

## **Operational Risk: A Practitioner's View**

By Silvan Ebnöther<sup>a</sup>, Paolo Vanini<sup>b</sup>  
Alexander McNeil<sup>c</sup>, and Pierre Antolinez<sup>d</sup>

<sup>a</sup> Corporate Risk Control, Zürcher Kantonalbank,  
Neue Hard 9, CH-8005 Zurich,  
e-mail: silvan.ebnoether@zkb.ch

<sup>b</sup> Corresponding author,  
Corporate Risk Control, Zürcher Kantonalbank,  
Neue Hard 9, CH-8005 Zurich,  
Institute of Finance, University of Southern Switzerland, CH-6900 Lugano,  
e-mail: paolo.vanini@zkb.ch

<sup>c</sup> Department of Mathematics,  
ETH Zurich, CH-8092 Zurich,  
e-mail: alexander.mcneil@math.ehtz.ch

<sup>d</sup> Corporate Risk Control, Zürcher Kantonalbank,  
Neue Hard 9, CH-8005 Zurich,  
e-mail: pierre.antolinez@zkb.ch

First version: June 2001, this version: October 11, 2002

### Abstract

The Basel Committee on Banking Supervision ("the Committee") released a consultative document that included a regulatory capital charge for operational risk. Since the release of the document, the complexity of the concept of "operational risk" has led to vigorous and recurring discussions. We show that for a production unit of a bank with well-defined workflows operational risk can be unambiguously defined and modelled. The results of this modelling exercise are relevant for the implementation of a risk management framework, and the pertinent risk factors can be identified. We emphasize that only a small share of all workflows make a significant contribution to the resulting VaR. This result is quite robust under stress testing. Since the definition and maintenance of processes is very costly, this last result is of major practical importance. Finally, the approach allows us to distinguish features of quality and risk management respectively.

**Keywords:** Operational Risk, Risk Management, Extreme Value Theory, VaR

**JEL Classification:** C19, C69, G18, G21

*Acknowledgement: We are especially grateful to Professor Embrechts (ETH) for his profound insights and numerous valuable suggestions. We are also grateful to A. Allemann, U. Amberg, R. Hottinger and P. Meier from Zürcher Kantonalbank for providing us with the data and relevant practical insight. We finally thank the participants of the seminars at the Universities of Pavia, Lugano and Zurich and the IBM 2002 Finance Forum, Zurich. Paolo Vanini would like to thank the Swiss National Science Foundation (NCCR FINRISK) for their financial support.*

## 1 Introduction

In June 1999, the Basel Committee on Banking Supervision ("the Committee") released its consultative document "The New Basel Capital Accord" ("The Accord") that included a proposed regulatory capital charge to cover "other risks". Operational risk (OR) is one such "other risk". From the time of the release of this document and its sequels (BIS (2001)), the industry and the regulatory authorities have been engaged in vigorous and recurring discussions. It is fair to say that at the moment, as far as operational risk is concerned the "Philosopher's Stone" is yet to be found.

Some of the discussions are on a rather general and abstract level. For example, there is still ongoing debate concerning a general definition of OR. The one adopted by the BIS Risk Management Group (2001) is "the risk of direct loss resulting from inadequate or failed internal processes, people and systems or from external events." How to translate the above definition into a capital charge for OR has not yet been fully resolved; see for instance Danielsson et al. (2001). For the moment, legal risk is included in the definition, whereas systemic, strategic and reputational risks are not.

The present paper contributes to these debates from a practitioner's point of view. To achieve this, we consider a number of issues of operational risk from a case study perspective. The case study is defined for a bank's production unit and factors in self-assessment as well as historical data. We try to answer the following questions quantitatively:

1. Can we define and model OR for the workflow processes of a bank's production unit (production processes)? A production process is roughly a sequence of business activities; a definition is given in the beginning of Section 2.
2. Is a portfolio view feasible and with what assets?
3. Which possible assessment errors matter?
4. Can we model OR such that both the risk exposure and the causes are identified? In other words, not only risk measurement but risk management is the ultimate goal.
5. Which are the crucial risk factors?
6. How important is comprehensiveness? Do all workflows in our data sample significantly contribute to the operational risk of the business unit?

The results show that we can give reasonable answers to all the questions raised above. More specifically, if operational risk is modelled on well-defined objects, all vagueness is dispelled although compared with market or credit risk, a different methodology and different statistical techniques are used. An important insight from a practitioner's point of view is that not all processes in an organization need to be equally considered for the purpose of accurately defining operational risk exposure. The management of operational risks can focus on key issues; a selection of the relevant processes significantly reduces the costs of defining and designing the workflow items. To achieve this goal, we construct the Risk Selection Curve (RiSC), which singles out the relevant workflows needed to estimate the risk figures. In a next step, the importance of the four risk factors considered is analyzed. As a first result, the importance of the risk factors depends non-linearly on the confidence level used in measuring risk. While for quality management all factors matter, fraud and system failure have a non-reliable impact on risk figures. Finally, with the proposed methodology we are able to link risk measurement to the needs of risk management: For each risk tolerance level of the management there exists an appropriate risk measure. Using this measure RiSC and the risk factor contribution analysis select the relevant workflows and risk factors.

The paper is organized as follows. In Section 2 we describe the case study. In Section 3 the results using the data available are discussed and compared for the two models. Further, some important issues raised by the case study are discussed. Section 4 concludes.

## 2 Case Study

The case study was carried out for Zürcher Kantonalbank's Production Unit. The study comprises 103 production processes.

### 2.1 Modelling Operational Risk: Framework

The most important and difficult task in the quantification of operational risk is to find a reasonable model for the business activities<sup>1</sup>. We found it useful, for both practical and theoretical reasons, to think of quantifiable operational risk in terms of directed graphs. Though this approach is not strictly essential in the present paper, for operational risk management full-fledged graph theory is crucial (see Ebnöther et al. (2002) for a theoretical approach). In this paper, the overall risk exposure is considered on an aggregated graph level solely for each process. This approach of considering first an aggregated level is essential from a practical feasibility point of view: Considering the costly nature of analyzing the operational risk of processes quantitatively on a "microscopic level", the important processes have to be selected first.

In summary, each workflow is modelled as a graph consisting of a set of nodes and a set of directed edges. Given this skeleton, we next attach risk information. To this end, we use the following facts: At each node (representing, say, a machine or a person) errors in the processing can occur (see Figure 1 for an example).

*Insert Figure 1 around here.*

The errors have both a cause and an effect on the performance of the process. More precisely, at each node there is a (random) input of information defining the performance. The errors then affect this input to produce a random output performance. The causes at a node are the risk factors, examples being fraud, theft or computer system failure. The primary objective is to model the link between effects and causes. There are, of course, numerous ways in which such a link can be defined. As operational risk management is basically loss management, our prime concern is finding out how causes, through the underlying risk factors, impact losses at individual edges.

We refer to the entire probability distribution associated with a graph as the *operations risk distribution*. In our modelling approach, we distinguish between this distribution and *operational risk distribution*. While the operations risk distribution is defined for all losses, the operational risk distribution considers only losses larger than a given *threshold*.

Operational risk modelling, as defined by the Accord, corresponds to the operations risk distribution in our setup. In practice, this identification is of little value as every bank distinguishes between small and large losses. While small losses are frequent, large losses are very seldom encountered. This implies that banks know a lot about the small losses and their causes but they have no experience with large losses. Hence, typically an efficient organization exists for small losses. The value added of quantitative operational risk management for banks thus lies in the domain of large losses (low intensity, high severity). This is the reason why we differentiate between operations risk and operational risk *if* quantitative modelling is considered. We summarize our definition of operational risk as follows:

**Definition 1** *Quantitative operational risk for a set of production processes are those operations risks which exceed a given threshold value.*

Whether or not we can use graph theory to calculate operational risk critically depends on the existence of *standardized and stable* workflows within the banking firm. The cost of defining

---

<sup>1</sup> Strictly speaking there are three different objects: Business activities, workflows, which are a first model of these activities, and graphs, which are a second model of business activities based on the workflows. Loosely speaking, graphs are mathematical models of workflows with attributed performance and risk information relevant to the business activities. In the sequel we use business activities and workflows as synonyms.

processes within a bank can be prohibitively large (i) if all processes need to be defined, (ii) if they are defined on a very deep level of aggregation, or (iii) if they are not stable over time.

## 2.2 Data

An important issue in operational risk is data availability. In our study we use both self-assessment and historical data. The former are based on *expert knowledge*. More precisely, the respective process owner valued the risk of each production process. To achieve this goal, standardized forms were used where all entries in the questionnaire were properly defined. The experts had to assess two random events:

1. The frequency of the random time of loss. For example, the occurrence probability of an event for a risk factor could be valued "high/medium/low" by the expert. By definition the "medium" class might, for example, comprise one-yearly events up to four-yearly events.
2. The experts had to estimate maximum and minimum possible losses in their respective processes. The assumed severity distribution derived from the self-assessment is calibrated using the loss history<sup>2</sup>. This procedure is explained in chapter 2.4.

If we use expert data, we usually possess sufficient data to fully specify the risk information. The disadvantage of such data concerns their quality. As Rabin (1998) lucidly demonstrates in his review article, people typically fail to apply the mathematical laws of probability correctly but instead create their own "laws" such as the "law of small numbers". An expert based database thus needs to be designed such that the most important and prominent biases are circumvented and a sensitivity analysis has to be done. We therefore represented probabilistic judgments in the case study unambiguously as a choice among real life situations.

We found three principles especially helpful in our data collection exercise:

1. *Principle I:* Avoid direct probabilistic judgments.
2. *Principle II:* Choose an optimal interplay between experts' know how and modelling. Hence the scope of the self-assessment has to be well defined. Consider for example the severity assessment: A possible malfunction in a process leads to losses in the process under consideration. The same malfunction can also affect other workflows within the bank. Experts have to be aware to whether they adopt a local point of view in their assessment or a more global one. In view of the pitfalls inherent in probabilistic judgments, experts should be given as narrow a scope as possible. They should focus on the simplest estimates, and model builders should perform more complicated relationships based on these estimates.
3. *Principle III:* Implement the right incentives. In order to produce the best result it is important not only to advise the experts on what information they have to deliver, but also to make it clear why it is beneficial for them and the whole institution to do so. A second incentive problem concerns accurate representation. Specifically, pooling behavior should be avoided. By and large, the process experts can be classified in three categories at the beginning of a self-assessment: Those who are satisfied with the functioning of their processes, those who are not satisfied with the status but have so far been unable to improve their performance and, finally, experts who well know that their processes should be redesigned but have no intention of doing so. For the first type, making an accurate representation would not appear to be a problem. The second group might well exaggerate the present status to be worse than it in fact is. The third group has an incentive to mimic the first type. Several measures are possible to avoid such pooling behavior, i.e. having other employees crosscheck the assessment values, and comparing with loss data where available. And ultimately, common sense on the part of the experts' superiors can reduce the extent of misspecified data due to pooling behavior.

---

<sup>2</sup> The loss history was not used in Ebnöther et al. (2002) because the required details were not available. The soundness of the results has been enhanced by the availability of this extended data.

The historical data are used for calibration of the severity distribution (see Section 2.4). At this stage, we restrict ourselves to noting that information regarding the severity of losses is confined to the minimum/maximum loss value derived from the self-assessment.

### 2.3 The Model

Within the above framework, the following steps summarize our quantitative approach to operational risk:

1. First, data are generated through simulation starting from expert knowledge.
2. To attain more stable results, the distribution for large losses is modelled using extreme value theory.
3. Key risk figures are calculated for the chosen risk measures. We calculate the VaR and the conditional VaR (CVaR)<sup>3</sup>.
4. A sensitivity analysis is performed.

Consider a business unit of a bank with a number of production processes. We assume that for workflow  $i$  there are 4 relevant risk factors  $R_{i,j}$ ,  $j = 1, \dots, 4$ , leading to a possible process malfunction such as system failure, theft, fraud, or error. Because we do not have any experience with the two additional risk factors external catastrophes and temporary loss of staff, we have not considered them in our model. In the present model we assume that all risk factors are *independent*.

To generate the data, we have to simulate two risk processes: The stochastic time of a loss event occurrence and the stochastic loss amount (the severity) of an event expressed in a given currency. The number  $N_{i,j}$  of workflow  $i$  malfunctions by risk factor  $j$  and the associated severity  $W_{i,j}(n)$ ,  $n = 1, \dots, N_{i,j}$ , are derived from expert knowledge.  $N_{i,j}$  is assumed to be a homogeneous Poisson process. Formally, the inter-arrival times between successive losses are i.i.d, exponentially distributed with finite mean  $1/\lambda_{i,j}$ . The parameters  $\lambda_{i,j}$  are calibrated to the expert knowledge database.

The severity distributions  $W_{i,j}(n) \sim F_{i,j}$ , for  $n=1, \dots, N_{i,j}$  are estimated in a second step. The distribution of severity  $W_{i,j}(n)$  is modeled in two different ways. First, we assume that the severity is a combined Beta and generalized Pareto distribution. In the second model, a lognormal distribution is used to replicate the severity.

If the  $(i,j)$ -th loss arrival process  $N_{i,j}(t)$ ,  $t \geq 0$ , is independent from the loss severity process  $\{W_{i,j}(n)\}_{n \in \mathbb{N}}$  and  $W_{i,j}(n)$  has the same distribution for each  $n$  and are independent, then the total loss experienced by process  $i$  due to risk type  $j$  up to time  $t$

$$S_{i,j}(t) = \sum_{n=1}^{N_{i,j}(t)} W_{i,j}(n)$$

is called a compound Poisson process. We always simulate 1 year. For example, 10,000 simulations of  $S(1)$  means that we simulate the total first years loss 10,000 times.

The next step is to specify the tail of the loss distribution as we are typically interested in heavy losses in operational risk management. We use extreme value theory to smooth the total loss distribution. This theory allows a categorization of the total loss distribution into different qualitative tail regions<sup>4</sup>.

In summary, Model 1 is specified by:

<sup>3</sup> VaR denotes the Value-at-Risk measure and CVaR denotes Conditional Value-at-Risk (CVaR is also called Expected Shortfall or Tail Value-at-Risk (See Tasche (2002)).

<sup>4</sup> We consider the mean excess function  $e_1(u) = E[S(1)-u | S(1) \geq u]$  for 1 year, which by our definition of operational risk is a useful measure of risk. The asymptotic behavior of the mean excess function can be captured by the generalized Pareto distribution (GPD)  $G$ . The GPD is a two-parameter distribution with distribution function

$$G_{\xi, \sigma}(x) = \begin{cases} 1 - \left(1 + \frac{\xi}{\sigma}x\right)^{-\frac{1}{\xi}} & \text{if } \xi \neq 0, \\ 1 - \exp\left(-\frac{x}{\sigma}\right) & \text{if } \xi = 0, \end{cases}$$

where  $\sigma > 0$  and the support is  $[0, \infty)$  when  $\xi \geq 0$  and  $[0, -\sigma/\xi]$  for  $\xi < 0$ . A good data fit is achieved which leads to stable results in the calculation of the conditional Value-at-Risk (see Section 3).

- Production processes which are represented as aggregated, directed graphs consisting of two nodes and a single edge,
- Four independent risk factors,
- A stochastic arrival time of loss events modelled by a homogeneous Poisson process and the severity of losses modeled by a Beta-GPD-mixture distribution. Assuming independence, this yields a compound Poisson model for the aggregated losses.
- It turns out that the generalized Pareto distribution, which is fitted by the POT<sup>5</sup> method, yields an excellent fit to the tail of the aggregate loss distribution.
- The distribution parameters are determined using maximum likelihood estimation techniques.

The generalized Pareto distribution is typically used in extreme value theory. It provides an excellent fit to the simulated data for large losses. Since the focus is not on choosing the most efficient statistical method, we content ourselves with the above choice while being very aware that other statistical procedures might work equally well.

## 2.4 Calibration

Our historical database<sup>6</sup> contains losses that can be allocated to the workflows in the production unit. We use this data to calibrate the severity distribution, noting that the historical data show an expected bias: Due to the relevance of operational risk in the last years, more small losses are seen in 2000 and 2001 than in previous years.

For the calibration of the severity distribution we use our loss history and the assessment of the maximum possible loss per risk factor and workflow. The data are processed in two respects. First, as the assessment of the minimum is not needed since it is used for accounting purposes only, we drop this number. Second, errors may well lead to losses instead of gains. In our database a small number of such gains occur. Since we are interested solely in losses, we do not consider events leading to gains.

Next we observe that the maximum severity assessed by the experts is exceeded in some processes. In our loss history, this effect occurs with an empirical conditional probability of 0.87% per event. In our two models, we factor this effect into the severity value by accepting losses higher than the maximum assessed losses.

Calibration is then performed as follows:

- We first allocate each loss to a risk factor and to a workflow.
- Then we normalize the allocated loss by the maximum assessed loss for its risk factor and workflow.
- Finally we fit our distribution to the generated set of normalized losses. It follows that the lognormal distribution and a mixture of the Beta and generalized Pareto distribution provide the best fits to the empirical data.

In the second simulation, we have to multiply the simulated normalized severity by the maximum assessed loss to generate the loss amount (reversion of the second calibration step).

### 2.4.1 Lognormal Model

In our first model of the severity distribution, we fit a lognormal distribution to the standardized losses.

The lognormal distribution seems to be a good fit for the systematic losses. However, we observe that the probability of occurrence for large losses is greater than the empirical data show.

---

<sup>5</sup> The Peaks-Over-Threshold (POT) method based on a GPD model allows construction of a tail fit above a certain threshold  $u$ ; for details of the method, see the papers in Embrechts (2000).

<sup>6</sup> The data range from 1997 to 2002 and contain 285 appropriate entries.

### 2.4.2 Beta-GPD-Mixture Model

We eliminate the drawbacks of the lognormal distribution by searching for a mixture of distributions which satisfies the following properties:

First, the distribution has to reliably approximate the normalized empirical distribution in the domain where the mass of the distribution is concentrated. The flexibility of the Beta distribution is used for fitting in the interval between 0 and the estimated maximum  $X_{\max}$ .

Second, large losses, which probably exceed the maximum of the self-assessment, are captured by the GPD with support the positive real numbers. The GPD distribution is estimated using all historical normalized losses higher than the 90% quantile. In our example, the relevant shape parameter  $\xi$  of the GPD fit is nearly zero, i.e. the distribution is medium tailed<sup>7</sup>. To generate the losses, we choose the exponential distribution which corresponds to a GPD with  $\xi=0$ .

Our Beta-GPD-mixture distribution is defined by a combination of the Beta- and the Exponential distribution. A Beta-GPD-distributed random variable  $X$  satisfies the following rules: With probability  $\pi$ ,  $X$  is a Beta random variable, and with probability  $(1-\pi)$ ,  $X$  is a GPD-distributed random variable. Since 0.87% of all historical data exceed the assessed maximum, the weight  $\pi$  is chosen such that  $P(X > X_{\max}) = 0.87\%$  holds.

The calibration procedure reveals an important issue if self-assessment and historical data are considered: Self-assessment data typically need to be processed if they are compared with historical data. This shows that the reliability of the self-assessment data is limited and that by processing this data, consistency between the two different data sets is restored.

## 3 Results

The data set for the application of the above approaches is based on 103 production processes at Zürcher Kantonalbank and self-assessment of the probability and severity of losses for four risk factors (see Section 2.1). The model is calibrated against an internal loss database. Since confidentiality prevents us from presenting real values, the absolute values of *all* results are fictitious but the relative magnitudes are real. The calculations are based on 10,000 simulations. Table 1 shows the results for the Beta-GPD-mixture model.

|                        | $\alpha = 95\%$ |      | $\alpha = 99\%$ |      | $\alpha = 99.9\%$ |      |
|------------------------|-----------------|------|-----------------|------|-------------------|------|
|                        | VaR             | CVaR | VaR             | CVaR | VaR               | CVaR |
| Empirical              | 17              | 41   | 60              | 92   | 134               | 161  |
| $u = 95\%$ -Quantile   | 17              | 55   | 52              | 129  | 167               | 253  |
| $u = 97.5\%$ -Quantile | -               | -    | 59              | 10   | 133               | 165  |
| $u = 99\%$ -Quantile   | -               | -    | 60              | 91   | 132               | 163  |
| $u = 99.5\%$ -Quantile | -               | -    | -               | -    | 134               | 161  |

**Table 1** Simulated data behavior of the tail distribution. "Empirical" denotes the results derived from 10,000 simulations for the Beta-mixture model. The other key figures are generated using the POT<sup>8</sup> model for the respective thresholds  $u$ .

Using the lognormal model to generate the severities the  $VaR_\alpha$  for  $\alpha = 95\%$  and  $99\%$  respectively are approximately the same. The lognormal distribution is more tailed than the Beta-GPD-mixture distribution that leads to higher key figures for the 99.9% quantile.

<sup>7</sup> The lognormal model bellows to the medium tailed distributions, too. But we observe that the tail behavior of the lognormal distribution converts very slowly to  $\xi = 0$ . For this reason, we anticipate that the resultant distribution on the yearly total loss will seem to be heavily tailed. Only a large-scale simulation could observe this fact.

<sup>8</sup> From Table 1 and 2 it follows that the POT model yields a reasonable tail fit. For further information on the underlying loss tail behavior and statistical uncertainty of the estimated parameters we refer to Ebnöther (2001).

|                        | $\alpha = 95\%$ |      | $\alpha = 99\%$ |      | $\alpha = 99.9\%$ |      |
|------------------------|-----------------|------|-----------------|------|-------------------|------|
|                        | VaR             | CVaR | VaR             | CVaR | VaR               | CVaR |
| Empirical              | 14              | 48   | 55              | 137  | 253               | 512  |
| $u = 95\%$ -Quantile   | 14              | 68   | 53              | 165  | 234               | 633  |
| $u = 97.5\%$ -Quantile | -               | -    | 53              | 195  | 236               | 706  |
| $u = 99\%$ -Quantile   | -               | -    | 55              | 277  | 232               | 911  |
| $u = 99.5\%$ -Quantile | -               | -    | -               | -    | 252               | 534  |

**Table 2** Simulated data behavior of the tail distribution. Instead of the Beta-mixture model of Table 1 the lognormal model is used for the severity.

We can observe that a robust approximation of the coherent risk measure CVaR is more sensitive to the underlying loss distribution. The Tables also confirm that the lognormal model is more heavily tailed than the Beta-mixture model.

### 3.1 Risk Selection Curve

A relevant question for practitioners is how much each of the processes contributes to the risk exposure. If it turns out that only a fraction of all processes significantly contribute to the risk exposure, then risk management needs only to be defined for these processes.

We therefore analyze how much each single process contributes to the total risk. We consider only VaR in the sequel as a risk measure. To split up the risk into its process components, we compare the risk contributions (RC) of the processes.

Let  $RC_\alpha(i)$  be the risk contribution of process  $i$  to VaR at the confidence level  $\alpha$

$$RC_\alpha(i) = VaR_\alpha(P) - VaR_\alpha(P \setminus \{i\}),$$

where  $P$  is the whole set of workflows.

Because the sum over all  $RC_\alpha$ 's is generally not equal to the VaR, the relative risk contribution ( $RRC_\alpha(i)$ ) of process  $i$  is defined as the  $RC_\alpha(i)$  normalized by the sum over all  $RC_\alpha$ , i.e.

$$RRC_\alpha(i) = \frac{RC_\alpha(i)}{\sum_j RC_\alpha(j)} = \frac{VaR_\alpha(P) - VaR_\alpha(P \setminus \{i\})}{\sum_j RC_\alpha(j)}.$$

As a further step, for each  $\alpha$ , we count the number of processes that exceed a relative risk contribution of 1%. We call the resulted curve with parameter  $\alpha$ , the Risk Selection Curve (RiSC).

*Insert Figure 2 around here.*

Figure 2 shows that on a reasonable confidence level only about 10 percent of all processes contribute to the risk exposure. Therefore only for this small number of processes is it worth developing a full graph theoretical model and analyzing this process in more detail. On lower or even low confidence levels, more processes contribute to the VaR. This indicates that there are a large number of processes of the high frequency/low impact type. These latter processes can be singled out for quality management, whereas processes of the low frequency/high impact type are under the responsibility of risk management. In summary, using RiSC graphs allows a bank to discriminate between quality and risk management in respect of the processes which matter. This reduces costs for both types of management significantly and indeed renders OR management feasible.

We finally note that the shape of the RiSC, i.e. not monotone decreasing, is not a product of modelling.

From a risk management point of view RiSC links the measurement of operational risk to its management as follows: Each parameter value  $\alpha$  represents a risk measure and therefore, in Figure 2

on the horizontal axes a family of risk measures is shown. The risk managers possess a risk tolerance that can be expressed with a specific value  $\alpha$ . Hence, RiSC provides the risk information managers are concerned with.

### 3.2 Risk Factor Contribution

The information concerning the most risky processes is important for splitting the Value at Risk into its risk factors. Therefore we determine the relative risk that a risk factor contributes to the  $\text{VaR}_\alpha$  in a similar manner to the former analysis. We define the relative risk factor contribution as

$$\text{RRFC}_\alpha(i) = \frac{\text{VaR}_\alpha - \text{VaR}_\alpha(P \setminus \{i\})}{\sum_{j=1}^4 (\text{VaR}_\alpha - \text{VaR}_\alpha(P \setminus \{j\}))},$$

with P now the whole set of risk factors.

The resultant graph clearly shows the importance of the risk factors.

*Insert Figure 3 around here.*

Figure 3 shows that the importance of the risk factors is not uniform and in linear proportion to the scale of confidence levels. For low levels, error is the most dominant factor, which again indicates that this domain is best covered by quality management. The higher the confidence level is, the more fraud becomes the dominant factor. The factor theft displays an interesting behavior too: It is the sole factor showing a virtually constant contribution in percentage terms at all confidence levels.

Finally, we note that both results, RiSC and the risk factor contribution, were not known to the experts in the business unit. These clear and neat results contrast with the diffuse and disperse knowledge within the unit about the risk inherent in their business.

### 3.3 Modelling Dependence

In the previous model we assumed the risk factors were independent. Dependence could be introduced though a so-called common shock model (see Bedford and Cooke (2001), Chapter 8, and Lindskog and McNeil (2001)). A natural approach to model dependence is to assume that all losses can be related to a series of underlying and independent shock processes. When a shock occurs, this may cause losses due to several risk factors triggered by that shock.

We did not implement dependence in our case study for the following reasons:

- The occurrence of losses which are caused by fraud, error and theft are independent.
- While we are aware of system failures dependencies, these are not the dominating risk factor. (See figure 3.) Hence, the costs for an assessment and calibration procedure are too large compared to the benefit of such an exercise.

### 3.4 Sensitivity Analysis

We assume that for each workflow and each risk factor the estimated maximum loss is twice the self-assessed value, and then twice that value again. In doing so, we also take into account that the calibration to the newly generated data has to be redone.

|                   | $\alpha = 95\%$ |      | $\alpha = 99\%$ |      | $\alpha = 99.9\%$ |      |
|-------------------|-----------------|------|-----------------|------|-------------------|------|
|                   | VaR             | CVaR | VaR             | CVaR | VaR               | CVaR |
| Empirical         | 17              | 41   | 60              | 92   | 134               | 161  |
| Stress scenario 1 | 22              | 45   | 57              | 92   | 129               | 178  |
| Stress scenario 2 | 21              | 48   | 65              | 103  | 149               | 186  |

**Table 3** Stress scenario 1 is a simulation using a maximum twice the self-assessed value. Stress scenario 2 is a simulation using a maximum four times the self-assessed value. A Beta-GPD-mixture distribution is chosen as severity model.

It follows that an overall underestimation of the estimated maximum loss does not have a significant effect on the risk figures since the simulation input is calibrated to the loss history.

Furthermore, the relative risk contributions of the risk factors and processes do not change significantly under these scenarios, i.e. the number of processes which significantly contribute to the VaR remains almost invariant and small compared to all processes.<sup>9</sup>

## 4 Conclusion

The scope of this paper was to show that quantification of operational risk (OR), adapted to the needs of business units, is feasible if data exist and if the modelling problem is seriously considered. This means that the solution of the problem is described with the appropriate tools and not by an ad hoc deployment of methods successfully developed for other risks.

It follows from the results presented that a quantification of OR and OR management must be based on well-defined objects (processes in our case). We do not see any possibility of quantifying OR if such a structure is not in place within a bank. It also follows that not all objects (processes for example) need to be defined; if the most important are selected, the costs of monitoring the objects can be kept at a reasonable level and the results will be sufficiently precise. The self-assessment and historical data used in the present paper proved to be useful: applying a sensitivity analysis, the results appear to be robust. In the derivation of risk figures we assumed that risk tolerance may be non-uniform in the management. Therefore, risk information is parameterized such that the appropriate level of confidence can be chosen.

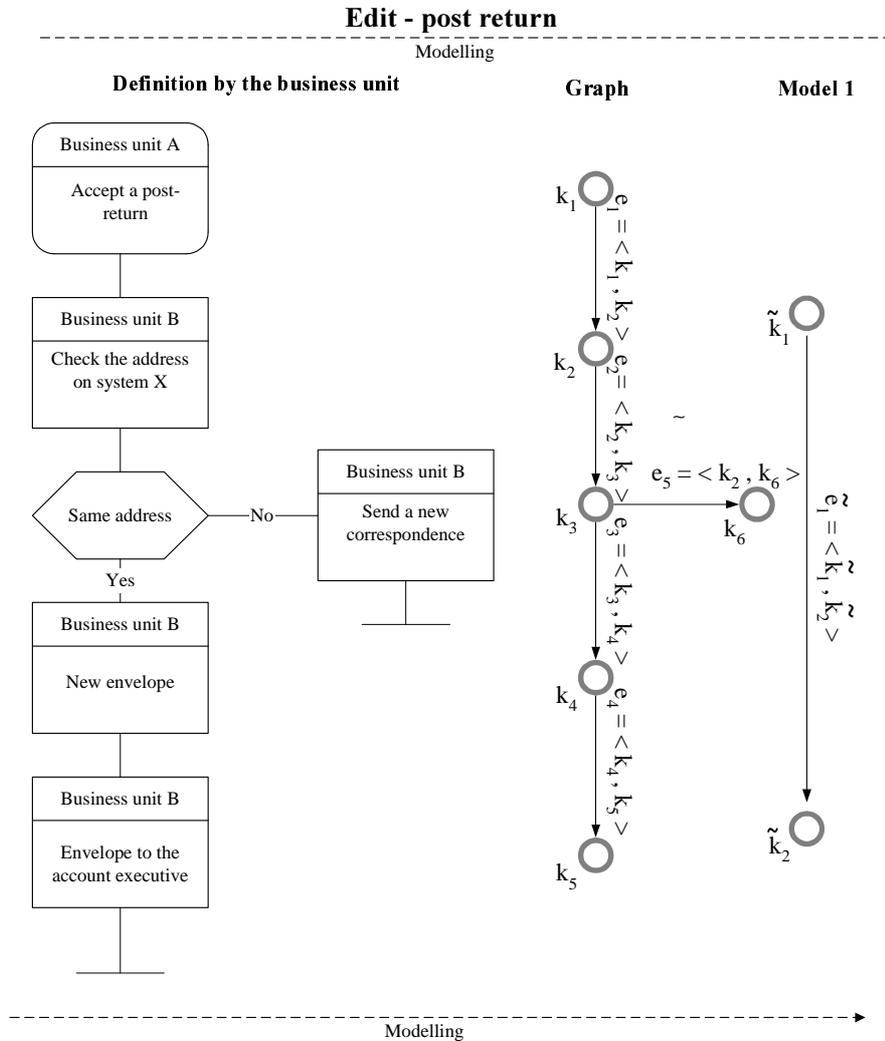
The models considered in this paper can be extended in various directions. First, if the Poisson models used are not appropriate, they can be replaced by a negative Binomial process (see Ebnöther (2001) for details). Second, production processes are only part of the total workflow processes defining business activities. Hence, other processes need to be modelled and using graph theory a comprehensive risk exposure for a large class of banking activities is derived.

---

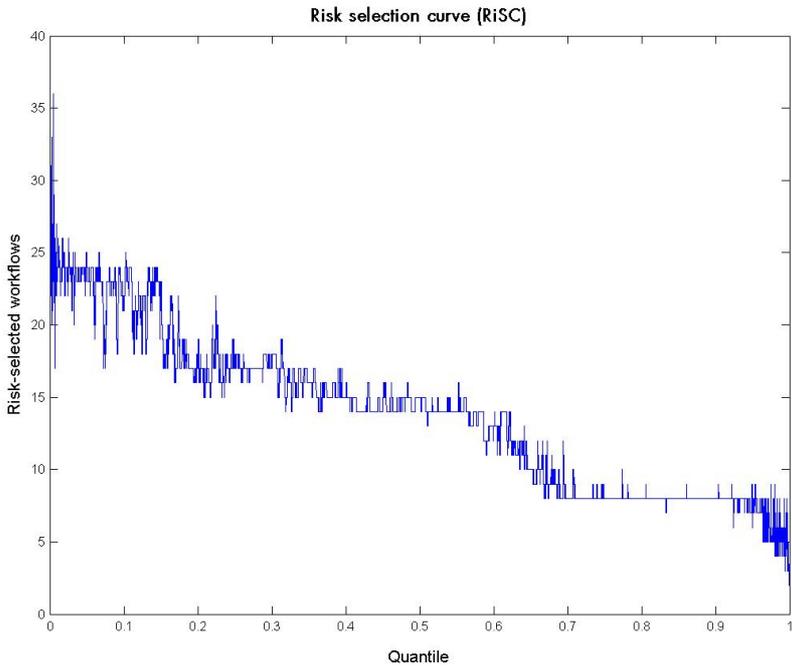
<sup>9</sup> At the 90% quantile, for both stress scenarios the number of “relevant “ workflows (8) remains constant whereas a small reduction from 15 to 14 (13) relevant workflows is observed at the median.

## 5 References

- BIS 2001, Basel Committee on Banking Supervision (2001), Consultative Document, The New Basel Capital Accord, <http://www.bis.org>.
- BIS, Risk Management Group of the Basel Committee on Banking Supervision (2001), Working Paper on the Regulatory Treatment of Operational Risk, <http://www.bis.org>.
- Bedford, T. and R Cooke (2001), Probabilistic Risk Analysis, Cambridge University Press, Cambridge.
- Danielsson J., P. Embrechts, C. Goodhart, C. Keating, F. Muenich, O. Renault and H. S. Shin (2001), An Academic Response to Basel II, Special Paper Series, No 130, London School of Economics Financial Markets Group and ESRC Research Center, May 2001
- Ebnöther, S. (2001), Quantitative Aspects of Operational Risk, Diploma Thesis, ETH Zurich.
- Ebnöther, S., M. Leippold and P. Vanini (2002), Modelling Operational Risk and Its Application to Bank's Business Activities, Preprint.
- Embrechts, P. (Ed.) (2000), Extremes and Integrated Risk Management, Risk Books, Risk Waters Group, London
- Embrechts, P., C. Klüppelberg and T. Mikosch (1997), Modelling Extremal Events for Insurance and Finance, Springer, Berlin.
- Lindskog, F. and A.J. McNeil (2001), Common Poisson Shock Models, Applications to Insurance and Credit Risk Modelling, Preprint, ETH Zürich.
- Medova, E. (2000), Measuring Risk by Extreme Values, Operational Risk Special Report, Risk, November 2000.
- Rabin, M. (1998), Psychology and Economics, Journal of Economic Literature, Vol. XXXVI, 11-46, March 1998.
- Tasche, D. (2002), Expected Shortfall and Beyond, Journal of Banking and Finance 26(7), 1523-1537.



**Figure 1** Example of a simple production process: The edit-post return. More complicated processes can contain several dozens of decision and control nodes. The graphs can also contain loops and vertices with several legs, i.e. topologically the edit - post process is of a particularly simple form. In the present paper, condensed graphs (Model 1) are only considered, while for risk management purposes the full graphs are needed.



**Figure 2** The risk selection curve (RiSC) of the Beta-GPD-mixture model.

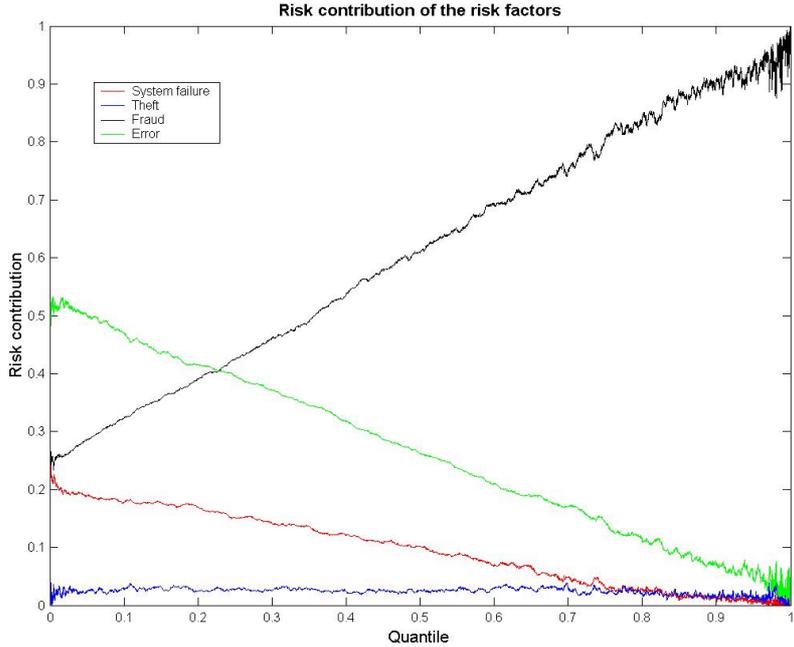


Figure 3 The segmentation of the VaR into its risk factors.