Sharpe Ratio Optimization
of an
Excess of Loss Reinsurance Contract

Sameer Nahal, FCAS

February 25, 2015

Abstract

Effects of the sharing and layering of losses by means of an excess of loss reinsurance contract are examined. In particular, the covariance of loss at any given position across the layers of an excess of loss reinsurance contract as a function of the first and second moments of the frequency and individual layer loss severity distributions is derived. Based on this result, a risk management application is presented and the Sharpe Ratio based investment equivalent paradigm for measuring profitability is explored.

1 Introduction

An excess of loss reinsurance contract facilitates the layering and sharing of losses by allowing individual reinsurers to take a position on the contract consisting of a set of shares across the various layers. In order to appropriately assess the profitability of a given reinsurer at its position on a variance based risk adjusted basis, it is necessary to take into account the interdependence between those layers. That is, evaluating any such measure of profitability for each layer individually will not take into consideration the correlation of loss between the layers at the reinsurers shares. Moreover, in the case that the cedent retains a portion of the reinsured layers then the measurement of the variance of loss to the cedent net of reinsurance is complicated for the same reasons also.

In the first part of this paper, a method for determining the variance of loss at a given position on an excess of loss reinsurance contract as a function of individual layer loss metrics is presented. In the second part, Kreps' investment equivalent approach is generalized to a multi-layer framework analogous to the Markowitz Portfolio Model whereby investors providing the capital required to back the reinsurer's position on the contract are concerned with the Sharpe Ratio of their investment. However, in this case it is the various layers of the reinsurance contract that serve as the individual risky assets for the investor to take a position across. In the third part of the paper, the Sharpe Ratio optimization paradigm will be examined through numerical examples exploring the view of the reinsurer as a price taker. This is intended as a thought experiment to examine the implications of layer risk load relativities on the behavior of market participants under this classical paradigm. Finally, an example is presented to demonstrate how the variance formula can be applied to the cedent from a risk management perspective.

2 Background and Methods

2.1 Variance of Shared and Layered Losses

While many more sophisticated measures are commonly used in practice, variance still remains at the least a benchmark measure of risk. For an excess of loss contract however, the effect of sharing and layering losses complicates the estimation of the variance at given shares of each layer. In particular, while if taking the same shares of each layer then the variance of loss at that position is just the square of the share times the total variance, at unequal shares of any two layers it is necessary to explicitly take into consideration the covariance of loss between those layers. This would also be the case from the perspective of the cedent if retaining some portion of one or more of the reinsured layers while of course retaining all of the loss beneath the attachment point of the contract which can be viewed as an underlying layer. And so, to appropriately account for the variance of loss to either a participating reinsurer or the cedent, it is necessary to take into consideration the individual layer variances as well as the covariances of loss between those layers.

Collective risk loss models separate the task of modelling losses into two distinct frequency and severity components by assuming that the number of losses and the severity of those losses are independent and that the loss severity distribution is the same for each loss. Unlike first dollar losses, for excess layers the loss severity can be defined based on large losses above some chosen threshold amount up to the attachment point of the contract. Then, the frequency is based on the number of losses in excess of the threshold and the severity is based on the amount of loss at a given position across the layers given a loss in excess of the threshold. That is, the random variable representing the severity here is the amount of loss at a given position as a function of the large loss random variable, the set of shares across the layers comprising the given position, and the various layer attachments and limits.

The collective risk loss model approach has the advantage of being able to capture certain loss characteristics entirely by either the frequency or the severity distribution resulting in a more accurate and flexible loss model [1]. For example, inflation can be explained by the severity distribution alone while exposure growth can be explained by that of the frequency. For an excess of loss contract, since the allocation of loss between the layers is a function of the loss severity only, the correlation of loss between those layers is explained entirely by the severity distribution. As will be shown, by taking a collective loss model approach it is possible to express the covariance of loss between any two layers of an excess of loss reinsurance contract at given shares of each layer as a function of individual layer loss metrics.

[1] Loss Models, Klugman et al. pg. 137

Let

$a_i =$ attachment point of the i_{th} layer
$b_i =$ upper bound of the i_{th} layer
$s_i =$ share of the i_{th} layer
$N =$ random variable representing the number of losses in excess of the threshold
$L =$ random variable representing the amount of loss given a loss in excess of the threshold
$L_i =$ loss to the i_{th} layer given a loss in excess of the threshold
$L_s =$ loss at a given position across the layers given a single loss in excess of the threshold
$f_L(l) =$ probability density function of L
$\mu_i =$ expected loss to the i_{th} layer given a single loss in excess of the threshold
$\sigma_i =$ standard deviation of the loss to the i_{th} layer given a single loss in excess of the threshold
$T =$ total loss at a given position across the layers, where T = 0 if N = 0
$\sigma_T =$ standard deviation of the total loss at a given position across the layers

For this collective risk model, by the law of total variance we have

$$\begin{aligned}
\sigma_T^2 &= E(N)Var(L_s) + E^2(L_s)Var(N) \\
&= E(N)Var(\sum_i s_i L_i) + E^2(\sum_i s_i L_i)Var(N) \\
&= E(N)[\sum_i Var(s_i L_i) + 2\sum_{j<k}\sum_k Cov(s_j L_j, s_k L_k)] + E^2(\sum_i s_i L_i)Var(N) \\
&= E(N)[\sum_i s_i^2\sigma_i^2 + 2\sum_{j<k}\sum_k s_j s_k Cov(L_j, L_k)] + (\sum_i s_i \mu_i)^2 Var(N)
\end{aligned}$$

Now to derive the covariance term, let any two layers be given and denoted by j and k for the lower and upper layers respectively. Then by definition, the covariance of L_j and L_k is given by

$$\begin{aligned}
Cov(L_j, L_k) &= E(L_j L_k) - E(L_j)E(L_k) \\
&= \int_{l=-\infty}^{l=\infty} L_j(l)L_k(l)f_L(l)dl - \mu_j\mu_k
\end{aligned}$$

But since L_j and L_k are both non-zero only when $L > a_k$, and since $L_j = (b_j - a_k)$ in that case, we have

$$\begin{aligned}
Cov(L_j, L_k) &= \int_{l=a_k}^{l=\infty} (b_j - a_j)L_k(l)f_{L(l)}dl - \mu_j\mu_k \\
&= (b_j - a_j)\mu_k - \mu_j\mu_k \qquad (1)
\end{aligned}$$

Plugging this into the equation for total variance we have

$$\sigma_T^2 = E(N)[\sum_i s_i^2\sigma_i^2 + 2\sum_{j<k}\sum_k s_js_k[(b_j - a_j)\mu_k - \mu_j\mu_k]] + (\sum_i s_i\mu_i)^2 Var(N) \tag{2}$$

And so we have the formula for the variance of the total loss at a given position across the layers of an excess of loss reinsurance contract as a function of the expected value and variance of the number of losses in excess of the threshold and the expected value and variance of the individual layer losses given a loss in excess of the threshold. Then if we estimate those individual layer metrics, the variance of the total loss at a given position across the layers can be determined formulaically.

A simplifying assumption that is often made when using a collective risk model is that the frequency is distributed Poisson. Under this assumption, and since for a Poisson distribution the mean and variance are equal, the total variance formula simplifies to

$$\begin{aligned}\sigma_T^2 &= \lambda[Var(L_s) + E^2(L_s)]\\ &= \lambda[E(L_s^2)]\end{aligned} \tag{3}$$

Where λ is the expected number of claims in excess of the given loss threshold. Now, since

$$\begin{aligned}Var(L_s) + E^2(L_s) &= \sum_i s_i^2\sigma_i^2 + 2\sum_{j<k}\sum_k s_js_k[(b_j - a_j)\mu_k - \mu_j\mu_k] + (\sum_i s_i\mu_i)^2\\ &= \sum_i s_i^2E(L_i^2) + 2\sum_{j<k}\sum_k s_js_k(b_j - a_j)\mu_k\end{aligned}$$

we have

$$E(L_s^2) = \sum_i s_i^2E(L_i^2) + 2\sum_{j<k}\sum_k s_js_k(b_j - a_j)\mu_k \tag{4}$$

and

$$\sigma_T^2 = \lambda[\sum_i s_i^2E(L_i^2) + 2\sum_{j<k}\sum_k s_js_k(b_j - a_j)\mu_k] \tag{5}$$

And so we have the formula for the variance of total loss at any given position under the Poisson assumption as a function of the expected number of losses in excess of the chosen loss threshold and the first and second moments of loss to the individual layers given a loss in excess of the threshold.

2.2 Investment Equivalent Reinsurance Pricing

2.2.1 Kreps' Paradigm

In Kreps' paper Investment Equivalent Reinsurance Pricing, he presents a paradigm for setting risk loads for layers of a single excess of loss reinsurance contract on a standalone basis by viewing the contract as an investment alternative to a given target investment. In particular, the risk load and capital allocated at contract inception by the investor to back the contract must be such that the expected return and risk as measured by variance of that return are at least as favorable as those of the target investment. Further, the amount of capital allocated to back the contract must be such that the probability of ruin not exceed some given loss safety level.

Kreps explicitly considers the reinsurance investment alternative as a combination of the reinsurance contract and a financial technique that is used to earn investment income on the total funds available to the reinsurer. Those funds consist of both the premium received and capital allocated at contract inception until losses are paid. And so, the risk and return to the investor providing the capital required to participate on the reinsurance contract are derived from both underwriting as well as investment income.

2.2.2 Dissimilarities of a Reinsurance Contract and Traditional Investments

A key difference between traditional investments such as stocks and bonds and a reinsurance contract is the determination of the amount of capital required to participate. While for stocks and bonds the capital required is the market price of the asset, for a reinsurer the amount of capital allocated at inception in order to back the contract must be determined. Further, while for stocks and bonds the greatest possible loss to the investor is the initial cost of acquiring the asset, for a reinsurer it is possible to lose more than the initial capital allocation requiring either additional capital to be provided by the investor to fund the additional losses or insolvency of the reinsurer. The greater the amount of capital initially allocated to back the contract, the less likely the investor will have to provide additional capital to cover losses and in turn the less likely that the cedent will not be reimbursed due to insolvency of the reinsurer. As such, the amount of capital allocated at inception can be viewed as a measure of underwriting conservatism [2] from the reinsurer's perspective and as a measure of security with respect to insolvency from the perspective of the cedent. Of course, credit risk to the cedent associated with collecting reinsurance recoverables still exists even in the case that the reinsurer remains solvent. Only in the extreme case where the investor collateralizes the reinsurer in the amount of the full limit provided is there no risk associated with the collectability of reinsurance recoverables for the cedent.

Another important difference between a reinsurance contract and more traditional investments is the liquidity of the investor's position. While with stocks and bonds an investor can usually liquidate their position at most any time, with a reinsurance contract the matter is more complicated. Firstly, the sale would occur in the private market and so would require the investor to find a buyer. Secondly, even if a buyer is found reinsurance contracts typically contain a special termination clause that requires the approval of the cedent in order to change control of the reinsurer. And so, in order to liquidate their position the investor would have to find an interested buyer and then get approval from the cedent in order to sell the reinsurer and liquidate their position. Assuming investors prefer liquidity, then this difference should necessitate a liquidity premium in the form of additional return to the investor.

[2] Kreps, Investment Equivalent Reinsurance Pricing

2.2.3 Sharpe Ratio of the Reinsurance Investment

While many measures for the performance of an investment are used in practice, a classical benchmark still used today is the Sharpe Ratio [3]. First introduced by William Sharpe in 1966, the ratio compares the expected excess return over the risk free rate to the standard deviation of that excess return. What this measure indicates is the risk premium provided to the investor per unit of risk taken as measured by standard deviation of the excess return, and so allows for a comparison of different investments on a risk adjusted basis. Under the assumption of unlimited borrowing and lending as is the case in the Markowitz Portfolio Selection Model,[4] by combining the risky asset with the risk free asset individual investors can take on a level of risk commensurate with their own degree of risk aversion based on individual utility and achieve a certain level of expected return. As such, when comparing two investments and given the individual risk appetite of the investor, the investment with the higher Sharpe Ratio will yield the higher expected return at that given level of risk chosen by the investor. And so, assuming that investors are risk averse, an investor will prefer the risky asset with the higher Sharpe Ratio as that will result in a higher expected return at the same level of risk.

In the case of a per risk excess of loss contract, if we assume that the investor capitalizing the reinsurer to write a single stand-alone contract has the ability to borrow or lend at the risk free rate, then a similar strategy as with the Markowitz Portfolio Selection Model can be taken. In particular, rather than target the individual expected return and variance of the target investment, the investor may instead consider the Sharpe Ratio of the return on the capital provided to capitalize the reinsurer. As is the case with the Markowitz portfolio selection model, the investor can then combine the reinsurance investment with the risk free asset in order to achieve a certain level of risk and corresponding expected return.

In order to derive the Sharpe Ratio of the reinsurance investment, the following assumptions are made:

1. The reinsurer writes a single standalone reinsurance contract
2. The term of the reinsurance contract is one year
3. There is a single loss payment at the end of the year
4. The ceded premium and allocated capital are invested at the risk free rate from contract inception until losses are paid

[3] Investments 9th Edition, Bodie, Kane, Marcus, pg. 133
[4] Investments 9th Edition, Bodie, Kane, Marcus, pg. 211

Let

$P =$ ceded premium at a given position across the layers of the reinsurance contract
$R_i =$ risk load defined as the difference between the ceded premium and the discounted expected loss for the i_{th} layer
$R =$ risk load at a given position across the layers of the reinsurance contract
$C =$ capital allocated to the reinsurer by the investor in order to take a given position on the reinsurance contract
$r =$ return on the capital provided by the investor
$\widehat{r} =$ expected return on capital
$\sigma_r =$ standard deviation of the return on capital
$r_f =$ risk free rate of return
$\mu_T =$ expected total loss at a given position across the layers of the reinsurance contract
$SR =$ Sharpe Ratio of the return on capital

Starting with the definition of the Sharpe Ratio we have

$$SR = \frac{\widehat{r} - r_f}{\sigma_r}$$

Now,

$$\begin{aligned} r &= \frac{P(1+r_f) + Cr_f - T}{C} \\ &= \frac{[R + \frac{\mu_T}{(1+r_f)}](1+r_f) + Cr_f - T}{C} \\ &= \frac{R(1+r_f) + r_f C - (T - \mu_T)}{C} \end{aligned}$$

And so,

$$\widehat{r} = \frac{R(1+r_f) + r_f C}{C} \text{ and } \sigma_r = \frac{\sigma_T}{C}$$

Which gives us

$$SR = \frac{R(1+r_f)}{\sigma_T}$$

If we assume a collective risk loss model with poisson(λ) frequency distribution, then substituting equation (3) for σ_T^2 we have

$$SR = \frac{R(1+r_f)}{[\lambda E(L_s^2)]^{\frac{1}{2}}} \tag{6}$$

$$= \frac{(\sum_i s_i R_i)(1+r_f)}{[\lambda[\sum_i s_i^2 E(L_i^2) + 2\sum_{j<k}\sum_k s_j s_k (b_j - a_j)\mu_k]]^{\frac{1}{2}}} \tag{7}$$

2.3 Reinsurance Market Dynamics

As mentioned by Kreps, reinsurance pricing is usually described as market driven. Prices are generally thought to be determined by the supply of capital available to provide reinsurance and the demand for coverage by cedents. Under the assumption that investors are concerned with the Sharpe Ratio of their investment, then the willingness of investors to allocate capital to provide coverage will depend on how the Sharpe Ratio of the reinsurance investment opportunity compares to that of the target investment. That is, taking into account the differences in the nature of the two investments, the Sharpe Ratio of the reinsurance investment opportunity must be sufficient in comparison to that of the target investment in order to entice investors to participate. The reinsurance supply curve under the current paradigm then is the aggregate of many reinsurers all concerned with the Sharpe Ratio of their investment but with differing views of the pertinent loss metrics.

2.3.1 Reinsurer View

Suppose that an investor is presented with a reinsurance investment opportunity consisting of a single two layer excess of loss reinsurance contract. Further, the investor here is a price taker in the sense that the ceded premium and contract terms are taken as given. The investor has the option to either capitalize a reinsurer to participate at shares of each layer of its own choosing such that a given internal ceded premium target is met or decline to participate entirely.

In order to assess the quality of the investment opportunity under the current paradigm, the investor must first determine the shares of each layer that will result in the highest Sharpe Ratio given the target ceded premium. For a two layer excess of loss reinsurance contract, assuming a poisson distribution with mean λ for the number of losses we have

$$SR = \frac{(s_1 R_1 + s_2 R_2)(1+r_f)}{[\lambda[s_1^2 E(L_1^2) + s_2^2 E(L_2^2) + 2s_1 s_2[(b_1 - a_1)\mu_2]^{\frac{1}{2}}} \tag{8}$$

In order to maximize the Sharpe Ratio of the investment, we look for points of extrema for the Sharpe Ratio formula. Let,

$$K = \text{reinsurer premium target}$$
$$P_i = \text{ceded premium for the } i_{th} \text{ layer}$$

Now, the target premium constraint given is $K = s_1 P_1 + s_2 P_2$ which reduces the optimization problem to one dimension.

If we substitute $s_2 = \frac{K - s_1 P_1}{P_2}$ and then take the derivative with respect to s_1, we get

$$SR' = \frac{(1 + r_f)}{f} \frac{[R' E(L_s^2)^{\frac{1}{2}} - \frac{1}{2} R E(L_s^2)^{\frac{-1}{2}} E(L_s^2)']}{E(L_s^2)}$$

which is zero when

$$\frac{R'}{R} = \frac{1}{2} \frac{E'(L_s^2)}{E(L_s^2)}$$

And so, positions of extrema occur when the rate of change in the percentage of the risk load equals half of the percentage change of the second moment of the loss severity distribution. However, to the author's knowledge no simple form of the solution for s_1 is available and so numerical methods will be relied upon in the following three examples that illustrate the effect of risk load sharing on the optimal position of the reinsurer.

Example 1

A reinsurer is presented with a two layer excess of loss reinsurance contract consisting of a 10M xs 10M and a 30M xs 20M layer. The reinsurer models losses to the cover using a collective risk model at a loss threshold of 5M and assumes that the severity of loss is distributed Pareto with alpha 1.8 and that the frequency is distributed Poisson with mean 2.2. The resulting individual layer loss metrics based on the loss distribution as estimated by the reinsurer along with the full coverage ceded premium amounts are summarized in the table below.

layer	a_i	b_i	μ_i	$E(L_i^2)$	P_i
1	10M	20M	1,527,951	12,143,363,968,303	4,463,585
2	20M	50M	1,071,173	23,499,405,810,277	4,087,942

Based on the reinsurers estimated individual layer loss metrics and frequency pick, using equation (8) the following relationship between the Sharpe Ratio of the reinsurance investment and the proportion of total shares in the first layer is obtained.

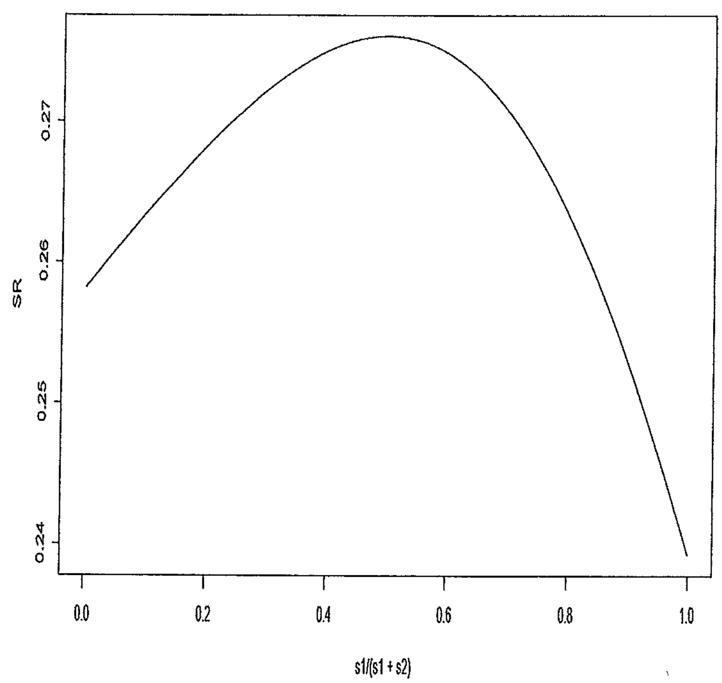

In this example, the risk load relativity between the layers is such that the optimal Sharpe Ratio occurs at equal shares of each layer and so the risk loading is balanced in the sense that the reinsurer does not have a preference to take more or less of any one of the layers. The optimal Sharpe Ratio here is 0.276 while the individual layer Sharpe Ratios are 0.239 and 0.258 for layers 1 and 2 respectively. It is interesting to note that setting the Sharpe Ratio's of the individual layers equal will not result in a balanced risk loading.

Example 2

Now, suppose instead that the total ceded premium amount of 8.6M is allocated as follows, all else the same as in the first example.

$P_1 = 4.4M, P_2 = 4.2M$

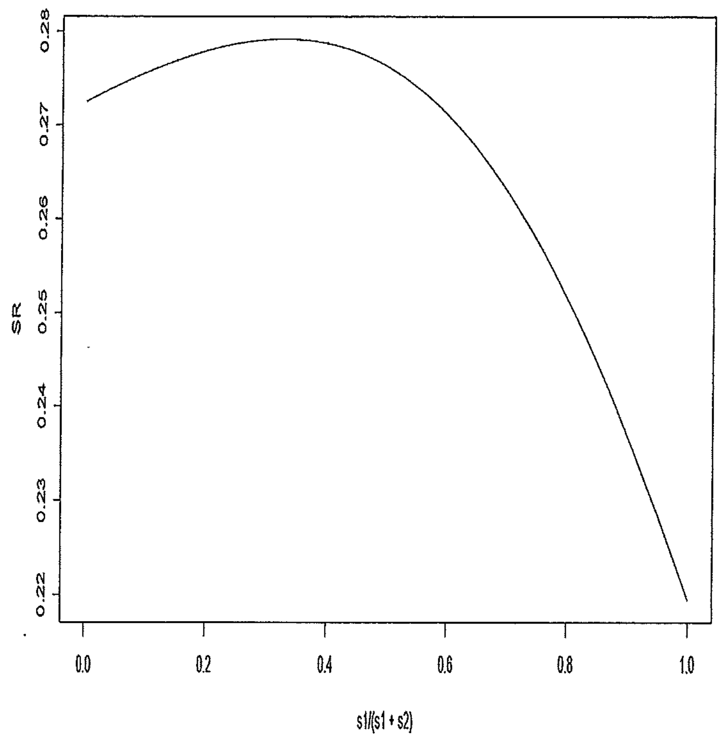

In this example, the second layer rates much better than the first layer with Sharpe Ratio's of 0.272 and 0.219 respectively. However, the Sharpe Ratio actually increases as shares are shifted from the second layer to the first and remains greater than the first layer alone until near the balance point. The optimal Sharpe Ratio is 0.279 and occurs around a two to one ratio of shares of layer 2 to layer 1.

Example 3

Again, all else the same only now let

$P_1 = 4.2M, P_2 = 4.4M$

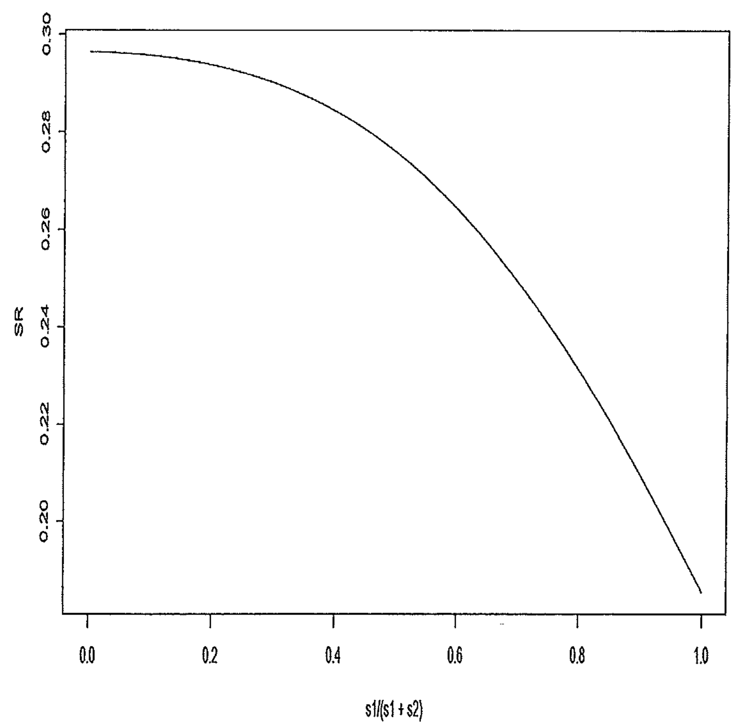

In this example, the first layer rates significantly better than the second with Sharpe Ratio's of 0.297 and 0.185 respectively. There is no benefit to diversify in this case and the optimal position for the reinsurer is to have the entire target premium from the first layer. It is interesting to note however, that even in this case there is enough of a diversification benefit such that a fairly flat region on the curve results as shares of the second layer are added to the first over which the reinsurer does not lose too greatly by taking up a small portion of shares of the less profitable layer.

2.3.2 Cedent View

While reinsurers are concerned primarily with the profitability of the reinsurance contract, cedents have other considerations to take into account such as capital management and regulation. From the view of the cedent, reinsurance functions more as a risk management tool and less as an investment and so decisions regarding the purchase of reinsurance must take into consideration regulatory and internal risk management objectives. That said, cedents are still concerned with giving away profits and so depending on their views of the pricing of their coverage relative to that of the market, the cedent may choose to cede more or less of a given layer taking into consideration any constraints on the reinsurance purchase regulatory or otherwise.

Example 4

An insurer would like to purchase a two layer excess of loss reinsurance treaty to manage the volatility of its book of business. Company management desires that the standard deviation of loss not exceed 1/3 of its current capital base of 10.7M. The cedent actuary estimates that the number of claims follows a Poisson distribution with mean 1200 and produces the following estimates for the first and second moments of the severity of loss for each of the reinsured layers as well as the underlying retained layer.

layer	a_i	b_i	μ_i	$E(L_i^2)$
0	0M	5M	14,544	423,049,790
1	5M	10M	1,956	19,138,644
2	10M	20M	1,500	58,223,477

Based on their internal analysis, the cedent believes that the market rate to fully place the first layer is too high and as such would rather cede less of that layer. At the rate that the cedent is willing to pay for the first layer, the supply of reinsurance will be sufficient for a 50 percent placement. In such a case where the insurer cedes all of the second layer while retaining half of the first layer and all of the underlying layer, the standard deviation of loss to the insurer is given by equation (5) as 3.5M which amounts to a minimum capital requirement of 10.5M. Since the cedent's current capital base exceeds the required amount, the cedent is able to meet its profitability objective within the minimum capital requirement constraint by ceding all of the second layer and half of the first layer.

3 Discussion

By taking a collective risk model approach it is possible to express the variance of loss at any combination of shares across the layers of an excess of loss reinsurance contract as a function of the first and second moments of the frequency and the individual layer loss severity distributions. While the presentation here was specifically in the context of an excess of loss reinsurance contract, the results are applicable to any setting involving the sharing and layering of losses including primary subscription policies or the combination of an excess of loss contract coupled with an underlying quota share.

The correlation of loss between the layers is dependent on the severity distribution only to the extent of the first and second moments and does not require any assumption about the distribution otherwise. And so, while it is common to assume a certain distribution for the severity of losses when modelling an excess of loss reinsurance contract, no such assumption is required in order to estimate the variance of loss at a given position. Moreover, the covariance term here is estimable to the same extent that the pertinent individual layer loss metrics upon which it is based are. While parameter risk was not taken into consideration and is beyond the scope of this paper, the applicability of the methods presented here is limited by the estimability of the pertinent loss metrics which may be difficult especially for more remote layers.

A key assumption underlying the collective risk model is the independence of the frequency and severity components. However, the validity of this underlying assumption may be questionable in certain cases. For example, for an excess of loss treaty covering a heavily catastrophe exposed property book, a large catastrophic event may impact several risks resulting in an unusually high frequency while the severity also is unique since all of those losses are caused by the same set of underlying extreme physical circumstances. And so there may be some dependence of the severity distribution on the number of losses in certain cases. Another common assumption made in collective risk modelling is that the frequency is distributed Poisson which simplifies the model to a more compact and mathematically tractable form. However, excess of loss reinsurance contracts typically have aggregate features such as occurrence limits and limits on reinstatements that may invalidate this assumption even as a reasonable approximation. Further, the type of cover may also invalidate the Poisson assumption as would be the case for an aggregate excess of loss cover for which the frequency is binary. Thus, caution should be exercised when employing the collective risk model approach in terms of both the underlying independence of frequency and severity assumption as well as the simplifying Poisson frequency distribution assumption.

When presented with the opportunity to participate on an excess of loss reinsurance contract, a reinsurer may choose to take a position consisting of unequal shares of the layers for various reasons including its views of the relative profitability of those layers. In such a case, the collective risk loss model approach can be employed to assess the profitability of different positions on a variance based risk adjusted basis. One such measure of profitability is the Sharpe Ratio of the reinsurance investment. In deriving the formula for the Sharpe Ratio, it was assumed that the funds available to the reinsurer over the duration of the contract are invested at the risk free rate, or as referred to by Kreps the 'Swap' financial technique. Under this assumption, the resulting formula for the Sharpe Ratio is independent of the capital allocated to back the contract and so this profitability measure does not require any assumption about how capital is allocated. For other financial techniques such as the 'Put' technique presented by Kreps however, the amount of capital allocated will need to be taken into consideration.

In order to further examine the implications of taking a multi-layer view, a thought experiment was presented analogous to the Markowitz Portfolio Selection Model whereby investors are concerned with maximizing the Sharpe Ratio of their investment by choosing shares of each layer that will result in an optimal portfolio while meeting a given ceded premium target. In essence, this amounts to examining the implications of risk load relativities across the layers of an excess of loss reinsurance contract on the behavior of reinsurers under the Markowitz paradigm. Three examples were presented that varied only in the allocation of risk load across the layers of a two layer contract in order to demonstrate the impact of the risk load relativity. It was interesting to note that only in the most extreme example where the two layers were most mispriced from the view of the reinsurer assessing the deal was there no benefit to take at least some of the less profitable layer. In the second example, even though there was a large difference in the Sharpe Ratio of the individual layers, the point of optimality is actually achieved by including some of the less profitable layer. And so, while one layer may rate better than the other on an individual basis, there may still be some benefit to diversify by taking at least some shares of the less profitable layer. This result is important from a practicality standpoint as cedents often restrict the choice of shares and may require reinsurers to take at least some shares of all layers in order to participate on the contract. As was the case in the third example where the disparity in profitability between the two layers was greatest, there was a fairly flat segment of the Sharpe Ratio curve as shares were shifted from the second to the first layer with the Sharpe Ratio holding fairly well until near to a two to one ratio of shares in the more profitable layer to the less profitable layers.

4 Conclusion

Under the framework of a collective risk loss model, it is possible to derive the covariance of loss at any given position across an excess of loss reinsurance contract as a function of the first and second moments of the frequency and the individual layer loss severity distributions. The covariance term facilitates the extension of a single layer variance based approach to a multi-layer setting formulaically and allows for a more holistic view of risk as variance to be taken. One such possible application of this result is the Markowitz Portfolio Selection Model applied to the various layers of a given excess of loss reinsurance contract. Under such a framework, the relativity of the risk load between the layers of the reinsurance contract will govern the behavior of reinsurers and ultimately determine the relative demand for the layers by the market. Even in the case where one layer rates poorly relative to others, there may still be a benefit to having some shares of that layer in the portfolio. The explicit consideration of the covariance of loss between layers is also important from the perspective of the cedent for variance based risk management measures in the case that their net position varies across the layers of their book.

Acknowledgements

The Author would like to give a special thanks to Raymond Nichols and Joseph Schreier for their support throughout the process of writing this paper. The author would also like to thank Dan Schlemmer, Scott Swaynay, and Rob Davies for their review and comments which greatly helped in creating this final version.