Stochastic Loss Reserving with the Collective Risk Model

Glenn Meyers, FCAS, MAAA, Ph.D.

Abstract
This paper presents a Bayesian stochastic loss reserve model with the following features.

1. The model for expected loss payments depends upon unknown parameters that determine the expected loss ratio for the given accident years and the expected payment for each settlement lag.
2. The distribution of outcomes is given by the collective risk model in which the expected claim severity increases with the settlement lag. The claim count distribution is given by a Poisson distribution with its mean determined by dividing the expected loss by the expected claim severity.
3. The parameter sets that describe the posterior distribution of the parameters in (1) above are calculated with the Gibbs sampler.
4. For each parameter set generated by the Gibbs sampler in (3), the predicted distribution of outcomes is calculated using a Fast Fourier Transform (FFT). The Bayesian predictive distribution of outcomes is a mixture of the distributions of outcomes over all the parameter sets produced by the Gibbs sampler.

This paper concludes by applying this model to the problem of calculating risk margins for loss reserves using a cost of capital formula.

Keywords
Reserving Methods, Reserve Variability, Uncertainty and Ranges, Collective Risk Model, Fourier Methods, Bayesian Estimation

1. Introduction

Over the years, there has been an increasing recognition that consideration of the random nature of the insurance loss process leads to better predictions of ultimate losses. Some of the papers that led to this recognition include Stanard [11] and Barnett and Zehnwirth [1]. Another thread in the loss reserve literature has been to recognize outside information in the formulas that predict ultimate losses. Bornhuetter and Ferguson [2] represents one of the early papers exemplifying this approach.

More recently, papers by Meyers [7] and Verrall [12] have combined these two approaches with a Bayesian methodology. This paper continues the development of the approach started by Meyers and draws from the methodology described by Verrall.

As the actuarial profession improves its ability to describe the variability of its ultimate loss projections, there arises the question on how one should take this variability into account when setting loss reserves. One proposal originated by the International Association of Insurance Supervisors (IAIS) calls for a risk margin to be added to the actuarial present value of the future loss payments. This paper applies its loss reserve model to the calculation of risk margins.

A significant accomplishment of the Meyers paper cited above was that it made predictions of the distribution of future losses of real insurers, and successfully validated these predictions on subsequent reported losses. To do this, it was necessary to draw upon data that, while generally available, comes at a price. While this made a good case that the underlying model is realistic, it tended to inhibit future research on this methodology. This paper uses simulated data so that readers can verify all calculations. In addition, this paper includes the code that produced all results and, with minor modifications, it should be possible to use this code for other loss reserving applications.

2. The Collective Risk Model

This paper analyzes a 10 x 10 triangle of incremental paid losses organized by rows for accident years 1, 2, …, 10 and by columns for development lags 1, 2, …, 10. We also have the premium associated with each accident year. Table 1 gives the triangle that underlies the examples in this paper.

Table 1 (000)

AY	Premium	*Lag* 1	*Lag* 2	*Lag* 3	*Lag* 4	*Lag* 5	*Lag* 6	*Lag* 7	*Lag* 8	*Lag* 9	*Lag* 10
1	50,000	7,168	11,190	12,432	7,856	3,502	1,286	334	216	190	0
2	50,000	4,770	8,726	9,150	5,728	2,459	2,864	715	219	0	
3	50,000	5,821	9,467	7,741	3,736	1,402	972	720	50		
4	50,000	5,228	7,050	6,577	2,890	1,600	2,156	592			
5	50,000	4,185	6,573	5,196	2,869	3,609	1,283				
6	50,000	4,930	8,034	5,315	5,549	1,891					
7	50,000	4,936	7,357	5,817	5,278						
8	50,000	4,762	8,383	6,568							
9	50,000	5,025	8,898								
10	50,000	4,824									

Our job is to predict the distribution of losses in the empty cells ($AY + Lag > 11$) and to predict the distribution of the sum of losses in the empty cells.

Let us start by considering two models for the expected loss.

Model 1 – The Cape Cod Model

$$\mathrm{E}\left[Loss_{AY,Lag}\right] = Premium_{AY} \cdot ELR_{AY} \cdot Dev_{Lag} \quad (1)$$

The unknown parameters in this model are ELR_{AY} (AY = 1, 2, …, 10) and Dev_{Lag} (Lag = 1, 2, …, 10). The structure of the parameters is similar to the "Cape Cod" method discussed in Stanard [11] but, as we shall see, this paper's method of parameterizing the model is different.

Model 2 – The Beta Model

In the Cape Cod model, set

$$Dev_{Lag} = \beta\left(Lag/10 \mid a,b\right) - \beta\left((Lag-1)/10 \mid a,b\right) \quad (2)$$

where $\beta\left(x \mid a,b\right)$ is the cumulative probability of a beta distribution with unknown parameters *a* and *b* as parameterized in Appendix A of Klugman, Panjer and Willmot [5].

The Beta model replaces the ten unknown Dev_{Lag} parameters in the Cape Cod model with the two unknown parameters *a* and *b*. I chose these models as representatives of a multitude of possible models that can be used in this approach. Other examples in this multitude include the models in Meyers [7], who uses a Cape Cod model with constraints on the Dev_{Lag} parameters, and Clark [3], who uses the Loglogistic and Weibull distributions to project Dev_{Lag} parameters into the future.

Let $X_{AY,Lag}$ be a random variable for the loss in the cell (AY,Lag). We describe the distribution of $X_{AY,Lag}$ by the collective risk model, which can be described by the following simulation algorithm.

Simulation Algorithm 1

1. Select a random claim count, $N_{AY,Lag}$ from a Poisson distribution with mean $\lambda_{AY,Lag}$.
2. For i = 1, 2, …, $N_{AY,Lag}$, select a random claim amount, $Z_{Lag,i}$.
3. Set $X_{AY,Lag} = \sum_{i=1}^{N_{AY,Lag}} Z_{Lag,i}$, or if $N_{AY,Lag} = 0$, then $X_{AY,Lag} = 0$.

This paper assumes that the claim severity distributions of Z_{Lag} are given. In our example, we use the Pareto distribution with the cumulative distribution function:

$$F(z) = 1 - \left(\frac{\theta}{z+\theta}\right)^{\alpha}. \quad (3)$$

We set α = 2 for all settlement lags. θ will vary by settlement lag as noted in the following table.

Table 2

Lag	1	2	3	4	5	6	7-10
θ (000)	10	25	50	75	100	125	150

Note that the average severity increases with the settlement lag, which is consistent with the common observation that larger claims tend to take longer to settle.

To summarize, we have two models (the Cape Cod and the Beta) that give E[$X_{AY,Lag}$] in terms of the unknown parameters {ELR_{AY}} and {Dev_{Lag}}. We also assume that the claim severity distributions of Z_{Lag} are known. Then for any selected {ELR_{AY}} and {Dev_{Lag}}, we can describe the distribution of $X_{AY,Lag}$ by the following steps.

1. Calculate $\lambda_{AY,Lag} = \frac{\mathrm{E}\left[X_{AY,Lag}\right]}{\mathrm{E}\left[Z_{Lag}\right]} = \frac{Premium_{AY} \cdot ELR_{AY} \cdot Dev_{Lag}}{\mathrm{E}\left[Z_{Lag}\right]}$.
2. Generate the distribution of $X_{AY,Lag}$ using Simulation Algorithm 1 above.

3. The Posterior Distribution of Model Parameters

Let **X** denote the data in Table 1. Let $\ell\left(\mathbf{X} \mid \left\{ELR_{AY}\right\}, \left\{Dev_{Lag}\right\}\right)$ be the likelihood (or probability) of **X** given the parameters $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$. Note that defining a distribution in terms of a simulation algorithm does not lend itself to calculating the likelihood. To do this we must resort to some math that is described in detail in Appendix B. At this point, the reader should know that we are approximating the likelihood with something called the overdispersed negative binomial distribution.

The maximum likelihood estimator has been historically important and, as we shall see, will also be important in this paper. Over the past decade or so, a number of popular software packages began to include flexible function-maximizing tools that will search over a space that includes a fairly large number of parameters. Excel™ Solver is one such tool. With such a tool, the software programs[1] that accompany this paper calculate the maximum likelihood estimates for the Cape Cod and the Beta models.

The Cape Cod program calculates the maximum likelihood estimate by searching over the space of $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$, subject to a constraint that $\sum_{Lag=1}^{10} Dev_{Lag} = 1$. The Beta program feeds the results of Equation 2 into the likelihood function used in the Cape Cod program as it searches over the space of $\{ELR_{AY}\}$, *a* and *b*. Table 3 gives the maximum likelihood estimates for each model.

[1] The programs are written in R, a freely downloadable statistical package. See Meyers [6] for a review of this package.

Table 3

	Cape Cod		Beta	
	ELR	*Dev*	*ELR*	*Dev*
AY/Lag				
1	0.89090	0.16948	0.89205	0.15991
2	0.65285	0.26864	0.65670	0.27295
3	0.64448	0.23763	0.69949	0.24156
4	0.55233	0.15539	0.51727	0.16661
5	0.48569	0.07865	0.51696	0.09488
6	0.57259	0.05524	0.53697	0.04410
7	0.56411	0.01771	0.60935	0.01576
8	0.58207	0.00581	0.53487	0.00378
9	0.61922	0.00654	0.68940	0.00044
10	0.52190	0.00491	0.63902	0.00001
			a =	1.90742
			b =	5.78613

Let us now develop the framework for a Bayesian analysis. The likelihood function $\ell\left(\mathbf{X}\,|\,\{ELR_{AY}\},\{Dev_{Lag}\}\right)$ is the probability of **X,** given the parameters $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$. Using Bayes' Theorem, one can calculate the probability of the parameters $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ given the data, **X**.

$$\Pr\left\{\{ELR_{AY}\},\{Dev_{Lag}\}\middle|\mathbf{X}\right\} \propto \ell\left(\mathbf{X}\middle|\{ELR_{AY}\},\{Dev_{Lag}\}\right)\cdot\Pr\left\{\{ELR_{AY}\},\{Dev_{Lag}\}\right\}. \quad (4)$$

A discussion of selecting the prior distribution $\Pr\left\{\{ELR_{AY}\},\{Dev_{Lag}\}\right\}$ is in order. This paper has the advantage that it is working with simulated (i.e., made up) "data" so it is editorially possible to select anything as a prior distribution. However, I would like to spend some time to illustrate one way to approach the problem of selecting the prior distribution when working with real data.

Actuaries always stress the importance of judgment in setting reserves. Actuarial consultants will stress the experience that they have gained by examining the losses of other insurers. Meyers [7] formalizes this by examining the maximum likelihood estimates of the $\{Dev_{Lag}\}$ parameters from the data of 40 large insurers. In an effort to keep the examples in this paper as realistic as possible, I looked at the same data and selected the prior distribution as follows.

Beta Model[2]:

$$a \sim \Gamma(\alpha,\theta) \text{ with } \alpha = 75 \text{ and } \theta = 0.02 \quad (5)$$

$$b \sim \Gamma(\alpha,\theta) \text{ with } \alpha = 25 \text{ and } \theta = 0.20 \quad (6)$$

Figure 1 shows the Dev_{Lag} paths generated from a sample of fifty (*a*,*b*) pairs sampled from the prior distribution.

Figure 1

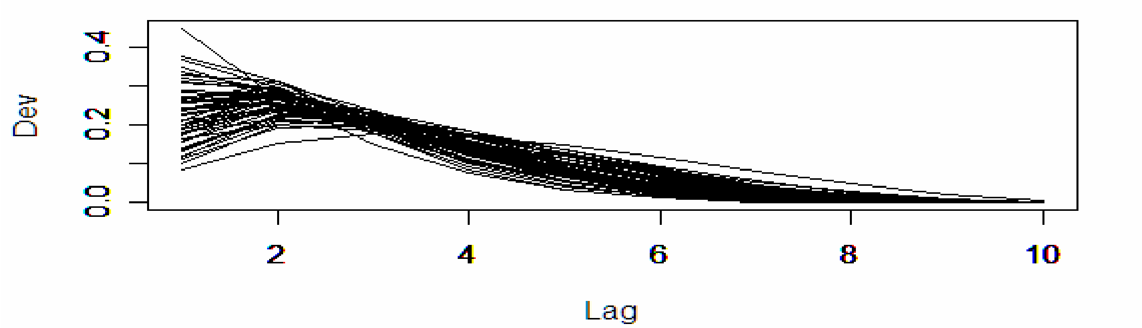

For the Cape Cod model, I calculated the mean and variance of the Dev_{Lag}s simulated from a large sample of (*a,b*) pairs and selected the following parameters for the gamma distribution for each Dev_{Lag}.

Table 4

$\Gamma \backslash Lag$	1	2	3	4	5	6	7	8	9	10
α	11.1010	64.6654	190.1538	34.9314	10.7284	4.4957	2.1298	1.0295	0.4574	0.1556
θ	0.0206	0.0041	0.0011	0.0040	0.0079	0.0101	0.0097	0.0073	0.0039	0.0009

[2] We will use the gamma (Γ) distribution as parameterized in Appendix A of Klugman, Panjer, and Willmot [5].

Before discussing the prior distribution of each ELR_{AY}, let us take a short side trip and look at the compound negative multinomial model[3], described by the following simulation algorithm.

Simulation Algorithm 2

1. For each accident year, select χ_{AY} at random from a gamma distribution with mean 1 and variance *c*.
2. For each accident year and settlement lag, select a claim count, $N_{AY,Lag}$, at random from a Poisson with mean $\chi_{AY}\cdot\lambda_{AY,Lag}$. (See the end of Section 2 for a description on how to determine the $\lambda_{AY,Lag}$s.)
3. For $i = 1, 2, \ldots, N_{AY,Lag}$, select a random claim amount, $Z_{Lag,i}$.
4. For each accident year and settlement lag, set $X_{AY,Lag} = \sum_{i=1}^{N_{AY,Lag}} Z_{Lag,i}$.

Note that for a given accident year, the $X_{AY,Lag}$s are correlated because of the common χ_{AY} that is in each *Lag*'s expected claim count.

This paper uses the compound negative multinomial model for the losses $X_{AY,Lag}$. At first glance, it might seem that this is different from the collective risk model described in Simulation Algorithm 1. But note that both the Cape Cod and the Beta models treat the ELR_{AY}s as unknown parameters. So by assigning a prior distribution to each ELR_{AY} so that its coefficient of variation squared is equal to the *c* in the negative multinomial model, we are explicitly modeling a random accident-year effect. With this in mind I selected each

$$ELR_{AY} \sim \Gamma(\alpha,\theta) \text{ with } \alpha = 100 \text{ and } \theta = 0.007. \quad (7)$$

Note that the expected value of each $ELR_{AY} = \alpha\cdot\theta = 0.70$ and the coefficient of variation of each $ELR_{AY} = \sqrt{1/\alpha} = 0.1$.

As we observe data points $x_{AY,Lag}$ in **X**, we gain information about the χ_{AY} in each accident year. As we shall see, treating each ELR_{AY} as an unknown parameter allows us to use this information in predicting the outcomes of future lags.

This paper uses the Gibbs sampler to generate random samples of the $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ parameters that represent the posterior distribution. Scollnik [10] introduced the Gibbs sampler to the CAS literature. Verrall [12] gives an application of it to a loss reserving problem.

For the Cape Cod model, this paper implements the Gibbs sampler as follows.

[3] The compound negative multinomial distribution was introduced to the CAS literature by Mildenhall [9].

Simulation Algorithm 3

1. Given the data triangle **X**, calculate the maximum likelihood estimates $Dev_{1,Lag}$ for Lag = 1, …,10 and $ELR_{1,AY}$ for AY = 1,…,10. Keep the maximum likelihood, ML, for future reference. Set i = 1.
2. Replace i by i+1, set $Dev_{i,Lag} = Dev_{i\text{-}1,Lag}$ and set $ELR_{i,AY} = ELR_{i\text{-}1,AY}$.
3. For Lag = 1 to 10:
    a. Replace $Dev_{i,Lag}$ with a random number taken from the prior distribution of Dev_{Lag} and calculate its likelihood L.
    b. Select a random number, u, from a uniform (0,1) distribution.
    c. If $L/ML < u$, then return to Step 3a, otherwise continue to the next step.
4. For AY = 1 to 10:
    a. Replace $ELR_{i,AY}$ with a random number taken from the prior distribution of ELR_{AY} and calculate its likelihood L.
    b. Select a random number, u, from a uniform (0,1) distribution.
    c. If $L/ML < u$, then return to Step 4a, otherwise continue to the next step.
5. Return to Step 2 until i is greater than a selected n.

The intuition behind this algorithm is that a parameter "applies" to be included in the Gibbs sample in proportion to its prior probability. Each applicant is "accepted" into the sample in proportion to its likelihood. So the probability of a parameter being included in the sample is the product of the probability of applying times its likelihood, which in turn is equal to its posterior probability. See Equation 4 above.

Figure 2 provides a graphic comparison between the prior distribution and the posterior distribution, as represented by the output of Simulation Algorithm 3. The upper histograms are random samples of ELR_1, taken from its prior and posterior distributions. The lower graphs represent the paths taken from the prior and posterior $\{Dev_{Lag}\}$ distributions.

Figure 2

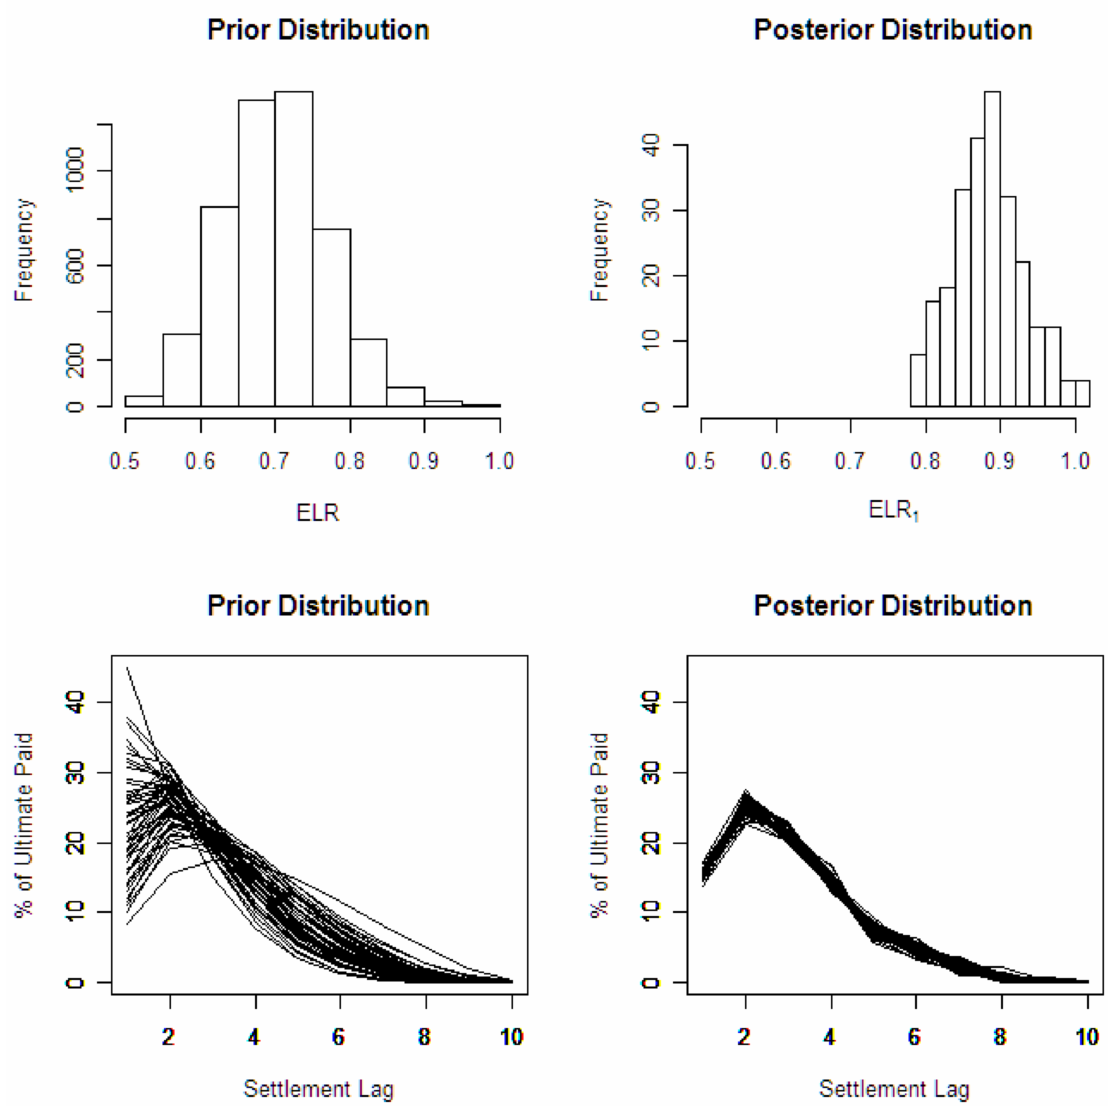

For the Beta model, we implement the Gibbs sampler as follows.

Simulation Algorithm 4

1. Given the data triangle **X**, calculate the maximum likelihood estimates a_1, b_1 and $ELR_{1,AY}$ for $AY = 1,\ldots,10$. Keep the maximum likelihood, ML, for future reference. Set $i = 1$.
2. Replace i by $i+1$, set $a_i = a_{i-1}$, set $b_i = b_{i-1}$ and set $ELR_{i,AY} = ELR_{i-1,AY}$.
3. For $p_i = a_i$ and then b_i:
   a. Replace p_i with a random number taken from the prior distribution of p, calculate the associated $Dev_{i,Lag}$s using Equation 2 and calculate the likelihood L.
   b. Select a random number, u, from a uniform (0,1) distribution.
   c. If $L/ML < u$, then return to Step 3a, otherwise continue to the next step.
4. For $AY = 1$ to 10:
   d. Replace $ELR_{i,AY}$ with a random number taken from the prior distribution of ELR_{AY} and calculate the likelihood L.
   e. Select a random number, u, from a uniform (0,1) distribution.
   f. If $L/ML < u$, then return to Step 4a, otherwise continue to the next step.
5. Return to Step 2 until i is greater than a selected n.

Each iteration is a step in a Markov chain of random transformations in the parameter space $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$. It is well know that Markov chains will converge to a limiting distribution and that, when executed as described in these simulation algorithms, the limiting distribution will be the posterior distribution.

The random parameters generated by the first several iterations of the Gibbs sampler may not be distributed as the limiting distribution. So it is a general practice to discard parameters that are generated early in the process. By examining successive blocks of parameters in the examples in this paper, I concluded that using parameters generated after 250 iterations[4] of Simulation Algorithms 3 and 4 was sufficiently accurate for our purposes. Table 5 shows some illustrative results that came out of Simulation Algorithm 4 being applied to the data in Table 1.

Table 5

[4] One may find other sources that recommend thousands of iterations. But these sources generally count one draw of a parameter from its prior distribution as one iteration. When counting that way, 250 iterations of Simulation Algorithms 3 and 4 represent 5,000 and 3,000 iterations respectively.

Iteration	ELR_1	ELR_2	ELR_3	ELR_4	ELR_5	ELR_6	ELR_7	ELR_8	ELR_9	ELR_{10}
251	0.75403	0.69942	0.62441	0.56447	0.51833	0.60362	0.61284	0.60487	0.66776	0.63434
252	0.84815	0.73598	0.64300	0.59662	0.50991	0.66961	0.63046	0.65861	0.78867	0.62436
253	0.82959	0.65535	0.62372	0.58285	0.54446	0.65939	0.62067	0.66337	0.77458	0.68514
254	0.82214	0.67833	0.72494	0.58108	0.59692	0.64186	0.64096	0.69660	0.63273	0.71884
255	0.85885	0.70338	0.65320	0.60643	0.57145	0.65768	0.74067	0.64207	0.61538	0.56273
256	0.82655	0.68402	0.71207	0.57685	0.50899	0.62291	0.68097	0.60459	0.73825	0.62867
257	0.86339	0.71486	0.62554	0.55949	0.54898	0.57494	0.63603	0.66952	0.68241	0.61616
258	0.81831	0.64761	0.73752	0.61186	0.63983	0.62646	0.61374	0.67133	0.64861	0.62245
259	0.80801	0.66089	0.70570	0.61823	0.57213	0.62688	0.58704	0.69212	0.62392	0.67231
260	0.81955	0.65917	0.61623	0.64292	0.56440	0.61969	0.61458	0.67270	0.74439	0.59132

Iteration	Dev_1	Dev_2	Dev_3	Dev_4	Dev_5	Dev_6	Dev_7	Dev_8	Dev_9	Dev_{10}
251	0.17353	0.26609	0.23075	0.16171	0.09592	0.04754	0.01863	0.00509	0.00072	0.00002
252	0.17373	0.26219	0.22815	0.16179	0.09773	0.04965	0.02012	0.00576	0.00087	0.00003
253	0.15662	0.25141	0.22857	0.16863	0.10601	0.05625	0.02396	0.00730	0.00119	0.00004
254	0.15514	0.24770	0.22656	0.16906	0.10796	0.05847	0.02559	0.00808	0.00139	0.00005
255	0.16275	0.25121	0.22557	0.16608	0.10487	0.05622	0.02435	0.00760	0.00130	0.00005
256	0.16274	0.24870	0.22378	0.16595	0.10596	0.05768	0.02550	0.00819	0.00145	0.00006
257	0.16549	0.25142	0.22449	0.16497	0.10422	0.05600	0.02436	0.00766	0.00132	0.00005
258	0.15983	0.24720	0.22401	0.16705	0.10721	0.05865	0.02607	0.00842	0.00151	0.00006
259	0.17049	0.25879	0.22734	0.16312	0.09993	0.05165	0.02138	0.00629	0.00099	0.00003
260	0.16584	0.26100	0.23092	0.16494	0.09979	0.05056	0.02034	0.00574	0.00085	0.00003

One can often find interesting information about the uncertainty in the parameter estimates by examining tables of parameters generated by the Gibbs sampler. Figure 3 below show the coefficients of variation (CV) of the loss ratio estimates taken from 2,500 additional iterations of the sample in Table 5. This illustrates how we gain information about the ultimate loss ratio as we get more data from each accident year. Wacek [13] gives another approach to estimating loss ratios as we gain information over time.

Figure 3

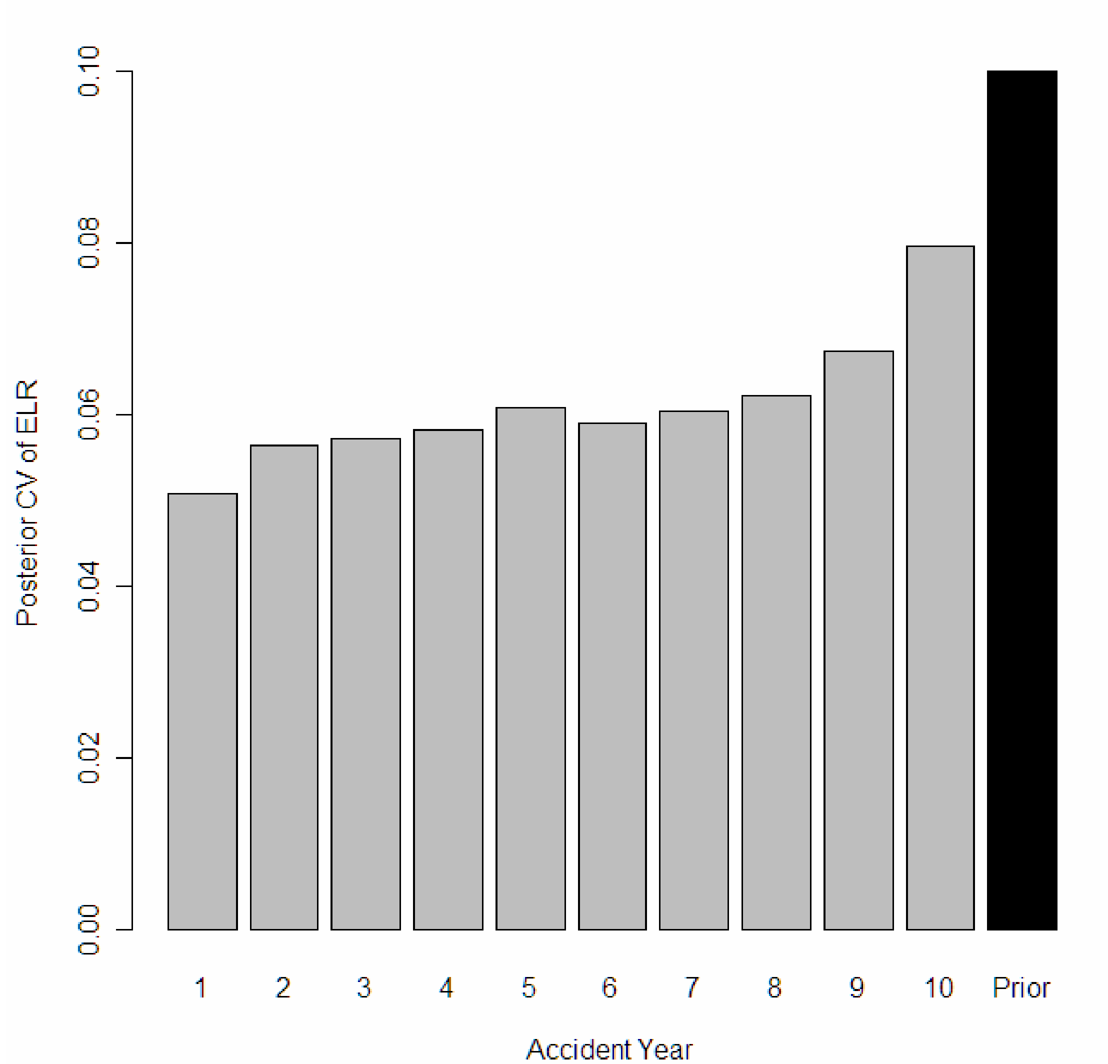

4. The Predictive Distribution of Outcomes

Now that we have the posterior distribution estimated from the data of Table 1, we now turn to the problem of predicting future outcomes, $X_{AY,Lag}$, when AY + Lag > 11.

Table 1 (000) (Repeated)

AY	Premium	*Lag* 1	*Lag* 2	*Lag* 3	*Lag* 4	*Lag* 5	*Lag* 6	*Lag* 7	*Lag* 8	*Lag* 9	*Lag* 10
1	50,000	7,168	11,190	12,432	7,856	3,502	1,286	334	216	190	0
2	50,000	4,770	8,726	9,150	5,728	2,459	2,864	715	219	0	$X_{2,10}$
3	50,000	5,821	9,467	7,741	3,736	1,402	972	720	50	$X_{3,9}$	$X_{3.10}$
4	50,000	5,228	7,050	6,577	2,890	1,600	2,156	592	$X_{4,8}$	$X_{4,9}$	$X_{4,10}$
5	50,000	4,185	6,573	5,196	2,869	3,609	1,283	$X_{5,7}$	$X_{5,8}$	$X_{5,9}$	$X_{5,10}$
6	50,000	4,930	8,034	5,315	5,549	1,891	$X_{6,6}$	$X_{6,7}$	$X_{6,8}$	$X_{6,9}$	$X_{6,10}$
7	50,000	4,936	7,357	5,817	5,278	$X_{7,5}$	$X_{7,6}$	$X_{7,7}$	$X_{7,8}$	$X_{7,9}$	$X_{7,10}$
8	50,000	4,762	8,383	6,568	$X_{8,4}$	$X_{8,5}$	$X_{8,6}$	$X_{8,7}$	$X_{8,8}$	$X_{8,9}$	$X_{8,10}$
9	50,000	5,025	8,898	$X_{9,3}$	$X_{9,4}$	$X_{9,5}$	$X_{9,6}$	$X_{9,7}$	$X_{9,8}$	$X_{9,9}$	$X_{9,10}$
10	50,000	4,824	$X_{10,2}$	$X_{10,3}$	$X_{10,4}$	$X_{10,5}$	$X_{10,6}$	$X_{10,7}$	$X_{10,8}$	$X_{10,9}$	$X_{10,10}$

While there are many statistics of interest that one could examine, I chose to examine the predictive distribution of the total reserve:

$$R = \sum_{AY=2}^{10} \sum_{Lag=12-AY}^{10} X_{AY,Lag} \,. \tag{8}$$

Suppose we have a set of parameters $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ calculated from several iterations of the Gibbs sampler. Conceptually, the easiest way to calculate the distribution of outcomes is by repeated use of the following simulation algorithm.

Simulation Algorithm 5

1. Select the parameters $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ from a randomly selected iteration.
2. For $AY = 2, \ldots, 10$, do:
    a. For $Lag = 12 - AY$ to 10, do:
        i. Set $\lambda_{AY,Lag} = \dfrac{Premium_{AY} \cdot ELR_{AY} \cdot Dev_{Lag}}{\mathrm{E}\left[Z_{Lag}\right]}$
        ii. Select N at random from a Poisson distribution with mean $\lambda_{AY,Lag}$.
        iii. If $N > 0$, for $i = 1, \ldots, N$ select claim amounts, $Z_{i,Lag}$, at random from the claim severity distribution for the Lag.
        iv. If $N > 0$, set $X_{AY,Lag} = \sum_{i=1}^{N} Z_{i,Lag}$, otherwise set $X_{AY,Lag} = 0$.
3. Set $R = \sum_{AY=2}^{10} \sum_{Lag=12-AY}^{10} X_{AY,Lag}$.

I expect that many actuaries will be satisfied with using this simulation algorithm to calculate the predictive distribution. However, this paper uses a Fast Fourier Transform (FFT) to calculate the predictive distribution. While it is very technical and harder to implement, it is faster and it produces more accurate results (relative to the model assumptions). Appendix A describes how to implement the FFT for this paper's application.

Figure 4 plots the density functions for the predictive distributions derived from the data in Table 1. For each model, I ran 500 iterations of the Gibbs sampler and discarded the first 250 because they are less likely to represent the posterior distributions.

Figure 4

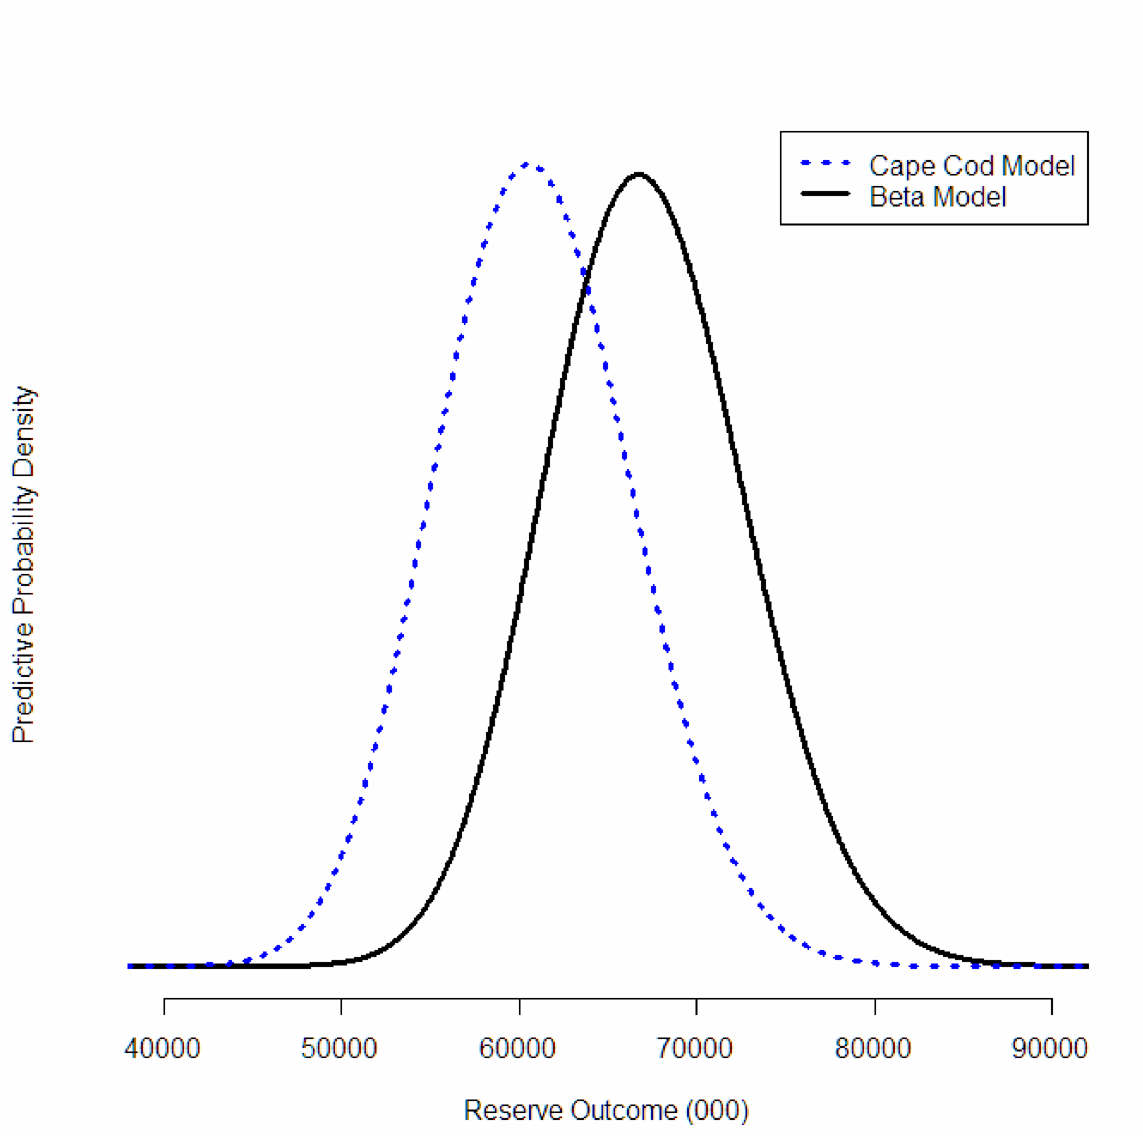

The predictive means and standard deviations are:

- 60,871,000 and 5,487,000 for the Cape Cod model; and
- 67,183,000 and 5,605,000 for the Beta model.

The difference in the predictive means for the two models is 5,982,000, illustrating the fact that we do face "model risk." If one wants to reflect model risk, one could modify Simulation Algorithm

5 by randomly selecting parameters from the $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ lists provided by the Gibbs samples for each model.

5. Risk Margins in Loss Reserves

Now that we have demonstrated a method that quantifies the uncertainty in the estimates of future loss payments, we now turn to exploring how this information might be used to post a loss reserve on a financial statement. The art of accounting has always had difficulty in dealing with uncertainty. A common practice, when possible, is to value a liability at its market value. When there is no active market, as is typically the case for loss reserves, the fallback position is to use a model to calculate the cost that an insurance market would "theoretically" charge to transfer the risky reserve.

As this paper is being written, there is still active debate on whether and how to do this. Meyers [8] provides some background and references on this subject. This section only addresses the "how."

I should add that in preparing this section I immeasurably benefited from the discussions that led to the paper jointly written by Kaufman, Broughton, Buchanan, and Meyers [4]. That paper discusses a variety of methods to calculate risk margins for loss reserves, whereas this paper illustrates only one of those methods.

The formula discussed here is called the Capital Cash Flow (CCF) risk margin. In words, this formula assumes that investors in a reinsurer would need to put up (or allocate) capital to take on the loss reserve risk by a ceding insurer. As claims are settled, the reinsurer expects to be able to release the capital over time. The CCF risk margin is the profit that the reinsurer would need to be persuaded to take on this risky venture.

We will now discuss the details. Let:

- i = Risk-free rate of return on investments.
- r = Total rate of return demanded by the reinsurer for taking additional insurance risk.
- C_t = Amount of capital required to (or allocated to) support an insurance portfolio at time t.

First look at the cash flow of the insurance transaction.

- At the beginning of the first year, at time $t = 0$, investors contribute a sum of C_0 to the reinsurer, which earns a risk-free rate of return, i, over the next year.
- At time $t = 0$, the reinsurer collects M_{CCF} from the ceding insurer and immediately transfers it to its investors. Equivalently, one could say that the investor contributes $C_0 - M_{CCF}$ to the reinsurer.
- At time $t = 1$, the investors expect to keep C_1 invested in the reinsurer, and they expect to receive a cash flow $C_0(1+i) - C_1$ at the end of year 1. Since the loss the reinsurer is required to pay and C_1 are uncertain, they discount the value of the amount returned at the risky rate of return $r > i$.
- Continuing on to time t, the investors expect to keep C_t invested in the reinsurer, and they expect a cash flow of $C_{t-1}(1+i) - C_t$ at the end of year t.

Since the cash flows are uncertain, it is appropriate to discount the cash flow at the risky rate of return, r. This leads to the following expression.

$$C_0 = M_{CCF} + \sum_{t=1}^{\infty} \frac{C_{t-1}(1+i) - C_t}{(1+r)^t}. \qquad (9)$$

This equation implies

$$\begin{aligned} M_{CCF} &= C_0 - \sum_{t=1}^{\infty} \frac{C_{t-1}(1+i) - C_t}{(1+r)^t} \\ &= \frac{C_0(1+r-1-i)}{1+r} + \frac{C_1(1+r-1-i)}{(1+r)^2} + \frac{C_2(1+r-1-i)}{(1+r)^3} + \ldots \\ &= (r-i)\sum_{t=0}^{\infty} \frac{C_t}{(1+r)^{t+1}}. \end{aligned} \qquad (10)$$

The following table shows how to calculate C_t for the example in this paper fit with the Beta model.

Table 6 (000)

(1) t	(2) L_t^{Nom}	(3) ΔL_t^{Nom}	(4) L_t^{Disc}	(5) TVaR_t^{Nom}	(6) $\Delta\text{TVaR}_t^{Nom}$	(7) TVaR_t^{Disc}	(8) C_t
0	67,183	27,103	61,224	80,617	28,086	72,373	11,149
1	40,080	18,847	36,993	52,531	21,984	47,799	10,805
2	21,233	11,391	19,809	30,547	14,167	28,033	8,224
3	9,843	5,978	9,270	16,380	8,224	15,129	5,859
4	3,864	2,653	3,671	8,156	4,315	7,570	3,899
5	1,211	940	1,160	3,841	2,075	3,581	2,422
6	271	237	261	1,766	856	1,659	1,398
7	34	33	33	909	803	877	845
8	1	1	1	106	106	103	102

(1) The time, *t*, after the liability is set.

(2) The nominal expected value of future payments, $L_t^{Nom} = \sum_{AY=2+t}^{10} \sum_{Lag=AY}^{10} \text{E}\left[X_{AY,Lag}\right]$.

(3) $\Delta L_t^{Nom} = L_t^{Nom} - L_{t+1}^{Nom}$.

(4) The discounted liability, $L_t^{Disc} = \sum_{k=t}^{8} \frac{\Delta L_k^{Nom}}{(1+i)^{k-t+0.5}}$, where i = 6%.

(5) The nominal Tail-Value-at-Risk, i.e., the conditional expected value of the nominal random losses, $\sum_{AY=2+t}^{10} \sum_{Lag=AY}^{10} X_{AY,Lag}$, given that they exceed their 99th percentile. The density functions for the nominal losses are plotted on Figure 5 for each *t*.

(6) $\Delta\text{TVaR}_t^{Nom} = \text{TVaR}_t^{Nom} - \text{TVaR}_{t+1}^{Nom}$.

(7) The discounted $\text{TVaR}_t^{Disc} = \sum_{k=t}^{8} \frac{\Delta\text{TVaR}_k^{Nom}}{(1+i)^{k-t+0.5}}$.

(8) The needed capital at time *t* is expected to be $C_t = \text{TVaR}_t^{Disc} - L_t^{Disc}$.

Now that we have the C_ts, we can then use Equation 10, with r = 10%, to calculate M_{CFF} = 1,368,000, which is 2.2% of the discounted liability, 61,224,000.

Figure 5 (000)

Density Functions for the Nominal Losses as They Run Off

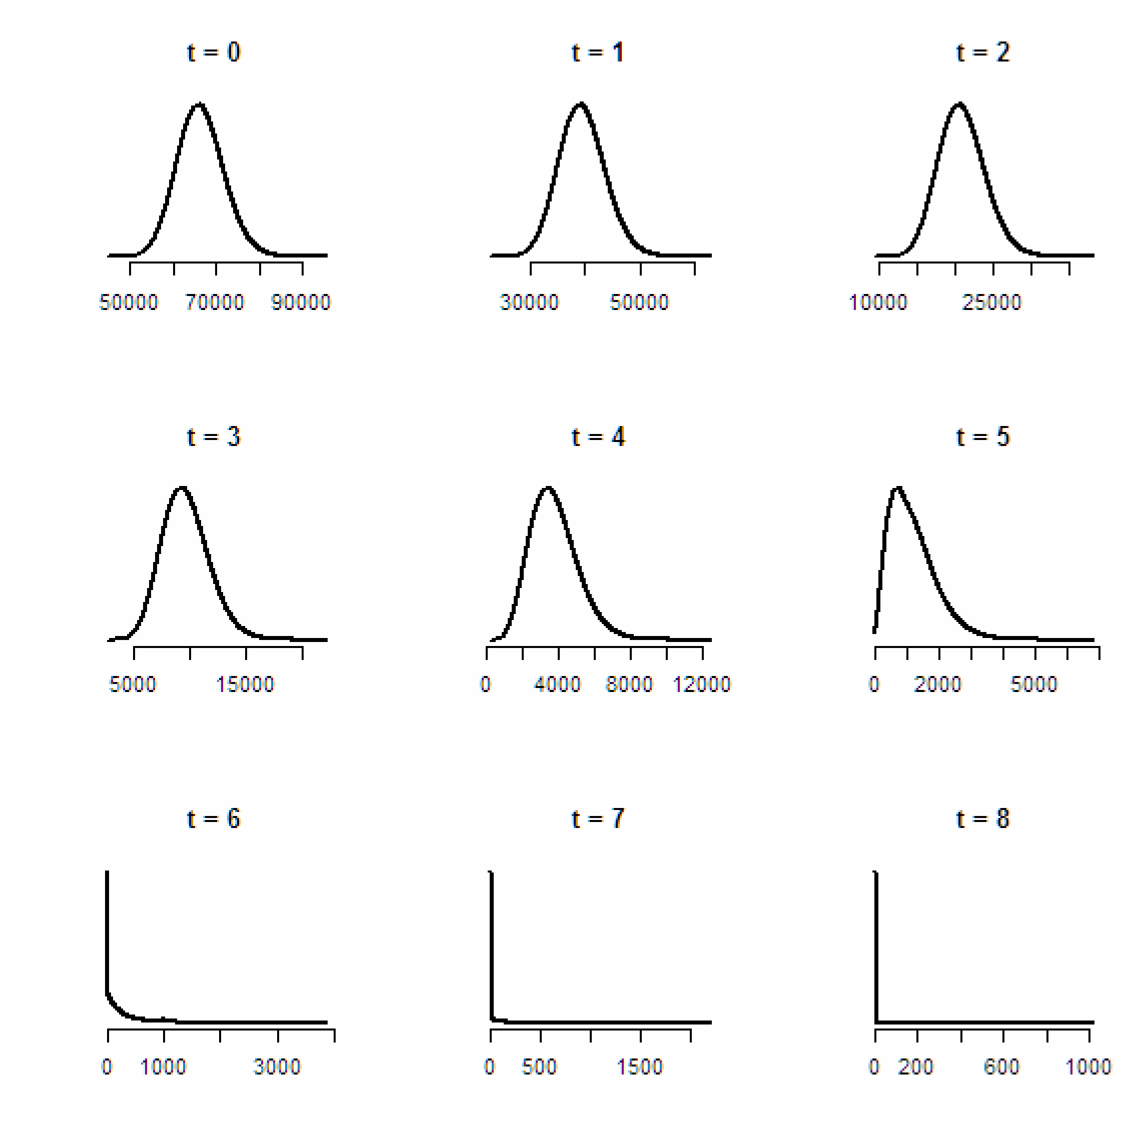

- In the latter stages of the runoff, there are a small number of potentially large claims (limited to 1,000,000) that occasionally are paid. Thus, you see the spikes at zero. The density function was plotted for those loss amounts for which the cumulative distribution function was less than 0.999999.

The risk margin calculation above was based on the nominal TVaR for the insurer's own losses. This is tantamount to assuming that the reinsurer has no other business to diversify the losses. If the liability is ever transferred, it will almost surely be transferred to a sizeable reinsurer with a diverse portfolio of losses. Rather than specify the characteristics of the reinsurer, a good approximation to the reinsurer's cost of capital would be to base the calculation of the distribution of the insurer's uncertainty in the expected values, as generated by the $\{ELR\}$ and the $\{Dev\}$ parameters in the Gibbs sampler. Table 7 calculates the risk margin under this assumption.

Table 7 (000)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t	L_t^{Nom}	ΔL_t^{Nom}	L_t^{Disc}	TVaR_t^{Nom}	$\Delta\text{TVaR}_t^{Nom}$	TVaR_t^{Disc}	C_t
0	67,183	27,103	61,224	76,583	29,581	69,488	8,264
1	40,080	18,847	36,993	47,002	21,079	43,202	6,208
2	21,233	11,391	19,809	25,923	13,294	24,092	4,283
3	9,843	5,978	9,270	12,629	7,270	11,850	2,580
4	3,864	2,653	3,671	5,359	3,514	5,076	1,405
5	1,211	940	1,160	1,845	1,381	1,763	603
6	271	237	261	464	397	447	186
7	34	33	33	67	65	65	33
8	1	1	1	3	3	3	2

The explanation of the columns is the same as for Table 6 except for Column 5.

(5) The nominal Tail-Value-at-Risk at the 99% level, where the random element is the expected value of the Gibbs sample, $\sum_{AY=2+t}^{10} \sum_{Lag=AY}^{10} Premium_{AY} \cdot ELR_{AY} \cdot Dev_{Lag}$. The histograms of the sums calculated from the Gibbs sample are plotted on Figure 6 for each t.

Now that we have the C_ts, we can then use Equation 10, with $r = 10\%$, to calculate M_{CFF} = 758,000 which is 1.2% of the discounted liability, 61,224,000.

Figure 6

Histograms of the Expected Runoff Scenarios Taken from the Gibbs Sample

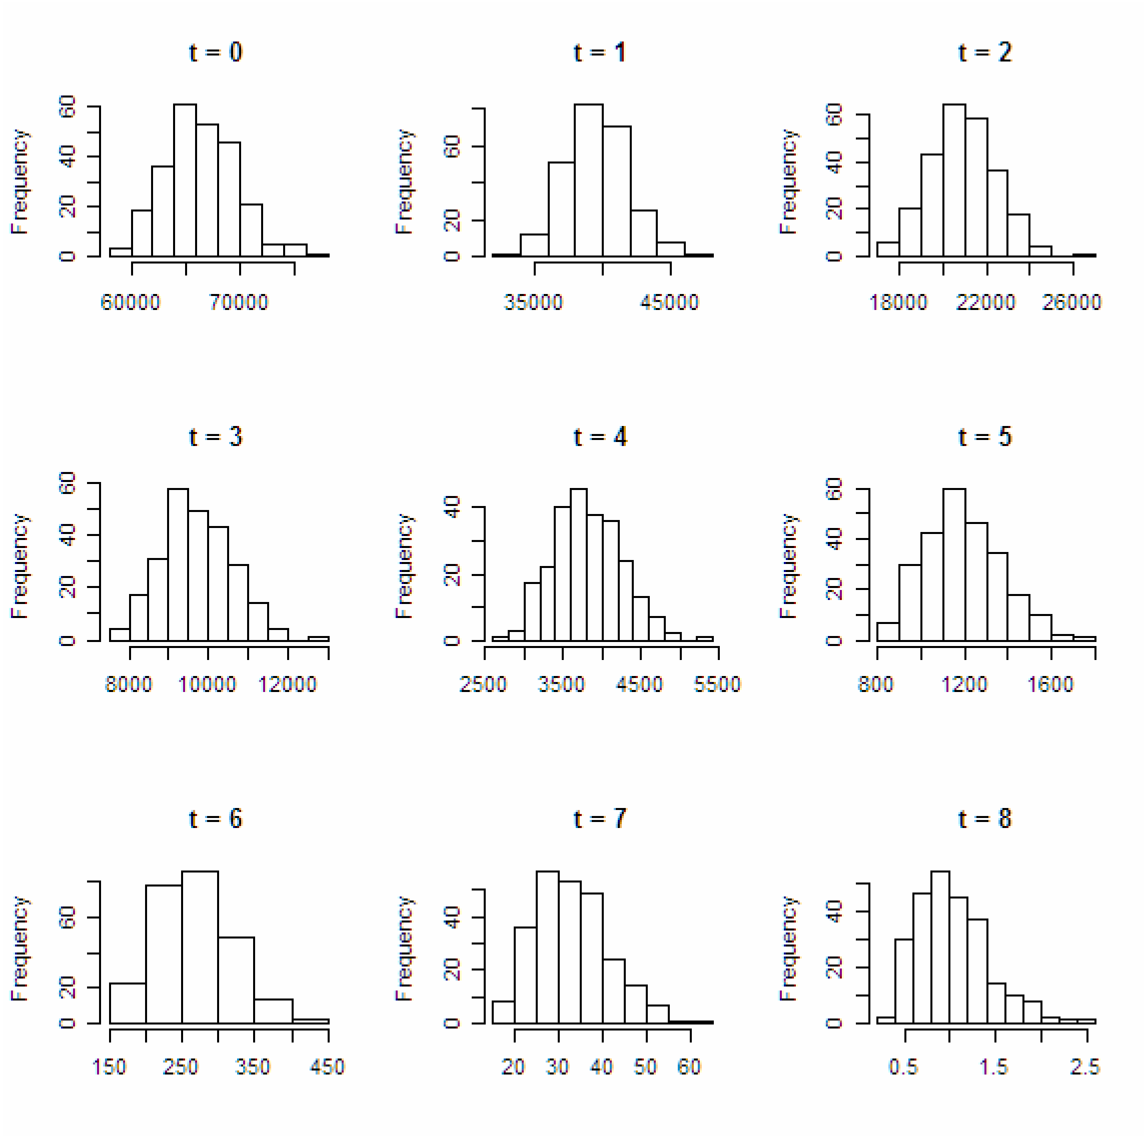

With the exception of workers compensation insurance, it is standard statutory accounting practice in the USA to post loss reserves at nominal, not discounted values. A common justification for this practice is that it provides a cushion for the risk in the posted reserve. In the above examples, the difference between the nominal and discounted expected values of the liability is 67,183,000 – 61,224,000 = 5,959,000. This difference is noticeably larger than the 1,368,000 and 758,000 risk margins calculated in the examples above.

Note that the CCF risk margin is sensitive to three factors that many consider when accessing risk:

1. The volatility of the future payouts as quantified by C_t. If desired, one can consider only parameter risk.
2. How long the insurer is exposed to the risk, as quantified by how C_t decreases over time.
3. The premium the market places on risk, as quantified by $r - i$.

Note that proposals for risk margins based solely on statistics taken from a predictive distribution, such as percentiles, do not address (2) and (3) above. The American practice of posting reserves at their nominal value does not address (1) and (3) above.

References

1. Barnett, Glen and Benjamin Zehnwirth, "Best Estimates for Reserves," *Proceedings of the Casualty Actuarial Society* LXXXVII, 2000, pp.245-321. http://www.casact.org/pubs/proceed/proceed00/00245.pdf

2. Bornhuetter, Ronald L. and Ronald E. Ferguson, "The Actuary and IBNR," *Proceedings of the Casualty Actuarial Society* LIX, 1972, pp.181-195. http://www.casact.org/pubs/proceed/proceed72/72181.pdf

3. Clark, David R., "LDF Curve Fitting and Stochastic Loss Reserving: A Maximum Likelihood Approach," *Casualty Actuarial Society Forum*, Fall 2003, pp. 41-92 http://www.casact.org/pubs/forum/03fforum/03ff041.pdf

4. Kaufman, Allan, Jonathan Broughton, Bob Buchanan, and Glenn Meyers, "Cost of Capital Method of Determining Risk Margins—Dealing with Neglected Issues: Observations from Three Empirical Analyses," 2007 GIRO Convention Proceedings.

5. Klugman, Stuart A., Harry A. Panjer, and Gordon E. Willmot, *Loss Models: From Data to Decisions*, Second Edition, John Wiley and Sons, 2004.

6. Meyers, Glenn G., "The R Programming Language," *Actuarial Review*, vol. 33, no. 4, 2006. http://www.casact.org/newsletter/index.cfm?fa=viewart&id=5311

7. Meyers, Glenn G., "Estimating Predictive Distributions for Loss Reserve Models," *Variance* Fall 2007, 1(2):248-272, http://www.variancejournal.org/issues/01-02/248.pdf.

8. Meyers, Glenn G., "Should Reserves Include Risk Margins?—International Developments," *Actuarial Review*, vol. 34, no. 4, 2007b. http://www.casact.org/newsletter/index.cfm?fa=viewart&id=5499

9. Mildenhall, Stephen J., "Correlations and Aggregate Loss Distributions With an Emphasis on the Iman-Conover Method." *Casualty Actuarial Society Forum,* Winter 2006, pp. 103-204. http://www.casact.org/pubs/forum/06wforum/06w107.pdf

10. Scollnik, David P. M., "Introduction to Markov Chain Monte Carlo Methods and Their Actuarial Applications," *Proceedings of the Casualty Actuarial Society* LXXXIII, 1996, pp. 114-165 http://www.casact.org/pubs/proceed/proceed96/96114.pdf

11. Stanard, James N., "A Simulation Test of Prediction Errors of Loss Reserve Estimation Techniques," *Proceedings of the Casualty Actuarial Society* LXXII, 1985, pp.124-148. http://www.casact.org/pubs/proceed/proceed85/85124.pdf

12. Verrall, Richard J., "Obtaining Predictive Distributions for Reserves Which Incorporate Expert Opinion," *Variance,* Spring 2007, 1(1):53-80. http://www.variancejournal.org/issues/01-01/053.pdf

13. Wacek, Michael G., "The Path of the Ultimate Loss Ratio Estimate," *Variance,* Fall 2007, 1(2):173-192. http://www.variancejournal.org/issues/01-02/173.pdf

Appendix A. Collective Risk Model Mathematics with Fast Fourier Transforms

This paper describes the collective risk model in terms of a simulation algorithm. Given the speed of today's personal computers, it is practical to actually do the simulations in a reasonable amount of time. This appendix describes how to do many of the calculations to a higher degree of accuracy in a significantly shorter time using FFTs.

The advantage to using FFTs is that the time-consuming task of calculating the distribution of the sum of random variables is transformed into the much faster task of multiplying the FFTs of the distributions. Simulation Algorithms 1, 2, and 5 show that the collective risk model requires the calculation of the distribution of the sum of random claim amounts. Furthermore, Simulation Algorithm 5 requires the calculation of the distribution of the sum of losses over different accident years and settlement lags.

This appendix has three sections. Since the FFTs work on discrete random variables, the first section shows how to discretize the claim severity distribution in such a way that the limited average severities of the continuous severity distribution are preserved. The second section will show how to calculate the probabilities associated with the collective risk model. The third section will show how to calculate the predictive distribution for the outstanding losses.

A.1 Discretizing the Claim Severity Distributions

The first step is to determine the discretization interval length h. Variable h, which depended on the size of the insurer, was chosen so the 2^{14} (16,384) values spanned the probable range of annual losses for the insurer. Specifically, let h_1 be the sum of the insurer's ten-year premium divided by 2^{14}. The h was set equal to 1,000 times the smallest number from the set $\{5, 10, 20, 25, 40, 50, 100, 125, 200, 250, 500, 1000\}$ that was greater than $h_1/1000$. This last step guarantees that a multiple, m, of h would be equal to the policy limit of 1,000,000.

The next step is to use the mean-preserving method (described in Klugman [5, p. 656] to discretize the claim severity distribution for each settlement lag. Let $p_{i,Lag}$ represent the probability of a claim with severity $h \cdot i$ for each settlement lag. Using the limited average severity (LAS_{Lag}) function determined from claim severity distributions, the method proceeds in the following steps.

1. $p_{0,Lag} = 1 - \text{LAS}_{Lag}(h)/h$.
2. $p_{i,Lag} = (2 \cdot \text{LAS}_{Lag}(h \cdot i) - \text{LAS}_{Lag}(h \cdot (i-1)) - \text{LAS}_{Lag}(h \cdot (i+1)))/h$ for i = 1, 2, ..., m-1.

3. $p_{m,Lag} = 1 - \sum_{i=0}^{m-1} p_{i,Lag}$.
4. $p_{ik} = 0$ for $i = m + 1, ..., 2^{14} - 1$.

A.2 Calculating Probabilities for the Compound Poisson Distribution

The purpose of this section is to show how to calculate the probabilities of losses defined by the collective risk model as defined in Simulation Algorithm 1. The math described in this section is derived in Klugman [5, Section 6.91]. The calculation proceeds in the following steps.

1. Set $\vec{\mathbf{p}}_{Lag} = \left\{ p_{0,Lag}, ...p_{2^{14}-1,Lag} \right\}$.
2. Calculate the expected claim count, $\lambda_{AY,Lag}$, for each accident year and settlement lag using Equation 2, $\lambda_{AY,Lag} \equiv \mathrm{E}\left[Paid\ Loss_{AY,Lag} \right] / \mathrm{E}\left[Z_{Lag} \right]$.
3. Calculate the Fast Fourier Transform (FFT) of $\vec{\mathbf{p}}_{Lag}$, $\Phi\left(\vec{\mathbf{p}}_{Lag}\right)$.
4. Calculate the FFT of each aggregate loss random variable, $X_{AY,Lag}$, using the formula

$$\Phi\left(\vec{\mathbf{q}}_{AY,Lag}\right) = e^{\left(\Phi\left(\vec{\mathbf{p}}_{Lag}\right)-1\right)}.$$

   This formula is derived in Klugman[5, Section 6.91].

5. Calculate $\vec{\mathbf{q}}_{AY,Lag} = \Phi^{-1}\left(\Phi\left(\vec{\mathbf{q}}_{AY,Lag}\right)\right)$, the inverse FFT of the expression in Step 4 above.

The vector ,$\vec{\mathbf{q}}_{AY,Lag}$, contains the probabilities of the discretized compound Poisson distribution defined by Simulation Algorithm 1.

A.3 Calculating Probabilities for the Predictive Distribution

To calculate the predictive distribution of the reserve outcomes by the methods in this paper, one needs the $\{ELR_{AY}, Dev_{Lag}\}$ parameter set that was simulated by the Gibbs sampler as described in Section 3 above.

1. For each parameter set, denoted by *i*, and $AY+Lag > 11$, do the following.

   a. Calculate the expected loss, $Premium_{AY,i} \cdot ELR_{AY,i} \cdot Dev_{Lag,i}$.

   b. Calculate the FFT of the aggregate loss $X_{AY,Lag,i}$, $\Phi\left(\vec{\mathbf{q}}_{AY,Lag,i}\right)$ as described in Step 4 in section A.2 above.

2. For each parameter set, *i*, calculate the product $\Phi\left(\vec{\mathbf{q}}_i\right) \equiv \prod_{AY=2}^{10} \prod_{Lag=12-AY}^{10} \Phi\left(\vec{\mathbf{q}}_{AY,Lag,i}\right)$.

3. Calculate the FFT of the mixture over all *i*, $\Phi\left(\vec{\mathbf{q}}\right) = \frac{\sum_i \Phi\left(\vec{\mathbf{q}}_i\right)}{n}$, where *n* is the number of Gibbs samples.

4. Invert the FFT, $\Phi\left(\vec{\mathbf{q}}\right)$, to obtain the vector, $\vec{\mathbf{q}}$, which describes the distribution of the of the reserve outcomes.

Here are the formulas to calculate the mean and standard deviation of the reserve outcomes:

- Expected Value = $h \cdot \sum_{j=0}^{2^{14}-1} j \cdot \vec{\mathbf{q}}_j$.
- Second Moment = $h^2 \cdot \sum_{j=0}^{2^{14}-1} j^2 \cdot \vec{\mathbf{q}}_j$.
- Standard Deviation = $\sqrt{\text{Second Moment} - \left(\text{First Moment}\right)^2}$.

Figure 4 has plots of the $\vec{\mathbf{q}}$'s for the Cape Cod and the Beta models.

Appendix B. An Approximate Likelihood Calculation for the Collective Risk Model

The goal of this appendix is to show how to calculate approximate likelihoods $\ell\left(\mathbf{X} \mid \left\{ELR_{AY}\right\}, \left\{Dev_{Lag}\right\}\right)$ for the Cape Cod model and $\ell\left(\mathbf{X} \mid \left\{ELR_{AY}\right\}, a, b\right)$ for the Beta Model, where the distribution of each $X_{AY,Lag}$ is defined by Simulation Algorithm 1 above.

This paper does not follow Meyers [7], which uses FFTs, as described in Appendix A to calculate the likelihood. The reason for this is the speed of calculation. While today's computers can calculate a likelihood with the FFT in a fraction of a second, the use of the Gibbs sampler can require the calculation of millions of likelihoods. My experience is that the approximate likelihood calculation described below cuts the computing time by a factor of 60.

The general strategy for calculating the likelihood is to start by calculating the first two moments of the aggregate loss for each accident year and settlement lag in terms of the expected loss and the first two moments of the claim severity distribution. The next step is to find an overdispersed negative binomial (ODNB) distribution that has the same first two moments. We then approximate the probability of the observed loss with its probability indicated by the ODNB distribution.

The log-likelihood for a given triangle of data is then given by:

$$\sum_{AY=1}^{10} \sum_{Lag=1}^{11-AY} \log\left(\text{ODNB}\left(x_{AY,Lag}\right)\right).$$

Here are the steps for calculating each $\log\left(\text{ODNB}\left(x_{AY,Lag}\right)\right)$:

Step 1 – Calculate the first two moments of $X_{AY,Lag}$.

Let μ_{Lag} and σ^2_{lag} be the mean and variance of claim severity distribution for the given settlement lag. Formulas for these moments are in Klugman [5]. Next calculate the expected claim count,

$$\lambda_{AY,Lag} = \frac{Premium_{AY} \cdot ELR_{AY} \cdot Dev_{Lag}}{\mathrm{E}\left[Z_{Lag} \right]}.$$

Then the variance of the compound Poisson distribution for $X_{AY,Lag}$ is given by

$$Var\left[X_{AY,Lag} \right] = \lambda_{AY,Lag} \cdot \left(\mu^2_{Lag} + \sigma^2_{Lag} \right).$$

Step 2 – Find an ODNB distribution with the same moments as that of $X_{AY,Lag}$.

We parameterize the negative binomial distribution so that the variance is equal to:

$$\lambda_{AY,Lag} + \frac{\lambda^2_{AY,Lag}}{\kappa_{AY,Lag}}.$$

If each claim has a constant size of $\mu_{AY,Lag}$, its variance is then equal to:

$$\mu^2_{AY,Lag} \left(\lambda_{AY,Lag} + \frac{\lambda^2_{AY,Lag}}{\kappa_{AY,Lag}} \right).$$

Equating the variance from Step 1 with the above variance and solving for κ yields:

$$\kappa_{AY,Lag} = \frac{\lambda_{AY,Lag} \cdot \mu^2_{Lag}}{\sigma^2_{Lag}}.$$

Given the parameters ELR_{AY} and Dev_{Lag}, we approximate the log-likelihood of an observation $x_{AY,Lag}$ follows.

1. Set $n_{AY,Lag} = x_{AY,Lag} / \mu_{Lag}$ rounded to the nearest integer.

2. Set $\log\left(\mathrm{ODNB}\left(x_{AY,Lag} \right) \right) = \log\left(\Pr\left(N = n_{AY,Lag} \mid \lambda_{AY,Lag}, \kappa_{AY,Lag} \right) \right)$.

Appendix C. Computer Code for the Algorithms.

This appendix describes the code that implements the algorithms in this paper. The code is written in R, a computer language that can be downloaded for free at www.R-Project.org[5]. The code itself will be posted in a zip folder that accompanies this paper on the CAS Web Site.

There is one feature of the code that is not described above. Occasionally the Gibbs sampler admits a set of parameters with low likelihood. The presence of such parameters causes subsequent parameters to have a high rejection rate with the result that the algorithm is "trapped." When this happens, the algorithm returns to a randomly selected parameter set that had been accepted earlier.

Here is a description of the files in the zip folder.

1. The Rectangle.csv – This is the triangle in Table 1 expressed in rectangular form so it fits into an R data frame.

2. CRM CCod Posterior.r – This code reads The Rectangle.csv and implements the Gibbs sampler to produce an output file containing sampled $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ parameters from the Cape Cod model.

3. CRM CCod Posterior.csv – The output from a run of CRM CCod Posterior.r

4. CRM Beta Posterior.r – This code reads The Rectangle.csv and implements the Gibbs sampler to produce an output file containing sampled $\{ELR_{AY}\}$ and $\{Dev_{Lag}\}$ parameters from the Beta model.

5. CRM Beta Posterior.csv – The output from a run of CRM Beta Posterior.r. Some of the records in this dataset are in Table 5.

6. Predict Outcomes.r – This code takes the output from Files 3 and 5 above and calculates the predictive distribution. It creates graphs like those in Figure 4.

7. Risk Margin.r – This code takes File 5 and calculates the expected losses and TVaRs needed for the risk margin calculation.

---

[5] Meyers [6] provides more information about the R programming language.

8. Risk Margin.xls – This spreadsheet takes the output of File 7 and produces Tables 6 and 7.

Biography of the Author

Glenn Meyers is Vice President and Chief Actuary for ISO Innovative Analytics, a division of ISO devoted to predictive modeling. He holds a bachelor's degree in mathematics and physics from Alma College, a master's degree in mathematics from Oakland University, and a Ph.D. in mathematics from the State University of New York at Albany. Glenn is a Fellow of the Casualty Actuarial Society and a member of the American Academy of Actuaries. Before joining ISO in 1988, Glenn worked at CNA Insurance Companies and the University of Iowa.

Glenn's current responsibilities at ISO include the development of scoring products. Prior responsibilities have included working on ISO Capital Management products, increased limits and catastrophe ratemaking, ISO's, and Property Size-of-Loss Database (PSOLD), ISO's model for commercial property size-of-loss distributions.

Glenn's work has been published in *Variance* and the *Proceedings of the Casualty Actuarial Society*. He is a three-time winner of the Woodward-Fondiller Prize, a two-time winner of the Dorweiller Prize, and a winner of the Dynamic Financial Analysis Prize. He is a frequent speaker at CAS meetings and seminars.

His service to the CAS has included membership on various education and research committees. He currently serves on the International Actuarial Association Solvency Committee and the CAS Board of Directors.