

A Cost-of-Capital Risk Margin Formula for Nonlife Insurance Liabilities

by Glenn Meyers

ABSTRACT

A Bayesian Markov chain Monte Carlo (MCMC) stochastic loss reserve model provides an arbitrarily large number of equally likely parameter sets that enable one to simulate future cash flows of the liability. Using these parameter sets to represent all future outcomes, it is possible to describe any future state in the model's time horizon including those states necessary to calculate a cost-of-capital risk margin. This paper shows how to use the MCMC output to (1) calculate the risk margin for an "ultimate" time horizon; (2) calculate the risk margin for a one-year time horizon; and (3) analyze the effect of diversification in a risk margin calculation for multiple lines of insurance.

KEYWORDS

Stochastic loss reserving, Bayesian MCMC, capital requirements, risk margins

1. Introduction

With the growing influence of Bayesian Markov chain Monte Carlo (MCMC) models in stochastic loss reserving (e.g., Meyers 2015), this paper will illustrate one way to use such a model to calculate a cost-of-capital risk margin for nonlife insurance liabilities. The need for such a calculation is found in the “technical provisions” specified in the European Union’s Solvency II act.¹

Those technical provisions refer to the insurer’s liability for unpaid losses. Specifically:

1. “The value of the technical provisions shall be equal to the sum of a best estimate and a risk margin.”
2. “The best estimate shall correspond to the probability-weighted average of future cash flows, taking account of the time value of money using the relevant risk-free interest rate term structure.”
3. “The risk margin shall be calculated by determining the cost of providing an amount of eligible own funds equal to the Solvency Capital Requirement necessary to support the insurance obligations over the lifetime thereof.”
4. “Insurance undertakings shall segment their insurance obligations into homogeneous risk groups, and as a minimum by lines of business, when calculating the technical provisions.”

This paper illustrates a way to implement the principles expressed in the above provisions of the act. Although the act goes on to provide some specific provisions on how to implement those principles, the purpose of this paper is more to show how to implement the principles underlying Solvency II using more theoretically sound risk management principles along with the Bayesian MCMC technol-

ogy. In the concluding remarks, we describe the main differences between the approach in this paper and that specified by Solvency II.

A Bayesian MCMC stochastic loss reserve model provides an arbitrarily large number (say 10,000) of equally likely parameter sets that enable one to simulate future cash flows of the liability. From those parameter sets, it is possible to describe any future state in the model’s time horizon including those states necessary to calculate the technical provisions. That is what this paper will do.

Here is a high-level description of that cash flow:

1. At the end of the current calendar year (call this time $t = 0$), the insurer posts its best estimate of the liability. The insurer also posts the amount of capital, C_0 , needed to contain the uncertainty in this estimate. It invests C_0 in a fund that earns income at the risk-free interest rate i .
2. At the end of the next calendar year, at time $t = 1$, the insurer uses its next year of loss experience to reevaluate its liability.² It then posts its updated estimate of the liability and the capital, C_1 , needed to contain the uncertainty in this estimate. The difference between $C_0 \cdot (1 + i)$ and C_1 is returned to the investor. If that difference is negative, as it occasionally will be, the investor is expected to contribute an amount to make up that difference.
3. The process continues for future calendar years, t , with the amount,

$$C_{t-1} \cdot (1 + i) - C_t,$$

being returned to (or being contributed by) the investor.

4. At some time $t = u$, the loss is deemed to be at the ultimate—that is, no significant changes in the loss are anticipated—and so we set $C_t = 0$ for $t > u$. For the examples in this paper, $u = 9$.

¹The provisions quoted here are stated in Section 2, Article 77 and Article 80, of Chapter VI of the act, p. 222, <http://register.consilium.europa.eu/pdf/en/09/st03/st03643-re01.en09.pdf>.

²As the risk margin is for the current liability, this paper does not consider new business in future calendar years.

The present value, discounted at the risky rate r , of the amount returned is equal to

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

Since $r > i$, this present value will be less than the initial capital investment of C_0 . To adequately compensate the investor for taking on the risk of insuring policyholder losses, the difference can be made up at time $t = 0$ by what we now define as the cost-of-capital risk margin, R_{COC} :

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

with the second equality coming after some algebraic manipulations.³

The problem that now needs to be addressed is the calculation of the C_t s. A straightforward way to project a future cash flow for this process would be to take a fitted Bayesian MCMC model and simulate an additional calendar year of losses for $t = 1$. Then fit another Bayesian MCMC model to the original data and the simulated data to get the loss estimate and capital requirements for $t = 1$. Then continue this process for $t = 2, \dots, u$.

While the execution speed of Bayesian MCMC software has significantly increased in recent years, repeating this for 10,000 simulated future cash flows would undoubtedly strain the patience of most practicing actuaries. This paper will propose a faster way to simulate the future cash flows to calculate the capital requirements for this process.

Now that we have defined the cost-of-capital risk margin, here is the route we will take to address the problems we must solve to calculate the risk margin:

- First we show how to use the Bayesian MCMC machinery to calculate the cash flows and corresponding loss estimates implied by the model.
- Then we show how to calculate the best estimate and the risk margins from the cash flows.
- Then we investigate the effect of insurer size and line of business on risk margins.
- Then we address the effect of diversification by line of business.

Whereas the examples in this paper focus on an “ultimate” time horizon, jurisdictions such as the European Union require insurers to calculate their capital requirements based on a one-year time horizon. The final section shows, with an example, how to adjust the models so that the one-year time horizon can be incorporated within the framework of this paper.

The data for the examples in this paper are taken from the Casualty Actuarial Society (CAS) Loss Reserve Database. The data consist of 50 loss triangles in the commercial auto (CA), personal auto (PA), workers’ compensation (WC), and other liability (OL) lines of insurance. The loss triangles used in this paper were selected from the list given in Appendix A of Meyers (2015).

The algorithms described in this paper are computationally intensive. The reader of this paper might question whether the computations can be done in a reasonable time. The answer is yes. The scripts that are included with the paper were run on the author’s standard issue laptop. The run times for the calculations are about two minutes per loss triangle for the model in Section 3 and about 17 minutes per triangle for the model in Section 5.

2. Cash flows and statistics of interest

This paper uses the changing settlement rate (CSR) model described in Meyers (2015) as modified in Meyers (2018). As those papers show, the model has been successfully validated on the lower triangle holdout data for a set of 200 loss triangles, 50 from

³Note that R_{COC} is similar to, but not identical to, the Solvency II risk margin: $R_{SII} \equiv (r-i) \cdot \sum_{t=0}^u \frac{C_t}{(1+i)^t}$.

each of four lines of business. The model describes the distribution of outcomes $X_{w,d} \sim \text{lognormal}(\mu_{w,d}, \sigma_d)$, where the accident year $w = 1, \dots, 10$ and the development year $d = 1, \dots, 11 - w$. It is fit to a cumulative paid loss triangle, $T_0 \equiv \{X_{w,d}\}$. This model allows for accident year effects, development year effects, and a variable claim settlement rate. The details of the model can be found in Section 3 of Meyers (2018). What is relevant for this paper is that given the loss triangle, T_0 , the model uses Bayesian MCMC to obtain a sample of 10,000 equally likely lognormal, $\{\mu_{w,d}^j, \sigma_d^j\}_{j=1}^{10,000}$, parameter sets from the posterior distribution, $\{\mu_{w,d}, \sigma_d | T_0\}$. This paper uses these parameter sets to describe the possible future cash flows by a simulation.

With these parameter sets we can calculate the best estimate of the liability, as specified by Solvency II, as the probability weighted average of the present value of expected future cash flows. This will be equal to the expected value of the differences in the cumulative payments between development years—that is,

$$E_{Best} = \frac{\sum_{j=1}^{10,000} \sum_{w=2}^{10} \sum_{d=12-w}^{10} \exp(\mu_{w,d}^j + (\sigma_d^j)^2/2) - \exp(\mu_{w,d-1}^j + (\sigma_{d-1}^j)^2/2)}{10,000 \cdot (1+i)^{w+d-11.5}} \quad (2)$$

This calculation assumes that the losses are paid one-half year before the end of future calendar year $t = w + d - 11$.

For the scope of this paper, let's also select the ultimate loss, U_j , associated with the j th parameter set to be the sum of the expected values of the losses for $d = 10$ over all the accident years—that is,

$$U_j = \sum_{w=1}^{10} \exp(\mu_{w,10}^j + (\sigma_{10}^j)^2/2) \quad (3)$$

For those wishing to allow for loss development after $d = 10$, we suggest that a Bayesian MCMC version of Clark (2003) would be a good place to begin.

For the lower triangle of $\{X_{w,d}^j\}_{j=1}^{10,000}$, define the simulated loss trapezoid for future calendar year t that includes the upper loss triangle, T_0 , and the first t diagonals of the lower loss triangle—that is,

$$T_t^j \equiv \begin{cases} X_{w,d} & \text{for } w = 1, \dots, 10 \text{ and } d = 1, \dots, 11 - w \\ X_{w,d}^j & \text{for } w = t + 1, \dots, 10 \text{ and} \\ & d = 12 - w, \dots, \min(11 - w + t, 10) \end{cases} \quad (4)$$

where $X_{w,d}^j$ is simulated from a lognormal distribution with parameters $\mu_{w,d}^j$ and σ_d^j .

Let's temporarily drop the assumption that we know the parameter set index j . All we have is an observed loss trapezoid, T_t . Then using Bayes' theorem and the fact that, initially, all j are equally likely, the probability that the parameter set index is equal to j given T_t , for $t > 0$, is given by

$$\Pr[J = j | T_t] = \frac{\prod_{X_{w,d} \in T_t} \phi(\log(X_{w,d}) | \mu_{w,d}^j \sigma_d^j)}{\sum_{k=1}^{10,000} \prod_{X_{w,d} \in T_t} \phi(\log(X_{w,d}) | \mu_{w,d}^k \sigma_d^k)} \quad (5)$$

where ϕ is the probability density function for the normal distribution.

At this point, one can choose from a number of options to calculate the various statistics that are of interest to insurer risk managers. For example, given T_t , one could calculate the ultimate loss estimate, E_t , as

$$E_t \equiv E \left[\sum_{w=1}^{10} X_{w,10} | T_t \right] = \sum_{j=1}^{10,000} \Pr[J = j | T_t] \cdot U_j. \quad (6)$$

If one accepts that the Bayesian MCMC output is representative of all future scenarios, Equation 6 is exactly the right calculation for the loss estimate given T_t . But let's consider what one should do to calculate, say, the 99.5th percentile. First one should sort the scenarios in order of increasing U_j . It is not uncommon to find a case where there is a scenario, j , with $\Pr[J \leq j | T_9] = 0.9900$ and $\Pr[J \leq j + 1 | T_9] = 0.9960$.

Algorithm 1. Calculate capital scenarios

```

1: for  $k = 1, \dots, 10,000$  do
2:   for  $t = 0, \dots, 9$  do
3:     Simulate cash flows  $\{T_t^k\}$  using the parameter set  $\{(\mu_{w,d}^k, \sigma_d^k)\}$ .
4:     Use Equation 5 to calculate  $\Pr[J = j|T_t^k]$  for each
        $j = 1, \dots, 10,000$ .
5:     Take a random sample of size 10,000 with replacement,
        $\{S_t^k\}$ , of  $\{U_j\}_{j=1}^{10,000}$  with sampling probabilities
        $\Pr[J = j|T_t^k]$ .
6:     Set  $E_t^k$  equal to the arithmetic average of  $\{S_t^k\}$ .
7:     Set  $C_t^k$  equal to the arithmetic average of the highest
        $(1 - \alpha) \cdot 10,000$  highest values of  $\{S_t^k\}$ , minus  $E_t^k$ .
8:   end for
9: end for
    
```

To deal with this, we decided to calculate the statistics of interest by first taking a random sample of size 10,000 (with replacement), $\{S_t\}$, of the U 's with sampling probabilities $\Pr[J = j|T_t]$. It was quite easy to implement and surprisingly fast in R. This is subject to an additional simulation error, but it should be small.

The statistics of interest for risk margin are, for $t = 0, \dots, 9$,

- the mean, E_t , which is equal to the arithmetic average of $\{S_t\}$; and
- the tail value at risk at the α level (TVaR@ α), which is equal to the arithmetic average of the $(1 - \alpha) \cdot 10,000$ highest values of $\{S_t\}$.⁴

Let's denote the total required capital by $C_t \equiv \text{TVaR@}\alpha - E_t$.


We summarize the above in Algorithm 1.

The examples in this paper use $\alpha = 97\%$. This selection is for illustrative purposes only.

Calculating E_t^k for $t = 0, \dots, 9$ yields the k th path that the loss estimate takes as it moves toward its ultimate value. Of interest for what follows is the set of possible paths that the loss estimate can take. Figure 2.1 shows those paths that contain the 100th, the 300th, . . . , and the 9,900th highest E_9^k s of Insurer #353 for commercial auto in the CAS Loss Reserve Database. This figure illustrates that the E_t^k s tend to become more certain over time.

⁴Although this paper does not use the value at risk (VaR) in its examples, one could calculate the VaR@ α as the $(1 - \alpha) \cdot 10,000$ th highest value of $\{S_t\}$.

Figure 2.1. Paths of ultimate loss estimates



Also of interest are the paths of the required capital, C_t^k , for $t = 0, \dots, 9$. Figure 2.2 shows the paths of C_t^k that correspond to the paths taken by E_t^k in Figure 2.1. This figure illustrates that as the estimates of the E_t^k s become more certain, the required capital, C_t^k , tends to decrease over time.

3. Risk margins

This section applies the cost-of-capital risk margin formula given by Equation 1 to the set of required capital paths, $\{C_0^k, \dots, C_9^k\}_{k=1}^{10,000}$. Recall that the formula defined the cost-of-capital risk margin as the present value of the capital released as the loss reserve liability becomes more certain. Figure 3.1 shows the paths of released capital that correspond to the paths taken by the C_t^k s in Figure 2.2. In general, this figure shows that most of the capital gets

Figure 2.2. Required capital by calendar year





Figure 3.1. Paths of released capital



released early on, and that occasionally it is necessary to add capital.

Applying Equation 1 we get for each k

$$R_{COC}^k \equiv C_0^k - \sum_{t=1}^u \frac{C_{t-1}^k \cdot (1+i) - C_t^k}{(1+r)^t} \quad (7)$$


Then the risk margin is given by

$$R_{COC} = \frac{1}{10,000} \sum_{k=1}^{10,000} R_{COC}^k \quad (8)$$

Figure 3.2 shows a histogram of the R_{COC}^k s for our example.

Of interest is the ratio of the risk margin and the size of the best estimate. To investigate, we calculated the risk margins for all 200 loss triangles in our data. After some exploratory analysis, we concluded that (1) significant differences exist by line of business; and (2) there is an approximate linear relationship

Figure 3.2. Risk margin



between the log of the risk margin and the log of the best estimate. Figure 3.3 shows the plots of the $\log(R_{COC})$ against $\log(E_{Best})$, along with the coefficients, along with their standard errors, of an ordinary linear regression of the form

$$\log(R_{COC}) = a + b \cdot \log(E_{Best}) \quad (9)$$

We can rewrite Equation 9 in the form

$$\frac{R_{COC}}{E_{Best}} = e^a \cdot (E_{Best})^{b-1} \quad (10)$$


Note from Figure 3.3 that $b < 1$ for all four lines of insurance. This implies that the risk margin-to-best estimate ratio decreases as the best estimate increases. As Figure 3.4 shows, the ratio can be quite high for insurers with small best estimates. It is easy to see where some insurers might object, especially if the line with the high ratio is a small part of the insurer’s book of business.

4. Diversification

As noted in the introduction, one provision of the European Union Solvency II act says explicitly, “Insurance undertakings shall segment their insurance obligations into homogeneous risk groups, and as a minimum by lines of business, when calculating the technical provisions.” This means that the total risk margin for a multiline insurer is the sum of the risk margins over its individual lines of business.

Longtime observers of the insurance business have recognized that multiline insurers benefit from the diversification of their risk of loss. This being the case, they might well want to reflect the benefits of diversification in their risk margins. The problem with a formal recognition of diversification is that the benefits have been difficult to quantify. What many are afraid of is the possibility that significant losses from the different lines of business could happen at the same time. This possibility is often referred to as the “dependency problem.”

Figure 3.3. $\log(\text{Risk Margin})$ vs. $\log(\text{Best Estimate})$ ⁵



As such, the Solvency II nonrecognition of diversification may appear to some to be prudent.

Mathematical tools that can be used to describe dependency have been available for quite some time. See, for example, Frees and Valdez (1998) and Wang (1998). The main tool described in those papers is called a copula, which is a multivariate distribution on an L -dimensional unit hypercube in which the marginal distributions have a uniform (0,1) distribution. Given a copula \mathcal{C} and samples $\{S_i^k\}$ (see Section 2), for each line l of L lines of busi-

ness one begins to calculate R_{COC} by first executing Algorithm 2.


Use the output of this algorithm to calculate $\{{}_T C_t^k\}_{k=1}^{10,000}$ for $t = 1, \dots, 9$ and Equations 7 and 8 to calculate ${}_T R_{COC}$.

Algorithm 2. Calculate samples for dependent lines

- 1: **for** $k = 1, \dots, 10,000$ **do**
- 2: **for** $t = 1, \dots, 9$ **do**
- 3: Simulate an L -tuple vector $\{P_i^k\}_{i=1}^L$ of uniform(0,1) numbers from the copula \mathcal{C} .
- 4: For each line of business, l , select Q_l^k to be the $P_l \cdot 10,000$ highest value of $\{S_i^k\}$.
- 5: **end for**
- 6: Set the total ultimate loss ${}_T S_t^k = Q_1^k + \dots + Q_L^k$.
- 7: **end for**

⁵Three small volatile insurers had negative best estimates and were excluded from the linear regression.

Figure 3.4. Risk margin ratio vs. best estimate⁶



So if one believes that the lines of business are correlated, it is possible to calculate the risk margin for the total liability that reflects whatever diversification one's choice of a dependency structure warrants. As it turns out, there has been some recent empirical work on determining that structure.

Let's first look at Avanzi, Taylor, and Wong (2016). The point of their paper is that correlations can arise from an inappropriate model. To quote their abstract: "We show with some real examples that, sometimes, most (if not all) of the correlation can be 'explained' by an appropriate methodology. Two major conclusions stem from our analysis.

1. "In any attempt to measure cross-LoB correlations, careful modeling of the data needs to be the order of the day. The exercise will not be well served by rough modeling, such as the use of simple chain ladders, and may indeed result in the prescription of excessive risk margins and/or capital margins.
2. "Such empirical evidence as examined in the paper reveals cross-LoB correlations that vary

only in the range zero to very modest. There is little evidence in favor of the high correlation assumed in some jurisdictions. The evidence suggests that these assumptions derived from either poor modeling or a misconception of the cross-LoB dependencies relevant to the purpose to which they are applied."

Meyers (2018) arrives at a similar conclusion. That paper first shows how to fit a bivariate CSR model, that allows for dependencies, to triangles for two lines of business from the same insurer. It then compares the fit of the bivariate model to a similar bivariate model that assumes independence for 102 within-insurer pairs. Taking into account the additional parameter introduced by the dependent model, it concludes that the model assuming independence has a better fit for all 102 pairs of triangles.

In other words, the appropriate dependency structure is to assume that the lines of business are independent. This assumes, as demonstrated in Meyers (2018) for the CSR model used in this paper, that careful modeling has been carried out.

The independence assumption allows us to simplify the procedure described at the beginning of this

⁶A small number of estimates fell outside the range of these figures.

section. Given the samples $\{S_l^k\}$, for each line l of L lines of business one begins to calculate R_{COC} by first executing Algorithm 3.

Use the sample, $\{S_l^k\}$, to obtain $\{C_l^k\}_{k=1}^{10,000}$ for $t = 1, \dots, 9$. Then use Equations 7 and 8 to calculate the combined risk margin, ${}_T R_{COC}$.

The combined risk margins in this paper were calculated using the independence assumption. This choice was *not* made for mathematical convenience. Meyers (2018) shows how to estimate the parameters of a model with dependency between the lines. The steps outlined at the beginning of this section show how to implement a dependency assumption if warranted.

From the loss triangles studied in Meyers (2015), five insurers had a loss triangle in all four lines. Table 4.1 gives the combined risk margin for those five insurers in the “Total” rows in the “Allocated Risk Margin” column. Over all five insurers, the diversification credit,

$$1 - \frac{\text{Combined Risk Margin}}{\text{Total Standalone Risk Margin}},$$

ranged from 30.3% to 48.3%.

Of interest, if not essential, is to see how this combined risk margin is allocated down to the individual lines of insurance. Allocating the cost of capital to individual lines is more important for pricing than for financial reporting as the former case requires an insurer to quote a price for an individual insurance contract. For the latter case, a risk margin need only apply to the total insurer liabilities.

Allocating the cost of capital has been debated in the actuarial profession for decades. About 15 years ago, a number of papers addressed the issue in a pricing context. Mango and Ruhm (2003) and Meyers

Algorithm 3. Calculate samples for independent lines

-
- 1: **for** $k = 1, \dots, 10,000$ **do**
 - 2: **for** $t = 1, \dots, 9$ **do**
 - 3: Set the total ultimate loss sample to be $\{S_l^k\} = \{S_1^k\} + \dots + \{S_L^k\}$.
 - 4: **end for**
 - 5: **end for**
-

(1999) are two of many papers that were published around then. Forgoing the seemingly endless discussion that accompanies this topic, this paper allocates combined capital to lines of insurance in proportion to each line’s marginal cost of capital.

Once one has done the coding necessary to calculate the combined risk margin, it takes only a little additional computer run time to allocate the combined risk margin to individual lines. So let’s proceed.

Given the samples, $\{S_l^k\}$, for each line l of L lines of business one begins to calculate marginal cost of capital for line l , ${}_{(l)}R_{COC}$, by first executing Algorithm 4 below. Then for each line l execute Algorithm 5.

The fourth column of Table 4.1 gives the marginal cost of capital, ${}_{(l)}R_{COC}$, by insurer for each line of insurance. Note that the sum of the marginal cost of capitals by line is less than the combined cost of capital in the “Total” column. We then allocate the cost of capital by line of insurance in proportion to the marginal capital by

$${}_{(l)}R_{ACOC} \equiv {}_{(l)}R_{COC} \cdot \frac{{}_T R_{COC}}{{}_{(1)}R_{COC} + \dots + {}_{(L)}R_{COC}} \quad (11)$$

Note that there are many instances where the diversification credit is in excess of 80%. This occurs

Algorithm 4. Calculate leave-line-out samples

-
- 1: **for** $k = 1, \dots, 10,000$ **do**
 - 2: Set the total ultimate loss sample to be $\{S_l^k\} = \{S_1^k\} + \dots + \{S_L^k\}$.
 - 3: **for** $t = 1, \dots, 9$ **do**
 - 4: Set the leave-line-out ultimate loss sample for line l to be $\{S_{(-l)}^k\} = \{S_l^k\} - \{S_l^k\}$.
 - 5: **end for**
 - 6: **end for**
-

Algorithm 5. Calculate marginal cost of capital

-
- 1: **for** $l = 1, \dots, L$ **do**
 - 2: **for** $t = 1, \dots, 9$ **do**
 - 3: Use the sample, $\{S_{(-l)}^k\}$, to calculate the leave-line-out capital, $\{C_{(-l)}^k\}_{k=1}^{10,000}$.
 - 4: **end for**
 - 5: Use Equations 7 and 8 to calculate the leave-line-out cost-of-capital risk margin, ${}_{(-l)}R_{COC}$.
 - 6: Calculate the marginal cost-of-capital risk margin, ${}_{(l)}R_{COC} \equiv {}_T R_{COC} - {}_{(-l)}R_{COC}$.
 - 7: **end for**
-

Table 4.1. Risk margins by type and line of business

Grp./Line	Estimated Ult. Loss	Best Estimate	Marginal Risk Margin	Allocated Risk Margin	Stand-Alone Risk Margin	Diver. Credit
1528/CA	88,756	13,822	464	852	1,447	41.1%
PA	311,659	36,507	542	996	1,519	34.4%
WC	129,762	13,207	61	111	550	79.8%
OL	19,143	4,697	243	447	1,065	58.0%
Total	549,320	68,233	1,310	2,406	4,581	47.5%
1767/CA	2,205,897	310,203	108	171	5,963	97.1%
PA	90,312,996	9,921,107	20,620	32,566	76,527	57.4%
WC	1,677,179	227,010	175	276	5,637	95.1%
OL	2,443,660	956,344	76,495	120,812	132,428	8.8%
Total	96,639,732	11,414,664	97,398	153,825	220,555	30.3%
3240/CA	97,298	18,684	346	554	1,653	66.5%
PA	1,092,757	136,373	1,862	2,983	3,208	7.0%
WC	38,960	4,155	42	67	467	85.7%
OL	13,774	2,638	36	58	459	87.4%
Total	1,242,789	161,850	2,286	3,663	5,787	36.7%
5185/CA	96,071	23,262	837	1,592	2,544	37.4%
PA	268,908	43,305	952	1,811	2,719	33.4%
WC	100,322	16,768	91	173	1,100	84.3%
OL	140,606	22,440	216	410	1,346	69.5%
Total	605,907	105,775	2,095	3,985	7,709	48.3%
14176/CA	28,929	11,759	982	1,716	2,191	21.7%
PA	144,563	26,494	380	663	1,439	53.9%
WC	111,498	20,075	263	460	1,229	62.6%
OL	5,290	1,287	35	61	326	81.3%
Total	290,280	59,615	1,660	2,900	5,185	44.1%

Note: CA = commercial auto; PA = personal auto; WC = workers' compensation; OL = other liability.

when a “small” line of insurance is part of the portfolio of a “large” insurer. Regardless of what one thinks of allocating the cost of capital, one cannot deny that a small line of insurance adds little to the risk of a large insurer. The insurer size effect illustrated in Figure 3.4 can be significantly reduced by taking diversification into account.

5. One-year time horizon

In the risk margin calculations above, we assumed an “ultimate” time horizon to establish the required capital. Some regulatory jurisdictions, such as Sol-

vency II, specify that the insurer should assume a one-year time horizon. This section extends the methodology of the previous sections to cover the one-year time horizon.

A high-level description of the methodology is to use a Bayesian MCMC model to obtain 10,000 equally likely scenarios that represent the future evolution of the line of business that produced the loss triangle. Then, as new losses come in, one uses Bayes' theorem to update the probability of each scenario. From those updated probabilities, one can then calculate the statistics that are needed to calculate the risk margin.

Under a one-year time horizon capital requirement, the capital is determined by the estimate of the ultimate losses after one more calendar year of loss experience. A key step in this methodology is to determine the ultimate loss estimate associated with each scenario. For the ultimate time horizon, it is simply U_j . However, as Figure 2.1 illustrates, with only one year of losses from the j th scenario, there may be several scenarios with a significant positive probability.

To get a good approximation, $O_{t,j}$, of the expected ultimate loss for the j th scenario, one can simulate future loss experience from the parameter set of that scenario and calculate the ultimate loss estimate, M times. Then set $O_{t,j}$ equal to the average of those estimates. Algorithm 6 shows the details.

Both the accuracy of the estimate of $O_{t,j}$ and the computer run time increase with M . After experimenting with different values of M , we found that $M = 12$ obtained results that were sufficiently accurate given the intrinsic variation of the underlying MCMC simulation.

Algorithm 7 is used to calculate the risk margin for the one-year time horizon. In this algorithm, one simply substitutes $O_{t+1,j}$ for U_j in the fifth step of Algorithm 1. Given the output of Algorithm 7, one then calculates risk margins using Equations 7 and 8.

Figures 5.1, 5.2, and 5.3 show the one-year time horizon capital paths, release paths, and risk margins of Insurer #353 for commercial auto that correspond

Algorithm 6. Calculate scenario estimates by calendar year

```

for m = 1, . . . , M do
  for j = 1, . . . , 10,000 do
    for t = 1, . . . , 9 do
      Simulate  $T_t$  using the parameters  $(\mu_{w,d}^k, \sigma_d^k)$ .
      Use Equation 5 to calculate  $\{\Pr [J = j|T_t]\}_{j=1}^{10,000}$ .
      Use Equation 6 to calculate the ultimate loss estimate,  $O_{t,k}^m$ .
    end for
    Set  $O_{10,j}^m = O_{9,k}^m$ 
  end for
end for
for j = 1, . . . , 10,000 do
  for t = 1, . . . , 10 do
    Set  $O_{t,j} = \text{mean}(O_{t,k}^m)$ .
  end for
end for

```

Algorithm 7. Calculate capital scenarios for a one-year time horizon

```

for k = 1, . . . , 10,000 do
  for t = 0, . . . , 9 do
    Simulate cash flows  $\{T_t^k\}$  using the parameter set  $\{(\mu_{w,d}^k, \sigma_d^k)\}$ 
    Use Equation 5 to calculate  $\Pr [J = j|T_t^k]$  for each  $j = 1, . . . , 10,000$ 
    Take a random sample of size 10,000 with replacement,  $\{S_t^k\}$ , of the  $\{O_{t+1,j}\}_{j=1}^{10,000}$  with sampling probabilities  $\Pr [J = j|T_t^k]$ .
    Set  $E_t^k$  equal to the arithmetic average of  $\{S_t^k\}$ .
    Set  $C_t^k$  equal to the arithmetic average of the highest  $(1 - \alpha) \cdot 10,000$  highest values of  $\{S_t^k\}$ , minus  $E_t^k$ .
  end for
end for

```

Figure 5.1. Required capital by calendar year






Figure 5.2. Paths of released capital



to Figures 2.2, 3.1, and 3.2, respectively, for the ultimate time horizon.

6. Concluding remarks

There has been no universal agreement on the assumptions underlying a cost-of-capital risk margin formula. Beyond the underlying Bayesian MCMC

Figure 5.3. Risk margin

stochastic loss reserve model, this paper makes the following key assumptions:

1. The required assets for an insurer are determined by the TVaR@ α measure of risk.
2. The required capital calculation assumes an “ultimate” time horizon.
3. The outcome distributions for the different lines of business are independent.

In numerous advisory committee meetings held at International Actuarial Association events, the author has heard the following argument supporting the one-year time horizon: Insolvency is usually not an instantaneous event. If the insurer finds itself under stress within a year, it will have time to make the necessary adjustments.

At the same meetings the following heuristic definition of a risk margin was also heard: The risk margin is to provide sufficient funds to transfer the insurer’s liability to another insurer. “Sufficient funds” should include the cost of capital.

The following considerations govern our approach to risk margins:

1. The term of such a portfolio risk transfer contract is unlikely to be for a single year, with the risk reverting back to the original insurer at the end of the year. This being the case, we used an ultimate time horizon to determine the capital requirements.

2. For a multiline insurer, the risk being transferred is unlikely to consist of a single line of insurance.
3. Dependency between lines is model dependent. In Meyers (2018) the author demonstrated that the independence assumption is warranted for the CSR model used in this paper.
4. The theoretical advantages of the TVaR@ α over the VaR@ α have been well documented by Artzner et al. (1999). Whatever computational difficulty there may have been with the TVaR is not an issue with the methodology used in this paper.

Recognizing that reasonable people may differ in their assumptions, this paper points the way to use alternative assumptions. The methodology described in this paper can be readily adopted for any Bayesian MCMC model.

Appendix

Included with this paper is a [zip archive](#) containing the following:

- RM 1Line.R—The script that produces the risk margin calculations in Sections 2 and 3
- RM 4Line.R—The script that produces the risk margin calculations in Section 4
- RM 1Line 1yr.R—The script that produces the risk margin calculations in Section 5
- Risk Margins for 200 Triangles.xlsx—Risk margin single-line calculations for all 200 triangles

The computer language for the scripts is R (<https://www.r-project.org>.) The computer language for the MCMC calculations is Stan (<http://mc-stan.org/interfaces/rstan.html>.)

References

- Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath, “Coherent Measures of Risk,” *Mathematical Finance* 9(3), 1999, pp. 203–228, <https://people.math.ethz.ch/~delbaen/ftp/preprints/CoherentMF.pdf>.
- Avanzi, B., G. Taylor, and B. Wong, “Correlations Between Insurance Lines of Business: An Illusion or a Real Phenomenon?”

- Some Methodological Considerations,” *ASTIN Bulletin* 46(2), 2016, pp. 225–263, doi: 10.1017/asb.2015.31, <https://www.cambridge.org/core/journals/astin-bulletin-journal-of-the-iaa/article/correlations-5BBB6C52A882572B4F4D587984ED10A7>.
- Clark, D. R., “LDF Curve-Fitting and Stochastic Reserving, A Maximum Likelihood Approach,” *Casualty Actuarial Society Forum*, Fall 2003, <http://www.casact.org/pubs/forum/03fforum/03ff041.pdf>.
- Frees, E. W., and E. A. Valdez, “Understanding Relationships Using Copulas,” *North American Actuarial Journal* 2(1), 1998, pp. 1–25, https://www.soa.org/Library/Journals/NAAJ/1998/january/naaj9801_1.aspx.
- Mango, D. F., and D. L. Ruhm, “A Method of Implementing Myers-Read Capital Allocation in Simulation,” *Casualty Actuarial Society Forum*, Fall 2003, <http://www.casact.org/pubs/forum/03fforum/03ff451.pdf>.
- Meyers, G. G., “Underwriting Risk,” *Casualty Actuarial Society Forum*, Spring 1999, pp. 185–220, <http://www.casact.org/pubs/forum/99spforum/99spf185.pdf>.
- Meyers, G. G., “Stochastic Loss Reserving Using Bayesian MCMC Models,” monograph no. 1, Arlington, VA: Casualty Actuarial Society, 2015, <http://www.casact.org/pubs/monographs/index.cfm?fa=meyers-monograph01>.
- Meyers, G. G., “Dependencies in Stochastic Loss Reserve Models,” *Variance* 11, 2018, pp. 74–94.
- Wang, S. S., “Aggregation of Correlated Risk Portfolios: Models and Algorithms,” *Proceedings of the Casualty Actuarial Society* 85, 1998, pp. 848–939, <http://www.casact.org/pubs/proceed/proceed98/980848.pdf>.