Portfolio modelling of counterparty reinsurance default risk

Modelling counterparty default risk of a reinsurance bouquet requires an appropriate underlying default correlation structure. A flexible model with two control parameters gives us an alternative of the current standard model in Solvency II.

BY PETER TER BERG

WE TREAT THE BOUQUET of reinsurance treaties and the induced counterparty default risk for a primary insurer as similar to an investment portfolio. Contrary to traditional mean-variance portfolio design, which implies a large number of realisations of random variables and comparable sized components, reinsurance default risk typically is a rare event process with potential large severity components. Moreover, the covariance structure is largely driven by a common shock. The probability law for this common shock, as well as the vulnerability of insurance companies to this shock, guide the derivation of the various key-quantities necessary for the determination of the solvency capital requirement.

The new risk-based solvency and supervisory standard for European (re)insurance companies, Solvency II, is planned to be introduced in 2012. In the process of completing it, insurance supervisors in Europe are developing and testing a set of risk modules together spanning the standard model in Solvency II for the calculation of the solvency capital requirement (SCR). Since this set of modules is expected to cover "all risks", not only traditional risk categories like market risk and underwriting risks are covered, but also more specialised risk categories. One of the more esoteric modules targets counterparty default risk, which focuses on the default risk of reinsurance and other risk mitigating instruments.

Through so-called quantitative impact studies these risk modules are tested by the European insurance industry. The technical specifications for the third quantitative impact study (QIS3) describe an approach for measuring counterparty default risk.

The current approach for default risk parallels credit risk modelling for banks under Basel II: risk is modelled using a Vasicek limiting probability distribution that is valid in case of large homogeneous populations. The main drawback of this approach for counterparty reinsurance default risk lies in a hidden assumption of a large homogeneous population. The rein-

Figure 1. Probability density function for shock size

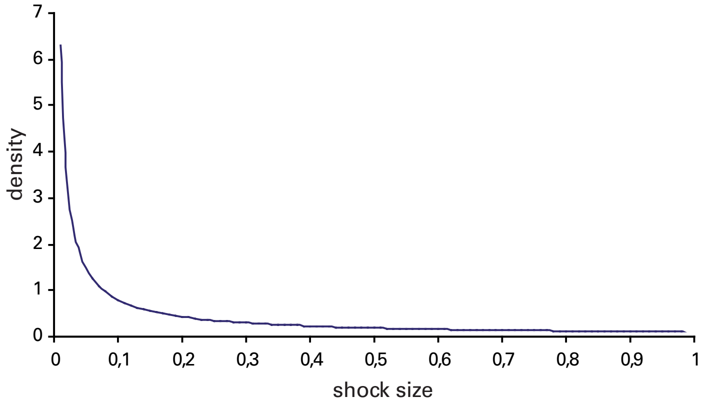

Figure 2. Default probability as function of shock size

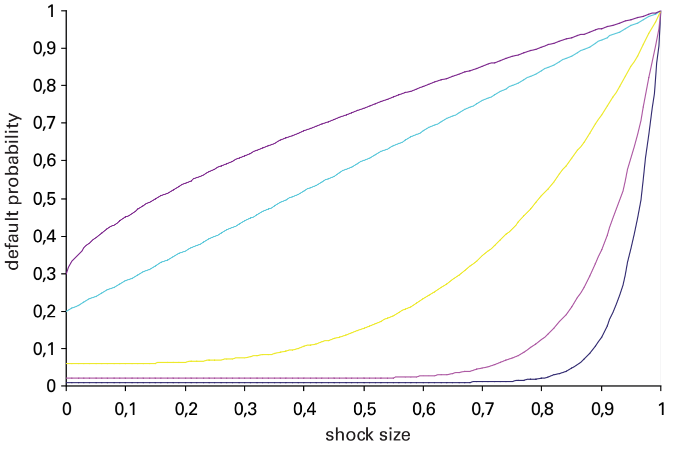

surance field typically is not only much smaller but also targets a heterogeneous population. Under the current version of Solvency II, correlation between reinsurance companies is approximated by (a function of) the concentration in the portfolio measured by the Herfindahl concentration measure. For such a heterogeneous environment, the use of a universal correlation between reinsurance companies is flawed. Furthermore, correlation is a property of a population and a reinsurance bouquet is just a finite sample of such a population. Any observed concentration in such a sample is a property of the sampling method. However, it does not change aspects of the population such as correlation. Relating correlation to a concentration measure such as Herfindahl consequently is conceptually flawed and is unlikely to yield correct outcomes in practice. The comparison in the numerical illustrations in this paper should support this point of view, in particular when the bouquet consists of a single reinsurance company.

As an alternative to the current approach we investigate the merits of default probabilities that are functions of a common shock size. This common shock size, a latent (unobservable) random variable, generates correlation between reinsurance companies in a natural way. The correlation structure follows from the deeper structural parameters that define the common shock model and the form of the shock dependent default probability. Also the implied correlations will be different for different pairs of reinsurers, reflecting the heterogeneity of the population.

The plan of this paper is as follows. First we will define the random nature of a world-wide shock which affects the (re)insurance industry. Next we model the default probability as a function of shock size and consider the group of reinsurance companies. Using the various losses given default we define the default risk for the primary insurance company. It will have a probability distribution with a large probability located at zero. Nevertheless the mean and variance of this reinsurance default risk (with a one-year horizon) can be calculated and used as an input for the Solvency II machinery. A numerical illustration displays mean, standard deviation and solvency capital required under varying reinsurance bouquets. Here it should become clear that the same input from insurance companies is needed as in QIS3: loss given default and the assignment of rating class to reinsurance companies. A discussion together with future vistas ends this article.

Common shock

Results in the insurance industry are driven by random mass phenomena as well as non-mass phenomena. The latter may correspond to slowly changing conditions related to business cycle and demography or sudden impacts induced by super-storms or stock exchange. In the same vein the reinsurance industry is vulnerable to such non-mass phenomena too. We imagine this to be captured by a yearly common shock which we denote by S and which takes values on the unit-interval $(0,1)$. Values close to zero will have minor impact on the industry, whereas values close to one should be expected to affect the whole industry in a worldwide catastrophic way. For its modelling we like to use the Beta probability density function which is a versatile function defined on the unit-interval. A special case of this function displays monotone decreasing probabilities for this shock size:

$$f(s\mid\alpha)=\alpha s^{\alpha-1} \qquad 0\le s\le 1 \qquad 0<\alpha<1$$

In Figure 1 such a density is displayed for the choice $\alpha=0.1$. The message of this figure is that small shocks are likely, but probabilities decline fast for large to huge shocks. However, the decline is such that even for an extreme large shock size close to 1, the density remains positive and does not collapse to 0. This should appeal to our subjective empirical perception of such shocks. This special density can also be interpreted in a different way. If we consider a sequence of n years the maximum of the n shocks will follow a similar Beta density with parameter α replaced by $n\alpha$. Considering $n=10$ the choice $\alpha=0.1$ implies $n\alpha=1$, which corresponds with a uniform density for this 10-year all-time high shock. Such considerations furnish a way to assign a numerical value to α. In this paper we will fix $\alpha=0.1$.

This probability density function will be fundamental for the derivation of first and second (cross) moments of default probabilities as functions of random shock size.

Default probability as function of shock size

After the determination of the probability density function for the common shock size, we have to specify a default probability that is a function of this particular common shock size. To this end we introduce a baseline for this

Table 1: Counterparty default risk with 7x2 reinsurers

	mean	0.0542	
	standard deviation	0.0631	
	SCR	0.1627	
	"observed" expected default probability	*implied base level default probability*	*loss given default*
reinsurer	p	b	y
1	0.00002	0.00001	0.07143
2	0.00002	0.00001	0.07143
3	0.00010	0.00007	0.07143
4	0.00010	0.00007	0.07143
5	0.00050	0.00033	0.07143
6	0.00050	0.00033	0.07143
7	0.00240	0.00160	0.07143
8	0.00240	0.00160	0.07143
9	0.01200	0.00800	0.07143
10	0.01200	0.00800	0.07143
11	0.06040	0.04027	0.07143
12	0.06040	0.04027	0.07143
13	0.30410	0.20273	0.07143
14	0.30410	0.20273	0.07143

Table 2: Counterparty default risk with 7 reinsurers

	mean	0.0542	
	standard dev	0.0810	
	SCR	0.2086	
	"observed" expected default probability	*implied base level default probability*	*loss given default*
reinsurer	p	b	y
1	0.00002	0.00001	0.14286
2	0.00002	0.00001	
3	0.00010	0.00007	0.14286
4	0.00010	0.00007	
5	0.00050	0.00033	0.14286
6	0.00050	0.00033	
7	0.00240	0.00160	0.14286
8	0.00240	0.00160	
9	0.01200	0.00800	0.14286
10	0.01200	0.00800	
11	0.06040	0.04027	0.14286
12	0.06040	0.04027	
13	0.30410	0.20273	0.14286
14	0.30410	0.20273	

probability. This probability will be denoted by *b* and will be specified for each reinsurance company. We propose the following function for the shock-modified default probability:

$$p(s) = b + (1-b)s^{\delta} \qquad 0 < b < 1 \qquad \delta > 0$$

This specification has two parameters: a baseline default probability *b* and a power parameter δ. This is a monotone increasing function in the shock size *s* from its baseline level *b* to the maximum level of 1.The use of the power parameter δ implies a convex curve if δ>1 and a concave curve if δ<1. The case δ=1 implies a straight line.

It is useful to make this power parameter δ a decreasing function of the baseline default probability *b*:

$$\delta = \tau b^{-1} \qquad \tau > 0$$

The idea behind this specification is that reinsurance companies with low baseline default probabilities will be rather immune to shocks as long as these remain non-extreme, whereas large baseline levels increase the sensitivity for shocks even if these are of modest size.

Figure 2 illustrates this behaviour for τ=0.2 and *b*=0.01,0.02,0.06,0.2,0.3. Baseline default probabilities *b*<τ are viewed as low. Clearly the choice τ=0.2 may need some sensitivity analysis. Also a quantitative impact study should give guidance on the effect of the value for τ on the measured size of this counterparty default risk, just the way as other risk modules are tested.

Derivation of baseline default probabilities

Using the probability density function *f*(*s*) for shock size and the shock-modified default probability function *p*(*s*), we derive the expected value *p* of this default probability as:

$$p = E\left[p(S)\right] = \int_0^1 p(s)f(s)\mathrm{d}s = \frac{(\tau+\alpha)b}{\tau+\alpha b}$$

This value for *p* can be observed indirectly, as under QIS3, as a default probability taken from rating agencies. So, for each reinsurance company we assign the value of its rating based default probability to the expected value *p*. From this we solve for the baseline default probability as:

$$b = \frac{p\tau}{\alpha(1-p)+\tau}$$

This shows that the baseline depends on α and τ as well as the observable *p*. So, the randomness of the reinsurance bouquet is governed by the two universal shape parameters α and τ.

Reinsurance bouquet

Until now we have specified the default probability for a specific reinsurance company as a function of shock-size that will be common to the whole industry. An insurance company itself may do business with more than a single reinsurance company, the so-called bouquet.

We consider an insurance company with a reinsurance bouquet consisting of *k* reinsurance companies indexed *i*=1,..., *k*. The default process itself is indicated through a set of Bernoulli random variables $(w_1,\ldots,w_k)$ where each component takes the value 0 in case of no default and 1 if default applies to the corresponding reinsurance company. Conditionally on the default of reinsurance company *i* there will arise a loss given default which we denote by y_i. For reasons of simplicity the loss given default is treated as a nonrandom variable, but an additional variance can be incorporated.

To emphasise the system structure, we express the various quantities involved as vectors:

$$\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_k \end{bmatrix} \quad \mathbf{p}(s) = \begin{bmatrix} p_1(s) \\ \vdots \\ p_k(s) \end{bmatrix} \quad \mathbf{p} = \begin{bmatrix} p_1 \\ \vdots \\ p_k \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_k \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_k \end{bmatrix}$$

The random variable for reinsurance default risk, denoted by Z, follows as the summation of the relevant loss given default quantities. With the variable w taking on the values 0 and 1 this can be expressed as:

$$Z = \sum_{i=1}^{k} w_i y_i = \mathbf{w}'\mathbf{y}$$

This random variable, conditional on an observed shock $S=s$, can be shown to have expectation (E), variance (V) and probability (P) of no default:

$$E(Z \mid S = s) = \sum_{i=1}^{k} p_i(s) y_i = \mathbf{p}'(s)\mathbf{y}$$

$$V(Z \mid S = s) = \sum_{i=1}^{k} p_i(s) y_i^2 - \sum_{i=1}^{k} \left(p_i(s) y_i\right)^2$$

$$P\left(Z = 0 \mid S = s\right) = \prod_{i=1}^{k} \left(1 - p_i(s)\right)$$

However, for the primary insurance company, the quantity of interest is the unconditional random variable Z. This takes the uncertainty on the universal shock size S into account. So we must average out over all possibilities for $S=s$. As Z is a linear in **w** and **y**, the mean and variance of Z are simply given as a linear and quadratic expression in the components of **y**:

$$E(Z) = \sum_{i=1}^{k} p_i y_i = \mathbf{p}'\mathbf{y}$$

$$V(Z) = \sum_{i,j=1}^{k} \omega_{ij} y_i y_j = \mathbf{y}'\Omega\mathbf{y}$$

where Ω denotes the covariance matrix of **w**. The derivation of this matrix is straightforward but tedious and not given here. The diagonal elements of Ω are given by:

$$\omega_{ii} = p_i(1 - p_i) \qquad i = 1, \ldots, k$$

and the off-diagonal elements

$$\omega_{ij} = \frac{\alpha(1 - b_i)(1 - b_j)}{\alpha + \tau b_i^{-1} + \tau b_j^{-1}} - (p_i - b_i)(p_j - b_j) \quad i \neq j; \quad i, j = 1, \ldots, k$$

In case the reinsurance bouquet consists of a single reinsurance company, all the above results apply with $k=1$. The variance of counterpart default risk boils down to:

$$V(Z) = p(1 - p)y^2$$

and all calculations can in principle be carried out using just pencil and paper. This is a common sense result. In case of a single reinsurance company, any reference to correlation evaporates and the reinsurance default risk process boils down to just two possibilities: no default with probability $(1-p)$ and default with probability p and y as loss given default. Such a risk has mean py and $p(1-p)y^2$ as variance, implying a standard deviation of $y\sqrt{p(1-p)}$. In Solvency II, the one-year horizon security level is set at 99.5%. The quantile factor corresponding to a standard normal distribution is 2.5758 and the SCR is calculated as 2.5758 times the standard deviation. Clearly, if we need the true quantile for counterparty default risk we should account for the firm probability at zero. However, as this SCR is processed with SCR's resulting from other risk modules through a square root of a quadratic form, this need is reduced in the Solvency II standard model. In contrast, for an internal model, explicit recognition of the probability at zero might be beneficial, just as this applies to the modelling of catastrophic risk.[1]

Numerical illustration

Rating agencies communicate their ratings in general as belonging to one out of seven rating classes. We will consider a bouquet consisting of fourteen reinsurance companies. Choosing these as pairs in the seven rating classes will enable us to view different rated reinsurance companies as well as similar rated ones in one stroke. For the loss given default we assume for simplicity an even distribution over all reinsurance companies in the bouquet. Clearly,

---

[1] Conditionally on the shock size the true probability distribution follows from convolution of the various independent risks, which are modelled as two-point distributions. Weighting these conditional (discrete) distributions with the (discretised) probability of shock-size gives the true probability distribution. It will have a firm probability mass at zero followed by a bumpy tail to the right. For Solvency II purposes such an approach seems a bridge too far. For internal models this might be a valid route.

Submission guidelines for technical articles

Life & Pensions welcomes the submission of technical articles on topics relevant to our practitioner readership. Core areas include solvency and economic capital modelling, the measurement and management of financial, biometric and operational risks, market-consistent valuation and financing of life and pension balance sheets and cashflows, and investment management. This list is not an exhaustive one.

The most important publication criteria are originality, exclusivity and relevance. In the interests of our readers, we attempt to strike a balance between these. Thus, while we will not publish executive summaries of longer papers (on the grounds of exclusivity), we may accept papers that draw partially but not completely on research submitted elsewhere should our referees recommend it on the grounds of originality and relevance to practitioners.

Given that *Life & Pensions* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the editor and his team, submissions are logged, and checked against the criteria above. Articles that fail to meet the criteria are rejected at this stage.Articles are then sent to one or more anonymous referees for peer review. Our referees are drawn from the actuarial, risk management, treasury and investment departments of major life and pensions companies, in addition to academia and regulatory bodies. Depending on the feedback from referees, the technical editor makes a decision to reject or accept the submitted article. His decision is final.

We also welcome the submission of brief communications. These are also peer-reviewed contributions to *Life & Pensions* but the process is less formal than for full-length technical articles. Typically, brief communications address an extension or implementation issue arising from a full-length article that while satisfying our originality, exclusivity and relevance requirements, does not deserve full-length treatment.

Submissions should be sent to the editor at *technical@incisivemedia.com*. The preferred format is MS Word, although Adobe PDFs are acceptable. The maximum recommended length for articles is 3,500 words, and for brief communications 1,000 words, with some allowance for charts and/or formulas. We expect all articles and communications to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations. Authors should allow four to eight weeks for the refereeing process.

this will be different in practice and also optimal portfolio theory would suggest something different. The summation of all losses given default in the bouquet is normalised at one. This is an innocuous normalisation. Using the set of default probabilities taken from rating agencies and documented in QIS3 technical specifications, we rank, just for ease of presentation, these reinsurance companies in ascending order of default probability p (or its implied b) and arrive at the results in Table 1 which uses $\alpha=0.1$ and $\tau=0.2$:

The present risk module in the standard module for Solvency II calculates the Herfindahl concentration measure H, which is the sum of squares of the portfolio shares. In this flat portfolio this results in $H=1/14$. From this a correlation parameter is specified as:

$$\rho = \tfrac{1}{2}(1+H)$$

A lower and upper bound for the (modified) default probability are given as:

$$\begin{aligned} p_L &= N\left(N^{-1}(p)\cdot\sqrt{2}+N^{-1}(0.995)\right) && \rho=\tfrac{1}{2} \\ p_U &= \min(1;100p) && \rho=1 \end{aligned}$$

where N denotes the cumulative standard normal distribution and N^{-1} its inverse known as the quantile function. Linear interpolation gives the final modified default probability as a linear function of ρ or the Herfindahl concentration measure H:

$$\begin{aligned} \tilde{p} &= (2-2\rho)p_L + (2\rho-1)p_U && \tfrac{1}{2}\le\rho\le 1 \\ &= (1-H)p_L + Hp_U && 0\le H\le 1 \end{aligned}$$

Application of this device gives a vector $\tilde{\mathbf{p}}$ of modified default probabilities and an SCR that results as $\tilde{\mathbf{p}}'\mathbf{y}$. In this illustration we get $SCR=0.2973$.

If instead of having two reinsurance companies in each rating class we decide to concentrate these in a single reinsurance company, the bouquet reduces to seven reinsurance companies. This can be arranged by nullifying appropriate elements of y and renormalising as in Table 2.

The mean of the counterparty default risk remains the same, as it should, but the standard deviation increases from 0.0632 to 0.0810, and the SCR increases from 0.1627 to 0.2086. So reducing this bouquet of 7×2=14 reinsurance companies to seven reinsurance companies increases the counterparty default risk, which should not come as a surprise. The SCR according to the present risk module of the standard model of Solvency II is calculated as 0.3108.

Finally, if we concentrate the bouquet into a single reinsurance company with $p=0.0542$ the SCR easily follows as

$$SCR = 2.5758\cdot y\sqrt{p(1-p)} = 2.5758\times 1\times\sqrt{0.0542\times(1-0.0542)} = 0.5832$$

Compared with the SCR of 0.1627 or 0.2086 corresponding with 14 or 7 reinsurance companies, this is a huge increase of the SCR and can be interpreted as the price of concentration.

The SCR according to the present risk module of the standard model of Solvency II is calculated as 1 which is due to an estimated $H=\rho=1$, which implies a modified default probability of 1.

Conclusion

Based on a flexible model with just two control parameters and using observable rating-based default probabilities, we have modelled mean and variance of counterparty reinsurance default risk. This gives an alternative for the current SCR-contribution of counterparty reinsurance default risk for the standard model of Solvency II. The model uses a common shock that affects default probabilities that are sensitive to a common shock. In the same vein however, other counterparty default risks, such as those originating from financial derivatives, can be incorporated. Obviously, an estimate for the mean default probability should become available then. A potential future refinement of this model will include randomness in the loss given default. Also a minimum variance bouquet could be derived and compared with the actual bouquet.

Peter ter Berg works for De Nederlandsche Bank and has written this article in a personal capacity. He is grateful to Marc Pröpper for useful remarks. Email: p.j.j.m.ter.berg@dnb.nl

References

Proposal for a Directive of the European Parliament and of the Council on the taking-up and pursuit of the business of Insurance and Reinsurance,
Solvency II, Brussels 10.7.2007

QIS3 Technical Specifications Part I: Instructions,
CEIOPS-FS-11/07 April 2007.