An introduction to non-life stochastic reserving using Markov Chain Monte Carlo (MCMC) models

Jean Rea
Director, Actuarial Services
KPMG Ireland

The Actuary as a Data Scientist – What, How and Why?
5 November 2018
Staple Inn Hall, London.

Reserve uncertainty – so what?

Reserve uncertainty

Solvency II

- Article 101 (3) - SCR shall correspond to the Value-at-Risk of the basic own funds of an insurance or reinsurance undertaking subject to a confidence level of 99.5 % over a one-year period.

IFRS 17

- The **risk adjustment** conveys information to users of financial statements about the amount the entity charged for bearing the uncertainty over the amount and timing of cash flows arising from non-financial risk.

Reserve uncertainty

IFRS 17 (continued)

- Principles based – no prescribed calculation technique.
- Insurers may prefer to adopt a methodology that allows management to use its own internal models to more accurately capture the specific and complex risks faced.
- Will form part of disclosures.
- Popular approaches may include:
  - Confidence level / VaR
  - Conditional Tail Expectation
  - Cost of Capital approach

A brief history of property casualty reserving

How reserves are established has changed little over the last century

1800

New York Insurance Law requires sufficient general reserves to pay all claims

1934

Tarbell paper in CAS Proceedings outlines a method of calculating one year runoff of pure IBNR, to add to case basis reserves

1972

Remarking how little has changed in reserving since 1934, Bornhuetter and Ferguson lay out methods of loss development, and the BF method that still underpins reserving techniques today

These are designed around batch computer printout reports, and green paper spreadsheets

1980

Electronic spreadsheets like Visicalc, and later Excel, are adapted to calculate the 1972 methodologies, replacing the paper greensheets

1990

Automated software tools are developed, such as ReservePro and ResQ that incorporate the 1972 methods

Advances like statistical ranges are introduced

2000

Modern version of SAS, R introduced – computing power increases exponentially

2018

Reserve modernization using detailed data, new tools, and computing power becomes practical

CAS introduced CSPA credential in 2016, showing the importance of Predictive Analytics in the profession

Technology

The emergence of new technology, coupled with enhanced computing power, has the potential to radically disrupt this historic approach.

Data preparation	R	HIVE	hadoop
Cognitive – machine learning	sas THE POWER TO KNOW.	R	python
Visualization	+ableau	Qlik	
Robotic process automation	AUTOMATION ANYWHERE Go be great.	blueprism	

Computing power has increased significantly over time

We have seen a **1 trillion-fold increase** in computer processing capabilities over the past **60 years**[1]

Today's smartphone has more computing power than the Apollo 11 Guidance Computer

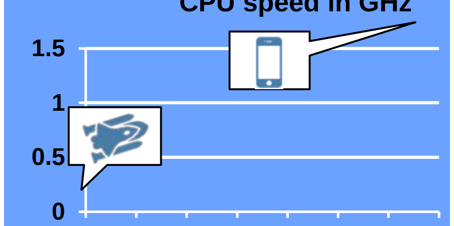

Source: [1]Experts Exchange, "Processing Power Compared"

Source: [2]Frost & Sullivan, "Addressing Mobile Cybersecurity"

Common approaches

Mack Method

Mack – my reference material

MEASURING THE VARIABILITY

OF CHAIN LADDER RESERVE ESTIMATES

Thomas Mack, Munich Re

Abstract:

The variability of chain ladder reserve estimates is quantified without assuming any specific claims amount distribution function. This is done by establishing a formula for the so-called standard error which is an estimate for the standard deviation of the outstanding claims reserve. The information necessary for this purpose is extracted only from the usual chain ladder formulae. With the standard error as decisive tool it is shown how a confidence interval for the outstanding claims reserve and for the ultimate claims amount can be constructed. Moreover, the analysis of the information extracted and of its implications shows when it is appropriate to apply the chain ladder method and when not.

Submitted to the 1993 CAS Prize Paper Competition on 'Variability of Loss Reserves'

Presented at the May, 1993 meeting of the Casualty Actuarial Society.

182

A Practitioner's Introduction to Stochastic Reserving

Alessandro Carrato MSc FIA IOA, Gráinne McGuire PhD FIAA, Robert Scarth PhD

2016-02-29

Mack – key assumptions

Mack, T (1993), *Distribution-free calculation of the standard error of chain-ladder reserve estimates*. ASTIN Bulletin, 22, 93-109

- Mack's model takes as input a triangle of cumulative claims. This could be a paid claims triangle or an incurred claims triangle.
- Cij is the cumulative claims in origin year i and development year j.

Key assumptions

1. For each $j = 1, \ldots, n-1$ there are development factors f_j such that $E[C_{i,j+1}|C_{i1}, \ldots, C_{ij}] = C_{ij} f_j$;

i.e. the expected value is proportional to the previous cumulative

2. For each $j = 1, \ldots, n-1$ there are parameters σ_j^2 such that $Var[C_{i,j+1}|C_{i,1}, \ldots, C_{ij}] = C_{ij}\sigma_j^2$.

i.e. the variance proportional to previous cumulative

3. Origin periods are independent.

Model validation

- Residual scatterplots
  - Should show an even scattering of residuals in a cloud centred around zero with no structure or discernible pattern.
  - Should be examined in all three dimensions of the triangle i.e. by origin, development and calendar (or payment) periods.
  - Include mean of residuals for each period/dimension – should have a mean of zero
  - Include standard deviation for each period/dimension – should have a constant standard deviation
- Residuals vs fitted values or log fitted values
- Heat maps

Mack specific validation

- Mack 1994 explains reasons why independence assumption could be violated in practice:
  - “The main reason why this independence can be violated in practice is the fact that we have certain calendar year effects such as major changes in claims handling or in case reserving or external influences such as substantial changes in court decisions or inflation”.
- Appendix H of Mack 1994 explains a procedure to test for calendar year influences
- Mack sigma graphs

R implementation code

```
library("ChainLadder")
Mack.example=MackChainLadder(MW2008,est.sigma = "Mack")
plot(Mack.example)
```

R implementation results

```
MackChainLadder(Triangle = MW2008, est.sigma = "Mack")

     Latest Dev.To.Date  Ultimate      IBNR Mack.S.E CV(IBNR)
1 3,678,633       1.000 3,678,633         0        0      NaN
2 3,902,425       0.999 3,906,803     4,378      566   0.1293
3 3,898,825       0.998 3,908,172     9,347    1,564   0.1673
4 3,548,422       0.992 3,576,814    28,392    4,157   0.1464
5 3,585,812       0.986 3,637,256    51,444   10,536   0.2048
6 3,641,036       0.970 3,752,847   111,811   30,319   0.2712
7 3,428,335       0.948 3,615,419   187,084   35,967   0.1923
8 3,158,581       0.885 3,570,445   411,864   45,090   0.1095
9 2,144,738       0.599 3,578,243 1,433,505   69,552   0.0485

                   Totals
Latest:    30,986,807.00
Dev:                0.93
Ultimate:  33,224,633.11
IBNR:       2,237,826.11
Mack.S.E      108,401.39
CV(IBNR):           0.05
```

R implementation results

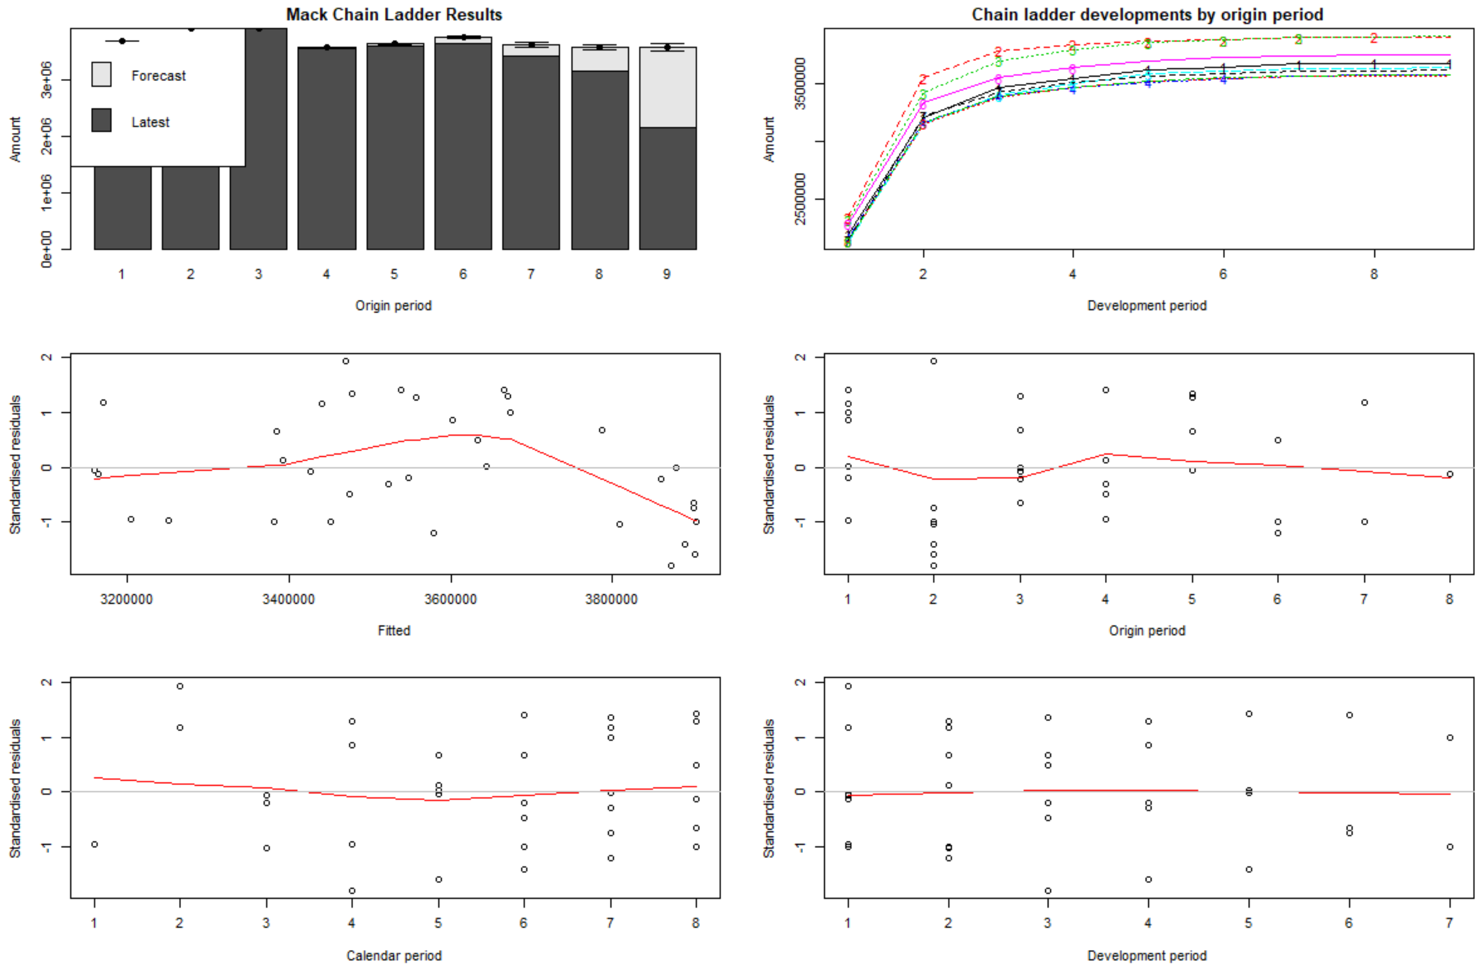

Tweedie Reserve

Tweedie – my reference material

Introduction

- Referenced in the context of model validation, for example in the case of internal model companies.
- Allows testing of different model structures.

Claims reserving with R:
ChainLadder-0.2.6 Package Vignette

Alessandro Carrato, Fabio Concina, Markus Gesmann, Dan Murphy, Mario Wüthrich and Wayne Zhang

May 29, 2018

Abstract

The ChainLadder package provides various statistical methods which are typically used for the estimation of outstanding claims reserves in general insurance, including those to estimate the claims development results as required under Solvency II.

1

Tweedie – introduction

- Over-Dispersed Poisson ("ODP") models have the following structure:

$Y \backsim as.factor(OY) + as.factor(DY)$ i.e. design.type=c(1,1,0)

- As noted earlier independence assumption could be violated in practice.
- For example, when the residuals plotted by calendar period show a pattern.
- This feature isn't being captured by the model structure above.
- With this model there is flexibility to change the regression structure to allow for these features.
- For example, a regression structure such as this could be tried:

$Y \backsim as.factor(DY) + as.factor(CY)$ i.e. design.type=c(0,1,1)

- Also flexibility to flex assumed underlying distribution i.e. p parameter. Noting that ODP is a special case of the Tweedie distribution with p equal to 1.

Tweedie – R help extracts

Tweedie Stochastic Reserving Model

“This function implements loss reserving models within the generalized linear model framework in order to generate the full predictive distribution for loss reserves. Besides, it generates also the one-year risk view useful to derive the reserve risk capital in a Solvency II framework. Finally, it allows the user to validate the model error while changing different model parameters, as the regression structure and diagnostics on the Tweedie p parameter.”

```
tweedieReserve(triangle, var.power = 1,
                    link.power = 0, design.type = c(1, 1, 0),
                    rereserving = FALSE, cum = TRUE, exposure = FALSE,
                    bootstrap = 1, boot.adj = 0, nsim = 1000,
                    proc.err = TRUE, p.optim = FALSE,
                    p.check = c(0, seq(1.1, 2.1, by = 0.1), 3),
                    progressBar = TRUE, ...)
```

https://www.rdocumentation.org/packages/ChainLadder/versions/0.2.5/topics/tweedieReserve

R implementation code – also see vignette

```
> p_profile <- tweedieReserve(MW2008, p.optim=TRUE,
+                             p.check=c(0,1.1,1.2,1.3,1.4,1.5,2,3),
+                             design.type=c(0,1,1),
+                             rereserving=FALSE,
+                             bootstrap=0,
+                             progressBar=FALSE)
0 1.1 1.2 1.3 1.4 1.5 2 3
........Done.
MLE of p is between 0 and 1, which is impossible. Instead, the MLE of p has been set
 to NA .  Please check your data and the call to tweedie.profile().
```

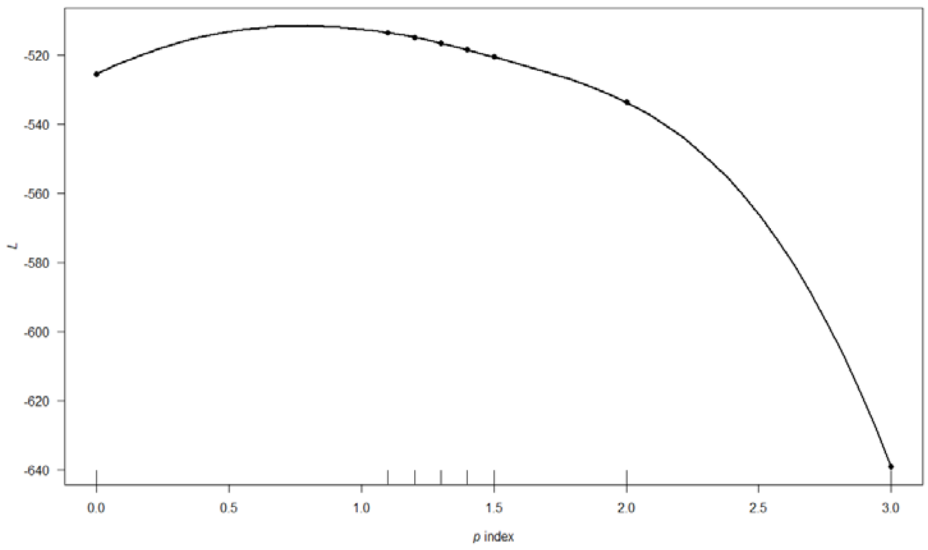

R implementation code and results

```
tweedie.res.a <- tweedieReserve(MW2008, var.power = 1, link.power = 0,
                                design.type=c(1,1,0),
                                rereserving=FALSE,
                                bootstrap=1,
                                progressBar=FALSE,cum=TRUE)
```

```
> tweedie.res.a$summary
        Latest    IBNR   IBNR.S.E        CoV Ultimate Det.IBNR Dev.To.Date
2      3902425    4288   5614.737 1.30940685  3906713     4378   0.9988794
3      3898825    9316   7272.266 0.78062113  3908141     9347   0.9976083
4      3548422   28629  12384.200 0.43257537  3577051    28392   0.9920622
5      3585812   51311  16178.072 0.31529443  3637123    51444   0.9858564
6      3641036  112120  21434.405 0.19117379  3753156   111811   0.9702064
7      3428335  187316  28947.913 0.15454052  3615651   187084   0.9482539
8      3158581  410528  41073.515 0.10005046  3569109   411864   0.8846463
9      2144738 1435228  94954.007 0.06615953  3579966  1433505   0.5993830
total 27308174 2238737 124826.879 0.05575772 29546911  2237826   0.9242596
```

R implementation code and results

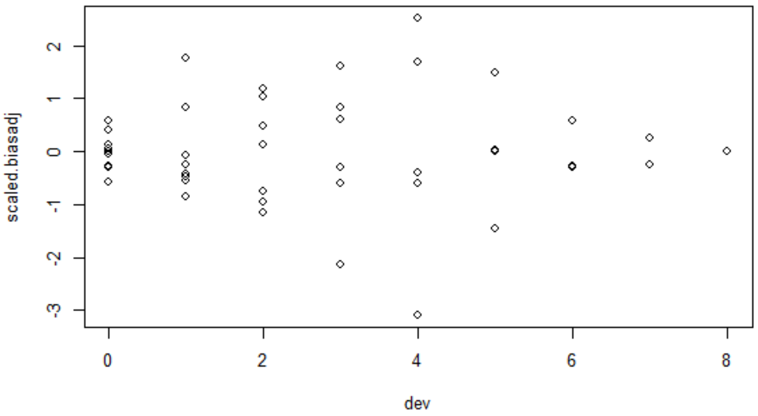

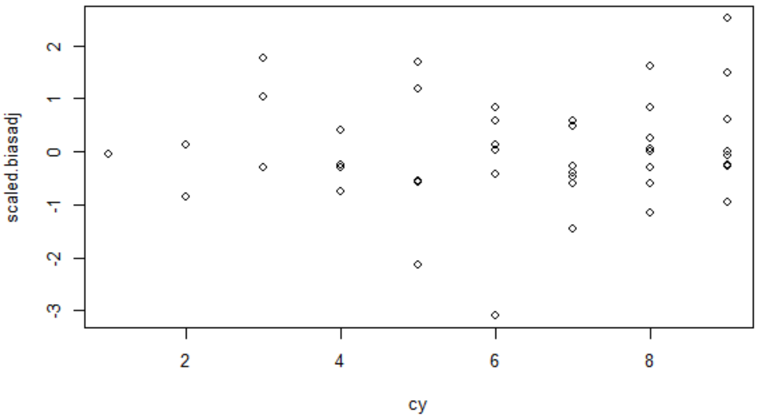

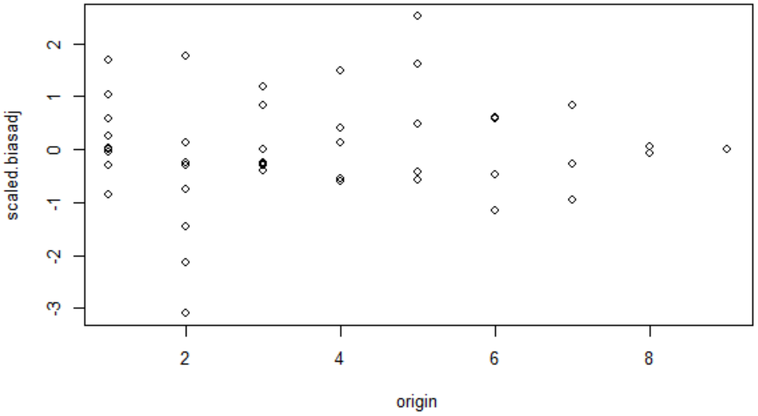

R implementation code and results

```
tweedie.res.b <- tweedieReserve(MW2008, var.power = 1, link.power = 0,
                                design.type=c(1,0,1),
                                rereserving=FALSE,
                                bootstrap=1,
                                progressBar=FALSE,cum=TRUE)
```

```
> tweedie.res.b$summary
       Latest    IBNR   IBNR.S.E       CoV Ultimate Det.IBNR Dev.To.Date
2     3902425     852   5102.055 5.9883273  3903277      955   0.9997553
3     3898825    3876  10616.182 2.7389529  3902701     3494   0.9991046
4     3548422    9602  17530.930 1.8257582  3558024     9712   0.9972705
5     3585812   27867  29414.195 1.0555207  3613679    27451   0.9924027
6     3641036   74145  49067.068 0.6617718  3715181    74841   0.9798591
7     3428335  194098  84557.879 0.4356453  3622433   194778   0.9462402
8     3158581  514345 145950.968 0.2837608  3672926   515572   0.8596760
9     2144738 1215652 278170.656 0.2288242  3360390  1220499   0.6373215
total 27308174 2040436 357299.360 0.1751093 29348610 2047300   0.9302583
```

R implementation code and results

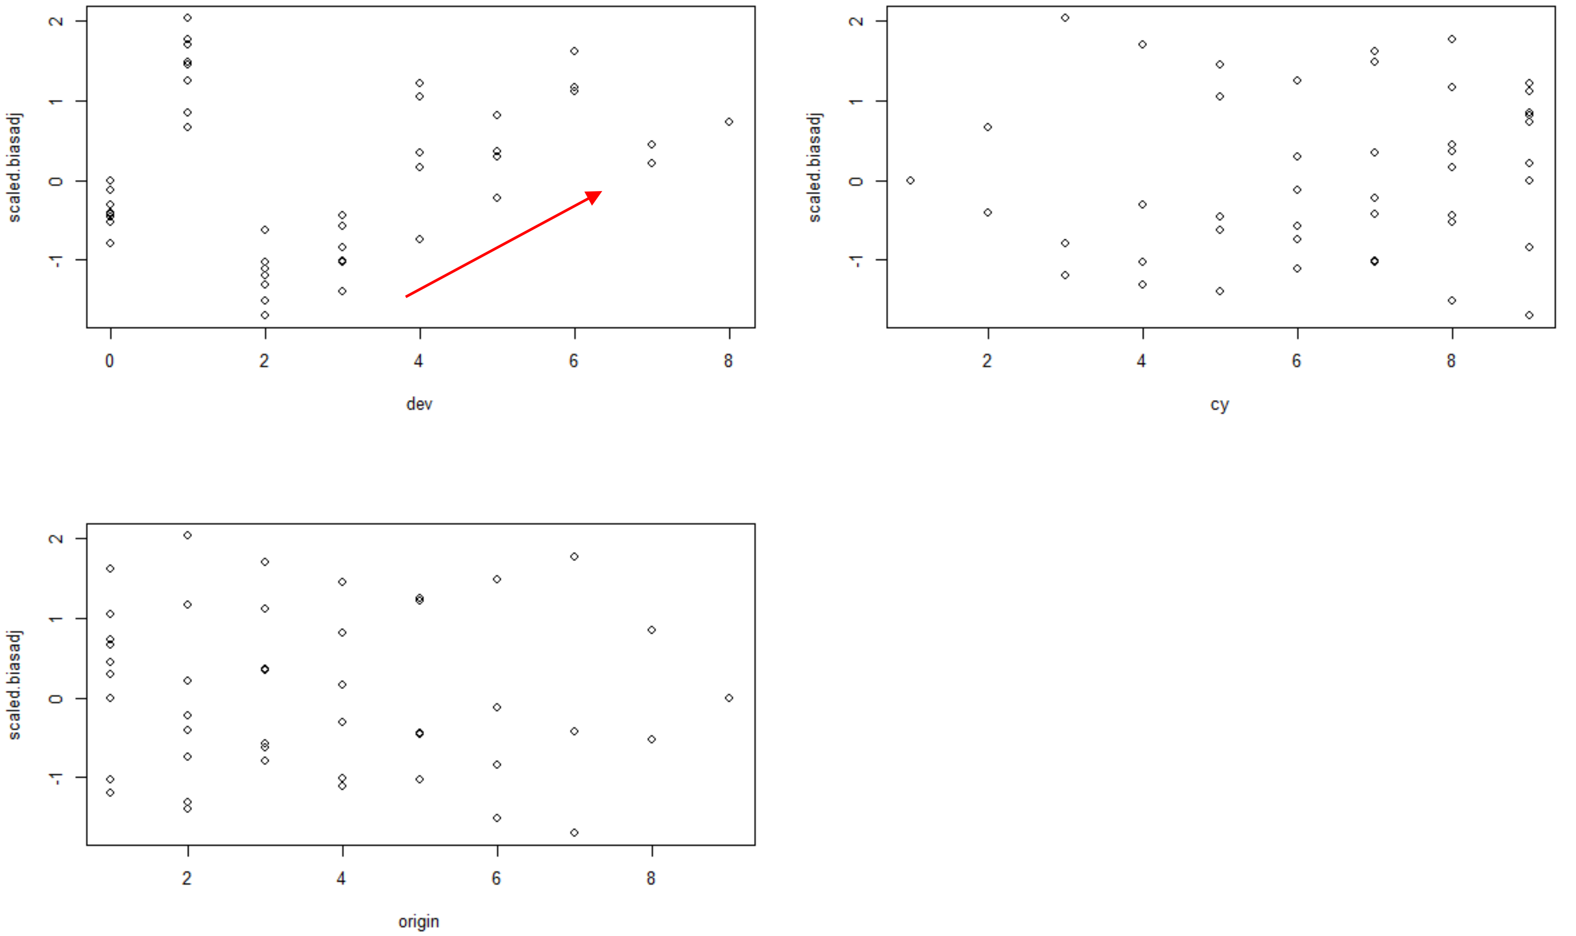

Markov Chain Monte Carlo (“MCMC”)

5 November 2018

MCMC – my reference material

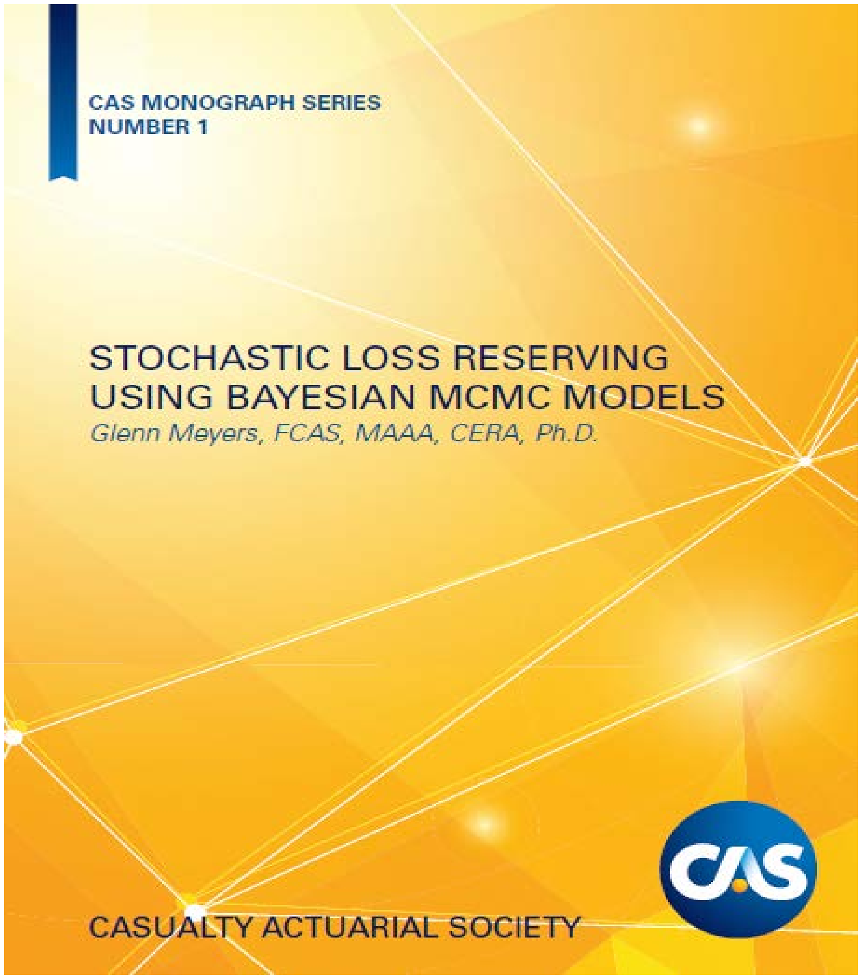

MCMC – what the paper did

- Notes a number of stochastic models have been developed
- However limited analysis had been done to retrospectively test, or validate, the performance of these models in an organised fashion on a large number of insurers.
- With the permission of NAIC the authors were able to build a database consisting of a large number of Schedule P triangles for six lines of insurance business and test the performance of various models based on retrospective testing.
- Basis of retrospective testing:
  - Specific models are build based on observed data.
  - The model is used to predict a distribution of outcomes that we will be observed in the future.
  - Since there is a reasonably large number of outcomes we expect the percentiles of outcomes vs. predictions to be uniformly distributed.
  - If they are not uniformly distributed it may suggest an issue with the model.

MCMC – what the paper found

- The variability predicted by Mack model was understated for the tested data, particularly in the tails.
- The paper notes that a key assumption of the model is that losses from different origin periods are independent.
- The paper proposed using a Correlated Chain Ladder (“CCL”) model as an alternative.
- The CCL allows for a form of dependency between the origin periods.
- The paper found that the CCL model predicted the distribution of outcomes within a confidence interval i.e. uniformity result held for this data.
- For paid data it found for the datasets considered that the Mack and Bootstrap ODP gave estimates that were high, suggesting change in environment not being captured by the models. Suggested alternatives such as Changing Settlement Rate models.
- It noted for one LOB that Mack and ODP validated better than the new models suggested.
- Concluded MCMC offers a flexible framework which can give improvements.

MCMC – a brief introduction

- Bayesian inference
  - Suppose we have a statistical model f(y I θ) (likelihood function), with unknown θ. We aim to infer θ conditional on observed data y.
  - Bayesian approach: specify a prior density g(θ) for θ, and then use Bayes' theorem which gives:
  - Posterior is proportional to the likelihood times the prior i.e. g(θ | y) $\propto$ f(y I θ) x g(θ)
- There are certain classes of Markov Chains that when given realisations from such a chain, asymptotic properties include distributional convergence of the realisations (Ergodic theory).
- In this set up the chains will converge to the posterior distribution of interest.
- Given that the chain has converged to the stationary distribution, realisations of the Markov chain can be regarded as a (dependent) sample from posterior distribution of interest.

https://www.youtube.com/watch?v=OTO1DygELpY

MCMC – example

EXPLANATION OF EXAMPLE CODE

To illustrate a simple Gibbs sampler in MATLAB and R, consider a data set $x = (x_1, \ldots, x_n)$ which each x_i is distributed as:

$$x_i \overset{iid}{\sim} N(\mu, \sigma^2)$$

for $i = 1, \ldots, n$ where μ and σ^2 are unknown parameters. The goal of this basic analysis is to estimate μ and σ^2. Assume the prior distributions for μ and σ^2 are:

$$\mu \sim N(m, s^2)$$

$$\sigma^{-2} = \phi \sim Gam(a, b)$$

where m, s^2, a, and b, are known hyperparameters specified by the researcher. Given this prior structure, the joint posterior distribution $p(\mu, \sigma^2|x)$ does not have closed form. Therefore, in order to conduct inference on μ and σ^2, we need to obtain samples $(\mu, \sigma^2)_j$ from $p(\mu, \sigma^2|x)$. To do so we can use the Gibbs sampler. Th Gibbs sampling algorithm for this example is:

(1) Sample μ from $p(\mu|\sigma^2, x)$
(2) Sample σ^2 from $p(\sigma^2|\mu, x)$

where $p(\mu|\sigma^2, x)$ is the "complete conditional" distibution of μ and $p(\sigma^2|\mu, x)$ is the complete conditional distribution of σ^2.

Under the prior distributions mentioned above and employing Bayes theorem, it can be shown that the complete conditional for μ is normal and the complete conditional for σ^2 is inverse gamma. Specifically,

$$\mu|\sigma^2, x \sim N(m^*, s_*^2)$$

$$\sigma^{-2}|\mu, x = \phi|\mu, x \sim Gam(a^*, b^*)$$

https://www2.stat.duke.edu/programs/gcc/ResourcesDocuments/CodeExplanation.pdf

MCMC – example (continued)

where,

$$m^* = \frac{\frac{1}{s^2}m + \frac{n}{\sigma^2}\bar{x}}{\frac{1}{s^2} + \frac{n}{\sigma^2}}$$

$$s_*^2 = \frac{1}{\frac{1}{s^2} + \frac{n}{\sigma^2}}$$

$$a^* = a + \frac{n}{2}$$

$$b^* = \frac{\sum_{i=1}^{n}(y_i - \mu)^2}{2} + b.$$

Given these complete conditional distributions, the Gibbs sampling algorithm will then proceed as follows:

(1) Sample μ from $N(\mu^*, s_*^2)$

(2) Sample σ^2 by sampling ϕ from $Gam(a^*, b^*)$ and setting $\sigma^2 = 1/\phi$.

The example code is this algorithm written in MATLAB and R. The code also provides examples of how to write data to a file, read data from a file, and plot some simple figures.

https://www2.stat.duke.edu/programs/gcc/ResourcesDocuments/CodeExplanation.pdf

MCMC – basic example with R code

```
#Simulate data of length 25 from N(10,20).
z=rnorm(25,10,sqrt(20))

mean(z)
sd(z)
sqrt(20)

gibbs2 = function(iters,y,mu=0,tau=1){
  # uniformative priors
  alpha0 = 0.00001
  beta0 = 0.00001
  mu0 = 0.0
  tau0 = sqrt(0.00001)

  x <-array(0,c(iters+1,2))
  x[1,1] = mu
  x[1,2] = tau
  n = length(y)
  ybar = mean(y)
  for(t in 2:(iters+1)){
    x[t,1] = rnorm(1,(n*ybar*x[t-1,2] + mu0*tau0)/
                 (n*x[t-1,2]+tau0), sqrt(1/(n*x[t-1,2]+tau0)))
    sn = sum((y-x[t,1])^2)
    x[t,2] = rgamma(1,alpha0+n/2)/(beta0+sn/2)
  }
  par(mfrow=c(1,2))
  plot(1:length(x[,1]),x[,1],type='l',lty=1,xlab='t',ylab='mu')
  plot(1:length(x[,2]),1/x[,2],type='l',lty=1,xlab='t',ylab='sigma^2')
  x
}
set.seed(1)
out2=gibbs2(1000,z,5,0.5)
out2
```

MCMC – basic example – trace plots

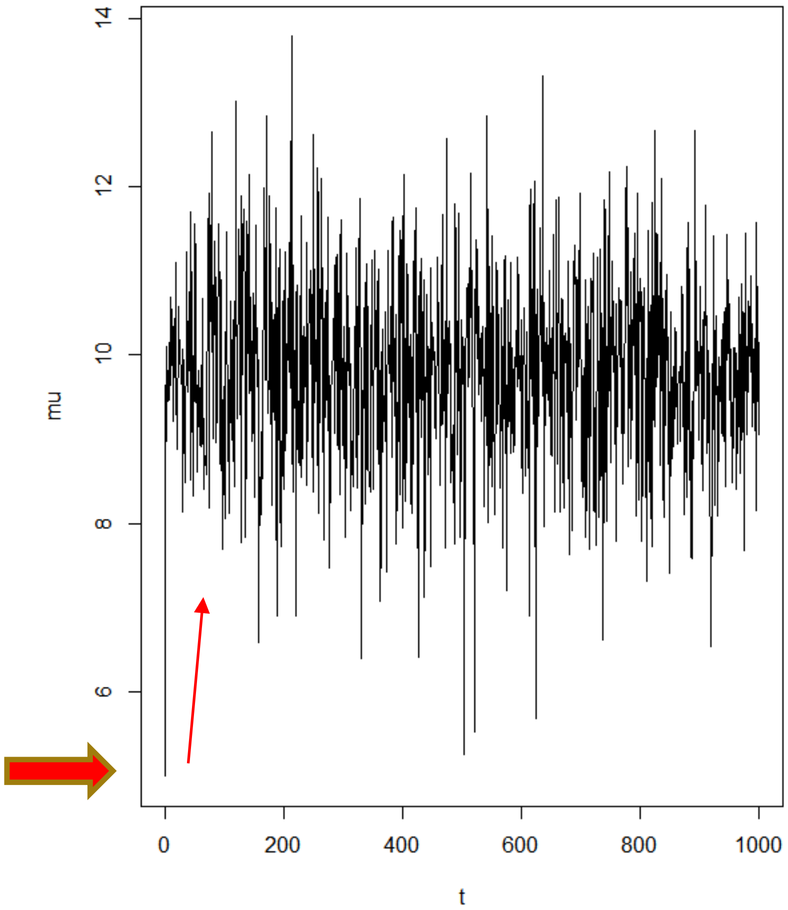

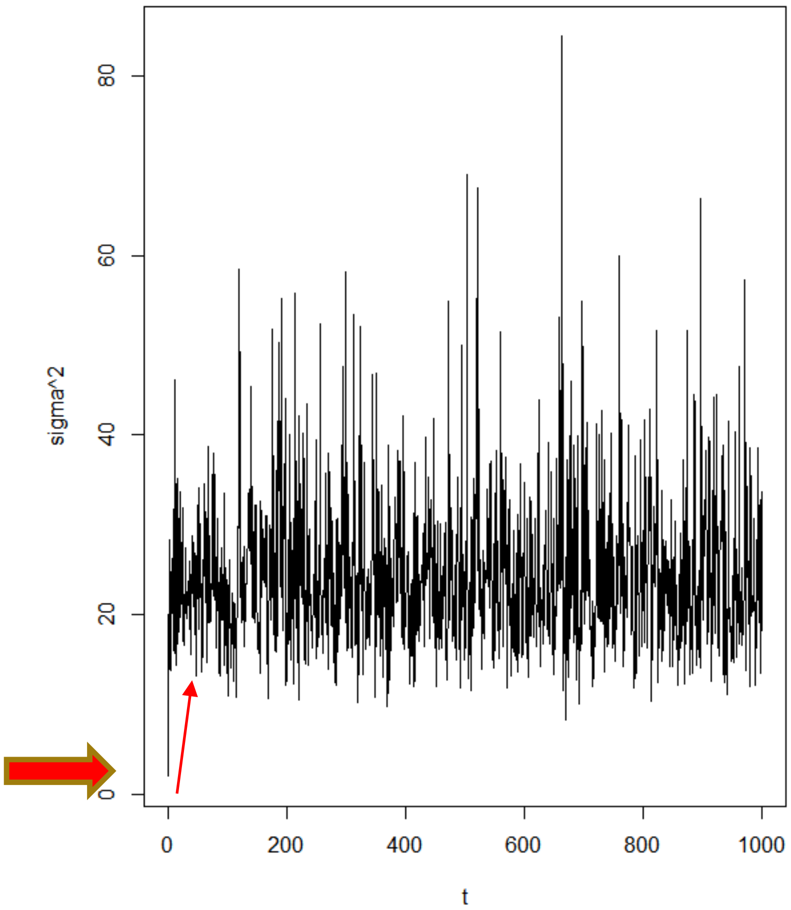

Questions Comments