar inside the box represents the median
- The whiskers at the end of the stems indicate the smallest (bottom) and largest (top) that is within 1.5 IQR from the box limits
- Observations beyond the whiskers are plotted as dots and constitute outliers as judged by the normal distribution