

Dependent Events and Operational Risk

Miro R. Powojowski, Diane Reynolds and Hans J.H. Tuenter

Most people acknowledge that some correlation between operational losses exists. Modelling it, however, is a complex issue. The questions arise: How can this be done? Is it beneficial? As is shown, including correlations in the calculation of operational capital can have a significant effect on the level of reserve capital required to cushion against the risk of operational losses. This article focuses on positively correlated events that arise from a single root cause. It also provides an example of how such events might be incorporated into operational risk measures.

Intuitively, operational loss events are correlated. Consider, for example, an electrical failure. This failure might affect all desks on the trading floor. Each desk would need to file a claim for lost income during the downtime and possibly for time required to ensure that all previous trades have been processed correctly. Clearly, this single event must be capitalized, but by which profit center? In attempting to allocate capital fairly among businesses, a more complex model is required—one that allows for common events to affect multiple business units.

In fact, as shall be seen, failure to acknowledge and account for this type of correlation between losses experienced in different businesses—and other positive correlations between events—leads to an understatement of the capital reserves that are required. One might argue that this increase in capital charge is undesirable. However, the goal of risk-sensitive capital reserves should be just that: to reflect properly all of the risk assumed. Also, it is entirely conceivable that charges related to correlations will become part

of the regulatory process, as is the case for credit risk.

This article presents a simple model that allows for the incorporation of positive correlations between operational units. However, before the model can be specified, some groundwork is required.

First, one must define a set of operational units to be modelled. For a complete definition of operational units, see Reynolds and Syer (2002a). Common examples include the event types and business lines—defined by the Basel Commission on Banking Supervision—actual organizational structures, and geographical locations. Each operational unit is allocated its related data. Typically, this would be loss data, but other data may also be used.

Second, an analysis of the data is done to determine whether one operational loss process per operational unit is sufficient, or whether several processes should be used to obtain a better fit to the data. Frequently, however, due to the paucity of data, this analysis results in each operational unit being modelled by a

single loss process, and independently of the other units.

As with many advanced approaches to the measurement of operational risk, an actuarial model is used. This is described in Reynolds and Syer (2002b). This approach divides the operational loss process into two components: frequency, which is the number of losses per period, and severity, which is the size of an individual loss. Each line of business, or event type, is assigned its own model for frequency and severity of losses that occur within an operational unit.

By correlating the frequency of loss across operational units, one effectively correlates the total loss across operational units. The model presented here focuses on the correlation of loss frequency across operational units through an underlying “common cause” approach.

The operational loss processes are modelled by representing the loss frequency as a combination of Poisson distributions, while allowing each severity distribution to follow any data-appropriate distribution. A discussion of the Poisson distribution and its principal properties provides essential background to understanding the overall model and its calibration.

Following an introduction to the Poisson distribution, the joint-frequency model is presented and discussed using a small example. A method of calibrating the joint model is then presented, and the example is extended to enrich the discussion. The paper concludes with some thoughts for future directions of research.

Poisson distribution

Named after the French mathematician Siméon Poisson (1781–1840), the Poisson distribution describes the number of events that occur in a given interval of time, when the probability of the event occurring is very small, but the number of trials is very large. The Poisson distribution is formally defined as

$$P(n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}, \quad n = 0, 1, 2, \dots,$$

where $P(n)$ is the probability that n events happen over a time interval of length t , and λ is the intensity or rate at which events happen per unit of time.

The distribution was studied more rigorously by the Russian-born statistician Ladislaus von Bortkiewicz (1868–1931), in his monograph, *Das Gesetz der kleinen Zahlen*. In his book (Bortkiewicz 1898), one can find an early, and perhaps the first, application of the Poisson distribution to operational risk. Von Bortkiewicz analysed data, collected on the number of soldiers kicked to death each year by horses in the Prussian army, and found an extremely good fit with the Poisson distribution.

Although today’s applications of the Poisson distribution to model the frequency of operational losses are less dramatic, it is still a viable and widely used model. One reason is that it is a simple model, since only one parameter must be estimated to specify the Poisson distribution—its mean. The property that the variance is equal to the mean is important for ease of calibration, as is seen in a subsequent section.

Another important property of the Poisson distribution is that it is a stable distribution: adding two Poisson processes with intensities λ_1 and λ_2 creates another Poisson process with intensity $\lambda_1 + \lambda_2$.

The Poisson distribution is well documented and, hence, one may draw on a large body of existing knowledge, see for instance, Haight (1967), Johnson and Kotz (1969) and Kingman (1993).

Dependent processes with Poisson marginals

Consider a set of operational events. Modelling many types of events simultaneously, to account for many types of operational loss processes and many business initiatives, is critical. This section provides a discussion of a bivariate model, and then generalizes it to a multivariate formulation. It concludes with a worked example based on the event types dictated by the Quantitative

Impact Study (QIS) of the Basel Committee on Banking Supervision (BCBS). See the study by the Bank of International Settlements (2001).

Two operational loss processes

First, consider the events originating from two operational loss processes Y_1 and Y_2 , and model these events by two separate Poisson distributions. For a fixed period of time, over which the events are recorded (typically one month or one year), let N_1 and N_2 be the number of events of the operational processes Y_1 and Y_2 , respectively.

Under these assumptions, N_1 will follow a Poisson distribution with intensity μ_1 , so that the mean and the variance of the number of events are both equal to μ_1 . The same applies to N_2 , with corresponding intensity μ_2 .

To complete the specification of the model, it is necessary to specify the dependence between N_1 and N_2 . The simplest assumption is, of course, to assume that they are independent. But, what if experience indicates that the events in Y_2 tend to coincide with the events in Y_1 ? From a risk management perspective, the ability to model simultaneous events in the processes Y_1 and Y_2 then becomes paramount.

There is a need for multivariate arrival processes that are capable of modelling the joint events in many processes. Individually, the events in each process remain Poisson distributed with a separate occurrence rate. One way to do this is to create three underlying loss processes. Assume X_1 , X_2 and X_3 are independent Poisson processes with intensities λ_1 , λ_2 and λ_3 , respectively. If

$$\begin{aligned} Y_1 &= X_1 + X_3, \\ Y_2 &= X_2 + X_3 \end{aligned} \tag{1}$$

is taken as the underlying relationship, a two-dimensional process that exhibits dependence is obtained. Denote the number of events of type X_1 , X_2 and X_3 by M_1 , M_2 and M_3 , respectively. The distributions of

$$\begin{aligned} N_1 &= M_1 + M_3, \\ N_2 &= M_2 + M_3 \end{aligned} \tag{2}$$

each have Poisson marginal distributions, with mean $\mu_1 = \lambda_1 + \lambda_3$ and $\mu_2 = \lambda_2 + \lambda_3$, respectively. However, the two components are now dependent. Their covariance and correlation coefficient are given by:

$$\text{Cov}(N_1, N_2) = \text{Var}(M_3) = \lambda_3,$$

and

$$\rho = \frac{\lambda_3}{\sqrt{(\lambda_1 + \lambda_3)(\lambda_2 + \lambda_3)}}.$$

From this one sees that, by introducing a set of underlying, abstract variables, each following a Poisson distribution, one can create dependent distributions for the frequencies of two loss processes, while maintaining their marginal distributions as Poisson.

Multiple operational loss processes

The example easily generalizes to the multi-dimensional case, where there are several processes. Suppose one has n observed processes Y_j , $j = 1, \dots, n$, for which one wants to model a dependent structure. To do this, m underlying Poisson processes are considered, with intensities λ_i , $i = 1, \dots, m$, and a corresponding number of events M_i . Each of these underlying processes can be assigned to one or more of the observed processes, which can be captured by introducing the indicator variables δ_{ij} :

$$N_j = \sum_{i=1}^m \delta_{ij} M_i.$$

The number of events N_j of process Y_j then follows a Poisson distribution with intensity

$$\mu_j = \sum_{i=1}^m \delta_{ij} \lambda_i. \tag{3}$$

This dependence model has an intuitive interpretation: it postulates the existence of

events that affect more than one operational unit. This type of event may be thought of as a common-cause event, affecting all processes Y_j . The covariance and correlation coefficients are readily determined as:

$$\text{Cov}(N_j, N_k) = \sum_{i=1}^m \delta_{ij} \lambda_i \delta_{ik}$$

and,

$$\rho_{jk} = \frac{\sum_{i=1}^m \delta_{ij} \lambda_i \delta_{ik}}{\sqrt{\sum_{i=1}^m \delta_{ij} \lambda_i \times \sum_{i=1}^m \delta_{ik} \lambda_i}} \quad (4)$$

The covariance structure is more easily described in matrix notation as

$$\mathbf{C} = \Delta \Lambda \Delta^T, \quad (5)$$

where \mathbf{C} is the covariance matrix, Δ is an $n \times m$ incidence matrix, describing the relationship between the observed and the underlying processes, and $\Lambda = \text{diag}(\lambda_i)$ is a diagonal matrix with the intensities λ_i as the elements. Note that all elements of the covariance matrix are nonnegative (and, hence, also all correlations), because all the elements of the matrices Δ and Λ are nonnegative.

Monte-Carlo simulations

To illustrate the model, we employ the same classification scheme used by the BCBS in collecting data for QIS. It divides risk into seven broad categories:

- Internal Fraud
- External Fraud
- Employment Practices
- Business Services
- Physical Assets
- Business Disruption, and
- Process Management.

For the sake of completeness, an eighth category is included, “Other Risks,” to allow for anything that does not fall into one of the QIS categories.

Independent events

First, assume that there are eight, independent, underlying loss processes: one specific to each risk category, and that all processes have a Poisson distribution with mean $\lambda = 0.5$. Thus, the expected frequency of events is once every two years. Given an event, let the resulting losses be independent and normally distributed with a mean of US \$4M and a standard deviation of US \$0.5M for all processes. This situation is depicted in Figure 1.

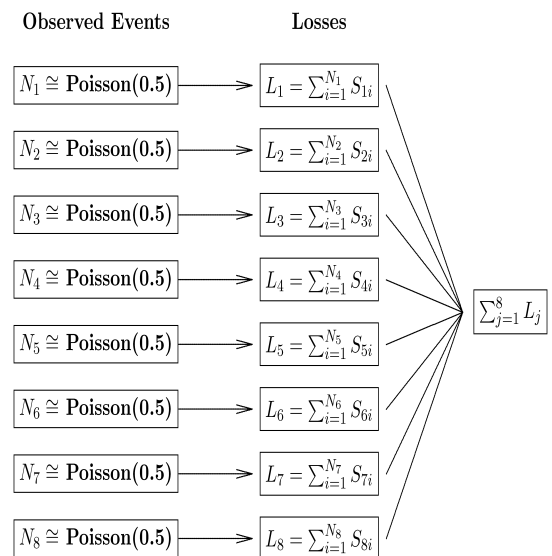


Figure 1: Independent loss processes; each process follows a Poisson distribution with intensity $\lambda = 0.5$

To gain insight into the shape of the distribution of the company-wide losses, a Monte-Carlo simulation with 10,000 scenarios was performed. The resulting, empirical distribution is depicted in Figure 2. The 99% quantile of this distribution is US \$37.82M.

Dependent events

To create dependencies between the processes, take the previous model, but with intensities 0.4

Independent Events

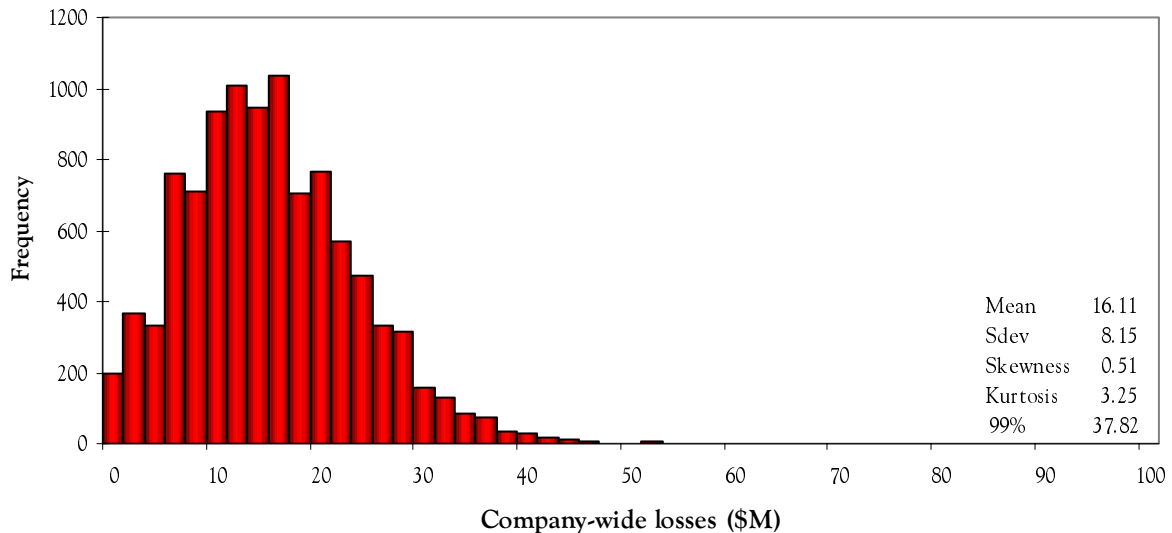


Figure 2: Empirical distribution of company-wide losses under the independent model, using a Monte Carlo simulation with 10,000 scenarios; the 99% quantile is \$37.82M

for all the independent loss processes, and introduce a single enterprise-wide source of loss, modelled by a Poisson distribution with intensity $\lambda_0 = 0.1$. The number of losses, experienced in each risk class, is now a combination of the overall number of losses experienced due to firm-wide issues and the losses specific to the risk class. This model is depicted in Figure 3.

Note that the marginal distributions of the events are still Poisson with one event expected every two years. However, the losses are now dependent with a pairwise correlation of 0.2, as calculated using Equation 4. This implies that differences in the loss distribution are due solely to the effect of correlations and not to changes in the marginal distribution.

Another Monte-Carlo simulation with 10,000 scenarios was performed, and the resulting empirical distribution is shown in Figure 4. The 99% quantile of this distribution is US \$59.19M, while the expected losses remain the same as in the independent case.

Comparison

It is evident that the overall budget requirements, as measured by the expected losses, are not affected by the correlation between events. However, the risk, as measured by the unexpected losses, has increased considerably by the inclusion of these firm-wide events. A cursory inspection of Figures 2 and 4 shows that the latter distribution has a considerably “fatter” tail.

The example also shows that the effect of dependence between processes on risk measures, such as Value-at-Risk (VaR), may be considerable even when the correlation between events is relatively low.

Parameter estimation

As with any model designed for practical applications, the value of the model is influenced significantly by one’s ability to calibrate and implement it as well as its ability to represent reality. In this case, the model is easily implemented, once the parameters are known, and it is possible to estimate the parameters using common methods.

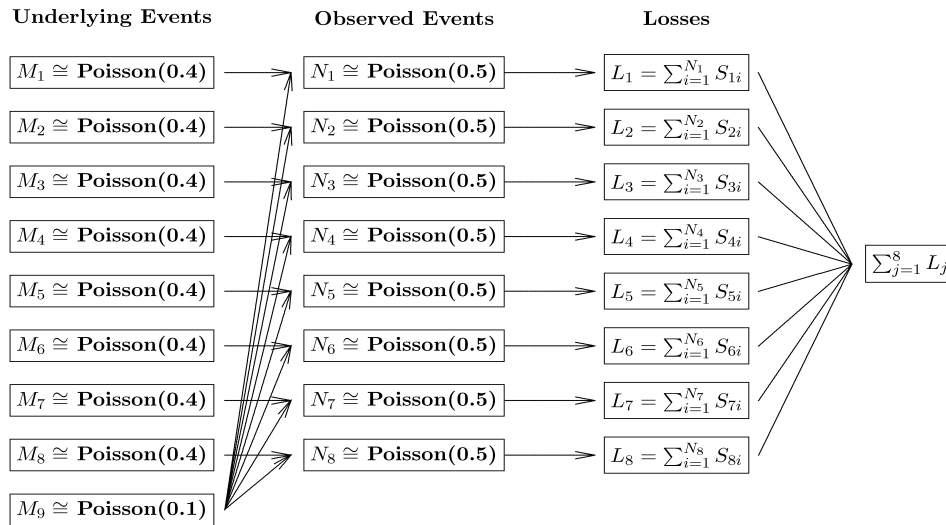


Figure 3: Dependent loss processes; loss processes specific to a risk class follow a Poisson distribution with intensity $\lambda = 0.4$. The firm-wide loss process follows a Poisson process with intensity $\lambda = 0.1$

In this section, the calibration method is discussed and illustrated by using the Monte-Carlo simulation results for the two models considered previously.

It is well known that the maximum likelihood estimate for the intensity of a Poisson process is the sample mean of the observed values. Ideally, one would like to use the maximum likelihood method to obtain estimates for the parameters in the multivariate processes. However, the expressions for the likelihood function are not always that easily derived, and the resulting expressions may be difficult to use in practice.

The approach used here is to utilize the maximum likelihood estimates for the intensities of the observed processes, and then incorporate these estimates into an optimization problem.

Let $\hat{\mu}_j$ be the sample mean for the events of the j -th observed process, and thus the maximum likelihood estimate for the intensity μ_j . For ease of exposition, assume, for the time being, that the system

$$\begin{aligned} \Delta\lambda &= \hat{\mu}, \\ \lambda &\geq 0 \end{aligned} \tag{6}$$

is solvable. Note that this approach ensures that the observed intensities are matched.

Given the above system of equations, one now wants to determine the parameter set that most closely matches the higher moments, as captured by the covariance matrix. This is done by formulating it as an optimization problem over a distance function

$$\begin{aligned} \min \quad & \|\hat{C} - \Delta\Lambda\Delta^T\| \\ \text{s.t.} \quad & \Delta\lambda = \hat{\mu}, \\ & \lambda \geq 0. \end{aligned} \tag{7}$$

If the Frobenius norm is chosen as the distance norm—the sum of squared distances of the matrix elements—a routine quadratic optimization problem is obtained, which is readily solved by standard optimization software.

Illustrative applications

Now the calibration method is applied to the models discussed previously.

Two processes

For the structure given in Equations 1 and 2, the incidence matrix is given by

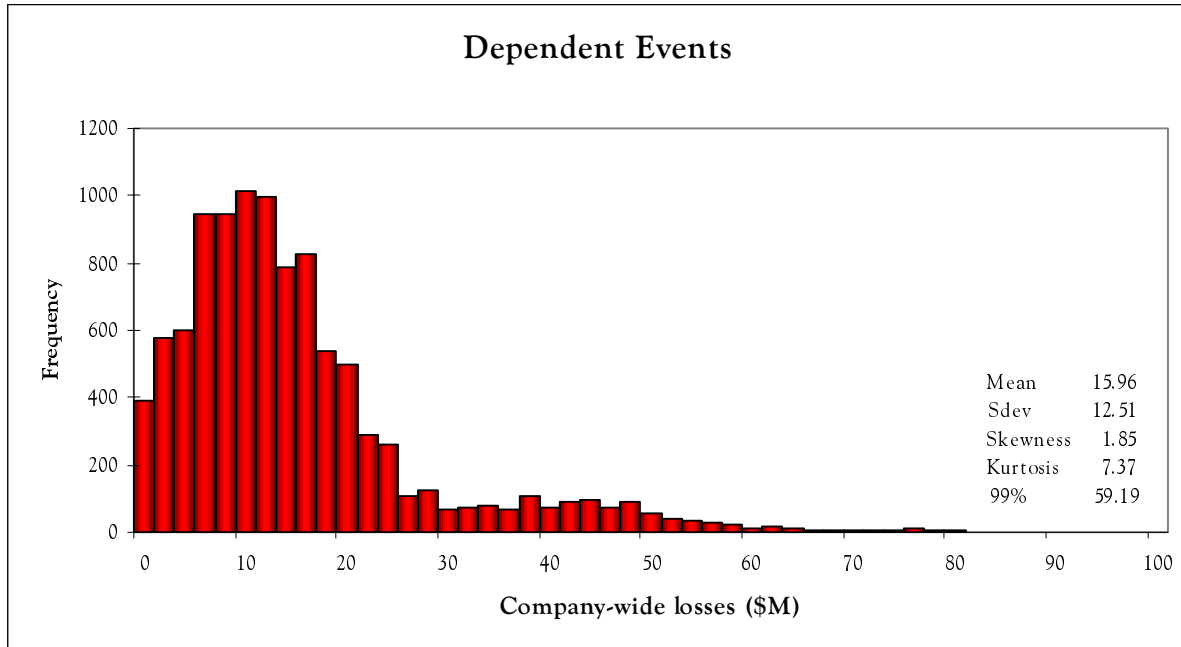


Figure 4: Empirical distribution of company-wide losses under the dependent model, using a Monte-Carlo simulation with 10,000 scenarios; the 99% quantile is \$59.19M

$$\Delta = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \tag{8}$$

and the covariance matrix by

$$C = \begin{bmatrix} \lambda_1 + \lambda_3 & \lambda_3 \\ \lambda_3 & \lambda_2 + \lambda_3 \end{bmatrix}. \tag{9}$$

The optimization problem becomes:

$$\begin{aligned} \min \quad & (\hat{C}_{11} - \lambda_1 - \lambda_3)^2 + 2(\hat{C}_{12} - \lambda_3)^2 \\ & + (\hat{C}_{22} - \lambda_2 - \lambda_3)^2 \\ \text{s.t.} \quad & \lambda_1 + \lambda_3 = \hat{\mu}_1, \\ & \lambda_2 + \lambda_3 = \hat{\mu}_2, \\ & \lambda_1, \lambda_2, \lambda_3 \geq 0. \end{aligned}$$

Because of the equality constraints, this particular instance is an optimization problem over λ_3 only, and is readily solved.

Independent example

In the independent model, shown in Figure 1, the incidence matrix is the identity matrix, and the solution to the optimization problem is given by the sample means of the observed events.

Dependent example

In the dependent model, shown in Figure 3, the incidence matrix is given by

$$\Delta = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}, \tag{10}$$

and the covariance matrix by

$$C = \text{diag}(\lambda_i) + \lambda_9 E,$$

where E is a matrix consisting of all ones.

The optimization problem becomes:

$$\begin{aligned} \min \quad & \sum_{i=1}^8 \sum_{j \neq i} (\hat{C}_{ij} - \lambda_9)^2 + \sum_{i=1}^8 (\hat{C}_{ii} - \lambda_i - \lambda_9)^2 \\ \text{s.t.} \quad & \lambda_i + \lambda_9 = \hat{\mu}_i, \quad i = 1, \dots, 8, \\ & \lambda_i \geq 0, \text{ for all } i. \end{aligned}$$

As before, by virtue of the equality constraints, this is an optimization problem over λ_9 only, and is readily solved numerically.

Numerical examples

To illustrate the optimization approach, 50 scenarios were taken that were generated under the “dependent event” model given in Figure 3. The sample covariance matrix was computed, and the optimization performed to determine estimates for the intensities. This resulted in

$$(\lambda_1, \dots, \lambda_9) = (0.38, 0.48, 0.40, 0.56, \\ 0.44, 0.38, 0.26, 0.38, 0.12).$$

The procedure was repeated with 50 scenarios generated under the “independent event” model given in Figure 1, and the following result was obtained:

$$(\lambda_1, \dots, \lambda_9) = (0.32, 0.32, 0.42, 0.58, \\ 0.50, 0.50, 0.44, 0.52, 0.00).$$

Note that, even though the dependent model was applied to scenarios generated from the independent model, the correct conclusion that there is no common cause of event is obtained.

Conclusions

Correlations have a significant impact on capital calculations. In order to ensure that risk-sensitive capital allocations are fair to all businesses, special models are required. One must be able to calculate correlations accurately and defensibly, and to use this information in determining and allocating capital. This paper presents a simple and easily calibrated model for including positive correlations. It demonstrates, by means of an example, the significant influence that positive correlations can have on required capital. Models such as these, which allow intuitive relationships to be modelled to provide more risk-sensitive capital allocation, need to be examined in greater detail if the quantitative measures of operational capital are to achieve credibility within the industry.

Acknowledgements

The authors would like to thank Thomas Eigl and Helmut Mausser for comments and helpful discussions.

References

- Bank of International Settlements, 2001, “Quantitative Impact Study of the Basel Committee on Banking Supervision,” <http://www.bis.org/publ/bcbsca.htm>.
- Bortkiewicz, L. von, 1898, *Das Gesetz der kleinen Zahlen*, Leipzig: B.G. Teubner.
- Haight, F.A. 1967, *Handbook of the Poisson distribution*, New York, John Wiley & Sons.
- Johnson, N.L. and S. Kotz, 1969, *Distributions in Statistics: Discrete Distributions*, Boston, MA: Houghton Mifflin Company.
- Kingman, J.F.C., 1993, “Poisson processes,” *Oxford Studies in Probability*, Vol. 3, New York, NY: Clarendon Press.
- Reynolds, D. and D. Syer, 2002a, “Mark to future and operational risk,” *Algo Research Quarterly*, 5(2): 39–51. (This issue)
- Reynolds, D. and D. Syer, 2002b, “Algo Academy Notes,” *Algo Research Quarterly*, 5(2): 31–37. (This issue)

Endnotes

1. The feasibility of the system of equations in Equation 6 can be verified by an application of Farkas’ lemma. In the case that this system does not have a solution, we can incorporate the first constraint in the objective function, using Lagrangian multipliers or a penalty method. These modifications also result in a quadratic optimization problem.
2. The optimization formulation easily allows one to incorporate preferences regarding the relative importance of elements of the sample covariance matrix by introducing weights into the elements of the objective function. Further, the optimization formulation allows one to examine the effect of different configurations of the underlying processes by taking different incidence matrices. This is a feature that a maximum likelihood method does not have.
3. To be able to use the model, one has to specify the number, m , of underlying processes, and decide how the underlying processes influence

the observed processes. Although potentially there are $m2^n$ possible ways of doing this, one needs to exercise caution. The paucity of data in the setting of operational risk does not allow for the estimation of too many parameters, and,

hence, there is the need to trim down the model. In most cases, the business analysis will help to determine the number of underlying processes and which processes they influence.