MOODY'S ANALYTICS

RESEARCH
AUGUST 2014
INSURANCE ERS

Steven Morrison PhD
Laura Tadrowski PhD
Moody's Analytics Research

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Capital Attribution Methods

Overview

Having calculated an overall capital requirement, insurers are often interested in quantifying how this overall capital can be attributed to:

- Sub-portfolios (for example business units, geographical locations or product types)
- Risk factors

This note describes some of the more popular methods for capital attribution by sub-portfolio, their estimation using Monte Carlo scenarios, and the statistical error in these estimates. While our previous paper on VaR estimation (Morrison and Tadrowski 2013) indicated that overall capital requirements can be relatively insensitive to choice of estimation scheme when using large scenario sets, this is not the case for one of the most popular attribution schemes (the continuous marginal method). In this case, empirical estimators suffer from a large amount of statistical error even using large scenarios sets, and there is a significant benefit in using smoothed estimates (e.g. Harrell-Davis).

While there is a large amount of literature on attribution by sub-portfolio, the question of attribution by risk factor appears to have received far less attention. In this note we describe a potential approach to this problem that has recently been proposed, and illustrated by applying a similar technique to attribution of capital by risk factor for a simple life insurance product.

Contents

1. Introduction

In a previous paper (Morrison and Tadrowski 2013) we discussed the estimation of insurers' capital requirements using Monte Carlo simulation, comparing performance of various statistical VaR estimators. In addition to estimating overall capital requirements, firms are often interested in attributing this to underlying sub-portfolios (e.g. business units) or across the different risk factors to which the business is exposed.

A number of different methods for capital attribution by sub-portfolio have been proposed by academics and practitioners. In this note we aim to demonstrate, via a life insurance case study, how some of the more common methods can be applied in a Monte Carlo framework. As for overall capital, different statistical estimators can be applied to the capital attribution problem. Our case study compares the performance of two different statistical estimators.

This note also introduces the problem of attribution by risk factor. While this question is also of great interest to insurers, it appears to have received relatively little attention in the academic literature. Here we present the ideas of a recent paper (Rosen and Saunders 2010) which provides a potential solution to this problem, and illustrate by applying a similar technique to attribution of capital by risk factor for a simple life insurance product.

The rest of the paper is structured as follows:

- Section 2 defines and compares some of the more popular methods for attributing capital to sub-portfolios.
- Section 3 discusses the estimation of contributions to VaR capital using Monte Carlo simulation.
- Section 4 presents a case study, in which various capital attribution methods are applied to a simple life insurance business, and statistical error quantified under different Monte Carlo estimation methods.
- Section 5 discusses the attribution of capital to risk factors.
- Section 6 concludes.

2. Capital attribution by sub-portfolio

Having calculated an overall capital requirement for some portfolio, insurers are often interested in quantifying how this overall capital can be attributed to sub-portfolios, for example business units, geographical locations or product types. The ability to attribute capital to individual sub-portfolios is useful for various management purposes, for example in the assessment of their relative risk-adjusted performance.

A number of different attribution methods have been proposed by academics and practitioners. In this section we describe some of the more popular methods.

Let $EC = \rho(L)$ denote the capital requirement of a portfolio with total loss L, where ρ is some risk-measure e.g. VaR or CTE.

We consider the situation where the total loss L can be broken down into a sum of losses over n sub-portfolios:

$$L = \sum_{i=1}^{n} L_i$$

Let EC_j denote the capital attributed to portfolio j. Figure 1 summarises some of the more popular definitions for EC_j.

Figure 1: Capital attribution measures

Measure	Description	Mathematical definition
Standalone capital	The capital requirement of an individual sub-portfolio in isolation.	$EC_j^{standalone} = \rho(L_j)$
Pro-rata contribution	The standalone capital scaled so that the sum of contributions equals the total capital requirement of the portfolio.	$EC_j^{pro-rata} = EC_j^{standalone} \times \frac{EC}{\sum_{j=1}^{n} EC_j^{standalone}}$
Discrete marginal contribution	The difference between the total capital requirement and the capital requirement after removing the sub-portfolio.	$EC_j^{discrete\ marginal} = \rho\left(\sum_{i=1}^{n} L_i\right) - \rho\left(\sum_{\substack{i=1 \\ (i \neq j)}}^{n} L_i\right)$
Scaled marginal contribution	The discrete marginal contribution scaled so that the sum of contributions equals the total capital requirement of the portfolio.	$EC_j^{scaled\ marginal} = EC_j^{discrete\ marginal} \times \frac{EC}{\sum_{j=1}^{n} EC_j^{discrete\ marginal}}$
Continuous marginal contribution	The sensitivity of the total capital requirement to a small additional exposure to the sub-portfolio.	$EC_j^{continuous\ marginal} = \frac{\partial}{\partial h} \rho\left(\sum_{i=1}^{n} L_i + hL_j\right)\Big\|_{h=0}$

In some applications it is desirable for the capital contributions to sum to the total capital requirement:

$$\sum_{j=1}^{n} EC_j = EC$$

This is sometimes known as the *full allocation property* (McNeil, Frey and Embrechts 2005).

Note that standalone capital requirements and discrete marginal contributions don't satisfy the full allocation property. In particular, the sum of standalone capital requirements doesn't (in general) equal the total capital requirement:

$$\sum_{j=1}^{n} EC_j^{standalone} \neq EC$$

This difference arises due to diversification of risk between different sub-portfolios[1]. The pro-rata contribution scales the standalone capital so that the full allocation property is met i.e. assumes that the 'diversification benefit' is attributed pro-rata according to standalone capital requirements. This is not necessarily a sensible allocation of diversification benefit – the amount of diversification benefit provided by a particular sub-portfolio intuitively depends on its dependency with the rest of the portfolio and not just its standalone risk.

Similarly, the discrete marginal method doesn't satisfy the full allocation property, though it can be scaled to achieve this (the scaled marginal contribution in Figure 1). Despite not satisfying full allocation, the discrete marginal contribution may be a useful measure in practice, indicating the reduction in capital that can be achieved by completely removing a particular sub-portfolio (or the additional capital created by adding it).

The method that appears to have achieved the greatest attention in the academic literature is the continuous marginal contribution. It has a number of desirable properties, in particular:

1. It naturally satisfies the full allocation principle[2]:

$$\sum_{j=1}^{n} EC_j^{continuous\ marginal} = EC$$

2. It is the only method satisfying the full allocation principle that is "suitable for performance measurement"(Tasche 1999), in the following sense: Using the continuous marginal contribution as a measure of risk contribution, if the Return on Risk Adjusted Capital (RORAC)[3] is better (worse) for a particular sub-portfolio than the overall portfolio then slightly increasing (decreasing) the size of that portfolio improves the overall performance of the portfolio. Other contribution methods (for example the scaled marginal method) do not satisfy this property.
3. In the specific case of the VaR risk measure, the continuous marginal contribution can be written as the expected loss on the sub-portfolio, conditional on the total loss equaling the total capital requirement[4]:

$$EC_j^{continuous\ marginal} = E\left[L_j | L = EC\right]$$

This alternative representation is useful in estimating the continuous marginal contribution by Monte Carlo simulation, as we will see below.

[1] Note also that VaR is not 'sub-additive' i.e. the sum of standalone capital requirements can in certain circumstances be *smaller* than the total capital requirement, implying that diversification creates a cost rather than benefit (though this is not the case for coherent risk measures like CTE).

[2] This is a result of Euler's Therorem. As a result, the continuous marginal contribution is sometimes referred to as the Euler allocation.

[3] The total portfolio RORAC is defined as the expected gain on the portfolio divided by eh total capital requirement, while the RORAC of a sub-portfolio is the expected gain on that sub-portfolio divided by its risk contribution.

[4] A similar representation exists for the CTE risk measure. In this case: $EC_j^{continuous\ marginal} = E\left[L_j | L \geq EC\right]$.

3. Estimation of VaR capital contributions using Monte Carlo simulation

In a previous paper (Morrison and Tadrowski 2013) we discussed the estimation of VaR using Monte Carlo simulation, comparing performance of various statistical VaR estimators. In addition to estimating overall capital requirements, these estimators can be used to measure the VaR of individual sub-portfolios or combinations of sub-portfolios, thus allowing estimation of standalone capital requirements and discrete marginal contributions (and hence their scaled versions).

To measure continuous marginal VaR contributions, we use the representation of the contribution as a conditional expectation, $EC_j^{continuous\ marginal} = E[L_j|L = EC]$. Suppose we have generated N for the subportfolio losses. We can estimate the conditional expectation as follows:

1. Measure the overall VaR capital requirement.
2. Select those scenarios with loss equal to the overall capital requirement.
3. Take the sample average of sub-portfolio losses on the selected scenarios.

The problem with such an approach is that losses are usually continuous so that, for any finite number of scenarios N, the number of scenarios with loss exactly equal to the estimated capital requirement is likely to be (at most) one. The resulting sample average is thus a very poor estimator of the conditional expectation.

For example, if the overall capital requirement is estimated using the 'empirical estimator', there is a single scenario with loss equal to the overall capital requirement, sometimes referred to as the 'biting scenario', and the resulting capital contribution estimates are therefore just the losses on individual sub-portfolios in the single biting scenario.

In order to increase the sample size, an alternative approach might be to select scenarios with losses lying in some range around the total VaR. By including a number of scenarios, we reduce the statistical noise in the estimate, at the potential expense of introducing bias (as we include scenarios with losses different from the total VaR). Selection of the size of the range thus involves a tradeoff between statistical variance and bias.

More generally, we can consider calculating a weighted sample average, with the weight put on individual scenarios depending on the size of the total loss relative to the total capital requirement. Such weighting schemes were discussed in our previous note on VaR estimation (Morrison and Tadrowski 2013) and their application to capital contribution estimation is discussed in (Epperlein and Smillie 2006).The method is illustrated in Figure 2, which plots scenarios for the joint loss on two sub-portfolios, under four different scenario sets each containing 10,000 scenarios. The two sub-portfolios here represent a Guaranteed Annuity Option portfolio and a Fixed Annuity portfolio, and are analysed further in Section 4.

For each of the four different scenario sets we highlight the 'biting scenario' in solid blue. The biting scenario is clearly very different for the four scenario sets – given our estimated total VaR of ~33,000, the biting scenario could lie anywhere on the straight line "*loss on fixed annuity portfolio + loss on GAO portfolio = 33,000*", with the exact location varying depending on the particular set of 10,000 scenarios generated. These biting scenarios provide estimates of the continuous capital contributions for the GAO and fixed annuity portfolios, albeit noisy estimates.

Rather than putting all weight on the single (noisy) biting scenario, suppose that we weight according to the Harrell-Davis weighting scheme. These weights are indicated by the size (area) of the circles representing the scenarios[5]. If we subsequently weight losses according to these weights, the weighted average losses are indicated by the solid red circles. These weighted average losses provide alternative estimates of the continuous capital contributions for the GAO and fixed annuity portfolios, which are more stable than the single biting scenarios as they are calculated as a weighted average over a number of scenarios.

[5] Note that 10,000 scenarios were used in each case here, but only a small number of scenarios have large enough Harrell-Davis weights to be visible in Figure 2.

Figure 2: Estimation of continuous marginal contributions for four different scenario sets

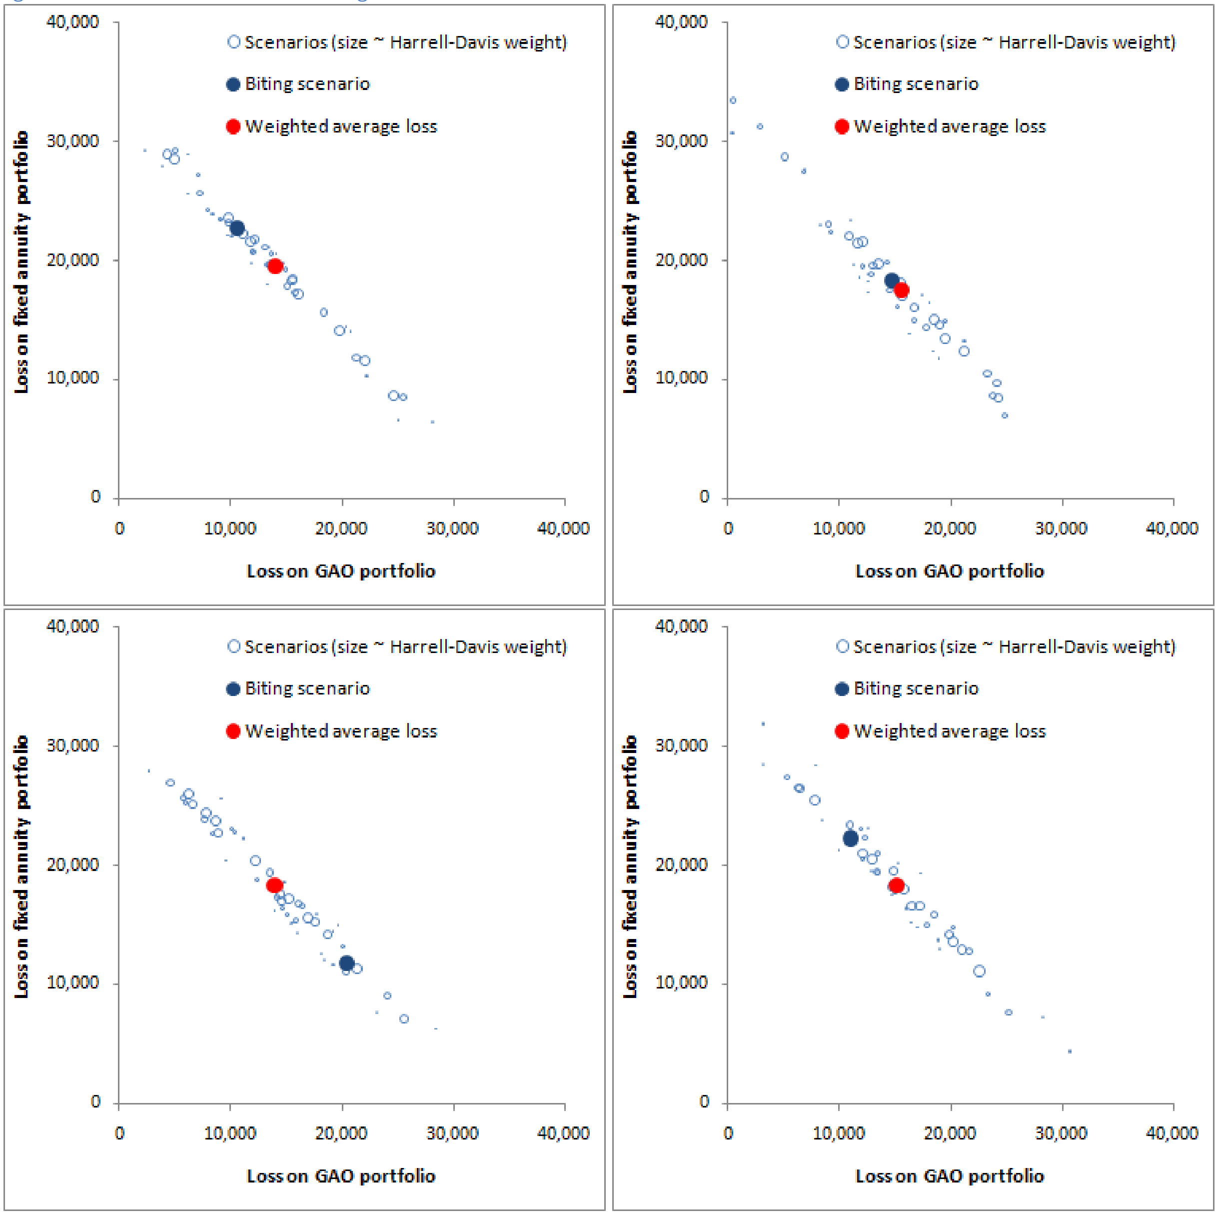

4. Case study

To illustrate, we consider a simple life insurance business with two sub-portfolios:

- With profits business with a Guaranteed Annuity Option, with underlying assets invested in a mix of equities and Government bonds. This business is sensitive to changes in risk-free interest-rates and equity returns.
- Non profit fixed annuity business, with backing assets in corporate bonds. This business is sensitive to longevity, credit returns and, to a lesser extent, changes in risk-free interest-rates (as we assume that backing assets are broadly cash flow matched).

Note that the losses on each of these sub-portfolios are approximately independent, as the only common risk factors are risk-free interest-rates, and the annuity business is relatively insensitive to these.

Figure 3 and Figure 4 show estimates of the total capital requirement and contributions of the two sub-portfolios using the different methods described in Section 2. All estimates were calculated using 100,000 scenarios, with empirical (i.e. biting scenario) estimators shown in Figure 3 and smoothed (Harrell-Davis) estimators shown in Figure 4.

The sizes of these portfolios were chosen such that their standalone capital requirements, and therefore their pro-rata contributions, are approximately equal. As a consequence, the discrete marginal contributions (and hence scaled marginal contributions) are also approximately equal[6]. The sum of standalone capital requirements exceeds the total capital requirement (and hence the sum of discrete marginal contributions is lower than the total capital requirement) by 12,571 (the 'diversification benefit'). Hence the standalone capital requirements need to be scaled down (and discrete marginal contributions scaled up) in order for them to sum to the total capital requirement. In contrast, the continuous marginal contributions always meet the full allocation property.

Estimates of total and standalone capital requirements (and hence pro-rata and discrete marginal contributions) are similar for both empirical and Harrell-Davis estimators. However continuous marginal contribution estimates vary considerably between empirical and Harrell-Davis estimators. Given the discussion in Section 3, this is not surprising – the empirical estimator is expected to be subject to a large amount of statistical error (while the Harrell-Davis estimate is expected to be far more accurate).

Figure 3: Estimated capital attributions using empirical estimators

	Standalone	Pro-rata contribution	Discrete marginal contribution	Scaled marginal contribution	Continuous marginal contribution
GAO portfolio	23,255	16,684	10,684	17,204	21,307
Fixed annuity portfolio	22,490	16,310	9,918	15,970	11,867
Sum	45,745	33,174	20,602	33,174	33,174
Total capital requirement	33,174				

Figure 4: Estimated capital attributions using smoothed (Harrell-Davis) estimators

	Standalone	Pro-rata contribution	Discrete marginal contribution	Scaled marginal contribution	Continuous marginal contribution
GAO portfolio	23,248	16,863	10,685	17,916	14,346
Fixed annuity portfolio	22,494	16,316	9,931	15,983	18,833
Sum	45,742	33,179	20,616	33,179	33,179
Total capital requirement	33,179				

[6] Note that since we just have two sub-portfolios, the discrete marginal contribution of the one portfolio is just the total capital requirement minus the standalone contribution of the other portfolio, and the sum of discrete marginal contributions is the total capital requirement minus the diversification benefit.

To quantify the amount of statistical error, we have repeated this exercise many times and measured the resulting variability in estimates. Figure 5 shows the average standalone capital requirements and continuous marginal contributions calculated using 80 independent sets of 100, 000 scenarios each, along with standard deviations (in brackets). The standalone capital requirements are relatively stable (with standard deviations < 1% of the average in all cases) and the standard deviations of a similar size for empirical and Harrell-Davis estimators. These results are consistent with the conclusions of our previous paper (Morrison and Tadrowski 2013), which concluded that empirical and Harrell-Davis estimators are likely to have similar levels of statistical accuracy when estimated using such large scenario sets.

However, when it comes to estimating continuous marginal contributions, the Harrell-Davis estimator is far more stable than the empirical estimator, as expected given the discussion in Section 3. The 80 different empirical and Harrell-Davis estimators of continuous marginal distributions are plotted in Figure 8, which further illustrates the increased stability of the Harrell-Davis estimates.

Note also that the Harrell-Davis estimates of continuous marginal contributions, although more stable than the corresponding empirical estimators, are still relatively 'noisy' in comparison with estimates of standalone capital requirements. The standard deviations of the continuous marginal contributions of the GAO and fixed annuity portfolios are 4.6% and 3.6% of the corresponding averages. This indicates that a far larger number of scenarios may be required in order to give a comparable level of statistical accuracy in estimates of continuous marginal contributions as for standalone capital measures.

Figure 5: Averages and standard deviations of estimated capital requirements and continuous marginal contributions (80 scenario sets)

	Standalone capital requirement		Continuous marginal contribution	
	Empirical	Harrell-Davis	Empirical	Harrell-Davis
GAO portfolio	23,318 (122)	23,323 (121)	14,978 (5,063)	14,351 (655)
Fixed annuity portfolio	22,601 (168)	22,609 (164)	18,113 (5,040)	18,742 (679)
Total capital requirement	33,091 (194)	33,095 (187)		

Figure 6: Continuous marginal contributions

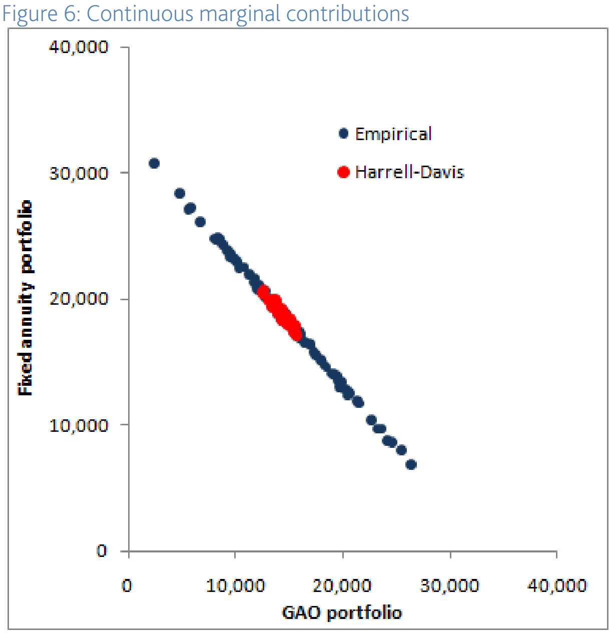

5. Capital attribution by risk factor

So far we have considered the problem of attributing capital requirements to sub-portfolios (for example individual business units). Another problem of interest to insurance groups is the attribution of capital based on an assessment of the contribution of individual *risk factors* (e.g. interest-rates, equity index levels, mortality rates etc). While there is a large amount of literature on the first of these problems, relatively little has been written on the second.

The methods for capital attribution defined above are based on the assumption that the total loss on a portfolio can be decomposed as a sum of losses on the sub-portfolios. All of the capital contribution definitions in Figure 1 involve increasing or decreasing the exposure to individual sub-portfolios while leaving exposure to the other sub-portfolios unchanged. Thus if the portfolio loss can be written as a sum of losses on sub-portfolios, each of which is exposed to a single risk factor, we could directly apply the above techniques to allocate capital requirements by risk factor. In practice things are not so simple, as the loss on a typical insurance portfolio contains *interactions* between individual risk factors i.e. the impact of a change in one of the risk factors will depend on the values of other risk factors.

As an illustrative example, consider an insurer who has sold a guaranteed equity bond, with a 10 year maturity and guarantee of 1.5, and invested assets in the underlying equity index. The net asset position is equivalent to a short vanilla put option, and the loss (over a one year horizon) is[7]:

$$L(S,R,\sigma) = V_{PUT}(S,R,\sigma) - V_{PUT}(S_0,R_0,\sigma_0)$$

where S, R, σ are the values of the three relevant risk factors at the one year horizon: the equity index, implied volatility and risk-free interest-rate respectively, S_0, R_0, σ_0 are their initial values, and V_{PUT} is the Black-Scholes formula for a put option.

This loss function contains interactions between all three risk factors. For example, Figure 7 shows the loss as a function of equity index for various values of interest-rate and fixed implied volatility (20%)[8].

Figure 7: Example loss function for guaranteed equity bond ($\boldsymbol{S_0 = 1, R_0 = 5\%, \sigma = \sigma_0 = 20\%}$)

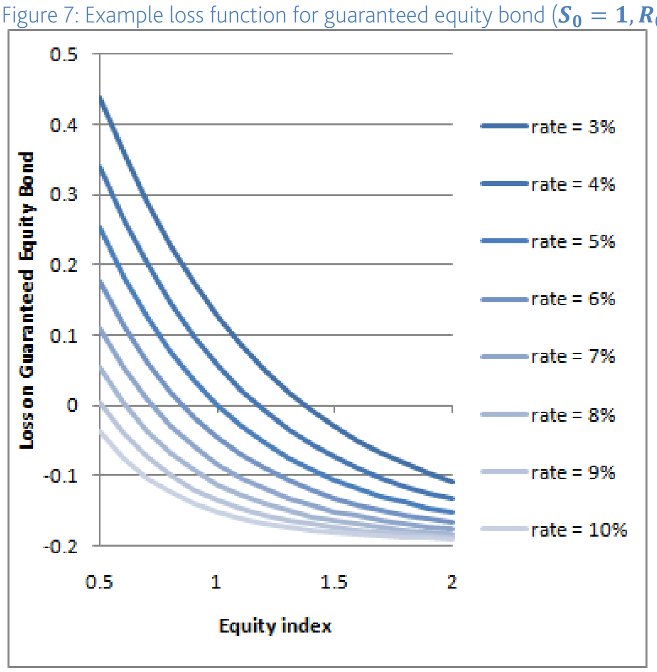

Clearly the loss function contains interaction between the equity index and risk-free rates: the price of a put option is more sensitive to the equity index when interest-rates are low than when they are high. As a result we cannot decompose the total loss as a sum over sub-portfolios that depend on single risk factors alone.

[7] Here we adopt a 'total balance sheet' definition of loss i.e. $Loss = -(Net\ Assets\ (1) - \ Net\ Assets\ (0))$, as in the Solvency 2 capital requirement for example.

[8] For simplicity, we assume that one year risk factor changes are applied *instantaneously*, which is consistent with the approach adopted by many insurance groups in practice.

In decomposing the loss as a sum over sub-portfolios, it can be useful to think about the sub-portfolios as representing hedges for individual risks or combinations of risks. For example, consider the equity sub-portfolio consisting of assets that only depend on the one year equity index (for example vanilla put options with a one year maturity). In the current problem, it is impossible to construct such a portfolio that perfectly hedges all equity risk. Any attempt to construct such a portfolio will always leave some exposure to equity risk due to its interaction with other risk factors.

Although we cannot construct a 'perfect' equity hedge using instruments that only depend on the equity index, we can construct equity hedges that are in some sense 'optimal'. Similarly we can construct optimal rate and implied volatility hedges. Such optimal hedge portfolios are natural sub-portfolios in a decomposition of the total loss though, as noted above, any such decomposition inevitably leaves some residual loss that depends on combinations of risk factors.

Below we consider two possible loss decompositions based on alternative definitions of the optimal hedge portfolios.

Optimal hedge portfolio definition (1): Optimal quadratic hedging

The idea of decomposing a portfolio in terms of optimal hedge portfolios for the purpose of risk decomposition was introduced in (Rosen and Saunders 2010). In this paper, a specific mathematical decomposition known as the Hoeffding decomposition is suggested. For the guaranteed equity bond example, the Hoeffding decomposition is as follows:

$$L(S,R,\sigma) = L_{cash} + L_S(S) + L_R(R) + L_\sigma(\sigma) + L_{S,R}(S,R) + L_{S,\sigma}(S,\sigma) + L_{R,\sigma}(R,\sigma) + L_{S,R,\sigma}(S,R,\sigma)$$

where:

- $L_{cash} = E[L(S,R,\sigma)]$

  This is a constant term, independent of all risk factors. It represents the 'optimal cash hedge' - the optimal hedge portfolio that can be constructed using cash alone.

- $L_S(S) = E[L(S,R,\sigma)|S] - L_{cash}$

  This term depends only on S. It represents the 'optimal equity only hedge' – the optimal hedge of remaining risk (after implementing the cash hedge) that can be constructed using instruments that only depend on the equity index S.

- $L_R(R) = E[L(S,R,\sigma)|R] - L_{cash}$

  This term depends only on R. It represents the 'optimal rate only hedge' – the optimal hedge of remaining risk (after implementing the cash hedge) that can be constructed using instruments that only depend on the risk-free rate R.

- $L_\sigma(\sigma) = E[L(S,R,\sigma)|\sigma] - L_{cash}$

  This term depends only on σ. It represents the 'optimal volatility only hedge' – the optimal hedge of remaining risk (after implementing the cash hedge) that can be constructed using instruments that only depend on the implied volatility σ.

- $L_{S,R}(S,R) = E[L(S,R,\sigma)|S,R] - L_S(S) - L_R(R) - L_{cash}$

  This term depends on S and R but not σ. It represents the 'optimal equity/rate interaction hedge' – the optimal hedge of remaining risk (after implementing the cash, equity only and rate only hedges) that can be constructed using instruments that only depend on the equity index S and risk-free rate R.

- etc.

Note that the final term[9] $L_{S,R,\sigma}(S,R,\sigma)$ is chosen so that the sum of all terms adds to the total loss i.e.:

$$L_{S,R,\sigma}(S,R,\sigma) = L(S,R,\sigma) - \left(L_{cash} + L_S(S) + L_R(R) + L_\sigma(\sigma) + L_{S,R}(S,R) + L_{S,\sigma}(S,\sigma) + L_{R,\sigma}(R,\sigma)\right)$$

In this decomposition, the sub-portfolios represent 'optimal' hedges in the sense of *quadratic hedging* i.e. they minimize expected squared hedging errors. For example, in constructing the 'optimal equity only hedge' our objective is to construct a portfolio containing cash and instruments depending on the equity index S alone, which minimizes the expected squared hedging error

[9] Note also that we could choose to terminate the decomposition at any point, for example we could alternatively write:

$$L(S,R,\sigma) = L_{cash} + L_S(S) + L_R(R) + L_\sigma(\sigma) + L_{S,R,\sigma}(S,R,\sigma)$$

where

$$L_{S,R,\sigma}(S,R,\sigma) = L(S,R,\sigma) - \left(L_{cash} + L_S(S) + L_R(R) + L_\sigma(\sigma)\right)$$

when averaged over all risks. Note that when risks are correlated, this hedge (which depends on the equity index S) partially hedges interest-rate risk and volatility in addition to equity risk. As noted in (Rosen and Saunders 2010), this results in some 'double counting', as the same risks are hedged by a number of different terms in the decomposition.

Another practical issue with the Hoeffding decomposition is that in general the conditional expectations need to be estimated numerically (for example using Monte Carlo simulation).

Optimal hedge portfolio definition (2): Perfect hedging if other risks are switched off

To avoid the problems with the Hoeffding decomposition mentioned above, as an alternative decomposition we can consider sub-portfolios that hedge individual risks (or combination of risks) under the assumption that other risks are 'switched off' (i.e. set equal to some constant values such as their initial values or their expected values). For example, freezing other risks at their initial values gives the following decomposition:

- $L_{cash} = L(S_0, R_0, \sigma_0) = 0$

  This 'optimal cash hedge' now represents the perfect cash hedge in the event that none of the risk factors change.

- $L_S(S) = L(S, R_0, \sigma_0) - L_{cash} = L(S, R_0, \sigma_0)$

  This 'optimal equity only hedge' now represents the perfect hedge of equity risk in the event that risk-free rates and implied volatility don't change.

- $L_R(R) = L(S_0, R, \sigma_0) - L_{cash} = L(S_0, R, \sigma_0)$

  This 'optimal rate only hedge' now represents the perfect hedge of rate risk in the event that the equity index and implied volatility don't change.

- $L_\sigma(\sigma) = L(S_0, R_0, \sigma) - L_{cash} = L(S_0, R_0, \sigma)$

  This 'optimal volatlity only hedge' now represents the perfect hedge of implied volatility risk in the event that the equity index and risk-free rates don't change.

- $L_{S,R}(S, R) = L(S, R, \sigma_0) - L_S(S) - L_R(R) - L_{cash}$

  This 'optimal equity/rate interaction hedge' now represents the perfect hedge of equity index and rate risk (after implementing the equity only and rate only hedges) in the event that the implied volatility doesn't change.

- etc.

For example, Figure 8 shows the loss on this 'optimal' equity portfolio as a function of equity index level and interest-rates (assuming implied volatility is fixed at 20%), along with the net loss on the guaranteed equity bond (post-hedging). By construction (for fixed implied volatility of 20%) the equity hedge portfolio provides a perfect hedge if rates stay at 5% but will result in a net loss (gain) if rates fall (rise). For comparison, we also show the corresponding loss function for a simpler delta hedge, which provides a linear approximation to the optimal hedge and thus gives a similar hedge performance for small movements in the index, but under-hedges for large movements due to convexity in the loss function. The 'optimal' hedge portfolio provides a perfect hedge of all (not just 'small') equity index changes, provided rates and implied volatility don't change. This 'optimal' hedge could be constructed dynamically or by static replication using 1-year put options of different strikes.

Figure 8: Example loss function for guaranteed equity bond index hedge portfolio and net position

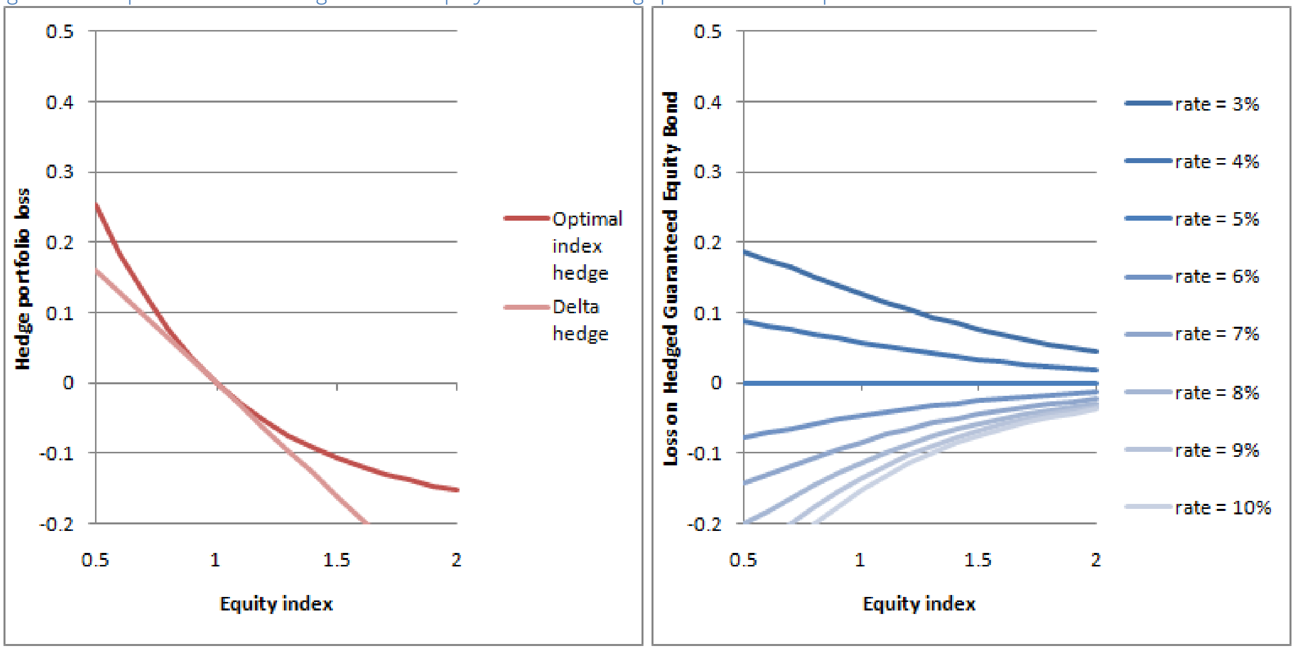

Figure 9 compares the loss on various hedge portfolios against the loss on the guaranteed equity bond, under 1,000 stochastic scenarios for the equity index, interest-rate and implied volatility. When all three 'optimal' standalone hedge portfolios (equity index, interest-rate and implied volatility) are combined we achieve a good but not perfect hedge, due to the interaction terms which remain unhedged.

Figure 9: Loss on hedge portfolios vs loss on guaranteed equity bond

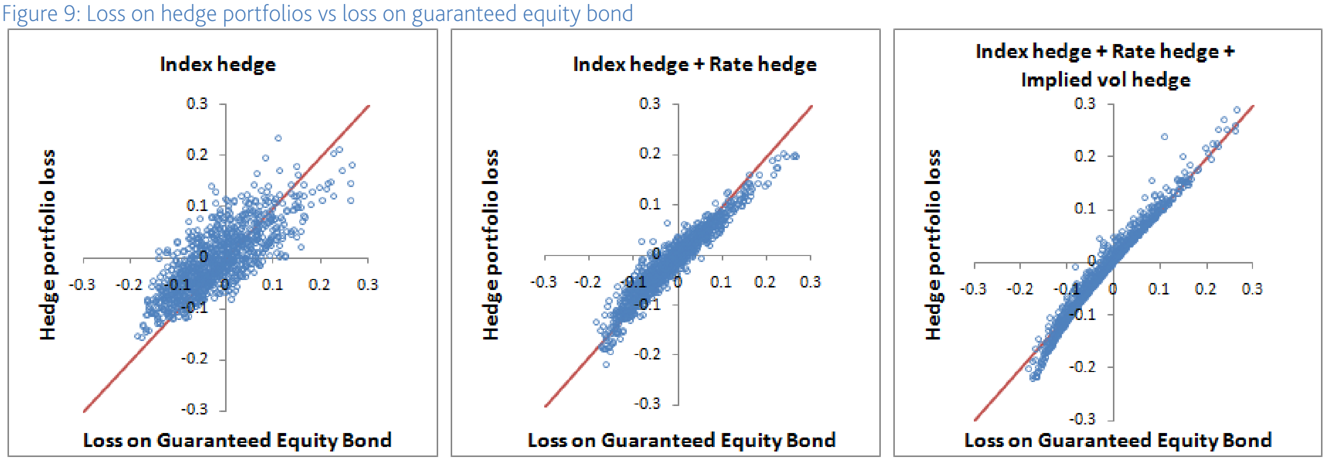

Having written the total portfolio loss as a sum of losses on sub-portfolios (each representing a hedge of an individual risk factors or combination of risk factors) we can use any of the techniques described in Section 2 to attribute capital requirements to these sub-portfolios. For example, Figure 10 shows standalone capital requirements for different hedge portfolios (with the 'residual' hedge portfolio representing the perfect hedge of all remaining risk after all individual hedges are in place), along with discrete marginal and continuous marginal contributions. All contributions are estimated using 10,000 scenarios. In this case, the (unscaled) discrete marginal contributions may be particularly useful as they represent the capital reduction that can be achieved for each individual risk using the optimal hedging strategies.

Figure 10: Capital attribution by risk factor for guaranteed equity bond

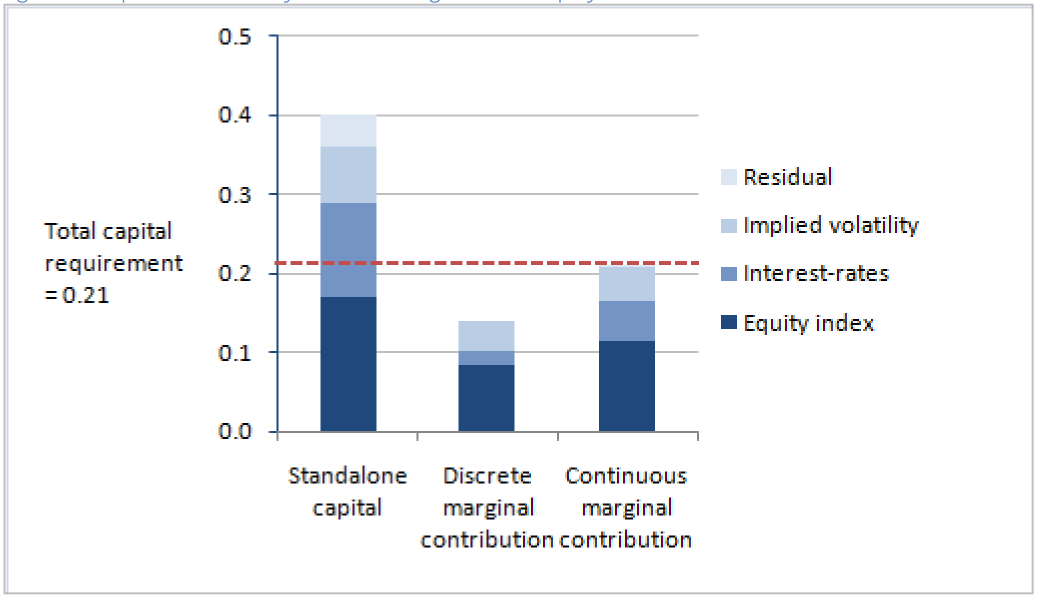

6. Summary

Having calculated an overall capital requirement, insurers are often interested in quantifying how this overall capital can be attributed to:

- Sub-portfolios (for example business units, geographical locations or product types)
- Risk factors

This note has described some of the more popular methods for capital attribution by sub-portfolio, their estimation using Monte Carlo scenarios, and the statistical error in these estimates. While our previous paper on VaR estimation (Morrison and Tadrowski 2013) indicated that overall capital requirements can be relatively insensitive to choice of estimation scheme when using large scenario sets, this is not the case for one of the most popular attribution schemes (the continuous marginal method). In this case, empirical estimators suffer from a large amount of statistical error even using large scenarios sets, and there is a significant benefit in using smoothed estimates (e.g. Harrell-Davis).

While there is a large amount of literature on attribution by sub-portfolio, the question of attribution by risk factor appears to have received far less attention. In this note we have described a potential approach to this problem that has recently been proposed, and illustrated by applying a similar technique to attribution of capital by risk factor for a simple life insurance product.

References

Epperlein, E, and A Smillie. "Cracking VAR with kernels." *RISK*, 2006.
McNeil, A J, R Frey, and P Embrechts. *Quantitative Risk Management.* Princeton University Press, 2005.
Morrison, S, and L Tadrowski. *Effiicient Statistical Estimation of 1-year VaR Economic Capital.* Moody's Analytics, 2013.
Rosen, Dan, and David Saunders. "Risk factor contributions in portfolio credit risk models." *Journal of Banking & Finance*, 2010.
Tasche, D. *Risk contributions and performance measurement.* Preprint, TU Munchen Department of Mathematics, 1999.