Operational risk modelled analytically

Regulators require banks to use an internal model to compute a capital charge for operational risk, which is thought to be sensitive to assumptions on dependence between losses that still remain a matter of debate. Vivien Brunel proposes an analytical way to quantify this risk, and shows that uniform correlation is a robust assumption for measuring capital charges

The current regulatory framework allows banks to compute their capital charge for operational risk under an internal model, which is often based on the loss distribution approach (LDA). Loss distributions are calibrated at the cell level (a cell is the elementary risk unit per business line and type of risk) and the bank's capital charge is estimated by aggregating cell loss distributions under some dependence assumption (Chernobai, Rachev & Fabozzi 2007).

The Basel Committee provides some guidelines about how banks should appropriately reflect the risk profile in their internal model (Bank for International Settlements 2011). However, banks benefit from some flexibility in their modelling choices and this may lead to some discrepancies in capital charges for similar risk profiles. The broad range of practices that are observed among banks is the result of different distributional or dependence assumptions in their models.

Many studies have focused on the modelling of the tails of the severity distributions (Dutta & Perry 2007; Moscadelli 2004), but the bulk of the correlation problem is still unsolved and controversial. There is much debate about the choice of copula function for losses across cells due to the scarcity of data, but the regulators advise banks to determine sound correlations and to retain conservative assumptions. Some institutions have selected the simplest option and use equal correlations between cell losses. This assumption is questionable, of course, and may embed some model risk, but both regulators and practitioners have great difficulty in agreeing on realistic and conservative correlation levels. Some authors believe that correlations between cell losses are as low as 4% (Frachot, Roncalli & Salomon 2004).

Most of the knowledge we have about operational risk quantification comes from complex models and heavy Monte Carlo simulations, and as far as we know there is no analytical model that takes into account risk and correlation dispersion among cells. This article fills this gap. Under the asymptotic single risk factor assumption, we obtain new results for the bank's capital charge sensitivity to the critical parameters of the model. In particular, we show the capital charge is not that sensitive to correlation dispersion, and that the constant correlation assumption is robust.

This new result is obtained with few specifications, and we conjecture that it remains valid, at least qualitatively, for real bank portfolios that have a finite number of cells. We believe that our approach also provides a way to pioneer a new method for computing capital charges and challenging internal model assumptions, as exemplified in this paper.

This paper is organised as follows. First, we provide some real data evidence about cell loss distributions and correlations. Second, we solve the asymptotic single risk factor model with lognormal losses at the cell level, even when individual cells have varying risk profiles. Third, we solve the case of non-equal correlations between cells and provide some key results about the sensitivity of capital charge to the main critical parameters of the model.

Some empirical facts about cell loss distributions and correlations

In the LDA framework, the aggregate operational loss for cell number i is equal to the sum of individual losses:

$$L_i = \sum_{n=1}^{N_i} X_n^i \tag{1}$$

where L_i is the aggregate loss of cell number i, N_i is the number of events over one year, and $(X_n^i)_{1 \leqslant n \leqslant M}$ is the sequence of individual loss severities for cell number i. The aggregate loss process is a compound Poisson process and, accordingly, the model is based on the following assumptions.

- The number of events and their severity are independent.
- Severities are independent and identically distributed random variables.

Cell loss distribution parameters. There are many studies in the literature about individual loss distributions (see, for example, Dutta & Perry 2004; Moscadelli 2004), but there are very few empirical studies about aggregate cell losses.

We have conducted such a study based on the SAS OpRisk Global Data database. As of November 2013, this database included 6,402 events that had occurred in financial firms since 2002, when financial institutions started to collect and report their operational losses systematically. We have calibrated the frequency of events and lognormal severity distributions for each of the 21 cells that have more than 30 losses. Direct calibration of the aggregate loss distribution from real data is, of course, impossible because there is only one observation per year. However, it is possible to assess the fit with the lognormal distribution of the aggregate loss distribution obtained through the LDA.

Let us assume that the loss distribution for cell i is lognormal with parameters μ_i and σ_i; the ratio between the expected value and any quantile depends only on the parameter σ_i:

$$\frac{\text{expected value } (i)}{\text{VaR}_q(i)} = \mathrm{e}^{\sigma_i^2/2+\sigma_i F_q} \tag{2}$$

$$\sigma_i = -F_q - \sqrt{F_q^2 + 2\ln\frac{\text{expected value } (i)}{\text{VaR}_q(i)}} \tag{3}$$

where VaR_q is the q-percentile of the lognormal distribution and $F_q = N^{-1}(1-q)$. Inverting (2) leads to two different solutions, and we have chosen the one with a minus sign in front of the square root in (3) because we require the parameters σ_i to decrease with the ratio of expected value to quantile for all cells. We observe that broader distribution assumptions for cell losses in the model can naturally be taken into account by choosing the plus sign solution in (3).

A. Implied value of σ from real data		
Confidence level (%)	Average (%)	Standard deviation (%)
95	98	41
97.5	99	39
99	107	44
99.5	112	46
99.9	124	48
All	107	42

1 Density of probability of the upper bound R for loss correlations

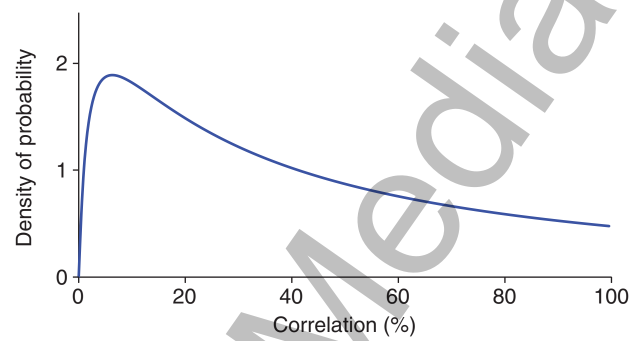

The LDA gives us these ratios for each cell in the tail of the loss distribution ($q \geqslant 95\%$). Table A provides the observed average value and the standard deviation of the parameters σ_i (implied from (3)) over all cells for several values of the confidence level.

The values of σ_i remain somewhat stable when the confidence level changes: the average value over all cells and confidence intervals is equal to 107%, and the observed standard deviation is equal to 42%. To assess the robustness of these estimates, we compute the median of the observed values of σ_i, which at 108.5% is very close to the average value. The med-med estimator (the median value of the spread with the median) is equal to 31%, which is lower than the measured standard deviation.

Cell loss correlations. In most studies (see, for example, Aue & Kalkbrener 2007; Frachot, Roncalli & Salomon 2004) cell loss correlations are calculated from the dependence of the number of events between cells rather than the dependence of severities. Under the assumption of lognormal severity distributions ($X^i \sim \mathrm{LN}(m_i, s_i)$), Frachot, Roncalli & Salomon show that the loss correlation between cell 1 and cell 2 is given by:

$$\mathrm{corr}(L_1, L_2) = \mathrm{corr}(N_1, N_2)\mathrm{e}^{-s_1^2/2 - s_2^2/2} \tag{4}$$

The correlation between the number of events, N_1 and N_2, is linked to the loss frequencies of cells 1 and 2. s_1 and s_2 are the standard deviation parameters of the two distributions. Bivariate Poisson variables are obtained by considering three independent Poisson variables Z, Y_1 and Y_2 with parameters r, $\lambda_1 - r$ and $\lambda_2 - r$, respectively; the variables $N_i = Z + Y_i$ are also Poisson with intensities λ_i, and their correlation is given by:

$$\mathrm{corr}(N_1, N_2) = \frac{r}{\sqrt{\lambda_1 \lambda_2}} \leqslant R = \sqrt{\frac{\min(\lambda_1, \lambda_2)}{\max(\lambda_1, \lambda_2)}} \tag{5}$$

The upper bound R for the correlation comes from the inequalities $\lambda_1 \geqslant r$ and $\lambda_2 \geqslant r$. Whenever the bank's portfolio includes a large number of cells, the intensities are distributed as a random variable. Internal data represents frequencies better than external data because it is specific to the bank, and it also includes the frequencies of rare but severe events that are taken into account by the scenario analyses in the model. Internal data and scenario analysis frequencies at Société Générale support the normal distribution assumption of the log-intensities of the Poisson processes (in particular, the skewness and normalised kurtosis are close to 0) with a standard deviation equal to $\gamma = 2.35 \pm 0.35$. Setting $\ln \lambda_i = \alpha + \gamma G_i$, where the $(G_i)_{i=1,2}$ are uncorrelated standard normal random variables, we obtain:

$$R = \mathrm{e}^{-\gamma |G_1 - G_2|/2} = \mathrm{e}^{-\gamma |X|/\sqrt{2}}$$

where X is a standard normal random variable. Under this assumption, the upper bound R follows a truncated lognormal law:

$$P[R \leqslant \rho] = P\left[|X| \geqslant -\sqrt{2}\frac{\ln \rho}{\gamma}\right] = 2N\left(\sqrt{2}\frac{\ln \rho}{\gamma}\right) \tag{6}$$

We plot the density function of the correlation upper bound of R corresponding to $\gamma = 2.35$ in figure 1.

The expected value of R is equal to $2\mathrm{e}^{\sigma^2/4} N(-\gamma/\sqrt{2})$, which works out to be 38.5% for $\gamma = 2.35$. This is in line with the findings of Aue & Kalkbrener (2007), who observed that frequency correlations were around 10%, with higher correlations being specific only to some pairs of cells. Frachot, Roncalli & Salomon (2004) claimed that loss correlations were as low as 4%: we recover this result when we take $\mathrm{corr}(N_1, N_2) = 38.5\%$ and $s_1 = s_2 = 1.5$ in (1), which is the lowest value observed by Frachot, Roncalli & Salomon for these parameters. From SAS OpRisk data, we observe that the parameters s_i have an average value of 2.03, a standard deviation of 0.42, and range between 1.34 and 2.90. Correlation upper bounds can be computed from this data with (4) and (5). We find an expected value of 1.33% and a standard deviation of 1.61%. All but a few correlation upper bounds lie in the range of 0–4%; the highest correlation has an upper bound at 11.27% and is found between the 'Execution, delivery and process management' and 'Internal fraud' cells of the retail brokerage business line. All of these studies confirm that we expect low levels of correlation between cells.

Correlation parameters in the Gaussian copula model. In the Gaussian copula framework with lognormal marginal cell losses, the correlation parameter ρ_{ij} between two cells is related to the cell loss correlation:

$$\mathrm{corr}(L_i, L_j) = \frac{\mathrm{e}^{\rho_{ij}\sigma_i\sigma_j} - 1}{\sqrt{(\mathrm{e}^{\sigma_i^2} - 1)(\mathrm{e}^{\sigma_j^2} - 1)}} \tag{7}$$

This formula with parameters $\sigma_i = \sigma_j = 107\%$ and a conservative assumption for loss correlations of $\mathrm{corr}(L_i, L_j) = 4\%$ leads to a correlation parameter of the Gaussian copula equal to $\rho_{ij} = 7.2\%$. External data supports the assumption of very low correlation parameters in the copula framework: much lower than 10%.

A class of solvable models with correlated risks

A simplified LDA model. The rest of this article is dedicated to building a simple portfolio model for operational risk. We assess that the bank's overall operational risk is composed of a portfolio of N operational risks at cell level. We make the following four assumptions.

- Lognormal distributions: the loss for cell number i is a lognormal random variable L_i with parameters μ_i and σ_i. As discussed earlier in this article, we assume that the σ_i have an expected value of $\sigma = 107\%$ and, unless otherwise stated, a variance of $v = 18\%$.
- Gaussian copula: pairwise correlations ρ_{ij} may be different to each other. For numerical estimations, we assume that the average correlation is equal to 10% (this is a conservative assumption, as mentioned earlier).
- One-factor model: cell losses are sensitive to the same systemic factor, which we denote by F. This factor is assumed to be a standard normal random variable. The specific (idiosyncratic) part of the risk is embedded in another independent normal random variable denoted by ϵ_i $(i = 1, \ldots, N)$. Systemic and specific factors are all assumed to be independent of each other.
- We assume that the parameters are not dependent on the number of cells N.

In this framework, the annual loss for a cell can be written as the exponential function of a normal random variable that is a linear combination of the systemic and specific factors. For cell number i $(i = 1, \ldots, N)$ we obtain:

$$L_i = \exp(\mu_i - \sigma_i(\beta_i F + \sqrt{1 - \beta_i^2}\epsilon_i)) \qquad (8)$$

The parameters β_i are linked to the pairwise correlations of the Gaussian copula: $\rho_{ij} = \beta_i \cdot \beta_j$. Because cells may have very different risk characteristics, and because correlations may be very different for different pairs of cells, we assume that the parameters μ_i, σ_i and β_i are the observations of independent and identically distributed random variables called M, Σ and B, respectively. In the limit $N \to \infty$, the bank's loss is equal to $N \cdot L(F)$ and is a function of the common factor F, as in Vasicek's model for granular homogeneous loan loss distributions (see Vasicek 2002):

$$\begin{aligned} L(F) &= \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} L_i \\ &= E[\mathrm{e}^{M - \Sigma(BF + \sqrt{1-B^2}\epsilon_i)} \mid F] \\ &= E[\mathrm{e}^{-\Sigma BF + \Sigma^2(1-B^2)/2} \mid F] \cdot E[\mathrm{e}^{M}] \end{aligned} \qquad (9)$$

Without loss of generality, we assume that $M = 0$ because this simply rescales the bank's loss by a constant factor $E[\mathrm{e}^M]$ in the $N \to \infty$ limit.

The stand-alone capital for cell number i, called KSA_i, is equal to the 99.9% percentile of the cell loss distribution and is given by $\mathrm{KSA}_i = \mathrm{e}^{-\sigma_i F_q}$, where $F_q = N^{-1}(0.1\%)$. The bank's capital charge is equal to $N \cdot L(F_q)$. The capital reduction coming from risk diversification is measured by the Diversification Index, which is defined as:

$$\text{Diversification Index} = \mathrm{DI} = \frac{N \cdot L(F_q)}{\sum_{i=1}^{N} \mathrm{KSA}_i} \xrightarrow{N\to\infty} \frac{L(F_q)}{E[\mathrm{e}^{-\Sigma F_q}]} \qquad (10)$$

Homogeneous risks. The simplest solvable model is obtained for homogeneous risks: the random variables Σ and B have constant values equal to σ and $\sqrt{\rho}$, respectively, for all cells, and $v = 0$. In the limit $N \to \infty$, the bank's loss distribution remains lognormal:

$$L(F) = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} L_i = \mathrm{e}^{-\sigma\sqrt{\rho}F + \sigma^2(1-\rho)/2} \qquad (11)$$

As in Markowitz's portfolio theory, risk is not diversified away because of loss correlations between cells. The correlation parameter determines DI:

$$\mathrm{DI} = \mathrm{e}^{\sigma(1-\sqrt{\rho})F_q + \sigma^2(1-\rho)/2}$$

For $\sigma = 107\%$ and $\rho = 10\%$, we get $\mathrm{DI} = 17.5\%$. We note that $\mathrm{DI} > 1$ when $\sigma > -2F_q((1 - \sqrt{\rho})/(1 - \rho)) = 4.70$ and capital charges are no longer sub-additive. However, super-additivity occurs for values of the parameter σ higher than the typical value of 107%. As explained earlier, broader distributions can easily be accommodated in our model by choosing the other solution of (3). This results in effects similar to those found in fat tail distributions (Neslehova, Embrechts & Chavez-Demoulin 2006).

The diversification ratio is particularly low because of the $N \to \infty$ limit. The perimeter and the number of cells is, however, a modelling choice and a convention. Choosing a very high number of cells would not necessarily result in regulatory arbitrage. To assess this, the scaling of the parameters of the model with the number of cells must be investigated.

Heterogeneous risk and identical correlations. In reality, cells have different risk characteristics. We assume that the parameters σ_i are normally distributed (that is, $\Sigma \equiv N(\sigma, v)$) and that correlations are constant (that is, $B = \sqrt{\rho}$). The loss at cell level is:

$$L_i = \mathrm{e}^{-\sigma_i(\sqrt{\rho}F + \sqrt{1-\rho}\epsilon_i)}$$

In the limit $N \to \infty$, the bank's loss derived from (9) as a Gaussian integral is:

$$\begin{aligned} L(F) &= \int_{-\infty}^{\infty} \frac{\mathrm{d}x}{\sqrt{v}} n\left(\frac{x-\sigma}{\sqrt{v}}\right) \mathrm{e}^{-x\sqrt{\rho}F + (1-\rho)x^2/2} \\ &= \frac{1}{\sqrt{1-(1-\rho)v}} \exp\left(-\sigma\sqrt{\rho}F + \frac{\sigma^2(1-\rho)}{2} \right. \\ &\qquad \left. + \frac{v}{2} \frac{((1-\rho)\sigma - \sqrt{\rho}F)^2}{1-(1-\rho)v}\right) \end{aligned} \qquad (12)$$

where $n(x) = \mathrm{e}^{-x^2/2}/\sqrt{2\pi}$. The bank's loss follows a g-and-h distribution. Because we assumed that the random variable Σ is normal and, strictly speaking, could take negative values, the resulting bank's loss is a non-decreasing function of the systemic factor. However, this occurs when the systemic factor F is larger than $F^* = \sigma/v\sqrt{\rho}$, which is very unlikely in practice (for instance, $F^* = 18.8$ when $\rho = 10\%$,

2 Diversification index impact as a function of $\sqrt{v}$ ($\sigma = 107\%$, $\rho = 10\%$)

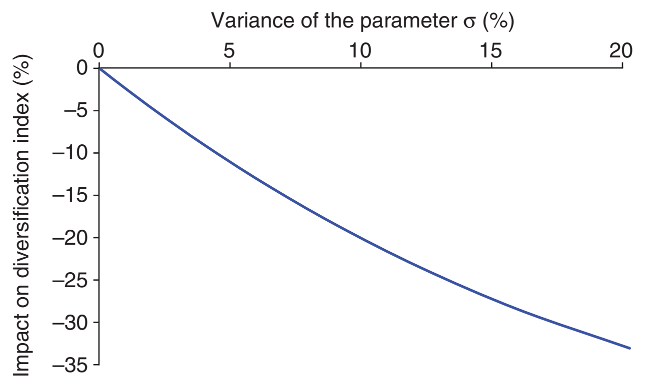

$\sigma = 107\%$ and $v = 18\%$). The normal law assumption for Σ is therefore suitable to model the tail of the loss distribution, whenever $F < 0$. We show later that the shape of the parameters' distribution function is not critical.

The bank's capital charge, $N \cdot L(F_q)$, increases with v. For $\rho = 10\%$, $\sigma = 107\%$ and $v = 18\%$, this increase is equal to $+62\%$. Unsurprisingly, the capital charge is very sensitive to risk dispersion measured by the parameter v, which is a critical parameter of the model. Changing the value of ρ from 10% to 20% leads to a capital charge increase of $+62\%$. The average correlation level is also a critical parameter of the model. As the sum of the stand-alone capital charges is equal to $N \cdot E[\mathrm{e}^{-\Sigma F_q}] = N \cdot \mathrm{e}^{-\sigma F_q + vF_q^2/2}$, the resulting DI is a decreasing function of v. Non-equal parameters σ_i, when included as uncorrelated additional risk, increase the capital charge but generate more diversification, as illustrated in figure 2.

Uncertain correlations

Correlations are not identical to each other, as discussed earlier, but estimating them from real data is a challenge from a statistical point of view. Data is scarce, limited as it is to only one observation per year for the aggregate loss. Estimation of the correlations between the number of events in each cell is no longer robust for the same reason, and severity correlations are only observable for cells that have a sufficient number of events per year. It is current practice to assume identical correlations among cells even if, in reality, correlations are unknown parameters. In what follows, we remain in the limit $N \to \infty$ and correlation uncertainty is included in the model by assuming that the random variable B has an expected value equal to $\beta = \sqrt{\rho}$ and a variance equal to w. For the sake of clarity, we assume that individual risks are all equal among cells: that is, Σ is a constant equal to σ. From (9) we obtain, in the limit $N \to \infty$:

$$\begin{aligned} L(F) &= \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} L_i \\ &= E[\mathrm{e}^{-\sigma BF + \sigma^2(1-B^2)/2} \mid F] \\ &= \int_{-\infty}^{\infty} \mathrm{d}x\, f(x)\mathrm{e}^{-x\sigma F + (1-x^2)\sigma^2/2} \end{aligned} \tag{13}$$

where the function $f(\cdot)$ is the density of the random variable B. If we assume that the variable B is normally distributed, ($B \equiv N(\beta, w)$), we obtain:

$$L(F) = \frac{1}{\sqrt{1+\sigma^2 w}} \exp\bigg(-\beta\sigma F + \frac{(1-\beta^2)\sigma^2}{2} + \frac{\sigma^2 w}{1+\sigma^2 w}\frac{(\beta\sigma + F)^2}{2}\bigg) \tag{14}$$

If the variable B is uniformly distributed between $\beta - \sqrt{3w}$ and $\beta + \sqrt{3w}$ (the bounds are chosen so that the expected value and the variance are equal to β and w, respectively), we obtain:

$$L(F) = \sqrt{\frac{\pi}{6w\sigma^2}}\mathrm{e}^{\sigma^2/2 + F^2/2} \times [N(\sigma(\beta + \sqrt{3w}) + F) - N(\sigma(\beta - \sqrt{3w}) + F)] \tag{15}$$

As pairwise correlations are equal to $\rho_{ij} = \beta_i \cdot \beta_j$, there is a direct link between the variance of ρ_{ij} and the variance of the sensitivity parameters β_i. Because of the independence of the β_i, we have:

$$\begin{aligned} \mathrm{var}(\rho_{ij}) &= E[\beta_i^2 \cdot \beta_j^2] - E[\beta_i]^2 \cdot E[\beta_j]^2 \\ &= w(w + 2\beta^2) \end{aligned}$$

We solve this second-order equation to get the value of w:

$$w = \sqrt{\beta^4 + \mathrm{var}(\rho_{ij})} - \beta^2 \tag{16}$$

For $\beta^2 = \rho = 10\%$ and $\sqrt{\mathrm{var}(\rho_{ij})} = 3\%$ (which is a conservative value compared with that measured from observed data: see above), we have $w = 0.44\%$; that is, the standard deviation of the parameter β is equal to 6.6%. The ratio of the capital charge including model risk ($w > 0$) to the capital charge without model risk ($w = 0$) measures the increase in capital due to dispersion or uncertainty on correlations. We plot this quantity in figure 3 as a function of the mean deviation of the correlation parameter ($\sqrt{w}$); we show that the impact of the mean deviation of the correlation parameter is lower than 2% for $\sqrt{w} = 6.6\%$ for both the normal and uniform assumptions. Additionally, as the curves are very close to each other, we conclude that the shape of the correlation distribution function is not a driver of the capital charge: this validates the choice of the normal law for distribution functions that we have made throughout this article.

Even with a much more conservative choice for the individual cell risk parameter $\sigma = 200\%$, the impact of correlation dispersion on the bank's capital charge would be around $+5\%$. Our conclusion is that correlation dispersion (measured by parameter w) is nowhere near as critical as the other parameters of operational risk models (average cell risk σ, cell risk dispersion v and average correlation level $\beta = \sqrt{\rho}$).

Conclusion

This article pioneers analytical models for computing banks' operational risk capital charges and provides some new results on the correlation problem. These simplified models are fairly realistic as they

3 Capital charge impact as a function of the correlation dispersion for normal and uniform laws ($\beta^2 = 10\%$ and $\sigma = 107\%$)

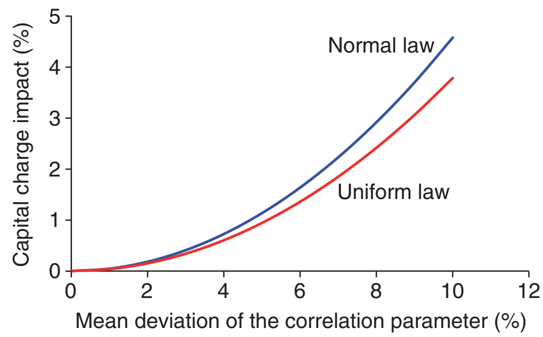

incorporate dispersion in individual cell risks and correlation levels. Our main result is that uniform correlation is a robust assumption for capital charge modelling. This result is important because it means that model risk associated with the value of correlations is not a major issue for capital measurement. The impact of the choice of the copula function and of the average correlation value is much more significant, albeit that calibration suffers because of the scarcity of observed data. At the end of the day, dependence appears to be a subjective choice that determines the diversification benefit at the bank level, and cell loss distribution functions remain the main driver of the capital charge. We emphasise that our approach can straightforwardly be extended to other cell loss distribution or copula functions (Student t for example) in the one-factor framework. The extensions of our analytical approach should, as a second step, focus on broad distribution functions for cell losses and a smaller number of cells. This is left for future research. **R**

Vivien Brunel is head of risk and capital modelling at Société Générale in Paris. **The author would like to thank two anonymous referees for their comments on this paper. Special thanks to Pavel Shevchenko for the very stimulating correspondence. This article reflects the author's opinions and not necessarily those of her employers.**
Email: vivien.brunel@socgen.com

REFERENCES

Aue F and M Kalkbrener, 2007
LDA at work: Deutsche Bank's approach to quantifying operational risk
Journal of Operational Risk 1, pages 49–93

Bank for International Settlements, 2011
Operational risk – supervisory guidelines for the advanced measurement approaches
Report 196, Basel Committee on Banking Supervision

Chernobai A, ST Rachev and F Fabozzi, 2007
Operational Risk: A Guide to Basel II Capital Requirements, Models and Analysis
Wiley

Dutta K and J Perry, 2007
A tale of tails: an empirical analysis of loss distribution models for estimating operational risk capital
Working Paper, Federal Reserve Bank of Boston

Frachot A, T Roncalli and E Salomon, 2004
The correlation problem in operational risk
Groupe de Recherche Opérationnelle, Crédit Agricole, France

Moscadelli M, 2004
The modelling of operational risk: experience with the analysis of the data collected by the Basel committee
Technical Report 517, Banca d'Italia

Neslehova J, P Embrechts and V Chavez-Demoulin, 2006
Infinite mean models and the LDA for operational risk
Journal of Operational Risk 1(1), pages 3–25

Vasicek O, 2002
Loan portfolio value
Risk December, pages 160–162

GUIDELINES FOR THE SUBMISSION OF TECHNICAL ARTICLES

Risk welcomes the submission of technical articles on topics relevant to our readership. Core areas include market and credit risk measurement and management, the pricing and hedging of derivatives and/or structured securities, and the theoretical modelling and empirical observation of markets and portfolios. This list is not exhaustive.

The most important publication criteria are originality, exclusivity and relevance – we attempt to strike a balance between these. Given that *Risk* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the technical editor and his team, submissions are logged and checked against these criteria. Articles that fail to meet the criteria are rejected at this stage.

Articles are then sent without author details to one or more anonymous referees for peer review. Our referees are drawn from the research groups, risk management departments and trading desks of major financial institutions, in addition to academia. Many have already published articles in *Risk*. Authors should allow four to eight weeks for the refereeing process. Depending on the feedback from referees, the author may attempt to revise the manuscript. Based on this process, the technical editor makes a decision to reject or accept the submitted article. His decision is final.

Submissions should be sent to the technical team, preferably by email addressed to technical@incisivemedia.com. LaTeX is the preferred format for submissions (your LaTeX file should be accompanied by a PDF of the output from it), although Word documents are also acceptable. It is helpful for graphs and figures to be submitted as separate Excel, postscript or EPS files. Vector graphics are strongly preferred, but rasterised images with a minimum resolution of 300dpi are sometimes acceptable.

The maximum recommended length for articles is 3,500 words, with some allowance for charts and/or formulas. We expect all articles to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations.