LDA at Work

Falko Aue
Risk Analytics & Instruments[1], Risk and Capital Management,
Deutsche Bank AG, Taunusanlage 12, 60325 Frankfurt, Germany

Michael Kalkbrener
Risk Analytics & Instruments, Risk and Capital Management,
Deutsche Bank AG, Taunusanlage 12, 60325 Frankfurt, Germany

February 2007

Abstract

The Advanced Measurement Approach in the Basel II Accord permits an unprecedented amount of flexibility in the methodology used to assess OR capital requirements. In this paper, we present the capital model developed at Deutsche Bank and implemented in its official EC process. The model follows the Loss Distribution Approach. Our presentation focuses on the main quantitative components, i.e. use of loss data and scenarios, frequency and severity modelling, dependence concepts, risk mitigation, and capital calculation and allocation. We conclude with a section on the analysis and validation of LDA models.

Keywords: Loss Distribution Approach, frequency distribution, severity distribution, Extreme Value Theory, copula, insurance, Monte Carlo, Economic Capital, model validation

[1]Deutsche Bank's LDA model has been developed by the AMA Project Task Force, a collaboration of Operational Risk Management, Risk Analytics and Instruments, and Risk Controlling.

Contents

1 Introduction

A key demand on a bank's Economic Capital methodology is to ensure that Economic Capital covers all material sources of risk. This requirement is a precondition for providing reliable risk estimates for capital management and risk-based performance measurement. Since operational losses are an important source of risk the quantification of operational risk has to be part of the calculation of a bank's Economic Capital. A strong additional incentive for the development of a quantitative OR methodology has been provided by the inclusion of operation risk into the Regulatory Capital requirements under Pillar I of the Basel II Accord (Basel Committee on Banking Supervision, 2006).

The Basel II Accord introduces three approaches to the quantification of operation risk. The most sophisticated option is the Advanced Measurement Approach. It requires the calculation of a capital measure to the 99.9%-ile confidence level over a one-year holding period.[2] The Advanced Measurement Approach permits an unprecedented amount of flexibility in the methodology used to assess OR capital requirements, albeit within the context of strict qualifying criteria. This flexibility sparked an intense discussion in the finance industry. Many quantitative and qualitative techniques for measuring operational risk have been proposed, most prominently different variants of the Loss Distribution Approach and techniques based on scenarios and risk indicators. In our opinion, the natural way to meet the soundness standards for Economic and Regulatory Capital is by explicitly modelling the OR loss distribution of the bank over a one-year period. In this sense, AMA models naturally follow the Loss Distribution Approach, differing only in how the loss distribution is modelled.

The application of the LDA to the quantification of operational risk is a difficult task. This is not only due to the ambitious soundness standards for risk capital but also to problems related to operational risk data and the definition of operational risk exposure, more precisely

1. the shortage of relevant operational risk data,
2. the context dependent nature of operational risk data, and
3. the current lack of a strongly risk sensitive exposure measure in operational risk modelling (cf market and credit risk).

The main objective of an LDA model is to provide realistic risk estimates for the bank and its business units based on loss distributions that accurately reflect the underlying data. Additionally, in order to support risk and capital management, the model has to be risk sensitive as well as sufficiently robust. It is a challenging

[2] Many banks derive Economic Capital estimates from even higher quantiles. For example, the 99.98% quantile is used at Deutsche Bank.

practical problem to find the right balance between these potentially conflicting goals. Finally, the model will only be accepted and implemented in the official processes of a bank if it is transparent and produces explainable results.

In this paper, we present the LDA model developed and implemented at Deutsche Bank. It is used for the quarterly calculation of OR Economic Capital since the second quarter of 2005. Subject to approval by regulatory authorities, the model will also be used for calculating Regulatory Capital.

The details of an LDA model are usually tailored to the specific requirements and limitations of a bank, e.g. the availability of data has an impact on the granularity of the model, the weights given to the different data sources, etc. However, the basic structure of LDA models as well as the fundamental modelling issues are rather similar across different banks. We therefore hope that the presentation of an LDA model that has been designed according to the Basel II guidelines and is part of the bank's official EC process is regarded as an interesting contribution to the current debate.

Section 2 outlines the Loss Distribution Approach implemented at Deutsche Bank and provides a summary of this document. Our presentation focuses on the quantitative aspects of the model and their validation. Qualitative aspects like generation of scenarios or development of a program for key risk indicators are beyond the scope of this paper.

2 Survey of the LDA model implemented at Deutsche Bank

Figure 1 provides the flowchart of the model. Each of the components will be discussed in the following sections.

The fundamental premise underlying LDA is that each firm's operational losses are a reflection of its underlying operational risk exposure (see subsection 3.1). We believe that loss data is the most objective risk indicator currently available. However, even with perfect data collection processes, there will be some areas of the business that will never generate sufficient internal data to permit a comprehensive understanding of the risk profile. This is the reason why internal data is supplemented by external data and generated scenarios: Deutsche Bank is a member of The Operational Riskdata eXchange Association, it has purchased a commercial loss database and has set up a scenario generation process. More information on the different data sources is provided in subsection 3.2.

The first step to generating meaningful loss distributions, is to organize loss data into categories of losses and business activities, which share the same basic risk profile or behaviour patterns. In subsection 3.3, we present the business line/event type matrix used in the model and discuss various criteria for merging cells. Subsection 3.4 focuses on the incorporation of external loss data and scenarios analysis.

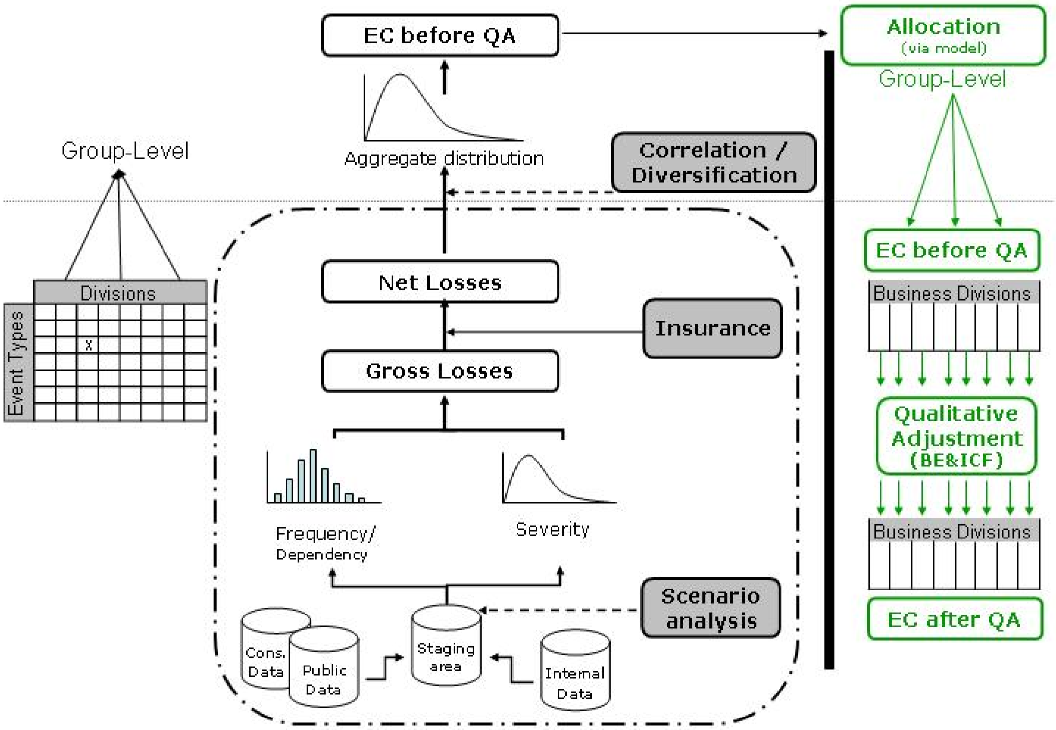

Figure 1: Flowchart of LDA model.

In general, all data points are regarded as a sample from an underlying distribution and therefore receive the same weight or probability in the statistical analysis. However, there are a number of exceptions: split losses, i.e. losses that are assigned to more than one business line, old losses, external losses in the commercial loss data base and scenarios. Section 4 presents algorithms for adjusting the weights of these data points.

Whereas sections 3 and 4 deal with the data sources that are used in the modelling process, sections 5 and 6 are devoted to the specification of loss distributions. More precisely, LDA involves modelling a loss distribution in each cell of the BL/ET matrix. The specification of these loss distributions follows an actuarial approach: separate distributions for event frequency and severity are derived from loss data and then combined by Monte Carlo simulation. In section 5, techniques are presented for calibrating frequency distributions and selecting the distribution that best fits the observed data.

OR capital requirements are mainly driven by individual high losses. Severity distributions specify the loss size and are therefore the most important component in quantitative OR models. Severity modelling is a difficult problem. In particular, tails of severity distributions are difficult to estimate due to the inherent scarcity of low frequency, high impact operational loss events. The methodology applied in DB's LDA model combines empirical distributions and parametric tail distributions which are derived with the Peaks-Over-Threshold method, a technique from Extreme Value Theory (EVT). The severity model is presented in section 6.

The overall capital charge for the firm is calculated by aggregating the loss distributions generated in the above fashion, ideally in a way that recognizes the risk-reducing impact of less than full correlation between the risks in each of the event type/business line combinations. In section 7, the most general mathematical concept for modelling dependence, so-called copulas, are applied to this aggregation problem. More precisely, the frequency distributions in the individual cells of the BL/ET matrix are correlated through a Gaussian copula in order to replicate observed correlations in the loss data.

A realistic quantification of operational risk has to take the risk reducing effect of insurance into account. Compared to other methodologies a bottom-up LDA has the benefit of allowing a fairly accurate modelling of insurance cover. Transferring risk to an insurer through insurance products alters the aggregate loss distribution by reducing the severity of losses that exceed the policy deductible amount. The frequency of loss is unaffected by insurance. More precisely, when frequency and severity distributions are combined through simulation, each individual loss point can be compared to the specific insurance policies purchased by the bank and the corresponding policy limits and deductibles. As a consequence, an insurance model in the Loss Distribution Approach consists of two main components: a quantitative model of the individual insurance policies and a mapping from the OR event types to the insurance policies. Both components are specified in section 8.

Section 9 focuses on the simulation of the aggregate loss distribution (including insurance) at Group level and on the calculation of Economic Capital and capital allocation. Risk measures are based on a one-year time horizon. At Deutsche Bank, Economic Capital for operational risk (before qualitative adjustments) is defined as the 99.98% quantile minus the Expected Loss. Expected Shortfall contributions are used for allocating capital to business lines, i.e. the contribution of a business line to the tail of the aggregate loss distribution. For business units at lower hierarchy levels that do not permit the specification of separate loss distributions the capital allocation is based on risk indicators instead.

Apart from generated loss scenarios LDA models mainly rely on loss data and are inherently backward looking. It is therefore important to incorporate a component that reflects changes in the business and control environment in a timely manner. In DB's LDA model, qualitative adjustments are applied to the contributory capital of business lines. The direct adjustment of EC reduces the complexity of the model and improves its transparency. However, it is difficult to justify with statistical means. Details of the incorporation of qualitative adjustments are given in section 10.

The final section of this paper deals with model analysis and validation. We present a sensitivity analysis of the model components frequencies, severities, dependence structure and insurance. The analysis uses basic properties of LDA models and is therefore not limited to the model implemented at Deutsche Bank. We briefly deal with the impact analysis of stress scenarios and outline the inherent problems with the application of backtesting techniques to OR models. However, the main focus of section 11 is on an approach for benchmarking quantiles in the tail of the aggregate loss distribution of the LDA model against individual data points from the underlying set of internal and external losses.

3 Loss data and scenarios

3.1 Importance of loss data

We believe that loss data is the most objective risk indicator currently available, which is also reflective of the unique risk profile of each financial institution. Loss data should therefore be the foundation of an Advanced Measurement Approach based on loss distributions (ITWG, 2003). This is one of the main reasons for undertaking OR loss data collection. It is not just to meet regulatory requirements, but also to develop one of the most important sources of operational risk management information. We acknowledge that internal loss data also has some inherent weaknesses as a foundation for risk exposure measurement, including:

1. Loss data is a "backward-looking" measure, which means it will not immediately capture changes to the risk and control environment.

2. Loss data is not available in sufficient quantities in any financial institution to permit a reasonable assessment of exposure, particularly in terms of assessing the risk of extreme losses.

These weaknesses can be addressed in a variety of ways, including the use of statistical modelling techniques, as well as the integration of the other AMA elements, i.e. external data, scenario analysis and factors reflective of the external risk and internal control environments, all of which are discussed in the next sections of this document.

3.2 Data sources

The following data sources are used in DB's LDA model:

- Internal loss data: Deutsche Bank started the collection of loss data in 1999. A loss history of more than five years is now available for all business lines in the bank.

- Consortium data: loss data from The Operational Riskdata eXchange Association (ORX).
- Commercial loss data base: data from OpVantage, a subsidiary of Fitch Risk.
- Generated scenarios: specified by experts in divisions, control & support functions and regions.

The process of selecting, enhancing and approving the loss data from all sources and finally feeding it into the LDA model is named Relevant Loss Data Process in Deutsche Bank. We will not provide details but list a few principles:

- As long as data is considered relevant according to defined criteria for business activity (firm type, product and region), it will be included in the capital calculations, no matter when the loss occurred (see section 4.2 for adjustments made to old losses). This ensures that the largest possible meaningful population of loss data is used, thus increasing the stability of the capital calculation.
- There is no adjustment to the size of the loss amount (scaling) in any data source except for inflation adjustment. However, the weights of data points from the public data source are adjusted as outlined in section 4.3.
- Gross losses after direct recoveries are used for capital purposes. Insurance recoveries are not subtracted at this stage because they are modelled separately.
- All losses are assigned to the current divisional structure. External data sources use different business lines and have to be mapped to the internal structure. If possible a 1:1 mapping is performed. However, the framework also allows mapping of one external business line to several internal business lines. In this case the weight of external data points is adjusted in order to reflect the mapping percentage (compare section 4.1).
- Boundary events are excluded from OR capital calculations, e.g. Business Risk, Credit Risk, Market Risk, Timing Losses.
- FX conversion into EUR generally takes place the date the event was booked. In order to report all cash flows of an event with multiple cash flows consistently, the FX rate of the first booking date is used.
- Procedures for avoiding double counting between data sources are in place.

3.3 Data classification and specification of BL/ET matrix

The first step to generating meaningful loss distributions, is to organize loss data into categories of losses and business activities, which share the same basic risk profile or behaviour patterns. For instance, we expect that fraud losses in Retail Banking will share a unique loss distribution, which may be quite different from employee claims in the Investment Banking business. If all losses are lumped together it may be difficult to discern a pattern, whereas if they are separated it becomes easier to describe a unique risk profile and be confident that it is a realistic picture of potential exposure. The Basel Committee has specified a standard matrix of risk types and business lines to facilitate data collection and validation across the various AMA approaches (Basel Committee, 2002a). Firms using LDA are required to map their loss data to the standard matrix, and prove that they have accounted for losses in all aspects of their operations, without being further restricted as to how they actually model the data. In other words, any given firm may choose to collapse or expand the cells in the matrix for purposes of building a specific loss distribution. Deutsche Bank's BL/ET matrix is specified according to

- the internal business lines represented in the Executive Committee of Deutsche Bank and
- Level 1 of the Basel II event type classification.[3]

The decision whether a specific cell is separately modelled or combined with other cells depends on several factors. The following criteria have been identified:

- comparable loss profile
- same insurance type
- same management responsibilities

Other important aspects are data availability and the relative importance of cells. Based on these criteria the seven Basel II event types have been merged into five event types:

Fraud	Internal Fraud
	External Fraud
Infrastructure	Damage to Physical Assets
	Business Disruption, System Failures
Clients, Products, Business Practices	
Execution, Delivery, Process Management	
Employment Practices, Workplace Safety	

Fraud and Clients, Products, Business Practices and Execution, Delivery, Process Management are the dominant risk types in terms of the number of losses as well as the width of the aggregated loss distribution. As a consequence, these event types are modelled separately by business line whereas Infrastructure and Employment are modelled across business lines. This results in the BL/ET matrix in Table 1.

There exist loss events that cannot be assigned to a single cell but

[3] We refer to Samad-Khan (2002) for a critical assessment of the Basel II event type classification.

Basel Level 1	Internal Event Types	Business Lines						
		BL1	BL2	BL3	BL4	B5	BL6	Group
Internal Fraud	Fraud	1	2	3	4	5	6	7
External Fraud								
Damage to physical assets	Infrastructure	8						
Business disruption ...								
Clients, Products, Business Practices	Clients, Products, Business Practices	9	10	11	12	13	14	15
Execution, delivery, process management	Execution, delivery, process management	16	17	18	19	20	21	22
Employment practices, workplace safety	Employment practices, workplace safety	23						

Table 1: BL/ET matrix.

- affect either the entire Group (Group losses)
- or more than one business line (split losses).

The cells 7, 15 and 22 are used for modelling Group losses. The modelling and allocation techniques applied in these Group cells are identical to the techniques in the divisional cells.

Some losses consist of several components that are assigned to different business lines but have the same underlying cause. For example, consider a penalty of 100m that has been split between business line A (70m) and business line B (30m). If the two components of 70m and 30m are modelled separately in the respective divisional cells, their dependence would not be appropriately reflected in the model. This would inevitably lead to an underestimation of the risk at Group level. This problem is avoided by aggregating the components of a split loss and assigning the total amount to each of the cells involved. However, the weights (or probabilities) of the components of a split loss are reduced accordingly: in the above example, the total amount of 100m is assigned to both business lines but the weight of this event is only 70% for business line A and 30% for business line B. We refer to section 4 for more information on adjustments to the weights of loss events and scenarios.

3.4 Use of data

3.4.1 Use of internal loss data

Deutsche Bank's internal losses are the most important data source in its model. Internal loss data is used for

1. modelling frequency distributions,
2. modelling severity distributions (together with external losses and scenarios),

3. analyzing the dependence structure of the model and calibrating frequency correlations.

3.4.2 Incorporation of external loss data

It seems to be generally accepted in the finance industry that internal loss data alone is not sufficient for obtaining a comprehensive understanding of the risk profile of a bank. This is the reason why additional data sources have to be used, in particular external losses (Basel Committee on Banking Supervision, 2006).

There are many ways to incorporate external data into the calculation of operational risk capital. External data can be used to supplement an internal loss data set, to modify parameters derived from the internal loss data, and to improve the quality and credibility of scenarios. External data can also be used to validate the results obtained from internal data or for benchmarking.

In DB's LDA model, external data is used as additional data source for modelling tails of severity distributions. The obvious reason is that extreme loss events at each bank are so rare that no reliable tail distribution can be constructed from internal data only. We are well aware that external losses do not reflect Deutsche Bank's risk profile as accurately as internal events but we still believe that they significantly improve the quality of the model.[4] In the words of Charles Babbage (1791-1871): "Errors using inadequate data are much less than those using no data at all."

3.4.3 Incorporation of scenario analysis

Scenario analysis is another important source of information. In this paper, we limit the discussion to the application of scenarios in DB's LDA model and refer to Scenario-Based AMA Working Group (2003), ITWG Scenario Analysis Working Group (2003), Anders and van den Brink (2004), Scandizzo (2006) and Alderweireld et al. (2006) for a more thorough discussion of scenario analysis including the design and implementation of a scenario generation process.

From a quantitative perspective, scenario analysis can be applied in several ways (ITWG Scenario Analysis Working Group, 2003):

- To provide data points for supplementing loss data, in particular for tail events
- To generate loss distributions from scenarios that can be combined with loss distributions from loss data
- To provide a basis for adjusting frequency and severity parameters derived from loss data
- To stress loss distributions derived from loss data

[4]The direct application of external loss data is a controversial issue. See, for example, Alvarez (2006) for a divergent opinion.

The main application of generated scenarios in DB's LDA model is to supplement loss data. More precisely, the objective of scenario analysis is to capture high impact events that are not already reflected in internal or external loss data. Starting point for the integration of scenarios is the set of relevant losses in OpVantage. These losses have been selected in the Relevant Loss Data Process and can therefore be regarded as one-event scenarios. In the next step, scenarios are generated as deemed necessary to fill in potential severe losses not yet experienced in the past. Each scenario contains a description and an estimate of the loss amount. The process for the generation of scenario descriptions and severities is driven by experts in divisions, control & support functions and regions and is followed by a validation process. The scenario data points are combined with the relevant OpVantage data and receive the same treatment, i.e. scaling of probabilities of individual data points. The combined data set of relevant OpVantage losses and generated scenarios is an important element in the calibration of the tails of severity distributions.

4 Weighting of loss data and scenarios

Loss data and scenarios are used for calibrating frequency and severity distributions. In general, all data points are regarded as a sample from an underlying distribution and therefore receive the same weight or probability in the statistical analysis. However, there are three exceptions:

1. split losses
2. old losses
3. external losses in the commercial loss data base and scenarios

4.1 Split losses

Split losses are loss events that cannot be assigned to a single cell but affect more than one business line. The treatment of split losses has already been discussed in section 3.3: the total amount of a split loss is assigned to each business line affected, the weight being set to the ratio of the partial amount of the respective business line divided by the aggregated loss amount. Note that the sum of the weights of a split loss equals one.

4.2 Old losses

Since the risk profile of a bank changes over time, old losses will be less representative. The impact of a given loss should therefore be reduced over an appropriate time period. In DB's LDA model, the phasing-out of old losses is implemented in the following way:

- For frequency calibration, only internal losses are used that occurred in the last five years.
- For severity calibration and scaling of OpVantage losses, a weighting by time is introduced. Losses that occurred within the last 5 years receive full weight. The weight of older losses is linearly decreased from one to zero over a period of 20 years.

4.3 Scaling of external data and scenarios

4.3.1 Characteristics of external data

External loss data is inherently biased. The following problems are typically associated with external data:

- *Scale bias* - Scalability refers to the fact that operational risk is dependent on the size of the bank, i.e. the scale of operations. A bigger bank is exposed to more opportunity for operational failures and therefore to a higher level of operational risk. The actual relationship between the size of the institution and the frequency and severity of losses is dependent on the measure of size and may be stronger or weaker depending on the particular operational risk category.
- *Truncation bias* - Banks collect data above certain thresholds. It is generally not possible to guarantee that these thresholds are uniform.
- *Data capture bias* - Data is usually captured with a systematic bias. This problem is particularly pronounced with publicly available data. More precisely, one would expect a positive relationship to exist between the loss amount and the probability that the loss is reported. If this relationship does exist, then the data is not a random sample from the population of all operational losses, but instead is a biased sample containing a disproportionate number of very large losses. Standard statistical inferences based on such samples can yield biased parameter estimates. In the present case, the disproportionate number of large losses could lead to an estimate that overstates a bank's exposure to operational risk (see de Fontnouvelle et al. (2003)).

4.3.2 Scaling algorithms

In the current version of the model, external loss data is not scaled with respect to size. The reason is that no significant relationship between the size of a bank and the severity of its losses has been found in a regression analysis of OpVantage data done at Deutsche Bank (compare to Shih et al. (2000)). This result is also supported by an analysis of internal loss data categorized according to business lines and regions.

General techniques for removing the truncation bias can be found in Baud et al. (2002) and Chernobai et al. (2006). In the frequency and severity models presented in this paper, the truncation bias does not pose a problem.

We scale publicly available loss data in order to remove the data collection bias.[5] The basic idea is to adjust the probabilities (and not the size) of the external loss events in order to reflect the unbiased loss profile, i.e. increase the probability of small losses and decrease the probability of large losses.[6] The crucial assumption in our approach is that ORX data and (unbiased) OpVantage data have the same risk profile, i.e. both reflect the generic OR profile of the finance industry. ORX data is assumed to be unbiased. As a consequence, the probabilities of the public loss events are adjusted in such a way that the OpVantage loss profile (after scaling) reflects the ORX profile. The scaling is performed at Group level, i.e. it is based on all OpVantage losses, scenarios and ORX events (including losses in DB) above 1m. The same scaling factors are applied across all business lines and event types.

The mathematical formalization of the scaling technique is based on stochastic thresholds. More precisely, following Baud et al. (2002) and de Fontnouvelle et al. (2003) we extract the underlying (unbiased) loss distribution by using a model in which the truncation point for each loss (i.e., the value below which the loss is not reported) is modelled as an unobservable random variable. As in de Fontnouvelle et al. (2003) we apply the model to log losses and assume that the distribution of the threshold is logistic. However, our model does not require any assumptions on the distribution of the losses.[7] We will now describe the model in more detail.

Let $X_1, \dots, X_m$ be independent samples of a random variable X. The variable X represents the "true" loss distribution and is identified with ORX data. Let $Y_1, \dots, Y_n$ be independent samples of the conditional random variable $Y := X|H \leq X$, where H is another random variable (independent of X) representing the stochastic threshold. We assume that the distribution function $F_\theta(x)$ of H belongs to a known distribution class with parameters $\theta = (\theta_1, \dots, \theta_r)$. The variable Y is identified with OpVantage data. The objective of scaling is to determine the threshold parameters $\theta = (\theta_1, \dots, \theta_r)$ from the data sets $\{X_1, \dots, X_m\}$ and $\{Y_1, \dots, Y_n\}$. The criterion for parameter calibration is to minimize

$$\sum_{i=1}^{k} \left(\mathbb{P}(H \leq X \leq S_i \,|\, H \leq X) - \mathbb{P}(Y \leq S_i)\right)^2,$$

where $S_1, \dots, S_k$ are positive real numbers, i.e. severities. The probabilities $\mathbb{P}(Y \leq S_i)$ are derived from the samples $Y_1, \dots, Y_n$. The probabilities $\mathbb{P}(H \leq X \leq S_i \,|\, H \leq$

[5] More precisely, the data set consisting of relevant OpVantage losses and generated scenarios (see section 3.4.3) is scaled.

[6] Since external data is not used for modelling frequencies the scaling methodology affects severity distributions only.

[7] In contrast, de Fontnouvelle et al. (2003) assume that the distribution of excesses of logarithms of reported losses converges to the exponential distribution which is a special case of the GPD.

X) are calculated as

$$\sum_{j=1}^{L_i} F_\theta(X_j) / \sum_{j=1}^{m} F_\theta(X_j),$$

where L_i is the highest index such that $X_{L_i} \leq S_i$ and w.l.o.g. $X_1 \leq \ldots \leq X_m$.

The following parameter setting is used in the current model: in order to specify $X_1, \ldots, X_m$, all ORX and internal losses above 1m EUR are selected. The samples $Y_1, \ldots, Y_n$ are the relevant OpVantage losses and scenarios. The threshold is assumed to follow a loglogistic distribution, i.e. it has the distribution function

$$F_{\mu,\beta}(x) = \frac{1}{1 + e^{-\frac{\log(x)-\mu}{\beta}}},$$

which is equivalent to applying a logistic threshold to log losses. This distribution has been chosen because it provides an excellent fit to our data.

Figure 2 displays the impact of scaling on the OpVantage profile by means of two QQ-plots: the unscaled OpVantage date is much heavier than the consortium data whereas the profiles match rather well after scaling.

5 Frequency distributions

The standard LDA uses actuarial techniques to model the behaviour of a bank's operational losses through frequency and severity estimation, i.e. the loss distribution in each cell of the BL/ET matrix is specified by separate distributions for event frequency and severity. Frequency refers to the number of events that occur within a given time period. Although any distribution on the set of non-negative integers can be chosen as frequency distribution, the following three distribution families are used most frequently in LDA models: the Poisson distribution, the negative binomial distribution, or the binomial distribution (see Johnson et al. (1993) or Klugman et al. (2004) for more information).

5.1 Data requirements for specifying frequency distributions

In DB's LDA model, the specification of frequency distributions is entirely based on internal loss data (in contrast to Frachot and Roncalli (2002) who suggest to use internal and external frequency data[8]). The main reasons for using only internal data are:

- Internal loss data reflects DB's loss profile most accurately.

[8]If external loss data is used for frequency calibration the external frequencies have to be scaled based on the relationship between the size of operations and frequency.

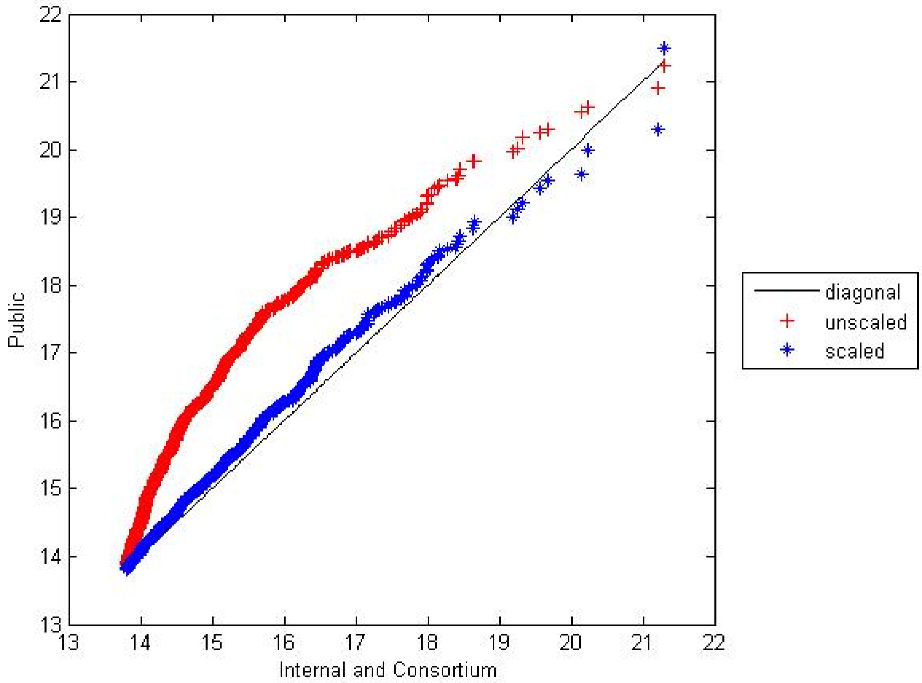

Figure 2: QQ-plots showing the log quantiles of internal and consortium losses above 1m on the x-axis and the log quantiles of public loss data and scenarios on the y-axis.

- It is difficult to ensure completeness of loss data from other financial institutions. However, data completeness is essential for frequency calibration.
- Data requirements are lower for calibrating frequency distributions than for calibrating severity distributions (in particular, if Poisson distributions are used).

For calibrating frequency distributions, time series of internal frequency data are used in each cell. Frequency data is separated into monthly buckets in order to ensure that the number of data points is sufficient for a statistical analysis.

5.2 Calibration algorithms

We have implemented calibration and simulation algorithms for Poisson, binomial and negative binomial distributions. In order to determine the appropriate distribution class for a particular cell we apply three different techniques to the corresponding time series.

- The dispersion of the time series is analyzed by comparing its mean and variance.[9] If the time series is equidispersed a Poisson distribution is used. In case of overdispersion (underdispersion) the frequency distribution is modelled as a negative binomial distribution (binomial distribution). Since it is not obvious which mean/variance combinations should be considered equidispersed the results of the dispersion analysis are compared with the following goodness-of-fit tests.

- We estimate the parameters of a Poisson distribution and a negative binomial distribution by matching the first two moments[10]. Then a goodness-of-fit test (to be more precise, a χ^2-test) is used to analyze the hypothesis of a Poisson and a negative binomial distribution respectively. The idea of the χ^2-test is based on comparisons between the empirically measured frequencies and the theoretically expected frequencies. More precisely, the frequencies of the observed data are aggregated on chosen intervals and compared to the theoretically expected frequencies. The sum over the (weighted) squared differences follows approximately a specific χ^2-distribution. This can be used to calculate for a given level (e.g. $\alpha = 0.05$) a theoretical "error". If the observed error is greater than the theoretical error, the hypothesis is rejected. The level α can be understood as the probability to reject a true hypothesis. In order to avoid the (subjective) choice of the level α, one can calculate a so called p-value, which is equal to the smallest α–value at which the hypothesis can be rejected, based on the observed data.

- Another test that we perform in this context analyzes the interarrival time of losses. If the data has been drawn from an independent random process (e.g. Poisson Process), the interarrival times of the losses follow an exponential distribution. The interarrival times are calculated over a particular time horizon and fitted by an exponential density, whose parameter is estimated with a ML-estimator. A χ^2-test is used to assess the quality of the fit.

Based on these tests we have developed an algorithm for selecting a Poisson or negative binomial distribution. Furthermore, we have analyzed the impact of different distribution assumptions on Economic Capital. It turned out that it is almost irrelevant for EC at Group and cell level whether Poisson or negative binomial distributions are used. This result agrees with the theoretical analysis in section 11 (see also Böcker and Klüppelberg (2005) and De Koker (2006)): in LDA models applied to OR data, the choice of severity distributions usually has a much more severe impact on capital than the choice of frequency distributions. It has therefore

[9] A distribution is called equidispersed (overdispersed; underdispersed) if the variance equals (exceeds; is lower than) the mean.

[10] Since no underdispersed time series were observed binomial distributions are not considered.

been decided to exclusively use Poisson distributions in the official capital calculations at Deutsche Bank. This decision reduces the complexity of the model since no statistical tests or decisions rules for frequency distributions are required.

6 Severity distributions

6.1 Complexity of severity modelling

OR capital requirements are mainly driven by individual high losses. Severity distributions specify the loss size and are therefore the most important component in quantitative OR models.

Severity modelling is a difficult task. One reason is the lack of data. Internal loss data covering the last 5 to 7 years is not sufficient for calibrating tails of severity distributions. It is obvious that additional data sources like external loss data and scenarios are needed to improve the reliability of the model. However, inclusion of this type of information immediately leads to additional problems, e.g. scaling of external loss data, combining data from different sources, etc.

Even if all available data sources are used it is necessary to extrapolate beyond the highest relevant losses in the data base. The standard technique is to fit a parametric distribution to the data and to assume that its parametric shape also provides a realistic model for potential losses beyond the current loss experience. The choice of the parametric distribution family is a non-trivial task and usually has a significant impact on model results (compare, for example, to Dutta and Perry (2006) or Mignola and Ugoccioni (2006)).

Our experience with internal and external loss data has shown that in many cells of the BL/ET matrix the body and tail of the severity distribution have different characteristics. As a consequence, we have not been able to identify parametric distribution families in these cells that provide acceptable fits to the loss data across the entire range. A natural remedy is to use different distribution assumptions for the body and the tail of these severity distributions. However, this strategy adds another layer of complexity to the severity model.

In summary, severity modelling comprises a number of difficult modelling questions including:

Treatment of internal and external loss data and scenarios

- How much weight is given to different data sources?
- How to combine internal and external data and scenarios?

Range of distribution

- One distribution for the entire severity range or different distributions for small, medium and high losses?

Choice of distribution family

- Two-parametric distributions like lognormal and GPD, more flexible parametric distribution families, i.e. three- or four-parametric, or even empirical distributions?
- One distribution family for all cells or selection of "best" distribution based on quality of fit?[11]

The main objective is to specify a realistic severity profile: severity distributions should provide a good fit to the available loss data over the entire range, in particular in the tail sections. However, fitting the data is not the only objective of severity modelling. Since the severity model is the key driver of OR capital requirements the sensitivity of severity distributions to changes in the input data (losses and scenarios) is of particular importance. While a capital model used for risk measurement and steering has to be risk sensitive wild swings in capital estimates are not acceptable for capital planning and performance measurement. It is a difficult task to find the right balance between these potentially conflicting goals. Another import requirement for a capital model used in practice is that its structure and results should be explainable to non-quants. Again, this is quite challenging for severity modelling: the severity model has to be sophisticated enough to capture complex severity profiles but it has to be kept as simple and transparent as possible in order to increase acceptance by business and management.

6.2 Modelling decisions

The availability of internal loss data differs significantly across financial institutions and there is no consensus about the application of external losses to severity modelling. It is therefore not surprising that there has not yet emerged a standard severity model that is generally accepted in the industry. In this subsection, we discuss the availability and characteristics of internal and external loss data at Deutsche Bank and present the basic structure of the severity model derived from the data.

6.2.1 Availability of data

In a typical cell of the BL/ET matrix, there are sufficient internal data points to specify a reliable severity profile between 10k (the internal data collection threshold) and 1m. However, the number of internal losses above 1m is rather limited. We therefore use all cell-specific external losses and scenarios as an additional data source. However, even if all data points in a cell are combined we do not have sufficient information on the extreme tail of the severity distribution, say beyond 50m.

[11] A description of goodness-of-fit tests like Kolmogorov-Smirnov test, Anderson-Darling test, QQ-plots, etc and their application to OR data can be found e.g. in Chernobai et al. (2005), Moscadelli (2004) or Dutta and Perry (2006).

This is the reason why a third data source is used: all internal losses, external losses and scenarios (across all cells) above 50m. This choice is based on the underlying assumption that the extreme tails of the severity distributions in the different cells have something in common. Of course, this assumption is debatable. However, we consider it the better option compared to extrapolating far beyond the highest loss that has occurred in a particular cell.

6.2.2 Characteristics of data

It seems to be generally accepted in the finance industry that OR capital requirements of large international banks are mainly driven by rare and extreme losses. This fact has a strong influence on the choice of the distribution families used for modelling severities in operational risk: it is quite natural to work with distributions that have been applied in insurance theory to model large claims. Many of these distributions belong to the class of subexponential distributions (see section 11 and Embrechts et al. (1997) for more information). Examples are Pareto, Weibull (with $\tau < 1$), Benktander-type-I and II, lognormal and loggamma distributions.[12]

We have experimented with a number of subexponential distributions, including truncated[13] lognormal, Weibull and Pareto. When we fitted these distributions to the internal and external data points we encountered two main problems:

1. In many cells of the BL/ET matrix, the body and tail of the severity distribution have different characteristics. As a consequence, we have not been able to identify parametric distribution families in these cells that provide acceptable fits to the loss data across the entire range. Typically, the calibrated distribution parameters were dominated by the large number of losses in the body which resulted in a poor fit in the tail.

2. Calibration results were rather unstable, i.e. for rather different pairs of distribution parameters the value of the maximum-likelihood function was close to its maximum. In other words, these different parametrizations provided a fit of comparable quality to the existing data points. Even different distribution families frequently resulted in a similar goodness-of-fit. In most cases, however, the calibrated distributions differed significantly in the extreme tails (compare to Mignola and Ugoccioni (2006)).

[12] In the literature, the calibration of various light and heavy-tailed distribution classes to operational risk data is analyzed. de Fontnouvelle and Rosengren (2004) discuss the properties of the most common severity distributions and fit them to OR data. A similar study can be found in Moscadelli (2004). Dutta and Perry (2006) also examine a variety of standard distributions as well as 4-parametric distributions like the g-and-h distributions and the Generalized Beta distribution of Second Kind. Alvarez (2006) suggests the 3-parametric lognormal-gamma mixture.

[13] The truncation point is 10k in order to reflect the internal data collection threshold. Chernobai et al. (2006) analyse the errors in loss measures when fitting non-truncated distributions to truncated data.

A potential solution is to apply more flexible parametric distribution families, e.g. distributions with more than two parameters. However, even if the additional flexibility improves the fit to the existing data points it seems doubtful whether these distributions provide a more reliable severity profile across the entire range.

Instead of applying high-parametric distribution families we have decided to model body and tail separately. Empirical distributions are used for modelling the body of severity distributions. This approach offers the advantages that no choice of a parametric distribution has to be made, the severity profile is reflected most accurately and high transparency is ensured.

For modelling severity tails, however, empirical distributions are clearly not sufficient. We combine empirical distributions with a parametric distribution in order to quantify the loss potential beyond the highest experienced losses. For the specification of the parametric distribution we have decided to apply Extreme Value Theory (EVT) or - more precisely - the Peaks-Over-Threshold method.

Extreme Value Theory is concerned with the analysis of rare and extreme events and therefore provides a natural framework for modelling OR losses. Most relevant for the application in operational risk is a theorem in EVT saying that, for a certain class of distributions, the generalized Pareto distribution (GPD) appears as limiting distribution for the distribution of the excesses $X_i - u$, as the threshold u becomes large. Hence, this theorem provides guidance for selecting an appropriate distribution family for modelling tails of severity distributions. Its algorithmic version is the Peaks-Over-Threshold method. We refer to Embrechts et al. (1997) for an excellent introduction to Extreme Value Theory and to Medova (2000), Cruz (2002), Embrechts et al. (2003), Moscadelli (2004) and Makarov (2006) for an application of EVT to OR data.

Unfortunately, the application of Extreme Value Theory in operational risk is not straightforward. In the words of Chavez-Demoulin et al. (2005): "Applying classical EVT to operational loss data raises some difficult issues. The obstacles are not really due to a technical justification of EVT, but more to the nature of the data". Depending on the data set used, the papers cited in this section (see, moreover, Nešlehová et al. (2006) and Mignola and Ugoccioni (2005)) come to different conclusions about the applicability of EVT to OR data. Our own experience is summarized in the following paragraph.

Generalized Pareto distributions are specified by two parameters, the shape and the scale parameter. Theory predicts that the calibrated shape parameter stabilizes as the threshold u becomes large and the distribution of excesses $X_i - u$ converges to a GPD. It depends on the underlying loss data at which threshold the corresponding shape becomes constant. However, when we apply the Peaks-Over-Threshold method to all losses and scenarios above 50m we do not observe stable shape parameters for large thresholds. On the contrary, shape parameters tend to decrease when thresholds are increased (compare to figure 3). This phenomenon is not necessarily contradicting theory but may be caused by the lack of loss data: additional extreme losses would enable us to increase thresholds further which might eventually lead to stable shape parameters. In the current setting, however, the decreasing shapes make the modelling of tails more difficult.[14] This problem will be addressed in more detail in the following subsections.

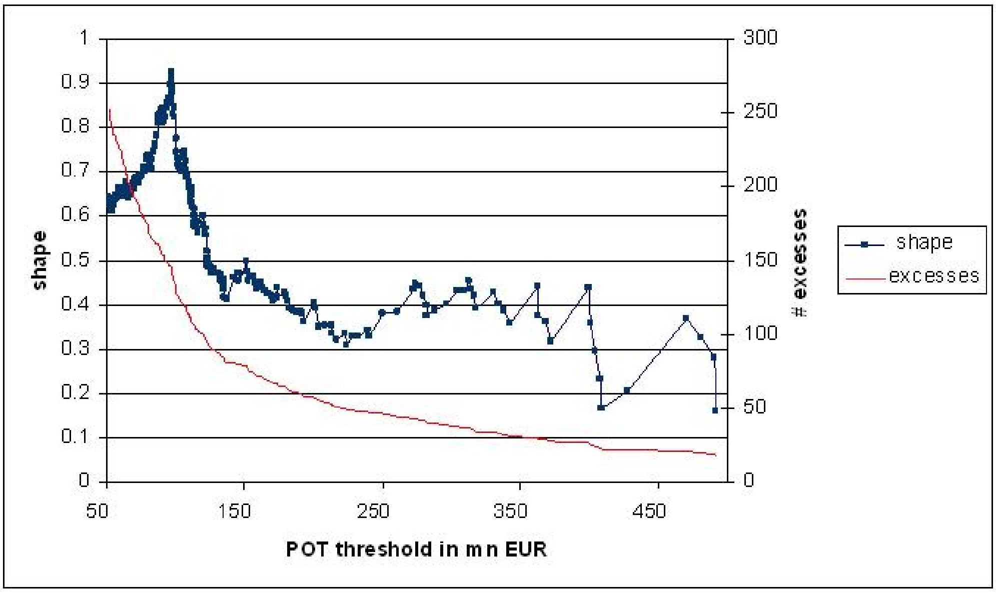

Figure 3: Estimated shape parameter and number of excesses for various POT thresholds.

Despite this problem we consider the POT method the most reliable technique for modelling tails of severity distributions. Other distributions, e.g. truncated lognormal and Weibull, used for modelling the extreme tail did not improve the quality of the fit and proved to be very unstable with respect to small deviations in the loss data.

[14]Note that we would observe decreasing shape parameters if the size of individual OR losses cannot be infinite but is capped at a certain level: in this case the shapes would eventually converge to a negative value.

6.2.3 Summary of modelling decisions

In each cell, the severity distribution is piece-wise defined with three disjunctive ranges:

1. In each cell, we derive separate severity distributions
   - from all internal losses >10k and
   - from all external losses and scenarios >1m respectively.

   Another severity distribution is calibrated from all losses >50m (across all cells) that provides additional information on the tail.

2. Both cell-specific distributions are directly derived from the historic losses, i.e. are modelled as empirical distributions.

3. The Peaks-Over-Threshold method is used to calibrate the parametric distribution >50m.

4. Empirical and parametric distributions are combined to a piece-wise defined severity distribution via a weighted sum applied to the cumulative distribution functions.

The following subsection provides details on the specification of the different distributions and the way they are combined.

6.3 Specification of severity distributions

6.3.1 Building piecewise-defined distributions

As outlined in section 6.2.3, different severity distributions are constructed at the thresholds 10k, 1m and 50m and are then aggregated to a piece-wise defined distribution with three disjunctive ranges. The cumulative distribution function F_1 of the first range (10k - 1m) is based on cell specific, internal losses. The CDF F_2 for the second range (1m - 50m) takes also cell specific, external data into account and the distribution F_3 for the tail (>50m) is derived from cell specific data and the data set consisting of all losses (internal, ORX, scaled OpVantage and scenarios) above 50m.

With $T_1 = 1\text{m}$ and $T_2 = 50\text{m}$, the combined severity distribution F is given by:

$$\overline{F}(x) = \begin{cases} \overline{F}_1(x) & x < T_1, \\ \overline{F}_1(T_1) \cdot \overline{F}_2(x) & T_1 \leq x < T_2, \\ \overline{F}_1(T_1) \cdot \overline{F}_2(T_2) \cdot \overline{F}_3(x) & T_2 \leq x, \end{cases}$$

where $\overline{F}(x) = 1 - F(x)$ denotes the tail of F.

In detail:

- F_1 is given by the empirical distribution of cell specific internal data.
- $F_2 = w_{\text{int}} F_{\text{int}}^{T_1} + w_{\text{ext}} F_{\text{ext}}^{T_1}$, with $F_{\text{int}}^{T_1}$ being the truncated empirical distribution of cell specific internal data above T_1 and $F_{\text{ext}}^{T_1}$ the empirical distribution of cell specific ORX, scaled OpVantage and scenario data above T_1.
- $F_3 = \bar{w}_{\text{int}} F_{\text{int}}^{T_2} + \bar{w}_{\text{ext}} F_{\text{ext}}^{T_2} + \bar{w}_{\text{tail}} F_{\text{tail}}$. The empirical distributions of internal and external data above T_2 are taken into account in order to capture the cell specific risk profile. F_{tail} is the parametric distribution calibrated on all losses above T_2.

The setting of the weights has not been determined by statistical techniques but is motivated by the following analyses and considerations:

- The weight of internal data increases with the number of internal data points: as internal data is considered most representative for the bank's risk profile it receives a high weight as long as sufficiently many internal losses are available to ensure stability.
- The weight for the parametric tail function reflects the contribution of the respective event type to losses >50m.
- Impact analyses have been performed to investigate the volatility due to inclusion or exclusion of data points.

6.3.2 Calibration of empirical distribution functions

In principle, the derivation of empirical distributions from loss data is straightforward. Note, however, that the data points in the loss sample receive different weights (compare section 4). Therefore the standard estimator for empirical distributions has to be modified and reads

$$\hat{F}^T(x) = \sum_{j=L_i}^{m} w(X_j) \cdot 1_{\{X_j \le x\}} / \sum_{j=L_i}^{m} w(X_j),$$

where $X_1, \ldots, X_m$ is the loss sample, L_i is the lowest index such that $X_{L_i} \ge T$ and w.l.o.g. $X_1 \le \ldots \le X_m$. The weights $w(X_1), \ldots, w(X_m)$ reflect the adjustments for the split ratios, time decay and data collection bias.

6.3.3 Calibration of parametric tail

The calibration of the parametric tail is done in two steps. First scale and shape parameters of GPDs are estimated at increasing excess thresholds $u_1, \ldots, u_K$ via the Peaks-Over-Threshold method. As outlined in section 6.2.3 we observe that the shape parameters do not become constant but decrease with increasing thresholds.

Thus, a modified GPD, which is a sequence of GPDs with varying scale and shape parameters, is constructed in the second step. This approach is a modification of the standard POT-method and provides the best and most stable fit to our data. Both steps are now described in detail.

The objective of the Peaks-Over-Threshold method is to fit a Generalized Pareto distribution to the excesses $X_1 - u, \ldots, X_n - u$, where $X_1, \ldots, X_n$ is the loss sample. For determining the scale and shape parameters we apply Maximum Likelihood Estimation (MLE).[15] The Maximum Likelihood Estimator calculates parameters of the distribution that maximize the "likelihood". In our application, the likelihood function is the density function of the Generalized Pareto distribution. Our loss sample consists of internal losses, generated scenarios and data points from ORX and OpVantage. A number of techniques for combining loss data from different sources have been proposed in recent years. Each of these techniques is based on specific assumptions on the relationship between the severity distributions of these sources.

The most straightforward approach assumes that all data sets follow the same severity distribution. Under this assumption direct mixing of data is possible, i.e. losses of all data sources are combined and used as input for the calibration of the joint severity distribution.

To derive the parameters of a GPD from unbiased and biased data sources we follow the methodology in Frachot and Roncalli (2002) and Baud et al. (2002). It is based on the assumption that biased data, i.e. OpVantage and scenarios, have the same distribution as unbiased data, i.e. internal and ORX, except that biased data are truncated below a stochastic threshold H_θ. The threshold is assumed to follow a loglogistic distribution with parameters $\theta = (\theta_1, \theta_2)$. Its purpose is to scale OpVantage to the generic OR profile of the finance industry represented by ORX and internal data (see subsection 4.3). Note, however, that for losses >50m the scaling function used to remove the data collection bias of the OpVantage data has insignificant influence and therefore does not affect the stability of the calibration. In this setting, MLE is applied to the log-likelihood

$$l_\theta(\delta) = \sum_{i=1}^{m} w(X_i) \cdot \ln(f_\delta(X_i - u)) + \sum_{j=1}^{n} w(Y_j) \cdot \ln(f_{\theta,u,\delta}(Y_j - u)),$$

where $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_n$ are series of internal/ORX and OpVantage/scenario data above a threshold u in a given cell, $f_\delta(x)$ is the density function of the underlying Generalized Pareto distribution with scale and shape parameters $\delta = (b, t)$ and $f_{\theta,u,\delta}(x)$ is the truncated density function of the severity distribution with stochastic truncation point H_θ and excess threshold u. The weights $w(\cdot)$ reflect the adjustments for time decay and the special treatment of split losses as outlined in section 4.

[15] We have compared Maximum Likelihood Estimation to Hill Estimators (Embrechts et al., 1997) and have obtained consistent results.

MLE is used to calculate the parameters $\delta = (b, t)$ of the severity distribution, the parameters $\theta = (\theta_1, \theta_2)$ have been calibrated in the scaling algorithm.[16] Applying this methodology to the set of excess thresholds $u_1, \ldots, u_K$ yields a set of scale and shape parameters $\delta_1 = (b_1, t_1), \ldots, \delta_K = (b_K, t_K)$.

In the second step we identify functions $b_\beta(u)$ and $t_\xi(u)$ with parameters $\beta = (\beta_1, \ldots, \beta_i)$ and $\xi = (\xi_1, \ldots, \xi_j)$ that describe the behaviour of the scale and shape parameters as functions of the threshold u. The precise form of $b_\beta(u)$ and $t_\xi(u)$ including the number of parameters i and j depends on the characteristics of the underlying data. Analysis on several data cycles showed that shapes decrease in the order of $1/u$ and scales increase in the order of $\ln(u)$. The parameters β and ξ are set such that the quadratic distance between estimated parameters $\delta_k = (b_k, t_k)$ and functions $b_\beta(u)$ and $t_\xi(u)$ taken at a selection of excess thresholds u_k is minimized. The two functions $b_\beta(u)$ and $t_\xi(u)$ are then used to define a sequence of n GPDs with varying scale and shape parameters for excesses x over $T_2 = 50$m.

An alternative approach to the piecewise specification of the parametric tail as sequence of n GPDs is based on the corresponding limit distribution: for $n \to \infty$, i.e. for infinite granularity, the tail distribution becomes

$$\tilde{G}_{\beta,\xi}(x) := 1 - \lim_{n\to\infty} \prod_{i=1}^{n} \left(1 - G_{b_\beta\left(\frac{x(i-1)}{n}+T_2\right), t_\xi\left(\frac{x(i-1)}{n}+T_2\right)} \left(\frac{x}{n}\right) \right).$$

This distribution does not depend on the parametric shape $t_\xi(u)$ and has the form[17]

$$\tilde{G}_\beta(x) = 1 - \exp\left(- \int_{T_2}^{T_2+x} \frac{1}{b_\beta(u)} du \right). \tag{1}$$

It can be calibrated from loss data using the standard maximum-likelihood estimator. Finally, it is interesting to note that if a linear scale function $b_\beta(u) = \beta_2 \cdot (u - T_2) + \beta_1$ is used, the distribution (1) becomes a GPD with scale and shape parameters β_1 and β_2.

7 Dependence

7.1 Types of dependence in LDA models

In a bottom-up Loss Distribution Approach, separate distributions for loss frequency and severity are modelled at "business line/event type" level and then combined.

[16]For extensions or alternatives to Maximum-Likehood-Estimation we refer to de Fontnouvelle and Rosengren (2004) who adapt a regression-based EVT technique that corrects for small-sample bias in the tail parameter estimate proposed by Huisman et al. (2001). See also Peng and Welsh (2001) for other robust estimators.

[17]We are grateful to Andreas Zapp for pointing out this fact to us.

More precisely, in each cell ($k = 1, \dots, m$) of the BL/ET matrix the loss distribution is specified by

$$L_k = \sum_{i=1}^{N_k} S_{ki},$$

where N_k is the frequency distribution and $S_{k1}, \dots, S_{kN_k}$ are random samples of the severity distribution S_k. The loss distributions in the cells are then aggregated to a loss distribution at Group level that determines the Economic Capital of the bank.

In this basic model, dependencies between the occurrence of loss events and the size of losses might occur in numerous ways:

Within a cell

- Dependence between the occurrence of loss events in a cell
- Dependence between the frequency distribution N_k and the severity distribution S_k in a cell
- Dependence between the severity samples $S_{k1}, \dots, S_{kN_k}$ in a cell

Between cells

- Dependence between the frequency distributions $N_1, \dots, N_m$ in different cells
- Dependence between the severity distributions $S_1, \dots, S_m$ in different cells

In this section, we analyze these types of dependence and discuss the consequences for modelling. Our analyses are based on time series of internal frequency and severity data covering 5 years. In order to increase the number of data buckets we use monthly intervals.

The statistical analysis of the dependence structure of OR losses does not only require information on the size of losses but also on the date of their occurrence.[18] The quality of the results crucially depends on the quality of the underlying data. Considering the inherent scarcity of OR loss data the results of the statistical correlation analyses have to be carefully assessed. Qualitative methods including expert judgement of reasonableness are important for validation. In particular, high (positive or negative) correlations have to be reviewed in order to determine whether they are justified or simply due to the lack of data.

Dependencies within a cell

[18] To complicate matters further, many operational events do not have a clearly defined date of occurrence but consist of a chain of incidents and are spread over a period of time.

If loss events in a particular cell do not occur independently the frequency distribution is not Poisson. In case of positive correlation it might be more appropriate to model the frequency distribution in this cell as Negative Binomial. We refer to section 6 for a detailed discussion.

The standard assumption in the LDA model is that

- frequency and severity distributions in a cell, i.e. N_k and S_k, are independent and
- the severity samples $S_{k1}, \ldots, S_{kN_k}$ are independent and identically distributed.

For validating the first assumption we have calculated rank correlations of frequency and severity time series. The independence of severity samples has been tested by calculating Durban-Watson autocorrelation coefficients of monthly severity data and sequential losses respectively. In summary, these statistical analyses support the LDA standard assumptions stated above. It also seems questionable whether any extension of the model is justified as long as there is not sufficient data available for its calibration. After all, the standard LDA model has been successfully applied in the insurance industry for many years.

Dependencies between cells

Intuitively, dependencies within a cell are higher than dependencies between different cells: all losses in a cell have the same event type and business line, they are often affected by the same internal processes, control and management structures and have occurred in the same business environment. On the other hand, losses in different cells, say fraud in retail banking and employee claims in investment banking, usually have nothing in common. The regulatory debate, however, is rather focused on modelling dependencies between cells, i.e. between different business lines or event types (see, for example, CEBS(2006)).

Central issues for modelling dependencies between cells are:

- What is the best level for specifying dependencies: frequencies, severities or loss distributions?
- What are appropriate mathematical models for dependencies in operational risk?

In DB's LDA model, dependencies of frequency distributions are specified. There are a couple of reasons for this choice.

- First of all, we think that frequencies are the natural concept for specifying dependence: dependencies of loss events are usually reflected in the date of occurrence of these losses. As a consequence, the frequencies in the respective cells will be positively correlated.

- The statistical analysis of frequency correlations shows a positive average correlation although significant frequency correlations are found between specific subsets of cells only.
- The third advantage is that the specification of dependencies at frequency level does not complicate the simulation of loss distributions (see section 9).

We do not assume positive correlations between severity distributions. More precisely, severity samples $S_{ki}, \ldots, S_{lj}$ are assumed independent for all i, j, k, l. The main reasons are:

- We have calculated linear and rank correlations of monthly severity time series and have not observed severity correlations in our loss data except for losses that are caused by the same event. These so-called split losses are not modelled separately but as one event with loss amount specified as the sum of the split loss components (see subsection 3.3).
- From a conceptual point of view, modelling severity correlations is more difficult than modelling frequency correlations (see Frachot et al., 2004): assume that the severity distributions in cells k and l are correlated, i.e. S_{ki}, S_{lj} have a positive correlation for all $i = 1, \ldots, N_k, j = 1, \ldots, N_l$. It immediately follows that also the severity samples within these cells are correlated, i.e. S_{ki}, S_{lj} have a positive correlation for all $i = 1, \ldots, N_k, j = 1, \ldots, N_l$. This is a contradiction to one of the underlying assumptions in the standard LDA.

  Another approach to modelling severity correlations is based on a weaker assumption that is consistent with the standard LDA model. It is used in common shock models (Lindskog and McNeil, 2001): instead of correlating severity samples S_{ki} in cell k to all samples $S_{l1}, \ldots, S_{lN_l}$ in cell l it is only assumed that S_{ki} is correlated to one specific sample S_{lj}. Hence, there exists a link between S_{ki} and S_{lj}, for example both losses are caused by a common shock. Since common shock events correspond to split losses in the internal data set this "weak" form of severity correlation is already incorporated in DB's LDA model.

Some authors propose to model dependencies between the aggregate losses L_k (see Di Clemente and Romano (2004), Chapelle et al. (2004), Nyström and Skoglund (2002)). We think, however, that it is more intuitive to directly induce dependencies between the components of the aggregate losses, i.e. frequencies or severities. Moreover, the dependence of aggregate losses has to be applied in the context of capital calculation after generation of the aggregate loss distribution and incorporation of insurance mitigation. As the insurance model adds another layer of dependence between cells it would be particularly hard to find a meaningful way to apply this approach.

7.2 Modelling frequency correlations

We model dependencies of frequency distributions with copula functions. Copulas are the most general mathematical concept for specifying dependence (see Nelson (1999) for an introduction). Their application to operational risk has been proposed in a number of papers, for instance in Ceske et al. (2000), Song (2000), Frachot et al. (2001), Frachot et al. (2004) and Bee (2005). For alternative mathematical concepts like common shock models and factor models we refer to Lindskog and McNeil (2001), Kühn and Neu (2002), Nyström and Skoglund (2002) and Powojowski et al. (2002).

In the current data environment, the selection of a copula family for modelling frequency correlations cannot be based on statistical tests. For simplicity reasons, we have decided to model dependencies of frequency distributions with Gaussian copulas. The correlation estimates are chosen very conservatively in order to take into account the uncertainty regarding the correct copula family.

Let m be the number of cells, $F_1, \ldots, F_m$ the distribution functions of the frequency distributions in the individual cells and C_Σ the Gaussian copula with correlation matrix Σ. Then the distribution function

$$C_\Sigma(F_1(x_1), \ldots, F_m(x_m))$$

specifies the m-variate frequency distribution.

It remains to specify the correlation matrix of the Gaussian copula. In principle, this matrix can be derived in different ways. There are two main approaches: multivariate and bivariate estimators. The bivariate estimators determine the correlation between two cells and use the results to construct the multivariate correlation matrix. Motivated by simulation results in Lindskog (2000) we apply a bivariate method.

Even though the Gaussian copula is specified by a linear correlation matrix (Pearsons Correlation Coefficient), we use a nonparametric rank estimator. The key concept of a rank estimator is to rank the observed data and derive the correlation coefficients from these ranks. The main advantage over the common sample (linear) correlation estimator lies in the fact that the distribution of the ranks is exactly known, namely a uniform distribution. Of course, some information is lost by replacing the observed values by their ranks, but the advantages over the standard linear estimator are significant. In particular, rank estimators are not sensitive to outliers and they are independent of the underlying distributions (compare to Embrechts et al. (1999)).

Having aggregated the frequency data on a monthly basis, we estimate the correlation between the cells by Spearman's Rho. This rank estimator is the common sample correlation estimator applied to the ranks of the data. See McNeil et al. (2005) and Chapelle et. al (2004) for its application and test results.

Our analysis indicates that significant frequency correlations are found between specific subsets of cells only: the average correlation is around 10%, higher correlations are only found between cells of the event type Execution, Delivery and Process Management. In the EC calculations in 2005 and 2006, a homogeneous correlation matrix with a correlation coefficient of 50% is used. In spite of this very conservative setting, the correlation between the aggregated loss distributions in the individual cells is rather low (due to independent severity distributions). As a consequence, the capital estimates are insensitive to frequency correlations or copula assumptions for frequencies. (We refer to Bee (2005) for a comparison between Gaussian and Gumbel copulas applied to frequencies and their impact on capital estimates.) In other words, we have experienced strong diversification effects, i.e. the 99.98%-quantile of the loss distribution at Group level is significantly lower than the sum of the quantiles of the loss distributions in the cells. This feature of the model is in line with the results reported in Frachot et al. (2004). Chapelle et al. (2004) also find low or even negative correlations between frequencies and observe a large diversification benefit.

8 Risk mitigation

8.1 Insurance models in the Loss Distribution Approach

Insurance is a well-established risk management tool that has been used by the banking sector for many decades. Critical to the viability of recognizing insurance as a risk mitigant in the Basel II Accord is a sound methodology for calculating the capital relief resulting from insurance.[19] A number of approaches based on insurance premia, limits or actual exposures have been proposed (see, for instance, Insurance Companies (2001)). The general set-up of DB's insurance model is tailored to the Loss Distribution Approach and follows the principles laid out in the working paper by the Insurance Companies (2001). Similar models are discussed in Brandts (2005) and Bazzarello et al. (2006).

Compared to other methodologies the Loss Distribution Approach has the benefit of allowing a fairly accurate replication of the risk profile of a bank, including the risk reducing effect of insurance. Transferring risk to an insurer through insurance products alters the aggregate loss distribution by reducing the severity of losses that exceed the policy deductible amount. The frequency of loss is unaffected by insurance. More precisely, when frequency and severity distributions are combined through simulation, each individual loss point can be compared to the specific insurance policies purchased by the bank and the corresponding policy limits and deductibles. As a consequence, an insurance model in the Loss Distribution Approach consists of two main components:

- a quantitative model of the individual insurance contracts,

[19] The recognition of insurance mitigation will be limited to 20% of the total operational risk capital charge for regulatory purposes.

- a mapping from the OR event types to the insurance contracts.

8.2 Modelling insurance contracts

Insurance contracts are characterized by certain specifications regarding the compensation of losses:

- A deductible d is defined to be the amount the bank has to cover by itself.
- The single limit l of an insurance policy determines the maximum amount of a single loss that is compensated by the insurer. In addition, there usually exists an aggregate limit l_{agg}.
- The recognition of the risk mitigating effects of insurance contracts within the AMA is subject to a number of prerequisites by the regulators. In particular, the methodology for recognizing insurance has to take into account the certainty of payment, the speed of payment, insurance through captives or affiliates and the rating of the insurer. We translate the rating into a default probability p and model defaults of the insurer by the binary random variable I^p, where $I^p = 0$ with probability p and $I^p = 1$ otherwise. The other limitations are transferred by expert judgement into a haircut H which is applied to the insurance payout.

In mathematical terms, the amount paid by the insurer is defined by

$$I^p \cdot \max(\min(l, x) - d, 0) \cdot H.$$

The actual specification of insurance parameters including haircuts has to be based on insurance data and expert knowledge and is beyond the scope of this paper.

8.3 Mapping OR event types to insurance policies

In the OR capital calculation, insurance contracts are applied to simulated OR losses. This requires a mapping from OR loss categories, e.g. OR event types, business lines or a combination of both, to insurance contracts. If all losses of one OR category are covered by the same insurance policy, this mapping is trivial. If the losses of one OR category can be covered by several insurance policies, the percentage of losses that fall into a specific insurance category need to be determined by expert assessment and historical data. For example, 70% of execution losses might be covered by General Liability, 20% by Professional Liability and 10% are not insured. A mapping from OR Level 3 Loss Events to standard insurance products can be found in Insurance Companies (2001).

9 Calculation of Economic Capital and capital allocation

In the previous sections the quantitative components of the loss distribution model have been presented. This section focuses on the simulation of the aggregate loss distribution (including insurance) at Group level and the calculation and allocation of Economic Capital.

9.1 Risk measures and allocation techniques

In recent years, a well-founded mathematical theory has been developed for measuring and allocating risk capital. The most important cornerstone is the formalization of the properties of a risk measure given in Artzner et al. (1999). Their axiomatization provides an appropriate framework for the analysis and development of risk measures, e.g. Expected Shortfall (Rockafellar and Uryasev, 2000; Acerbi and Tasche, 2002). General principles for capital allocation can be found in a number of papers, for example in Kalkbrener (2005).

The general theory of measuring and allocating risk is independent of specific risk types. In fact, it is important that the same measurement and allocation techniques are applied in the Economic Capital calculation for market, credit and operational risk. This is a prerequisite for a consistent framework for measuring risk and performance across the bank.

Economic Capital is a measure designed to state with a high degree of certainty the amount of capital needed to absorb unexpected losses. The exact meaning of "high degree of certainty" depends on the risk tolerance of the bank. A popular rule is to choose the tolerance level that reflects the average default rate associated with the rating of the bank. This rule immediately translates into a definition of Economic Capital based on Value-at-Risk (defined as a quantile of the loss distribution): for example, Deutsche Bank's EC is determined by the Value-at-Risk at confidence level 99.98% reflecting the AA+ target rating of the bank. The simple link between ratings and Value-at-Risk based Economic Capital is one reason why this specific Economic Capital methodology has become so popular and has even achieved the high status of being written into industry regulations.

The confidence level of 99.98% requires a high precision in modelling tails of loss distributions. It is obvious that this problem is particularly challenging for operational risk. In fact, this very high confidence level is one of the reasons why internal data has to be supplemented by external losses. Even if external losses and generated scenarios are used as additional information the 99.98% quantile (and to a lesser degree the 99.9% quantile) has to be based on extrapolation techniques beyond the highest losses in the data base. It is inevitable that there is some degree of uncertainty about the correct value of these high quantiles due to the scarcity of data in the extreme tail.[20] We refer to Mignola and Ugoccioni (2006) and to section 11 for more information.

The standard procedure for credit and operational risk is to specify the Economic Capital as Value-at-Risk minus Expected Loss. In credit risk, Expected Loss denotes the mean of the portfolio loss distribution. If the same definition is applied in operational risk the Expected Loss significantly exceeds the annual loss in an average year: due to the fat-tailed severities the mean of the aggregate OR loss distribution is usually much higher than its median. This is the reason why alternative definitions of Expected Loss are analyzed, e.g. replacing the mean of the aggregate loss distribution by the median (see Kaiser and Kriele (2006) for approaches to the calculation of EL estimates). In Deutsche Bank, expert opinion is used to transform an initial EL estimate based on historic internal loss data into an EL forecast for the one-year time horizon. Using raw internal losses as starting point for an EL estimate has the advantage that the Expected Loss methodology is based on internal data only and that no sophisticated statistical techniques are required. Furthermore, the involvement of OR experts from the different divisions of the bank ensures that the quantification of the Expected Loss is closely linked to the budgeting and capital planning process of the bank. This feature of the EL quantification process is particularly important for meeting the regulatory requirement that the bank has measured and accounted for its EL exposure. This requirement is a precondition for subtracting EL from the 99.9% Value-at-Risk in the calculation of Regulatory Capital.

The allocation of risk capital to BL/ET cells is based on Expected Shortfall contributions. Intuitively, the Expected Shortfall contribution of a BL/ET cell is defined as its contribution to the tail of the aggregate loss distribution:

$$\mathrm{E}(L_k \,|\, L > \mathrm{VaR}_\alpha(L)),$$

where L_k denotes the loss distribution in cell k, L is the aggregate loss distribution at Group level and α has been chosen such that

$$\mathrm{E}(L \,|\, L > \mathrm{VaR}_\alpha(L)) = \mathrm{VaR}_{0.9998}(L).$$

Expected Shortfall allocation has a number of theoretical and practical advantages compared to classical allocation techniques like volatility contributions (Kalkbrener et al., 2004). Furthermore, it is consistent with the methodology used in DB's EC methodology for credit risk and risk type diversification.

In the following two subsections, the Monte Carlo based EC calculation and allocation is described in more detail.

[20] An alternative approach would be the replacement of the direct calculation of the 99.98% quantile by the calculation of a lower quantile and a subsequent scaling. However, the scaling factor for the 99.98% quantile is as uncertain as the direct calculation of the 99.98% quantile and therefore does not improve the reliability of the EC estimate.

9.2 Simulation of the aggregate loss distribution

The aggregate loss distribution at Group level cannot be represented in analytic form. In order to overcome this problem, different approximation techniques have been proposed in the literature including Panjer recursion, fast Fourier transform methods and approximations based on properties of subexponential distributions (Klugman et al., 2004; McNeil et al., 2005; Böcker and Klüppelberg, 2005). However, all these techniques limit the functionality of the underlying model. In particular, it becomes more difficult to incorporate advanced modelling features like insurance or more sophisticated allocation techniques. We have therefore chosen Monte Carlo simulation for calculating and allocating Economic Capital. This decision has been made despite the fact that Monte Carlo estimates of high quantiles of the aggregate loss distribution and therefore MC estimates of EC are numerically unstable. In order to reduce the number of required MC samples we have developed a semi-analytic technique for Expected Shortfall allocation. Furthermore, we are working on the adaptation of importance sampling techniques to LDA models. Importance sampling has been successfully applied in market and credit risk (see Glassermann et al. (1999), Kalkbrener et al. (2004) and Glasserman (2003)). The adaptation of this variance reduction technique to operational risk requires the introduction of rare event simulation techniques for heavy-tailed distributions (see, for example, Asmussen et al., 2000). We will report on the variance reduction techniques used in a forthcoming paper.

These are the main steps in generating one MC sample of the aggregate loss distribution at Group level:

Simulation of loss distributions in individual cells

- Generate a sample $(n_1, \ldots, n_m)$ of the m-dimensional frequency distribution specified by the Gaussian copula C_Σ. This can easily be done by a standard software package or by implementing a simulation algorithm as proposed in e.g. McNeil et al. (2005).
- Compute n_k samples $s_{k1}, \ldots, s_{k\,n_K}$ of the severity distribution in cell $k, k = 1, \ldots, m$.

Incorporation of insurance

- Take all losses in all cells and order them randomly.
- Determine the insurance policy (if available) for each loss.
- Apply the insurance model to the ordered loss vector and compute the corresponding vector of net losses (loss minus compensation).

Sample of aggregate loss distribution

- The sum of the net losses is a sample l of the aggregate loss distribution at Group level.

9.3 Monte Carlo estimates for Economic Capital

The simulation of the aggregate loss distribution is used as input for the calculation and allocation of Economic Capital: assume that an ordered list $l^1 > \ldots > l^n$ of samples of the aggregate loss distribution at Group level has been computed. The EC computation is done in several steps:

EC calculation

The Economic Capital (before qualitative adjustment) is calculated as the 99.98% quantile of the aggregate loss distribution at Group level minus the Expected Loss. Note that the 99.98% quantile is represented by l^q, where $q := n/5000$.

EC allocation to BL/ET cells

The Economic Capital is then allocated to the cells of the BL/ET matrix using Expected Shortfall techniques. More precisely, define an index r and consider the r highest losses $l^1 > \ldots > l^r$ as a representation of the tail. In DB's LDA model, r is chosen such that the sum $\sum_{i=1}^{r} l^i/r$ is approximately the 99.98% quantile l^q . Note that it is easy to compute a decomposition $l^i = \sum l_k^i$, where l_k^i denotes the contribution of the cell k to loss l^i: since each loss l^i is a sum of losses in individual cells the l_k^i are obtained as byproducts in the simulation process. The capital allocated to cell k is the Expected Shortfall contribution defined by $\sum_{i=1}^{r} l_k^i/r$.

EC allocation to business lines

The EC of the individual business lines (before qualitative adjustment) is obtained by aggregation of the Economic Capital allocated to the BL/ET cells: the capital of each business line is the sum of

- the capital allocated to its three cells covering the event types Fraud and Clients, Products, Business Practices and Execution, Delivery, Process Management and
- a proportional contribution from the Group cells and the cells for the event types Infrastructure and Employment Practices, Workplace Safety.

The above procedure ensures that the entire amount of EC for the Group is allocated to the business lines.

The sub-allocation of EC to business units within business lines cannot be based on Expected Shortfall techniques since separate loss distributions are not specified for these units. Deutsche Bank applies risk indicators instead.

Qualitative adjustment of EC

The EC of each business line is scaled by a multiplicative factor that reflects qualitative adjustments. The Economic Capital at Group level has to be modified accordingly in order to ensure additivity of EC figures, i.e. the adjusted EC of the Group is obtained by adding up the adjusted EC figures of the individual business lines. More information on qualitative adjustments is given in the following section.

10 Incorporation of business environment and internal control factors

Apart from generated loss scenarios LDA models mainly rely on loss data and are inherently backward looking. It is therefore important to incorporate a component that reflects changes in the business and control environment in a timely manner.

There have been intense discussions in the industry regarding different strategies a bank could follow in order to perform qualitative adjustments. The ways that adjustments are performed vary in form and complexity and the processes for collecting relevant information differ across financial institutions. Common industry practice is to compile the data into a scoring mechanism which translates qualitative information into numerical values. The most prevalent forms of qualitative adjustment utilize data from key risk indicators (KRIs) and risk self-assessment.

We will not discuss these general topics in this paper but focus on the incorporation of qualitative adjustments into LDA models.[21] Qualitative adjustments might occur in different components of an LDA model, e.g. adjustments of

- the parameters of frequency distributions,
- the parameters of severity distributions, or
- the Economic Capital of business lines.

We have decided to directly adjust the allocated Economic Capital of a business line (third option). The main reasons for this decision are simplicity and transparency:

- The application of qualitative adjustments to the Economic Capital separates the quantitative EC calculation from the qualitative adjustments and therefore reduces the complexity of the model.
- It is important for the acceptance of the methodology by the business that the impact of qualitative adjustments on allocated Economic Capital is as transparent as possible. This transparency is only guaranteed by applying adjustments directly to the allocated EC.

[21] For more information on self-assessment and KRIs we refer to Anders and Sandstedt (2003) and Davies et al. (2006)

- A further consequence of a direct EC adjustment is that QA for one business unit does not affect the other business units.

The direct application of qualitative adjustments to Economic Capital is difficult to justify with statistical means. However, we think that the advantages listed above outweigh this disadvantage at the current stage of the model development. We also believe that the statistical validation of QA techniques is one of the main challenges for the development of the next generation of LDA models.

11 Analysis and validation of LDA models

The final section of this paper deals with model analysis and validation. The techniques and results in this section are not restricted to Deutsche Bank's LDA model but can be applied to a more general class of models (see section 11.1.1 for a definition).

Basel II requires that an operational risk AMA model generates a capital measure to the 99.9%-ile confidence level over a one-year holding period. The Economic Capital of Deutsche Bank is derived from an even higher quantile, i.e. the 99.98% quantile of the aggregate loss distribution for one year. The objectives of model validation are to show that the estimates for Regulatory and Economic Capital are reasonable. Given these high quantiles it is obvious that a validation of OR capital requirements is a very ambitious task. But not only demanding confidence levels complicate the application of statistical validation techniques in operational risk, in particular compared to market risk: there are a number of additional problems including

1. the shortage of relevant operational risk data,
2. the context dependent nature of operational risk data, and
3. the current lack of a strongly risk sensitive exposure measure in operational risk modelling (cf market and credit risk).

As a consequence, operational risk measurement and statistical validation techniques inevitably suffer from a higher degree of uncertainty than market risk validation and have to be combined with qualitative methods including expert judgement of reasonableness.

The validation of an operational risk model involves a review of data inputs, model methodology (including any assumptions contained within the model such as correlations), and model outputs (AMA Working Group, 2004). Since the specific methodology implemented in Deutsche Bank's LDA model has been extensively discussed in previous sections we now focus on sensitivity analysis and output validation. More precisely, we present a sensitivity analysis of the model components frequencies, severities, dependence structure and insurance. This analysis comprises the assessment of

1. the capital impact of different modelling assumptions and

2. the sensitivity of capital requirements to parameter changes

for each model component (section 11.2). The analysis uses basic properties of LDA models presented in sections 11.1.2 and 11.1.3 and is therefore not limited to the model implemented at Deutsche Bank. Section 11.3 deals with the impact analysis of stress scenarios. Section 11.4 briefly outlines the inherent problems with the application of backtesting techniques to OR models and reviews some of the benchmarks for AMA capital charges suggested in the literature. However, the main focus of section 11.4 is on an approach for benchmarking quantiles in the tail of the aggregate loss distribution of the LDA model against individual data points from the underlying set of internal and external losses.

11.1 Model characteristics

11.1.1 Definition of a general class of LDA models

The formal framework for our analysis is a model which is based on m distinct (aggregate) loss distributions. These loss distributions may correspond to combinations of business lines and event types, i.e. cells in the BL/ET matrix. The specification of each loss distribution follows an actuarial approach: separate distributions for loss frequency and severity are derived and then combined to produce the aggregate loss distribution. More precisely, let $N_1, \dots, N_m$ be random variables which represent the loss frequencies, i.e. the values of N_i are non-negative integers. For each $i \in \{1, \dots, m\}$, let $S_{i1}, S_{i2}, \dots$ be iid random variables independent of N_i. These variables are independent samples of the severity distribution S_i of loss type i. The loss variable L_i in this cell is defined by

$$L_i := \sum_{r=1}^{N_i} S_{ir}.$$

The aggregate loss distribution at Group level is defined by

$$L := \sum_{i=1}^{m} L_i.$$

11.1.2 Variance analysis

Variance analysis does not provide information on quantiles of loss distributions but

1. it quantifies the impact of frequencies, severities and their dependence structure on the volatility of aggregate losses and

2. it is independent of specific distribution assumptions, i.e. the first two moments of the aggregate loss distribution only depend on the first two moments of the frequency and severity distributions.

Variance analysis therefore is a rather general technique that provides insight into the significance of the individual model components.

For calculating moments we make the additional assumption that severity samples in different cells are independent, i.e. S_{ik} and S_{jl} are independent for $i, j \in \{1, \dots, m\}$ and $k, l \in \mathbb{N}$. However, we allow for dependence between the frequency variables $N_1, \dots, N_m$.

Under the assumption that the first and second moments of frequency distributions N_i and severity distributions S_i exist we obtain the following formulae for the expectation and the variance of the aggregate loss distribution at cell and Group level:

$$\mathrm{E}(L_i) = \mathrm{E}(N_i) \cdot \mathrm{E}(S_i) \tag{2}$$

$$\mathrm{E}(L) = \sum_{i=1}^{m} \mathrm{E}(N_i) \cdot \mathrm{E}(S_i) \tag{3}$$

$$\mathrm{Var}(L_i) = \mathrm{E}(N_i)\mathrm{Var}(S_i) + \mathrm{Var}(N_i)\mathrm{E}(S_i)^2 \tag{4}$$

$$\begin{aligned}\mathrm{Var}(L) = & \sum_{i=1}^{m} \mathrm{E}(N_i)\mathrm{Var}(S_i) + \mathrm{Var}(N_i)\mathrm{E}(S_i)^2 + \\ & \sum_{i,j=1,\ i\neq j}^{m} \mathrm{Cov}(N_i, N_j)\mathrm{E}(S_i)\mathrm{E}(S_j)\end{aligned} \tag{5}$$

The derivation of these formulae is a rather straightforward generalization of the proof of Proposition 10.7 in McNeil et al. (2005)).

Formula (5) provides the general representation of $\mathrm{Var}(L)$ for the model class specified in section 11.1.1. In order to better understand the impact of the different model components, i.e. frequencies, severities and frequency correlations, it is useful to simplify formula (5) in an idealistic homogeneous model. More precisely, we assume that

1. frequencies in all cells follow the same distribution N,

2. severities in all cells follow the same distribution S, and

3. all frequency correlations equal a fixed c, i.e. for $i, j \in \{1, \dots, m\}$

$$c = \mathrm{Corr}(N_i, N_j) = \mathrm{Cov}(N_i, N_j)/\mathrm{Var}(N).$$

Then the variance $\mathrm{Var}(L)$ of the aggregate loss distribution can be written in the form

$$\mathrm{Var}(L) = m \cdot (\mathrm{E}(N) \cdot \mathrm{Var}(S) + \mathrm{Var}(N) \cdot \mathrm{E}(S)^2 \cdot (c \cdot (m-1) + 1)). \tag{6}$$

11.1.3 Loss distributions for heavy-tailed severities

Since OR capital requirements of large international banks are mainly driven by rare and extreme losses it is quite natural to work with severity distributions that have been applied in insurance theory to model large claims. These distributions typically belong to the class of subexponential distributions.

The defining property of subexponential distributions is that the tail of the sum of n subexponential random variables has the same order of magnitude as the tail of the maximum variable. More precisely, let $(X_k)_{k\in\mathbb{N}}$ be positive iid random variables with distribution function F and support $(0,\infty)$. Then F is subexponential if for all $n \geq 2$,

$$\lim_{x\to\infty} \frac{\mathbb{P}(X_1 + \ldots + X_n > x)}{\mathbb{P}(\max(X_1, \ldots, X_n) > x)} = 1. \tag{7}$$

The following property justifies the name subexponential for distributions satisfying (7): if the distribution function F is subexponential then for all $\epsilon > 0$,

$$e^{\epsilon x}\overline{F}(x) \to \infty, \qquad x \to \infty$$

where $\overline{F}(x) = 1 - F(x)$ denotes the tail of F. Hence, $\overline{F}(x)$ decays to 0 slower than any exponential $e^{\epsilon x}$. Examples of subexponential distributions include the families Pareto, Weibull ($\tau < 1$), lognormal, Benktander-type-I and -II, Burr, loggamma.

Under certain conditions, random sums of subexponential variables are subexponential. Let $(X_k)_{k\in\mathbb{N}}$ be iid subexponential random variables with distribution function F. Let N be a random variable with values in $\mathbb{N}_0$ and define

$$p(n) := \mathbb{P}(N = n), \qquad n \in \mathbb{N}_0.$$

Denote the distribution function of the random sum $\sum_{i=1}^N X_i$ by G, i.e.

$$G(x) = \mathbb{P}(\sum_{i=1}^{N} X_i \leq x), \qquad x \geq 0.$$

Suppose that $p(n)$ satisfies

$$\sum_{n=0}^{\infty}(1+\epsilon)^n p(n) < \infty \tag{8}$$

for some $\epsilon > 0$. Then

$$\text{the distribution function } G \text{ is subexponential and } \lim_{x\to\infty} \frac{\overline{G}(x)}{\overline{F}(x)} = \mathrm{E}(N). \qquad (9)$$

If N follows a Poisson or negative binomial distribution then (8) is satisfied (see examples 1.3.10 and 1.3.11 in Embrechts et al. (1997)).

The above result states that the tail of the aggregate loss distribution only depends on the tail of the subexponential severity distribution. An explicit approximation formula for $\overline{G}(x)$ can be found in Böcker and Klüppelberg (2005).

For the total sum $L := \sum_{i=1}^{m} L_i$, results similar to (9) hold under extra conditions. We refer to Embrechts et al. (1997) and McNeil et al. (2005).

11.2 Sensitivity analysis of LDA models

In the previous subsection, basic facts on frequencies, severities and aggregate loss distributions have been presented. These results will now be used to analyze for each model component:

1. the capital impact of different modelling assumptions and
2. the sensitivity of capital requirements to parameter changes.

Sensitivity analysis gives an impression about which inputs have a significant impact upon the risk estimates and by implication, need prioritizing in terms of ensuring their reasonableness.

11.2.1 Frequencies

The impact of the shape of frequency distributions on capital requirements is rather limited. The reason is that OR capital is mainly driven by cells with fat-tailed severities. In those cells, the distribution assumptions for frequencies have no significant impact on the aggregate loss distributions. This is a consequence of (9): for subexponential severities, the tail of the aggregate loss distribution is determined by the tail of the severity distribution and the expected frequency (but not its precise shape). This result is in line with formula (4) which says that the impact of the frequency distribution on the variance of the aggregate loss distribution depends on the relationship of the frequency vol $\mathrm{Var}(N)/\mathrm{E}(N)$ and the severity vol $\mathrm{Var}(S)/\mathrm{E}(S)^2$. In high impact cells, the volatility of severities dominates and the actual form of the frequency distribution is of minor importance. These theoretical results are confirmed by our own test calculations discussed in section 6 and by a number of papers including Böcker and Klüppelberg (2005) and De Koker (2006).

11.2.2 Severities

OR capital requirements are mainly driven by individual high losses. Severity distributions specify the loss size and are therefore the most important component in quantitative OR models. This claim is again supported by formula (4) and the result in (9): in a cell with fat-tailed severity distribution the volatility and the shape of the aggregate loss distribution are mainly determined by the severity distribution. All relevant aspects of severity models therefore have an impact on the aggregate loss distribution and on capital estimates. In particular, OR capital requirements depend on

1. weights and techniques for combining different data sources for severity modelling and

2. distribution assumptions for severities, in particular for severity tails.

11.2.3 Insurance

The impact of insurance on risk capital depends on the insurance coverage of the dominant OR event types. In Deutsche Bank's LDA model, most of the OR capital is allocated to the event type Clients, Products and Business Practices. The high losses in this event type are predominantly classified as insurance type Professional Liability (PL). Since individual extreme losses are the main driver for capital the single limit for Professional Liability is a very significant insurance parameter in the current model setup. Another important parameter is the actual coverage of Professional Liability, i.e. the proportions between insured and uninsured PL losses.

11.2.4 Dependence

In a bottom-up Loss Distribution Approach, separate distributions for loss frequency and severity are modelled at "business line/event type" level and are then combined. In this basic LDA model, dependencies between the occurrence of loss events and the size of losses might occur in a cell or between different cells.

If loss events in a particular cell do not occur independently the frequency distribution is not Poisson. In case of positive correlation it might be more appropriate to model the frequency distribution in this cell as Negative Binomial. The impact of these distribution assumptions for frequencies has already been discussed in section 11.2.1.

The impact of dependencies between cells depends on the level where these dependencies are modelled, e.g. frequencies, severities or aggregate losses. Equality (6) provides information on the impact of frequency correlations on capital requirements: it depends on

1. the number of (relevant) cells m and

2. the relationship of $Var(F)/E(F)$ (frequency vol) and $Var(S)/E(S)^2$ (severity vol).

In Deutsche Bank's LDA model, for example, the number of relevant (= high impact) cells is relatively small and $Var(S)/E(S)^2$ clearly dominates $Var(F)/E(F)$ in these cells. The impact of frequency dependence is therefore insignificant. In contrast, dependence between severity distributions or loss distributions has a significant impact on capital. For example, assume that loss distributions in different cells are comonotonic. Then every α-quantile of the aggregate loss distribution (Group level) equals the sum of the α-quantiles of the loss distributions in the cells (see, for example, McNeil et al. (2005)). For $\alpha = 0.9998$, this sum would clearly exceed the Economic Capital calculated in Deutsche Bank's LDA model.

11.3 Impact analysis of stress scenarios

Stress testing has been adopted as a generic term describing various techniques used by financial firms to gauge their potential vulnerability to exceptional but plausible events. In a typical OR stress scenario, internal and/or external losses are added or removed and the impact on capital is analyzed. The scenarios are often provided by business and OR management to quantify

1. potential future risks,
2. the impact of business strategies, and
3. risk reduction by OR management.

In this context, stress scenarios have a wider scope than scenario analysis which is limited to the generation of single loss data points (compare section 3.4.3).

Stress tests also play an important role in model analysis and validation. The objective is to study the response of the model to changes in model assumptions, model parameters and input data. In particular, we have analyzed the sensitivity of the model to changes in the loss data set. These stress tests are an important part of the model development.

A different application of stress scenarios is discussed in the next section: LDA model outputs are tested by comparison with adverse extreme, but realistic, scenarios.

11.4 Backtesting and benchmarking

Backtesting, as the name suggests, is the sequential testing of a model against reality to check the accuracy of the predictions. More precisely, the model outcomes are compared with the actual results during a certain period. Backtesting is frequently used for the validation of market risk models. In credit and operational risk, the inherent shortage of loss data severely restricts the application of backtesting techniques to capital models.

The most useful validation of an LDA model would be the backtesting of capital estimates against actual annual losses. However, it is obvious that capital requirements derived from high quantiles of annual loss distributions cannot be tested against the actual loss history of a single bank in a reasonable way.[22] In this paper we propose a different approach. It is based on the assessment that OR capital requirements are mainly driven by individual high losses. We argue that individual data points from the set of internal and external losses can therefore be used as benchmarks for quantiles in the tail of the aggregate loss distribution of the LDA model. More precisely, the basic idea is to compare

1. the tail of the aggregate loss distribution calculated in a bottom-up LDA model, e.g. Deutsche Bank's LDA model, and
2. the tail of an aggregate loss distribution directly specified at Group level.

An approximation of the tail specified at Group level can be directly derived from the loss data, i.e. each quantile of the aggregate loss distribution is represented by a single data point. Hence, this method allows to compare quantiles of the aggregate loss distribution in a bottom-up LDA model to individual losses from the underlying set of internal and external loss data.

We will now describe the proposed technique in more detail. The first step is the specification of an aggregate loss distribution at Group level, i.e. all internal and external losses across business lines and event types are considered. However, since we are only interested in the tail of the distribution we take only losses above a high threshold, say 1m, into account. Let n denote the bank's average annual loss frequency above 1m. We assume that the frequency distribution satisfies the condition in (8) and the severity distribution S at Group level is subexponential. Let $S_1, \ldots, S_n$ be independent copies of S. It follows from (7) and (9) that the α-quantiles of the aggregate loss distribution and the maximum distribution $\max(S_1, \ldots, S_n)$ converge for $\alpha \to 1$. The α-quantile of the maximum distribution is now identified with the $1 - ((1 - \alpha)/n)$-quantile of the severity distribution. The latter is directly derived from the data set for appropriate α's. Note that the amount of loss data provides a limit for the confidence level α, i.e. depending on the annual loss frequency n the $1 - ((1 - \alpha)/n)$-quantile of the empirical severity distribution is only specified up to a certain α.

In summary, we propose to compare the α-quantile of the aggregate loss distribution calculated in an LDA model and the $1 - ((1 - \alpha)/n)$-quantile of the empiric

[22] It should be noted that this situation will not improve over time, since context dependency of OR loss data will mean that 5-10 years of data will probably be the maximum useful data set against which banks can validate their capital models (depending on the risk type and the speed of change in environment and control processes).

severity distribution derived from the complete set of internal and relevant external losses above 1m. Application of this method to DB's LDA model yields the results shown in figure 4: the tail of the aggregate loss distribution has the same magnitude as the corresponding high impact losses in the underlying data set.

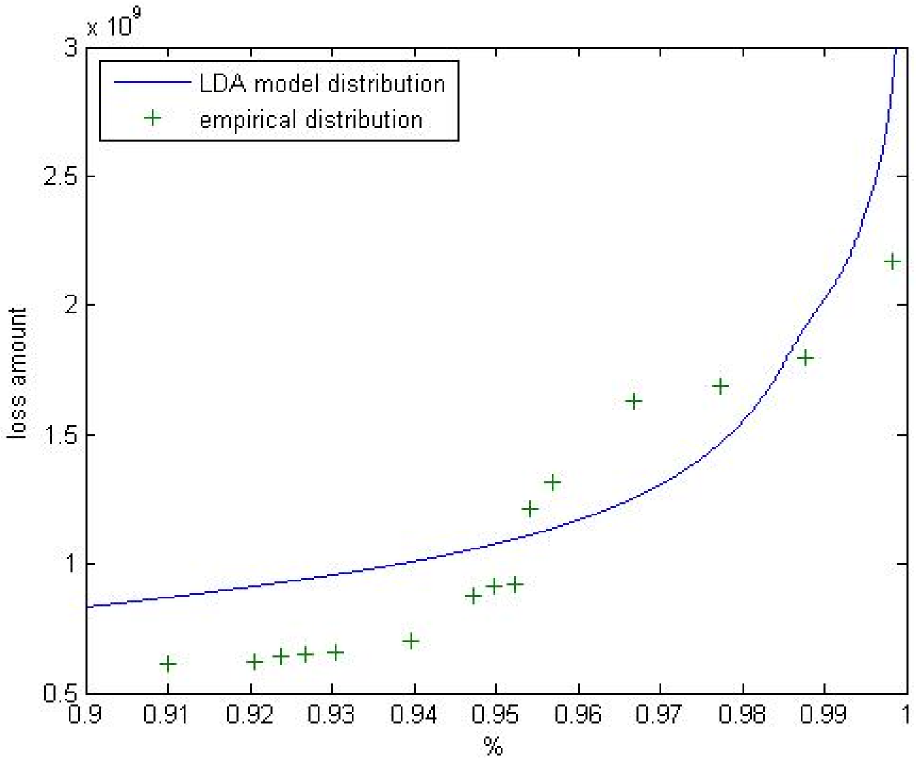

Figure 4: 1y aggregate loss distribution from AMA model vs. corresponding high impact losses in the underlying data set.

Our approach has a lot in common with the methodology proposed in De Koker (2006). First of all, both approaches are based on the assumption that the size of a single (super-sized) loss event provides a meaningful capital estimate. The likelihood of these events is derived from the frequency above a high threshold and the shape of the severity distribution. The main difference between both approaches is that De Koker (2006) approximates the severity distribution above the threshold by a power law whereas we use an empirical distribution, i.e. raw loss data. Empirical distributions have the advantage that no assumptions on the shape of the severity distribution have to be made but, in contrast to the technique in De Koker (2006), only quantiles up to a certain confidence level can be benchmarked.

Other benchmarks for OR risk estimates have been proposed in the literature:

1. comparison of a bank's operational risk capital charge against a bank's close peers,

2. comparison of the AMA capital charge against the BIA or TSA capital charges,

3. comparison of the LDA model outputs against adverse extreme, but realistic, scenarios.

Although these tests help to provide assurance over the appropriateness of the level of capital there are obvious limitations. Peer comparison will not indicate whether banks are calculating the right level of risk capital, just whether they are consistent with each other. Given that the TSA/BIA are non-risk sensitive methodologies, and the AMA is risk-sensitive, the value of any comparison between AMA and TSA/BIA capital charges is questionable. Validation through scenarios is based on the identification of an appropriate range of specific OR events and the estimation of loss amounts in each of these scenarios. This process is highly subjective and requires substantial resources for specifying a comprehensive operational risk profile of a bank.

Given the limitations of OR validation techniques it is particularly important that a model meets the "Use Test", i.e. it is incorporated within the bank's operational risk management framework and used to manage the risks facing the bank. Although the Use Test does not validate the capital estimates calculated in the model it provides assurance that the model is working appropriately.

12 References

Acerbi, C., and Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking and Finance* 26, 1487-1503.

Alderweireld, T., Garcia, J., and Léonard, L. (2006). A practical operational risk scenario analysis quantification. *RISK* (Feb. 2006), 93-95.

Alvarez, G. (2006): Operational risk Economic Capital measurement: mathematical models for analysing loss data. Pages 59-117 in Davis (2006): *The Advanced Measurement Approach to Operational Risk*, Risk Books, London.

AMA Working Group (2004). Operational risk validation. London, August 2004.

Anders, U., and van den Brink, G.J. (2004). Implementing a Basel II scenario-based AMA for operational risk. Pages 343-368 in Ong (2004): *The Basel Handbook*, Risk Books, London.

Anders, U., and Sandstedt, M. (2003). Self-assessment for scorecards. *RISK* (Feb. 2003), 39-42.

Artzner, P., Delbaen F., Eber J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance* 9/3, 203-228.

Asmussen, S., Binswanger, K., and Hoegaard, B. (2000). Rare events simulation for heavy-tailed distributions. *Bernoulli* 6/2, 303-322.

Basel Committee on Banking Supervision (2002a). Sound Practices for the Management and Supervision of Operational Risk. Bank for International Settlements, July 2002. Available at www.bis.org.

Basel Committee on Banking Supervision (2002b). Operational Risk Rules Language. Bank for International Settlements, August 2002. Available at www.bis.org.

Basel Committee on Banking Supervision (2006). International Convergence of Capital Measurement and Capital Standards: A Revised Framework - Comprehensive Version. Bank for International Settlements, June 2006. Available at www.bis.org.

Baud, N., Frachot, A., and Roncalli, T. (2002). Internal data, external data and consortium data for Operational Risk measurement: How to pool data properly? Group de Recherche Operationnelle, Credit Agricole, France. Available at www.gloriamundi.org.

Bazzarello, D., Crielaard B., Piacenza F., and Soprano A. (2006). Modeling insurance mitigation on operational risk capital. *Journal of Operational Risk* 1(1), 57-65.

Bee, M. (2005). Copula-based multivariate models with applications to risk management and insurance. University of Trento. Available at www.gloriamundi.org.

Böcker, K., and Klüppelberg, C. (2005). Operational VAR: a closed-form approximation. *RISK* (Dec. 2005), 90-93.

Brandts, S. (2005). Reducing risk through insurance. Pages 305-314 in Davis (2005): *Operational Risk - Practical Approaches to Implementation*, Risk Books, London.

CEBS (2006). Guidelines on the implementation, validation and assessment of Advanced Measurement (AMA) and Internal Ratings Based (IRB) Approaches. Committee of European Banking Supervisors. Available at www.c-ebs.org.

Ceske, R., Hernández, J.V., and Sánchez, L.M. (2000). Quantifying event risk: the next convergence. *Journal of Risk Finance*, Spring 2000.

Chavez-Demoulin, V., Embrechts, P., and Nešlehová, J. (2005). Quantitative models for operational risk: extremes, dependence and aggregation. Swiss Federal Institute of Technology, Zurich. Available at www.gloriamundi.org.

Chapelle, A., Crama Y., Huebner G., Peters J.-P. (2004). Basel II and operational risk: implications for risk measurement and management in the financial sector. Working Paper, National Bank of Belgium. Available at www.bnb.be.

Chernobai, A., Menn, C., Rachev, S.T., Trück, S., and Moscadelli M. (2006). Treatment of incomplete data: the effects on parameter estimates, EL and UL figures. Pages 145-168 in Davis (2006): *The Advanced Measurement Approach to Operational Risk*, Risk Books, London.

Chernobai, A., Rachev, S., Fabozzi, F. (2005): Composite goodness-of-fit tests for left-truncated loss samples. University of California, Santa Barbara. Available at www.gloriamundi.org.

Cruz, M. (2002). *Modelling, Measuring and Hedging Operational Risk.* Wiley, Chichster, UK.

Embrechts, P., Furrer, H., and Kaufmann, R. (2003). Quantifying regulatory capital for operational risk. *Derivatives Use, Trading and Regulation* 9(3), 217-233.

Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events.* Springer, Berlin Heidelberg.

Embrechts, P., McNeil, A., and Straumann, D. (1999). Correlation and dependence in risk management: properties and pitfalls. ETH Zurich. Available at www.math.ethz.ch.

Davies, J., Finlay, M., McLenaghen, T., Wilson, D. (2006). Key risk indicators – their role in operational risk management and measurement. Pages 215-246 in Davis (2006): *The Advanced Measurement Approach to Operational Risk*, Risk Books, London.

de Fontnouvelle, P., DeJesus-Rueff, V., Jordan, J., and Rosengren, E. (2003). Capital and risk - new evidence on implications of large operational losses. Federal Reserve Bank of Boston, Boston, MA. Available at www.gloriamundi.org.

de Fontnouvelle, P., and Rosengren, E. (2004). Implications of alternative operational risk modeling techniques. Federal Reserve Bank of Boston, Boston, MA. Available at www.gloriamundi.org.

De Koker, R. (2006): Operational risk modelling: where do we go from here. Pages 37-58 in Davis (2006): *The Advanced Measurement Approach to Operational Risk*, Risk Books, London.

Di Clemente, A. and Romano, C. (2004). A copula-Extreme Value Theory approach for modelling operational risk. Pages 189-208 in Cruz (2004): *Operational Risk Modelling and Analysis*, Risk Books, London.

Dutta K. and Perry, J. (2006). A Tail of Tails: An Empirical Analysis of Loss Distribution Models for Estimating Operational Risk Capital. Working Paper, Federal Reserve Bank of Boston, Boston, MA. Available at www.gloriamundi.org.

Frachot, A., Georges, P., and Roncalli, T. (2001). Loss Distribution Approach for Operational Risk. Group de Recherche Operationnelle, Credit Agricole, France.

Available at www.gloriamundi.org.

Frachot, A., and Roncalli, T. (2002). Mixing internal and external data for managing Operational Risk. Group de Recherche Operationnelle, Credit Agricole, France. Available at www.gloriamundi.org.

Frachot, A., Roncalli, T. and Salomon, E. (2004). The Correlation Problem in Operational Risk. Group de Recherche Operationnelle, Credit Agricole, France. Available at www.gloriamundi.org.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering.* Applications of Mathematics, Vol. 53, Springer.

Glasserman, P., Heidelberger, P., and Shahabuddin, P. (1999). Importance sampling and stratification for value-at-risk. In Y.S. Abu-Mostafa, B. LeBaron, A.W. Lo and A.S. Weigand (eds.): *Computational Finance* 1999 (*Proceedings of the 6th International Conference on Computational Finance*), MIT Press, Cambridge, Mass.

Huisman, R., Koedijk, K., Kool, C., and Palm, F. (2001), Tail-index estimates in small samples. *Journal of Business and Economic Statistics* 19, 208-216.

Insurance Companies (2001). Insurance of Operational Risk under the new Basel Capital Accord. Working paper submitted by Insurance Companies, November 2001.

Industry Technical Working Group on Operational Risk (2003). An LDA-Based Advanced Measurement Approach for the Measurement of Operational Risk - Ideas, Issues and Emerging Practices. Draft Document for Discussion Purposes, May 2003.

ITWG Scenario Analysis Working Group (2003). Stylized AMA. Draft Document for Discussion Purposes, Feb 2003.

Johnson, N.L., Kotz, S. and Kemp, A.W. (1993). *Univariate Discrete Distributions (Second Edition).* Wiley, New York.

Kaiser, T., and Kriele, M. (2006). The big unknown. *OpRisk and Compliance* (April 2006), 33-35.

Kalkbrener, M. (2005). An axiomatic approach to capital allocation. *Mathematical Finance* 15/3, 425-437.

Kalkbrener, M, Lotter, H. and Overbeck, L. (2004). Sensible and efficient capital allocation for credit portfolios. *RISK* (Jan. 2004), 19-24.

Klugman, S., Panjer, H., and Wilmot, G. (2004). *Loss models - from data to decisions.* Wiley, Hoboken, New Jersey.

Kühn, R., and Neu, P. (2002). Functional Correlation Approach to Operational Risk in Banking Organizations. Dresdner Bank, Frankfurt. Available at www.gloriamundi.org.

Lindskog F. (2000). Linear Correlation Estimation. ETH Zurich. Available at www.gloriamundi.org.

Lindskog, F., and McNeil, A.J. (2001). Common Poisson Shock Models: Applications to Insurance and Credit Risk Modelling. ETH Zurich. Available at www.gloriamundi.org.

Makarov, M. (2006). Extreme value theory and high quantile convergence. *Journal of Operational Risk* 1(2), 51-57.

McNeil, A.J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management*. Princeton University Press, Oxfordshire, UK.

Medova, E. (2000). Measuring risk by extreme values. *RISK* (Nov. 2000), 20-26.

Mignola G., and Ugoccioni, R. (2005). Tests of extreme value theory applied to Operational Risk data. *OpRisk and Compliance* (Oct 2005).

Mignola G. and Ugoccioni, R. (2006). Sources of uncertainties in modelling operational risk losses. *Journal of Operational Risk* 1(2), 33-50.

Moscadelli, M. (2004). The modelling of operational risk: experience with the analysis of the data collected by the Basel Committee. Working paper, Banca d'Italia. Available at www.gloriamundi.org.

Nelson, R.B (1999). *An Introduction to Copulas*. Lecture Notes in Statistics, Springer, New York.

Nešlehová, J., Embrechts P., and Chaves-Demoulin, V. (2006). Infinite-mean models and the LDA for operational risk. *Journal of Operational Risk* 1(1), 3-25.

Nyström K., and Skoglund, J. (2002). Quantitative Operational Risk Management. Swedbank. Available at www.gloriamundi.org.

Peng L., and Welsh, A. (2001). Robust Estimation of the Generalized Pareto Distribution, *Extremes* 4, 53-65.

Powojowski, M.R., Reynolds, D., and Tuenter, H.J.H (2002). Dependent Events and Operational Risk. *Algo Research Quarterly* 5/2, 65-73.

Rockafellar, R.T., and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk* 2, 21-41.

Samad-Khan, A. (2002). How to categorize operational losses. OpRisk Analytics, Stamford, CT.

Scandizzo, S. (2006). Scenario Analysis in Operational Risk Management. Pages 187-214 in Davis (2006): *The Advanced Measurement Approach to Operational Risk*, Risk Books, London.

Scenario-Based AMA Working Group (2003). Scenario-based AMA.

Available at www.newyorkfed.org/newsevents/events/banking/2003/con0529d.pdf.

Shih, J., Samad-Khan, A., and Medapa, P. (2000). Is the size of an operational loss related to firm size. *Operational Risk* (Jan 2000).

Song, P. (2000). Multivariate dispersion models generated from Gaussian copula. *Scandinavian Journal of Statistics*, 27/2, 305-320.