Integrated Risk Modelling

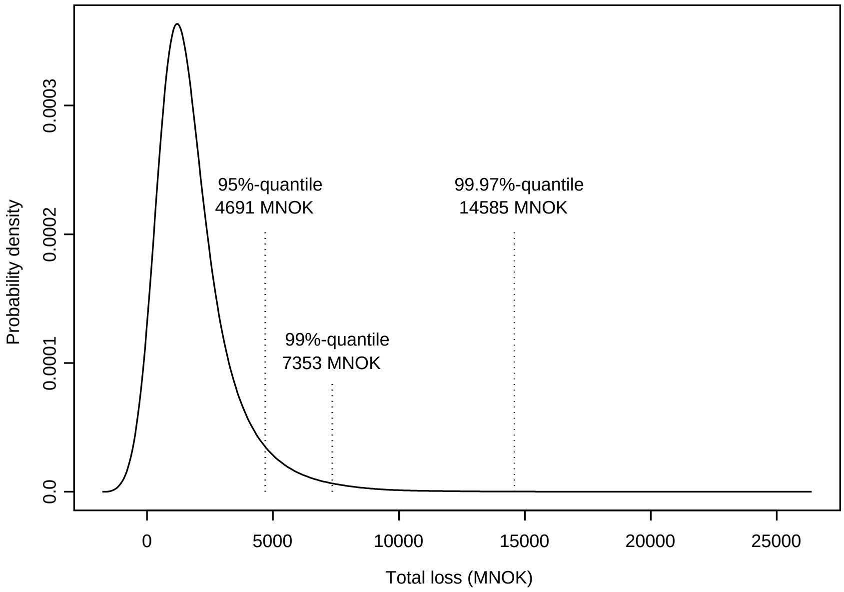

ISBN 82-539-0506-8

Xeni K. Dimakos
Kjersti Aas

18th December 2003

NR Norwegian Computing Center
APPLIED RESEARCH AND DEVELOPMENT

NR Report

Title: Integrated Risk Modelling

Author: Xeni K. Dimakos
Kjersti Aas

Date: 18th December 2003
Year: 2003
ISBN 82-539-0506-8
Publication no.: 998

Abstract: In this paper we present a new approach to modelling the total economic capital required to protect a financial institution against possible losses. The approach takes into account the correlation between risk types and in this respect it improves upon the conventional practice that assumes perfectly correlated risks. A statistical model is built and Monte Carlo simulation is used to estimate the total loss distribution. The methodology has been implemented in the Norwegian financial group DnB's system for risk management. Incorporating current expert knowledge of relationships between risks, rather than taking the most conservative stand, gives a 20 % reduction in the total economic capital for a one-year time horizon.

Keywords: Total risk, credit risk, market risk, operational risk, correlation, simulation

Target group: Financial risk managers

Availability: Open

Project: StAR

Project no.: 830105

Research field: Financial and Insurance Risk

No. of pages: 13

Norwegian Computing Center
Gaustadalléen 23, P.O. Box 114 Blindern, NO-0314 Oslo, Norway
Telephone: +47 2285 2500, telefax: +47 2269 7660, `http://www.nr.no`

Contents

1 Introduction

Many financial institutions and also larger companies, have developed formal methods to deal with risk. Typically capital reserving is carried out for different risk types individually and then added to the buffer for the whole corporation (Bock, 2000*a*). Such an approach implicitly assumes that all risks are perfectly correlated, i.e. that great losses occur simultaneously. Generally, the buffer capital (or economic capital) to meet a certain solvency standard is smaller than calculated from the perfect correlation hypothesis. A sharper evaluation would lower the required buffer capital and increase corporation competitiveness and profit.

The purpose of this paper is to present a framework for these more realistic assumptions through joint modelling of risk types and their correlations. Central ingredients of our approach are simple correlations between pairs of risks, empirical modelling, and Monte Carlo simulations as the technical tool. Risk is in our setting defined as potential large losses, and our main interest is in the upper tail of the loss distributions. Using Monte Carlo methods we are not confined to single quantities, such as VaR, and may support our analysis with a variety of measures. Sensitivity analysis and stress tests may easily be performed within the suggested framework by varying the parameters to the model. The methodology was developed in close cooperation with the Norwegian financial group DnB. It is implemented in their system for risk management, which was used by Moody's as an example of a sophisticated risk modelling system at the Nordic and Baltic Banking, Monthly Bank Forum London, April 30, 2001.

Integrated risk modelling is also the topic in Kuritzkes et al. (2002), Medova & Smith (2003) and Ward & Lee (2002). Kuritzkes et al. (2002) discuss different levels of complexity of integrated models. Medova & Smith (2003) focus on credit and market risk only, while Ward & Lee (2002) also include operational risk and use a copula approach.

DnB is like other financial institutions exposed to credit risk, market risk, and operational risk. In addition, DnB faces the risk that stems from its ownership in the life insurance company Vital. This is due to external parameters for life insurance operations in Norway, i.e. regulations of risk and profit sharing between policyholders and owner. As this ownership risk is special to DnB and of less general interest, we have omitted the description of this risk and the corresponding methodology from this paper. However, our numerical results include also this risk component.

The focus of this paper is the integration of different risks to assess the total risk. For marginal evaluation of credit and market risk most banks are equipped with advanced risk assessment software. Systems for dealing with operational risk, are less common. A risk manager concerned with the total risk needs to determine whether a system for total risk is to be complex to accommodate the different marginal systems or simple and based on the essential features of the existing systems. The model we suggest uses the latter principle. It is not the aim of this paper to present detailed models for marginal risks, but to integrate the essential features of marginal models to model the total risk.

Each of the risks has distinct statistical properties. However, by using a common confidence interval and common time horizon, the economic capital requirements for different risk factors can be made equivalent and it is possible to assess the risk for a financial institution. For banking, the convention is to adapt to a one-year time horizon, and the economic capital requirement is set to protect losses over one year at the 99.97% level, which is roughly equivalent to an "AA" rating from leading rating agencies.

The paper is organized as a case study of the risk faced by DnB, and the modelling is tailor-made to reflect the history of this particular institution, the data it has access to and how it perceives dependency between risks. However, the general approach is applicable in a broader context. The model is presented in Section 2 and the implementation and numerical results in Sections 3 and 4.

2 The Model

2.1 General

One of the fundamental technical hurdles in risk management for a financial institution is the aggregation of a diverse set of risks. The total loss T of a financial institution in one year is

$$T = C + O + M \tag{1}$$

where C, O, and M are the yearly losses due to credit, operational, and market risk. Typically the risk manager has some knowledge of the marginal distribution of each risk type and is interested in the joint distribution, i.e. the distribution of all risks together. Aggregating these risks is technically quite challenging. Since the underlying risk distributions for each risk type not necessarily follow the same distributional form, it is necessary to do a numerical integration or simulation to convolve them. Only a few approaches for aggregating credit, market and operational risk have been proposed in the literature. In Kuritzkes et al. (2002) a highly simplified approach is given, in which one assumes that all risks are jointly normally distributed. Another method that recently has become very popular in finance, is using copulas to link the marginals to the joint distribution. Embrechts et al. (1999) were among the first to introduce this toolkit to the finance literature. To our knowledge the only present use of copula to aggregate credit, market, and operational risk can be found in Ward & Lee (2002) who use a multivariate Gaussian copula to aggregate marginal distributions.

Another challenge in integrated risk management is specifying a common time horizon for all the risk types. Market risk is typically measured on a daily basis. Credit risk is typically calibrated to a one-year horizon, as is operational risk. We use the convention for modelling risks and assessing capital in banks, which is to adopt to a one-year horizon, see Kuritzkes et al. (2002). A one-year time horizon is reasonable as it corresponds to the internal capital allocation and budgeting cycle, it is a period in which an institution can access the markets for addition capital, and it is also the horizon used in the New Basel Accord.

As far as the aggregation of risk types is concerned, we use simulation to convolve the marginal distributions. Upon sampling C, O, M from the simultaneous distribution $P(C, O, M)$

we construct the total loss T according to (1). The multiplication law of probability from standard statistical theory states that the joint probability of the credit, market, and operational losses $P(C, O, M)$ can be decomposed as follows

$$P(C, O, M) = P(C)\, P(O|C)\, P(M|C, O). \tag{2}$$

The motivation for taking the credit risk as the base risk to which all the other risk types are linked, is that it is the dominant risk for the financial institution we consider. Note that (2) does not imply that we try to predict the operational loss and market loss from the credit loss. To simplify (2), we assume that

$$P(M|C, O) = P(M|C), \tag{3}$$

which expresses that market and operational losses are *conditionally* independent. That is, once the credit loss is known, any evidence about the operational loss does not change our belief about the market loss. Combining (3) with (2) yields

$$P(C, O, M) = P(C)\, P(O|C)\, P(M|C). \tag{4}$$

If we are able to simulate from the three probability distributions on the right, we will be able to obtain the joint distribution of risk types. In Sections 2.2–2.4 we will describe how $P(C)$, $P(O|C)$, and $P(M|C)$ are obtained.

2.2 Credit Risk

The credit risk faced by DnB is defined as the risk of losses resulting from failure by its financial counterparties to meet their obligations. We use the existing credit model of DnB, which is similar to the approach used in Ward & Lee (2002), but different from the approach used by other credit risk systems (Crosbie, 1997; Goupton et al., 1997). It can be summarized as follows. First, credit commitments are divided into ten categories based on expected default frequency. The classification of commitments is based on both financial factors and non-financial factors such as managerial aspects. The next step is to estimate exposure, i.e. the size of the claim at the time of a potential default and how much would be repaid to the bank. The expected loss for each commitment is calculated as the product of the expected default frequency, exposure at default and the loss ratio. The commitments are then added to obtain the expected loss μ for the portfolio.

The risk level in the portfolio depends on how the losses occur relative to each other. The credit model specifies how each loan contributes to the overall risk through an assumed correlation between the individual commitment and the total credit losses. Based on these correlations it is possible to compute the standard deviation σ of the credit portfolio. In our setting, the credit part of the total risk takes as input μ and σ which are outputs from DnB's credit management systems.

To be able to simulate from our model we need the whole distribution and not only the mean and standard deviation. In the DnB model, we work with the credit loss rate R, which is the total credit loss C of the institution divided by the total exposure e. Following the same reasoning as presented in Ward & Lee (2002), we have chosen to use a beta distribution to model the portfolio of correlated loans.

More specifically, the probability density of R is:

$$b(r) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} r^{\alpha-1}(1-r)^{\beta-1}, \quad 0 < r < 1, \tag{5}$$

where the Gamma function $\Gamma(\alpha)$ is defined by

$$\Gamma(\alpha) = \int_0^\infty t^{\alpha-1}\, e^{-t}\, dt$$

for $\alpha > 0$.

The beta distribution is fully determined by the two parameters α and β. These can be obtained from the expectation $\mu' = \mu/e$ and standard deviation $\sigma' = \sigma/e$ of the loss ratio R using the following relationships

$$\alpha = (1-\mu')\left(\frac{\mu'}{\sigma'}\right)^2 - \mu'$$

and

$$\beta = \frac{\alpha}{\mu'} - \alpha.$$

2.3 Operational Risk

The operational risk of DnB is due to direct and indirect losses caused by internal factors such as inadequate or ineffective internal processes and systems, and by external events such as natural disasters and criminal acts. Some of these losses occur frequently, but are of moderate value, whereas others are rare, but very large. This suggests heavy-tailed models, and extreme value theory (EVT) has been proposed in Cruz et al. (1998); Medova (2000) and Kyriacou & Medova (2000).

As the size and quality of DnB's database on operational losses does not allow reliable estimation of the parameters in an EVT-model, we had to rely on expert opinions and subjective choices. The risk managers felt that they had a relatively clear opinion on the following three quantities; the size of the most frequent loss m, the risk-adjusted capital o_p needed to cover operational risk (here the institution follows international industry benchmark, which has also been acknowledged by the Basel Committee (BIS (2001)), and let operational risk represent around 20% of overall capital requirement), and the correlation ν between operational and credit risk. The latter is assumed to be positive, since operational errors linked to credit activity will often not appear until the credit risk materializes.

The most frequent loss m can be related to the mode of the operational loss distribution, and the risk-adjusted capital o_p to the 99.97% quantile. The latter is linked to the rating of the institution, e.g. an S&P rating of "AA" corresponds to an average default probability of the institution of 0.03%. Hence, a convenient distribution for operational loss is one determined by two parameters, the mode and the 99.97% quantile, that is heavy-tailed and that may easily be linked to the beta distribution for the credit risk. The lognormal distribution was found to be a reasonable choice. One of the approaches suggested by the Basel Committee BIS (2001) is to simulate the number of operational loss events one year from a Poisson distribution and the severity of these events from a lognormal distribution and compute the total operational loss as the sum of the individual events. Using this approach the total operational losses is distributed as a sum of lognormals. Approximating this distribution with another lognormal distribution, like we do, is a crude approximation.

The dependency between operational and credit losses is achieved by transforming C and O to standard normal variables X and Y, specifying a correlation ν between these variables, and retransforming to the correct distributions. As noted by Ward & Lee (2002) this is a copula approach. Even if a normal copula implies independence (correlation breakdown) in the tails (see Bradley & Taqqu (2003)), this is further out in the tail than the 99.97% quantile that we are interested in and does not prevent it from being used in the calculation of capital at risk. Using our approach, the actual correlation between C and O will be smaller than ν (Kendall & Stuart, 1979) (see p. 600), but this can be compensated for by adjusting the correlation between X and Y upwards.

In detail the modelling is as follows. The credit loss C, as described in Section 2.2, can be represented as

$$C = e\, B^{-1}\{\Phi(X)\} \tag{6}$$

where $B^{-1}(\cdot)$ is the inverse cumulative beta distribution, $\Phi(\cdot)$ is the cumulative standard normal distribution function, and X is a standard normal variable. Moreover, the lognormally distributed operational loss O can be written

$$O = \exp(\xi + \tau\, Y), \tag{7}$$

where Y is another standard normal variable and ξ and τ are the parameters of the lognormal distribution. By specifying a correlation between X and Y, O and C become dependent.

It remains to specify the parameters ξ and τ in (7). In the absence of appropriate historical data the parameters are determined from a subjective assessment of the 99.97% quantile o_p and the mode m by solving the equations

$$m = \exp(\xi - \tau^2) \tag{8}$$

and

$$o_p = \exp\{\xi + \tau\, \Phi^{-1}(p)\}, \tag{9}$$

where $\Phi^{-1}(p)$ is the 99.97% quantile of the standard normal distribution.

Due to limited information on historical losses, the approach for modelling operational risk presented in this paper must be considered as preliminary. DnB has started the work of establishing a database for operational losses (possibly in collaboration with other financial institutions), and to develop risk indicators for use in monitoring operational risk. The plan is to refine the current model as soon as the data material is considered to be sufficiently large and of satisfactory quality.

2.4 Market Risk

Market risk is a consequence of the open positions of the financial institution in the foreign exchange, interest rate and capital markets, and the risk is linked to fluctuations in market prices and exchange rates. The total market risk of DnB is composed of risks associated with 15 different instruments representing different types of market risk, e.g. the equity, foreign exchange, interest rate and commodities markets and it is managed through limits for the various types. DnB makes a distinction between market risks stemming from the brokerage activities of the financial institution, mainly in the interest rate and currency markets, and from banking activities, where investments have a longer-term perspective, e.g. in equity instruments. The trading books' part of the portfolio only constitutes only 50 MNOK out of a total of approximately 3000 MNOK for the whole market portfolio.

Market risk is typically measured with VaR on a short time horizon such as 10 days, see BIS (1995), assuming that market liquidity will always be sufficient to allow positions to be closed out at minimal losses. While effective in short-term trading environment, the usefulness of VaR begins to break down when it is employed to measure the market risk generated by long-term operations. Hickman et al. (2002) gives four reasons why VaR falls short as a long-term metric. One of the reasons is that VaR fails to reflect the effect of management intervention policies (i.e. stop-loss limits) that can substantially limit the cumulative effect of losses in a severe downside scenario.

In our model we incorporate the fact that an intermediate loss is likely to be realized to avoid the risk of large losses by fixing a holding period for each market instrument. The worst-case change occurred during such holding periods is taken as the yearly loss. The holding periods vary from 250 days for equity investments (as the vast majority of the financial institution's stock investments are long-term) to two days for positions in the most commonly traded currencies.

Let Δ_i be the holding period for instrument i and P_t^i the value of instrument i on day t. Then the change in instrument i on day t is defined as

$$L_t^i = E^i \, \frac{P_t^i - P_{t+\Delta_i}^i}{P_t^i},$$

where E^i is the position limit of instrument i. The change in the whole market portfolio on day t is defined as the sum of the changes in all instruments, i.e.

$$L_t^M = \sum_i L_t^i.$$

Finally, the market loss over one year is defined as the worst daily change, i.e.

$$M = \max_t L_t^M.$$

Note however, that the approach described below for linking market and credit risk goes through with any definition of a market loss that is based on the values of the instruments.

To model the daily fluctuations in the prices P_t^i we use the conventional choice, the geometric Brownian motion model (Jorion, 1997),

$$P_t^i = P_t^{i-1} \exp(d_t^i),$$

where the innovations d_t^i are Gaussian and uncorrelated over time.

The innovations at time t for different instruments are assumed to be correlated. In our application, we have used estimated average daily correlations. However, this may lead to a downward bias in the estimated economic capital. Recently there has been a number of studies, see for instance Longin & Solnik (2001), that have suggested that the correlations between market returns increase in periods of global turbulence. Hence, one may use a model in which the correlation is time-varying, e.g. a multivariate GARCH-model Bollerslev et al. (1998). Such models can be unstable and difficult to use in an operative system and as an alternative one may use fixed extreme correlations, that is the correlation between returns far out in the tail of the multivariate market distribution. To incorporate extreme correlations in our model, the average correlations could be replaced by extreme correlations found by using the parts of the period 1983-2000 that are characterized by a 5% or more weekly decrease in the Standard & Poor 500 Index. In these sub-periods the correlation between the Norwegian Stock Market Index (TOTX) and the Standard & Poor 500 Index was almost 50% higher than the average correlation during the whole time period. We are aware of the fact that by always using extreme correlations, we overestimate the returns in "good times", but as our focus is the adverse events, this is not problematic.

As shown in (4), our model assumes a link between a market loss M and a credit loss C. This does *not* mean that we predict the market losses from the credit losses. We could equally well have modelled $P(C|M)$ instead of $P(M|C)$. Our model links a given credit loss to the observed situation in the market the same year. The credit and market losses could be linked to the same macroeconomic factors, thereby inducing a correlation between them (Iscoe et al. (2002), Medova & Smith (2003)). This approach is often difficult to implement in practice. First, one has to identify the appropriate macroeconomic factors, and then one has to determine in what way the credit and the market losses depend on these factors. To avoid these problems we have chosen a simple approach in which we let the expectation and standard deviation in the distribution of each market instrument depend on the size of the credit losses.

More specifically, we let the expected daily geometric return μ^i and volatility σ_i of each instrument i be dependent of the yearly credit loss C through

$$\mu^i = \alpha^i C + \beta^i \tag{10}$$

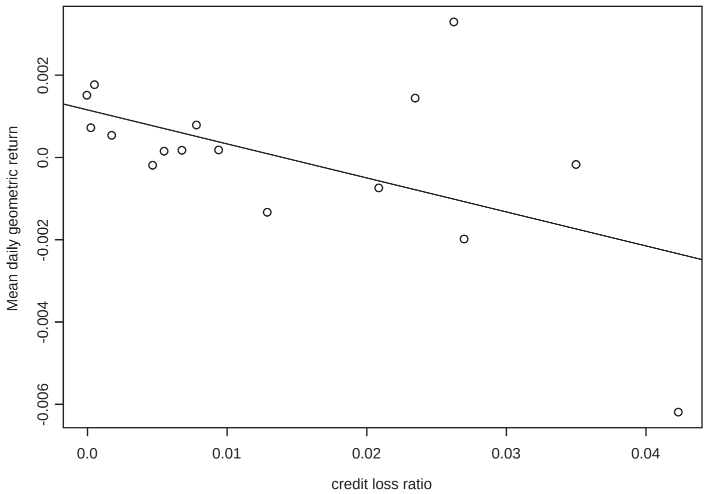

Figure 1: *DnB's credit loss ratios in the period 1984–1999, plotted against the mean daily geometric return of FINX the same years. Regression lines are superimposed.*

and

$$\sigma^i = \gamma^i C + \delta^i. \tag{11}$$

The parameters α^i, β^i, γ^i and δ^i are determined empirically. The historical data available was the annual geometric returns and the corresponding standard deviations for selected stocks, currencies, interest rates and the oil price in the period 1984–1999[1], and the yearly credit loss ratios for the same period. Examples are given in Figure 1 where we have plotted DnB's credit loss ratios in the period 1984–1999 against the mean daily geometric return of FINX the same years, and in Figure 2 which shows a plot of DnB's credit loss ratios in the period 1984–1999 against the standard deviation of the daily geometric returns the same years. The mean returns and volatilities were used as response variables and the credit loss ratios as explanatory variables in a linear least squares regression analysis to estimate the parameters α^i, β^i, γ^i and δ^i in (10) and (11). The resulting lines have been superimposed on the plots. As can be seen from Figure 1, the fit is not very good, but there is a tendency of large credit losses occurring together with low returns of FINX, while from Figure 2, we see that for the years with small credit losses, the volatility in the market seem to be low, while it is higher for years with larger credit losses.

[1]Shorter period for some of the instruments

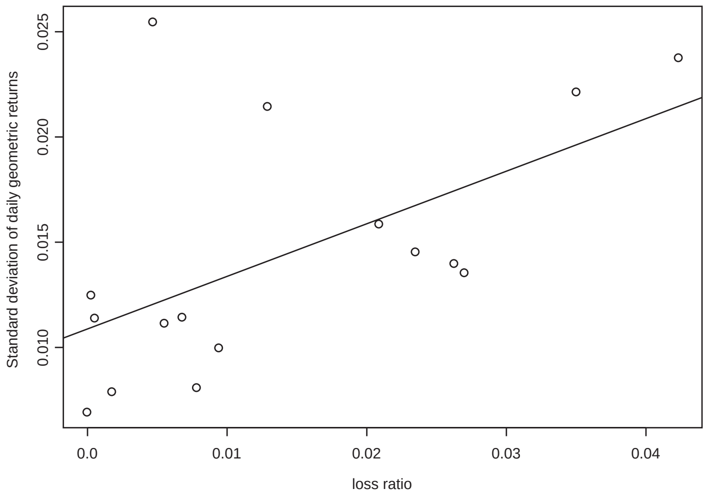

Figure 2: *DnB's credit loss ratios in the period 1984–1999, plotted against the standard deviation of the daily geometric returns of FINX the same years. Regression lines are superimposed.*

3 Implementation

3.1 Simulation Procedure

The distribution of the total loss is complicated and may only be calculated by simulation. Realizations $T_1, ..., T_N$ of the total loss are generated by sampling from the model (4), following the decomposition on the right of (4). First, credit losses C_j are sampled by drawing the credit loss ratios R_j from the beta distribution as described in Section 2.2. Second, we use the procedure described in Section 2.3 to simulate operational losses O_j from the distribution $P(O|C)$, and third, the market losses M_j are drawn dependent on C_j as outlined in Section 2.4.

Having sampled all the marginal distributions, we finally compute the total losses $T_j = C_j + M_j + O_j$. Figure 3 shows an example of the total loss distribution of the of the DnB group.

3.2 Simulation Variability

As described in Section 2.1 our target is the economic capital K_p for small percentiles p. We use a Monte Carlo estimate $\hat{K}_p$ based on N simulations. An approximate formula for the

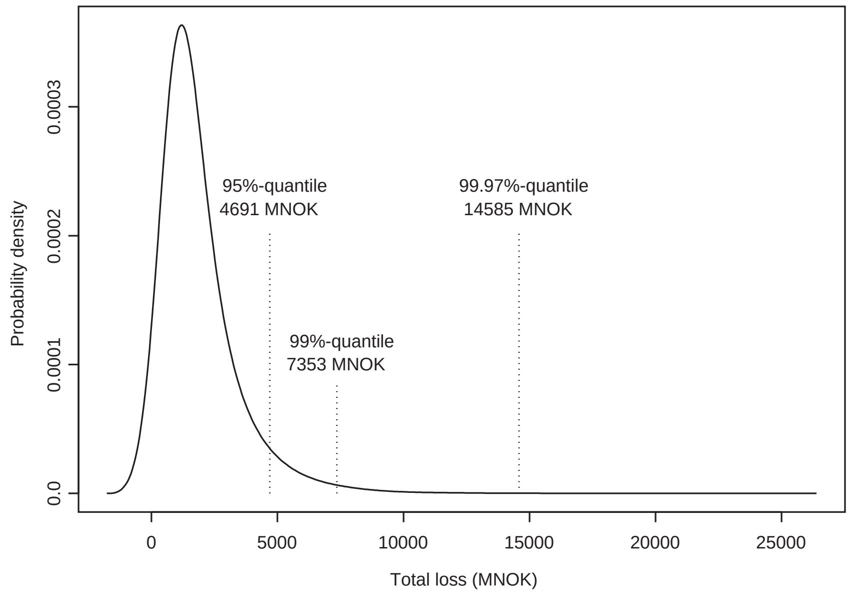

Figure 3: *The estimated total loss distribution for the DnB group.*

standard error due to sampling (Jorion, 1997) is

$$se(\hat{K}_p) = \frac{1}{f(K_p)}\sqrt{\frac{p(1-p)}{N}}, \tag{12}$$

where $f(\cdot)$ is the probability density function of $\hat{K}_p$ and is N the number of simulations. The factor $f(K_p)$ is unknown, but can be estimated from the simulations by using a density estimation method. We used a kernel-density smoother (Silverman, 1986).

Using (12) one may evaluate the Monte Carlo error in the reported estimates of the economic capital and judge how many simulations are required. For the DnB group, the 99.97% quantile (i.e. $p = 0.9997$) is of particular interest as it is related to the official rating of the financial institution. DnB's defined goal is to achieve an "AA" rating from leading rating agencies. The institution has stipulated that risk-adjusted capital should cover 99.97% of potential losses within a one-year horizon, which is in accordance with an "AA" rating. Table 1 shows $se(\hat{K}_p)/\hat{K}_p$ (i.e. the coefficient of variation) for different sample sizes for this quantile as well as the 95% and 99% quantiles.

Based on the results in Table 1, DnB found 500,000 simulations to be a reasonable compromise between accuracy and computational cost (approximately one hour on a standard PC). With this number of simulations, Table 1 shows that the upper and lower 95 % confidence bonds ($\pm 2se(\hat{K}_p)$) do not differ more than 2 % from the estimated 99.97 % quantile.

	Number of simulations				
quantile	*50,000*	*100,000*	*200,000*	*500,000*	*1,000,000*
95.00 %	0.005	0.004	0.003	0.0017	0.0001
99.00 %	0.009	0.006	0.004	0.0028	0.0020
99.97 %	0.029	0.021	0.015	0.0092	0.0065

Table 1: *The coefficient of variation associated with quantiles of the total loss distribution for different number of simulations.*

4 Empirical results

The modelling approach described in this paper gives large reductions in the economic capital when compared to methods following the hypothesis of perfect correlation. The 99.97% quantile becomes 20% smaller than it would have been if the separate capital requirements had been added, while the reductions at the 95 % and 99 % quantiles are 11% and 12%, respectively. This shows that considerable savings and increased competitiveness may be obtained modelling the correlations between the risk types more realistically.

It is interesting to compare the diversification results with the actual correlations between the risk types. The correlations in our model are shown in Table 2. We see that all the correlations are positive, but are markedly smaller than 1.

	Credit	Market	Operational
Credit	1.00	0.30	0.44
Market	0.30	1.00	0.13
Operational	0.44	0.13	1.00

Table 2: *Correlations between different risk sources.*

An important question is how to attribute the total diversified economic capital value to the different risk types. The treatment of this problem is outside the scope of this paper. Readers interested in this issue may for instance see Ward & Lee (2002), Denault (2001) and Bock (2000*b*).

5 Summary and Discussion

In this paper we have presented a new approach to modelling the required total economic capital needed to protect a financial institution against possible losses. The approach takes into account the correlation between several risk types and in this respect it improves upon the conventional practice that assumes perfectly correlated risks. We use Monte Carlo simulation to estimate the loss distribution of the marginal risk types as well as the total risk.

This allows us to choose risk measures freely and not be restricted to single measures such as for instance VaR.

The system we present uses output from the available risk systems for the individual risks and no costly data collection is required. Our motivation has been to construct a simple model to correlate the different risk types, that could be easily used in practice. There are shortcomings of the model and in particular it would be of interest to do model validation, for instance by varying the distributions used for the different risk types. We also believe that more work should be done on the modelling of credit losses. The beta distribution is a crude approximation, and alternatives like the Vasicek distribution (see Bluhm et al. (2002) and Overbeck & Wagner (2000)) should be explored.

The system has been implemented in the statistical software package S-Plus, and is currently being used by DnB for computing the bank's economic capital. We wish to demonstrate that in some cases a financial institution may find a simple model to be a very useful tool, the alternatives being no model or a complex model that is difficult to establish as an everyday tool in the organization.

Acknowledgments

The authors would like to thank Roar Hoff (head of Group of risk management), Ingvill Storli and Øistein Garsjømoen at DnB for introducing us to the risk modelling area and for many interesting discussions. We will also thank Erlend Berg and Jacob K. Laading for their support and cooperation in the early stages of this work and Professor Erik Bølviken and Assistant professor Fred Espen Benth at Department of Mathematics, University of Oslo for valuable comments that contributed to improving this paper.

References

BIS (1995), 'An internal model-based approach to market risk capital requirements'. Basle Commitee on Banking Supervision.

BIS (2001), 'Working paper on the regulatory treatment of operational risk'. Basle Commitee on Banking Supervision.

Bluhm, C., Overbeck, L. & Wagner, C. (2002), *An Introduction to Credit Risk Modeling*, CRC Press.

Bock, J. (2000*a*), 'A capital idea', *Risk Professional* **2**(9), xx–xx.

Bock, J. T. (2000*b*), 'Efficient allocation of economic capital', *MKIRisk, discussion paper.*

Bollerslev, T., Engle, R. F. & Wooldridge, J. M. (1998), 'A capital asset pricing model with time varying covariances', *Journal of Political Economy* **96**, 116–131.

Bradley, B. & Taqqu, M. (2003), Finacial risk and heavy tails, *in* S.T.Rachev, ed., 'Handbook of Heavy Tailed Distributions in Finance', Elsevier, North-Holland.

Crosbie, P. (1997), 'Modeling default risk'. Technical Document, KMV Corporation.

Cruz, M., Coleman, R. & Salkin, G. (1998), 'Modeling and measuring operational risk', *The Journal of Risk* **1**(1), 63–72.

Denault, M. (2001), 'Coherent allocation of risk capital', *Journal of Risk* **4**(1), 1–34.

Embrechts, P., McNeil, A. J. & Straumann, D. (1999), 'Correlation: Pitfalls and alternatives', *Risk* **12**, 69–71.

Goupton, G. M., Finger, C. C. & Bhatia, M. (1997), 'Creditmetrics'. Technical document.

Hickman, A., J., R. & Tange, C. (2002), 'A consistent framework for market and credit risk management', *The Risk Desk* **2**(11), xx–xx.

Iscoe, I., A., K. & Rosen, D. (2002), 'An integrated market and credit risk portfolio model', *Algo Research Quarterly* **2**(3), 21–37.

Jorion, P. (1997), *Value at Risk*, McGraw-Hill.

Kendall, M. & Stuart, A. (1979), *The advanced theory of statistics, Vol 2., ed. 4*, Charles Griffin, London.

Kuritzkes, A., Schuermann, T. & Weiner, S. M. (2002), 'Risk measurement, risk management and capital adequacy in financial conglomerates', *Wharton Working Paper.*

Kyriacou, M. N. & Medova, E. A. (2000), 'Extreme values and the measurement of operational risk – Part 2', *Operational Risk* **1**(8), xx–xx.

Longin, F. & Solnik, B. (2001), 'Extreme correlation of international equity markets', *The Journal of Finance* **56**(2), 649–676.

Medova, E. A. (2000), 'Extreme values and the measurement of operational risk – Part 1', *Operational Risk* **1**(7), xx–xx.

Medova, E. A. & Smith, R. G. (2003), 'A framework to measure integrated risk', *Judge Institute of Management Working paper 09/2003.*

Overbeck, L. & Wagner, C. (2000), 'Term Structure of Loss Cascades in Portfolio Securitisation'. Deutsche Bang AG - Risk Management, November 17,2000.

Silverman, B. W. (1986), *Density estimation for Statistics and Data Analysis*, Chapman & Hall, London.

Ward, L. & Lee, D. (2002), 'Practical application of risk-adjusted return on capital framework', *CAS Forum Summer 2002, Dynamic Financial Analysis Discussion Paper.*