Measuring Operational Risk in Financial Institutions

Séverine Plunus[a] · Georges Hübner[b] · Jean-Philippe Peters[c]

Abstract The scarcity of internal loss databases tends to hinder the use of the advanced approaches for operational risk measurement (AMA) in financial institutions. As there is a greater variety in credit risk modelling, this paper explores the applicability of a modified version of CreditRisk+ to operational loss data. Our adapted model, OpRisk+, works out very satisfying Values-at-Risk at 95% level as compared with estimates drawn from sophisticated AMA models. OpRisk+ proves to be especially worthy in the case of small samples, where more complex methods cannot be applied. OpRisk+ could therefore be used to fit the body of the distribution of operational losses up to the 95%-percentile, while Extreme Value Theory, external databases or scenario analysis should be used beyond this quantile.

Keywords: Operational Risk · Basel II · Modelling · CreditRisk+.

JEL: G28, G21

[a] Corresponding Author. HEC-Management School, University of Liège. Rue Louvrex, 14, B-4000 Liège, Belgium. E-mail: splunus@ulg.ac.be. Tel. (+32)485 944 370. Gambit Financial Solutions, Belgium.

[b] HEC Management School, University of Liège; Faculty of Economics and Business Administration, Maastricht University ; Gambit Financial Solutions, Belgium.

[c] Deloitte Luxembourg, Advisory and Consulting Group (Risk Management Unit), Luxembourg.

1 Introduction

Over the past decades, financial institutions have experienced several large operational loss events leading to big banking failures. Memorable examples include the Barings bank losing 1.4 billion USD from rogue trading in his branch in Singapore leading to the failure of the whole institution[1]; Allied Irish Banks losing 750 MM USD in rogue trading[2], or Prudential Insurance incurring 2 billion USD settlement in class action lawsuit[3], to name a few. These events, as well as developments such as the growth of e-commerce, changes in banks' risks management or the use of more highly automated technology, have led regulators and the banking industry to recognize the importance of operational risk in shaping the risk profiles of financial institutions.

Reflecting this recognition, regulatory frameworks such as the New Capital Accord of the Basel Committee on Banking Supervision ("Basel II") have introduced explicit capital requirements for operational risk. Similar to credit risk, Basel II does not impose a "one-size-fits-all" approach to capital adequacy and proposes three distinct options for the calculation of the capital charge for operational risk: the Basic Indicator Approach, the Standardized Approach and the Advanced Measurement Approaches (AMA). The use of these approaches of increasing risk sensitivity is determined according to the risk management systems of the banks. The first two methods are a function of gross income, while the advanced methods are based on internal loss data, external loss data, scenario analysis, business environment and internal control factors.

In 2001, the Basel Committee was encouraging two specific AMA methods: (i) the Loss Distribution Approach (LDA) and (ii) an Internal Measurement Approach (IMA) developing a linear relationship between unexpected loss and expected loss to extrapolate credit-risk's internal rating based (IRB) approach to operational risk. While the Basel Committee dropped formal mention of the IMA in favour of Value-at-Risk approaches in the final version of the Accord, it is still legitimate to be inspired by modelling approaches for credit risk in order to model the distribution of operational loss data. Indeed, both risk measurement frameworks have similar features, such as their focus on a one-year measurement horizon or their use of an aggregate loss distribution skewed towards zero with a long right-tail.

This paper explores the possibility of adapting one of the current proposed industry credit-risk models to perform much of the functionality of an actuarial LDA model (see Crouhy *et al.* (2000) or Gordy (2000) for a comparative analysis of the main credit risk models). We identify CreditRisk+, the model developed by Credit Suisse, as an actuarial-based model, whose characteristics can be adapted to fit the Loss Distribution Approach (LDA). The LDA is explicitly mentioned in the Basel II Accord as eligible among the Advanced Measurement Approaches (AMA) to estimate risk capital, and has unambiguously emerged as the standard industry practice[4]. After some adjustment, we construct a distribution of operational losses through an adapted CreditRisk+model, that we name "OpRisk+"[5]. As this model calibrates the whole distribution, not only can we retrieve the quantiles of the operational loss distribution, but also an estimate of its expectation, needed for the computation of the economic capital.

Our research is aimed at answering the following questions: (i) How would the adaptation of CreditRisk+ model perform compared to sophisticated models such as the approach developed by Chapelle *et al.* (2008) (henceforth CCHP) or Moscadelli (2004) among others? (ii) Does OpRisk+ provide a reasonable assessment of the body of the distribution of the operational losses? (iii) Are the VaR computed with OpRisk+ more conservative than the lower bound of Alexander, (2003), an extended IMA approach?

We address the questions with an experiment based on generated databases using three different Pareto distributions, proven to be appropriate to model operational loss data by Moscadelli (2004) and de Fontnouvelle and Rosengren (2004). The knowledge of the true distribution of losses is necessary to assess the quality of the different fitting methods. Had a real data set been used instead of controlled numerical simulations as proposed by McNeil and Saladin (1997), we would not be able to benchmark the observed results against the true loss distribution and to evaluate the performance of OpRisk+ for different loss generating processes and sample sizes.

We assess the influence of the number of losses recorded in the database on the quality of the estimation. Indeed, Carrillo-Menéndez and Suárez (2012) have shown the difficulty of selecting the correct model from the data when only small samples are available. We also test our new adapted internal rating based model against Alexander's existing improvement to the basic IMA formula. Alexander's Value-at-Risk on operational loss data (OpVaR) is effectively a quantile value from a normal distribution table which allows identification of the unexpected loss if one knows the mean and variance of the loss severity distribution and the mean of the frequency distribution.

Our main findings are twofold. First, we note that the precision of OpRisk+ is not satisfactory to estimate the very far end of the loss distribution, such as the Value-at-Risk (VaR)[6] at the 99.9% confidence level ($VaR_{99.9}$). Yet, our model works out very satisfying quantile estimates, especially for thin-tailed Pareto-distributions, up to a 95% confidence level for the computation of the VaR. Secondly, the simplicity of our model makes it applicable to "problematic" business lines, that is, with very few occurrences of events, or with limited history of data. Procedures that rely on extreme-value theory, by contrast, are very data-consuming, and yield very poor results when used with small databases. Moreover, as argued by Malevergne et al. (2006), when there is a lack of data, non-parametric methods are useful to assess risk at probability level 95% but fail at high probability level such as 99% or larger.

These findings make the OpRisk+ approach clearly not an effective substitute, but indeed a very useful complement to approaches that specifically target the extreme tail of the loss distribution. In particular, the body of the loss distribution can be safely assessed with our method, while external data or scenario analysis, as specifically mentioned in the Accord, can be used to estimate the tail. This model could also represent for a bank a very good cross-check method for scenario-based approaches and for regulators who would want to challenge banks using internal approaches that might seems too aggressive on small samples. Moreover, being able to simultaneously rely on the body and the tail of the distribution is crucial for the operational risk capital estimation, because one needs the full distribution of losses in order to capture the expected loss that enters the regulatory capital estimate.

Next section describes the adjustment needed in order to apply CreditRisk+ model to operational loss data and presents two alternative methods to calibrate a VaR on operational loss data (OpVaR). We then describes our database, presents our results and compares them to the other approaches' results.

2 Alternative Approaches for the Measurement of Operational Risk

This section presents three alternative ways to calibrate a Value-at-Risk on operational loss data. The first one represents an adaptation of the CreditRisk+ framework, while the second one proposes an adaptation of the Loss Distribution Approach (LDA) in the context of operational losses with the use of Extreme Value Theory (EVT). Finally, we introduce an IMA approach developed by Alexander (2003).

2.1 OpRisk+: Application of CreditRisk+ to Operational Loss Data

CreditRisk+ developed by Credit Suisse First Boston is an actuarial model derived from insurance losses models. It models the default risk of a bond portfolio through the Poisson distribution. Its basic building block is simply the probability of default of a counterparty. In this model, no assumptions are made about the causes of default: an obligor is either in default with a probability P_A, or not in default with a probability 1-P_A. Although operational losses do not depend on a particular counterparty, this characteristic already simplifies the adaptation of our model, as we do not need to make assumptions on the causes of the loss.

CreditRisk+ determines the distribution of default losses in three steps: the determination of the frequency of defaults, approximated by a standard Poisson distribution, the determination of the severity of the losses, and the determination of the distribution of default losses.

The determination of the frequency of events leading to operational losses can be modelled through the Poisson distribution as for the probability of default in CreditRisk+:

$$P(N = n) = \frac{\mu^n e^{-\mu}}{n!} \text{ for } n = 0,1,2, \ldots \quad (1)$$

where μ is the average number of defaults per period, and N is a stochastic variable with mean μ, and standard deviation $\sqrt{\mu}$.

CreditRisk+ computes the parameter μ by adding the probability of default of each obligor, supplied, for instance, by rating agencies. However, operational losses do not depend on a particular obligor. Therefore, instead of being defined as a sum of probabilities of default depending on the characteristics of a counterpart, μ can be interpreted as the average number of loss events of one type occurring in a specific business line during one period.

CreditRisk+ adds the assumption that the mean default rate is itself stochastic in order to take into account the fat right tail of the distribution of defaults. Nevertheless, the Poisson distribution being one of the most popular in operational risk frequency estimation, according to Cruz (2002)[7] and Basel Committee on Banking Supervision (2009), we keep on assuming that the number of operational loss events follows a Poisson distribution with a fixed mean μ.

In order to perform its calculations, CreditRisk+ proposes to express the exposure (here, the losses) in a unit amount of exposure L.[8] The key step is then to round up each exposure size to the nearest whole number, in order to reduce the number of possible values and to distribute them into different bands. Each band is characterized by an average exposure, v_j, and an expected loss, ε_j, equal to the sum of the expected losses of all the obligors belonging to the band. Table 1 shows an example of this procedure.

Insert Table 1

CreditRisk+ posits that

$$\varepsilon_j = \nu_j \mu_j \tag{2}$$

where ε_j is the expected loss in band j, ν_j is the common exposure in band j, and μ_j is the expected number of defaults in band j.

As the operational losses do not depend on a particular transaction, we slightly modify the definition of these variables. The aim is to calculate the expected aggregate loss. We will therefore keep the definition of ε_j unchanged. However, as noted earlier, μ_j is not an aggregate expected number of defaults anymore but simply the (observed[9]) average number of operational loss events of size j occurring in one year. Consequently, in order to satisfy equation (2), ν_j must be defined as the average loss amount per event for band j.

Table 2 illustrates the reprocessing of the data.

Insert Table 2

Each band is viewed as a portfolio of exposures by itself. Because some defaults lead to larger losses than others through the variation in exposure amounts, the loss given default involves a second element of randomness, which is mathematically described through its probability generating function.

Thus, let $G(z)$ be the probability generating function for losses expressed in multiples of the unit L of exposure:

$$G_j(z) = \sum_{n=0}^{\infty} P(\text{loss} = nL) z^n = \sum_{n=0}^{\infty} P(n \text{ defaults}) z^{n\nu_j} \tag{3}$$

As the number of defaults follows a Poisson distribution, this is equal to:

$$G_j(z) = \sum_{n=0}^{\infty} \frac{e^{-\mu_j} \mu_j^n}{n!} z^{n\nu_j} = e^{-\mu_j + \mu_j z^{\nu_j}} \tag{4}$$

As far as operational losses are concerned, we can no more consider a band as a portfolio but simply as a category of loss size. This also simplifies the model, as we do not distinguish exposure and expected loss anymore. For credit losses, one first sorts exposures, and then calculates the expected loss, by multiplying the exposures by their probability of default. As far as operational losses are concerned, the loss amounts are directly sorted by size. Consequently, the second element of randomness is not necessary anymore. This has no consequences on the following results except simplifying the model.

Whereas CreditRisk+ assumes the exposures in the portfolio to be independent, OpRisk+ assumes the independence of the different loss amounts. Thanks to this assumption, the probability generating function for losses of one type for a specific business line is given by the product of the probability generating function for each band:

$$G(z) = \prod_{j=1}^{m} e^{-\mu_j + \mu_j z^{\nu_j}} = e^{-\sum_{j=1}^{m} \mu_j + \sum_{j=1}^{m} \mu_j z^{\nu_j}} \tag{5}$$

Finally, the loss distribution of the entire portfolio is given by:

$$P(\text{loss of } nL) = \frac{1}{n!} \left. \frac{d^n G(z)}{dz^n} \right|_{z=0} \quad \text{for } n = 1,2,... \tag{6}$$

Note that this equation allows only computing the probability of losses of size 0, L, $2L$ and so on. This probability of loss of nL will further be denoted A_n.

Then, under the simplified assumption of fixed default rates, Credit Suisse has developed the following recursive equation:[10]

$$A_n = \sum_{j:\nu_j \leq n} \frac{\varepsilon_j}{n} A_{n-\nu_j} \tag{7}$$

where $A_0 = G(0) = e^{-\mu} = e^{-\sum_{j=1}^{m} \mu_j} = e^{-\sum_{j=1}^{m} \frac{\varepsilon_j}{\nu_j}}$.

The calculation depends only on 2 sets of parameters: ν_j and ε_j, derived from μ_j , the number of events of each range, *j*, observed. With operational data, A_0 is derived directly from $A_0 = e^{-\mu}$.

To illustrate this recurrence, suppose the database contains 20 losses, 3 (resp. 2) of which having a rounded-off size of 1*L* (resp. 2*L*):

$$A_0 = e^{-20} = 2.06.10^{-9}$$

$$A_1 = \sum_{j:\nu_j \leq 1} \frac{\varepsilon_j}{1} A_{1-\nu_j} = \varepsilon_1 A_0 = 3 \text{ x } 2.06.10^{-9} = 6.18.10^{-9}$$

$$A_2 = \sum_{j:\nu_j \leq 2} \frac{\varepsilon_j}{2} A_{2-\nu_j} = \frac{1}{2}\left(\varepsilon_1 A_1 + \varepsilon_2 A_0\right) = \frac{1}{2}\left(3\text{x } 6.18.10^{-9} + 2 \text{ x } 2.06.10^{-9}\right) = 1.13.10^{-8}$$

Therefore, the probability of having a loss of size resp. 0, 1*L* and 2*L* is resp. $2.06.10^{-9}$, $6.18.10^{-9}$ and $1.13.10^{-8}$, and so on.

From there, one can re-construct the distribution of the loss of size *nL*.

2.2 The Loss Distribution Approach adapted to Operational Risk

Among the Advanced Measurement Approaches (AMA) developed over the recent years to model operational risk, the most common one is the Loss Distribution Approach (), which is derived from actuarial techniques (see Frachot *et al.* (2001; Cruz (2004); Chavez-Demoulin *et al.*, 2006); Peters *et al.* (2011)for an introduction).

By means of convolution, this technique derives the aggregate loss distribution (ALD) through the combination of the frequency distribution of loss events and the severity distribution of a loss given event.11 The operational Value-at-Risk is then simply the 99.9th percentile of the ALD. As an analytical solution is very difficult to compute with this type of convolution, Monte Carlo simulations are usually used to do the job. Using the CCHP procedure with a Poisson distribution with a parameter μ equal to the number of observed losses during the whole period to model the frequency,12 we generate a large number M of Poisson(μ) random variables (say, 100,000). These M values represent the number of events for each of the M simulated periods. For each period, generate the required number of severity random variables (that is, if the simulated number of events for period m is x, then simulate x severity losses) and add them to get the aggregate loss for the period. The obtained vector represents M simulated periods and OpVaRs are then readily obtained (e.g. the OpVaR at 99.99% confidence interval is the 10th lowest value of the M sorted aggregate losses)

Examples of empirical studies using this technique for operational risk include Moscadelli (2004) on loss data collected from the Quantitative Impact Study (QIS) of the Basel Committee, de Fontnouvelle and Rosengren (2004) on loss data from the 2002 Risk Loss Data Collection Exercise initiated by the Risk Management Group of the Basel Committee or Chapelle *et al.* (2008) with loss data coming from a large European bank.

In the latter case, mixing two distributions fits more adequately the empirical severity distribution than a single distribution. Therefore, the authors divide the sample into two parts: a first one with losses below a selected threshold, considered as the “normal” losses, and a second one, including the “large” losses. To model the “normal” losses, CCHP compare several classic continuous distributions such as gamma, lognormal or Pareto. In our example, we will use the lognormal distribution.

To take extreme and very rare losses into account (i.e. the “large” losses), CCHP apply the Extreme Value Theory (EVT) on their results.[13] The advantage of EVT is that it provides a tool to estimate rare and not-yet-recorded events for a given database,[14] hence providing an attractive solution for loss databases with limited collection history that are used to reach a very high confidence levels like the one required by Basel II (i.e. 99.9%).

2.3 Alexander's Internal Measurement Approach

The basic formula of the Internal Measurement Approach (IMA) included in the Advanced Measurement Approaches of Basel II is:

$$UL = \gamma EL \tag{8}$$

where UL = unexpected loss, determining the operational risk requirement,[15] γ is a multiplier, and EL is the expected loss.

Gamma factors are not easy to evaluate as no indication of their possible range has been given by the Basel Committee. Therefore, Alexander (2003) suggests that instead of writing the unexpected loss as a multiple (γ) of expected loss, one writes unexpected loss as a multiple (Φ) of the loss standard deviation (σ). Using the definition of the expected loss, she gets the expression for Φ:

$$\Phi = \frac{VaR_{99.9} - EL}{\sigma} \tag{9}$$

The advantage of this parameter is that it can be easily calibrated.

The basic IMA formula is based on the binomial loss frequency distribution, with no variability in loss severity. For very high-frequency risks, Alexander notes that the normal distribution could be used as an approximation of the binomial loss distribution, providing for Φ a lower bound equal to 3.1 (as can be found from standard normal tables when the number of losses goes to infinity). She also suggests that the Poisson distribution should be preferred to the binomial as the number of transactions is generally difficult to quantify.

Alexander (2003) shows that Φ, as a function of the parameter μ of the Poisson distribution, must be in a fairly narrow range: from about 3.2 for medium-to high frequency risks (20 to 100 loss events per year) to about 3.9 for low frequency risks (one loss event every one or two years) and only above 4 for very rare events that may happen only once every five years or so. Table 3 illustrates the wide range for the gammas by opposition to the narrow range of the phi's values.

Insert Table 3

Then, assuming the loss severity to be random, i.e. with mean μ_L and standard deviation σ_L, and independent of the loss frequency, Alexander writes the Φ parameter as:

$$\Phi = \frac{VaR_{99.9} - \lambda\mu_L}{\sqrt{\left[\lambda\left(\mu_L^2 + \sigma_L^2\right)\right]}} \quad (10)$$

where λ is the average number of losses.

For $\sigma_L > 0$, this formula produces slightly lower Φ than with no severity uncertainty, but it is still bounded below by the value 3.1.

For a comparison purpose, we will use the following value for the needed OpVaRs, derived from equation (10) in which we replace Φ by a value corresponding to the selected level of confidence:

$$OpVaR = \Phi\sqrt{\left[\lambda\left(\mu_L^2 + \sigma_L^2\right)\right]} + \lambda\mu_L \quad (11)$$

3 An Experiment on Simulated Losses

3.1 Data

OpRisk+ makes the traditional statistical tests impossible, as it uses no parametric form but a purely numerical procedure. Therefore, as proposed in McNeil and Saladin (1997), in order to perform tests of the calibrating performance of OpRisk+ on any distribution of loss severity, we simulate databases to obtain an exhaustive picture of the capabilities of the approach. Moscadelli (2004) and de Fontnouvelle and Rosengren (2004) having shown that loss data for most business lines and event types may be well modelled by a Pareto-type distribution, we simulated our data on the basis of three different kinds of Pareto distributions: a heavy-tail, a medium-tail and a thin-tail Pareto distribution.

A Pareto distribution is a right-skewed distribution parameterized by two quantities: a minimum possible value or location parameter, x_m, and a tail index or shape parameter, ξ. Therefore, if X is a random variable with a Pareto distribution, the probability that X is greater than some number x is given by:

$$\Pr(X > x) = \left(\frac{x}{x_m}\right)^{-k} \tag{12}$$

for all $x \geq x_{m.}$, and for x_m and $k = 1/\xi > 0$.

The parameters of our distributions are Pareto(100;0.3), Pareto(100;0.5) and Pareto(100;0.7): the larger the value of the tail index, the fatter the tail of the distribution. The choice of these functions has been found to be reasonable with a sample of real data obtained from a large European institution.

We run three simulations: one for the thin-tailed Pareto severity distribution case, one for the medium-tailed Pareto severity distribution case and one for the fat-tailed Pareto severity distribution case. For each of these cases, we simulate two sets of 1000 years of 20 and 50 operational losses respectively and two sets of 100 series of 200 and 300 operational losses respectively.[16] 1

Table 4 gives the characteristics of each of the twelve databases (each thus comprising 1000 or 100 simulated loss distributions) constructed in order to implement OpRisk+. For each series of operational losses we compute the expected loss, that is, the mean loss multiplied by the number of losses, as well as the standard deviation, median, maximum and minimum of these expected losses.

Insert Table 4

These results clearly show that data generated with a thin-tailed Pareto-distribution exhibit characteristics that make the samples quite reliable. The mean loss is very close to its theoretical level even for 20 draws. Furthermore, we observe a standard deviation of aggregate loss that is very limited, from less than 10% of the average for N=20 to less than 3% for N=200. The median loss is also close to the theoretical value. For a tail index of 0.5 (medium-tailed), the mean loss still stays close to the theoretical value but the standard deviation increases. Thus, we can start to question the stability of the loss estimate.

When the tail index increases, the mean aggregate loss becomes systematically lower than the theoretical mean, and this effect aggravates when one takes a lower number of simulations (100 drawings) with a larger sample. The standard deviation and range become extremely large, thereby weakening inference based on a given set of loss observations.

This highlights the difficulty of modelling operational risk losses (which often exhibit this type of tail behaviour) using classical distribution fitting methods when only a limited number of loss data points are available.

3.2 Application of OpRisk+

To apply OpRisk+ to these data, the first step consists of computing $A_0 = e^{-\mu}$, where μ is the average number of loss events. For instance, for N=200, this gives the following value: $A_0 = e^{-200} = 1.38 \cdot 10^{-87}$. Then, in order to assess the loss distribution of the entire population of operational risk events, we use the recursive equation (7) to compute A_1, A_2 etc.

Once the different probabilities A_n for the different sizes of losses are computed, we can plot the aggregate loss distribution as illustrated in Figure 1-1[17].

Insert Figure 1

With this information, we can compute the different Operational Values-at-Risk (OpVaRs). This is done by calculating the cumulated probabilities for each amount of loss. The loss for which the cumulated probability is equal to *p*% gives us the OpVaR at percentile *p*.

We repeat the procedure for each year of losses and report the average values of the different yearly OpVaRs in Tables 5 and 6. Even though this procedure is likely to underestimate the true quantiles (see Section 3.4), we view this setup as more realistic than merely computing a single OpVaR on the whole number of years. Indeed, the operational risk manager is likely to be confronted with a few years of limited data, which is consistent with our simulation procedure.

Table 5 compares the OpVaRs obtained using OpRisk+ with the simulated data for the small databases. The first column represent the average observed quantiles of the aggregate distribution when simulating 100,000 years with a Poisson(μ) distribution for the frequency and a Pareto(100, ξ) for the severity. The tables also gives the minimum, maximum and standard deviation of the 100(0) OpVaRs produced by OpRisk+.

Insert Table 5

Panel A of Table 5 shows that OpRisk+ achieves very satisfactory OpVaRs for the Pareto-distribution with thin tail. The mean OpVaRs obtained for both the samples of 20 and 50 observations stays within a 3% distance from the true value. Even at the level of 99.9% required by Basel II, the OpRisk+ values remain within a very narrow range, while the root mean square error (RMSE) of the estimates is kept within 13% of the true value.

The results obtained with the OpRisk+ procedure with medium and fat tails tend to deteriorate, which is actually not surprising as the adaptation of the credit risk model strictly uses observed data and does necessarily underestimate the fatness of the tails. However, we still have very good estimation for $OpVaR_{95}$. It mismatches the true 95% quantile by 2% to

7% for the medium and fat tailed Pareto-distribution, while the RMSE tends – naturally – to increase very fast.

The bad news is that the procedure alone is not sufficient to provide the $OpVaR_{99.9}$ required by Basel II. It severely underestimates the true quantile, even though this true value is included in the range of the observed values of the loss estimates, mainly because the support of the distribution generated by the OpRisk+ method is finite and thus truncates the true loss distribution. This issue had been pointed out by Mignola and Ugoccioni (2006) who propose to reduce the sources of uncertainty in modelling the operational risk losses, by lowering the percentile at which the risk measure is calculated and finding some other mechanism to reach the 99.9% percentile.

Further reasons for this systematic underestimation can be found in the setup of the simulations. The procedure averages the individual yearly OpVaRs, each of them being computed using a very small number of losses. This modelling choice mimics a realistic situation as closely as possible. There is thus a small likelihood of observing extreme losses over a particular year, and the averaging process tends to lead to the dominance of too small OpVaR estimates for the extreme quantiles. Table 6 displays the results of the simulations when a large sample size is used.

Insert Table 6

Table 6, Panel A already delivers some rather surprising results. The OpRisk+ procedure seems to overestimate the true operational risk exposure for all confidence levels. This effect aggravates for a high number of losses in the database. This phenomenon may be due to an intervalling effect, where losses belonging to a given band are assigned the value of the band's upper bound. Given that extreme losses are likely to occur in the lower part of the band, as the distribution is characterized by a thin tail Pareto-distribution, taking the upper bound limit value for aggregation seems to deteriorate the estimation, making it too conservative. Nevertheless, the bias is almost constant in relative terms, indicating that its seriousness does not aggravate as the estimation gets far in the tail of the distribution. Sub-section 3.4 investigates further this issue.

This intervalling phenomenon explains the behaviour of the estimation for larger values of the tail index. In Panel B, the adapted credit risk model still overestimates the distribution of losses up to a confidence level of 99%, while in Panel C, the underestimation starts earlier, around the 95% percentile of the distribution. In both cases, the process does not capture to distribution at the extreme end of the tail (99.9%), similar to what we observed for smaller sample sizes.

Nevertheless, from panels B and C altogether, the performance of OpRisk+ still stays honourable when the confidence level of 95% is adopted. The RMSE of the estimates also remains within 20% (with the tail index of 0.5) and 32% of the mean (with a tail index of 0.7), which is fairly large but mostly driven by large outliers as witnessed in the last column of each panel.

A correct mean estimate of the $OpVaR_{95}$ would apply to a tail index between 0.5 and 0.7, which corresponds to a distribution with a fairly large tail index. Only when the tail of the Pareto-distribution is actually thin, one observes that the intervalling effect induces a large discrepancy between the theoretical and observed values.

It remains to be mentioned that the good empirical application of OpRisk+ does not depend on the number of observed losses as it only affects the first term of the recurrence, namely A_0.

3.3 Comparison with the CCHP and the Alexander's approaches

These results, if their economic and statistical significance have to be assessed, have to be compared with a method that aims at specifically addressing the issue of operational losses in the Advanced Measurement Approaches setup. We choose the CCHP approach, which is by definition more sensitive to extreme events than OpRisk+, but has the drawback of requiring a large number of events to properly derive the severity distributions of "normal" and "large" losses. For low frequency database, the optimization processes used by this type of approaches (e.g. Maximum Likelihood Estimation) might not converge to stable parameters estimates.

The graphs from Figure 1-2 display the OpVaRs (with confidence levels of 90, 95, 99 and 99.9%) generated from three different kind of approaches, that is the sophisticated CCHP approach, OpRisk+ and the simpler Alexander (2003) approach (see Section 2.3) for each of the three tail index values (0.3, 0.5 and 0.7) and for each of the four sample size (20, 50, 200 and 300 loss events).

Insert Figure 2

From the graphs in Figure 1-2, we can see that for most databases, OpRisk+ is working out a capital requirement higher than the Alexander's IMA, but smaller than the CCHP approach. This last result could be expected as CCHP is more sensitive to extreme events. In next sub-section, we will discuss the fact that the database with 300 observations shows higher OpVaRs for OpRisk+ than CCHP. However, we can already conclude that our model is more risk sensitive than a simple IMA approach.

Considering the thin tailed Pareto-distribution in Panel A, we can observe that OpRisk+ produces the best estimations for the small database. Indeed, those are very close to the theoretical OpVaRs for all confidence level. However, for the large database, it is producing too cautious (large) OpVaRs. The comparison with other methods sheds new light on the results obtained with Panel A of Table 6: OpRisk+ overestimates the true VaR, but the CCHP model, especially dedicated to the measurement of operational risk, does frequently worse. Actually, Alexander (2003) approach, also using observed data but not suffering from an intervalling effect, works out very satisfactory results when the standard deviation of loss is a good proxy of the variability of the distribution.

For the medium and fat tailed Pareto-distributions, neither of the models is sensitive enough for OpVaRs of 99% and more. This could raise some questions on the feasibility or appropriateness of a requirement of a 99.9% VaR by the Basel Accord, where it appears that even an LDA model is far from being able to estimating economic capital with such a high level of confidence. Nevertheless, as far as the small databases are concerned, it is interesting to note that OpRisk+ is producing the best estimations for $OpVaR_{95}$.

While none of these approach seems good enough for the level of confidence required by Basel II, we would first recommend OpRisk+ or Alexander's for low frequency databases, as none of these needs the pre-determination of the shape of the distribution. Then, although Alexander's approach is simpler and provides as good OpVaRs as our model for the thin-tail Pareto distribution, this method has the drawback of deteriorating much faster than OpRisk+ for larger tails. Unfortunately, risk managers usually do not know the type of distribution they are dealing with, and in this case, we would recommend the OpRisk+ method that seems a bit more complicated but yields more consistent results.

3.4 Comparison with OpRisk+ taking an average value of loss for each band

As shown above, taking the upper bound limit value for aggregation as described in the CreditRisk+ model tends to overestimate the true operational risk exposure for all confidence levels; especially with larger databases. A solution could be to take the average value of losses for each band.[18] Table 7 displays the results of the simulations when a relatively large sample size is used.

Insert Table 7

Panel A of Table 7 shows that OpRisk+ achieves very good results for the Pareto-distribution characterized by a thin tail when using an average value for each band ("round" column). The OpVaR values obtained for the sample of 200 observations is very close to the theoretical value, whereas it stays within a 6% range from the "true" value with a 300 observations sample, including at the Basel II level of 99.9%.

When the loss Pareto-distributions are medium-tailed, the results obtained with the OpRisk+ procedure with the databases are very good for quantiles up to 95% but deteriorate for more sensitive OpVaRs. OpRisk+ is still totally unable to capture the tailedness of the distribution of aggregate losses for very high confidence interval, such as the Basel II requirement.

Table 8 compares the two methods when applied to small databases of 20 and 50 observations. In such cases, OpRisk+ provides better results with the "round up" solution than with the "round-off" one. This bias could be due to the fact that with the second method we tend to loosen the "extreme value theory" aspect of the model. Small databases tend indeed to lack extreme losses and taking the upper bound limit value for the aggregation makes the resulting distribution's tail fatter.

Insert Table 8

4 An application to real loss data

As an illustration of the application of our model, we applied the three models on operational loss data provided by a large European bank. As this bank required confidentiality, we will not publish our results. Given that we only had a collection of one year of data, we could not apply the CCHP model on low frequency data.

We first applied the three models on two cells characterized by more than 100 losses, and noticed that OpRisk+ VaRs were systematically higher than the lower bound of

Alexander and lower than the CCHP VaRs at the 99.9% level of confidence. However, we found close VaRs for the 90th and the 95th percentiles.

We then applied OpRisk+ to two lower frequency cells, with about 20 losses, and were able to computed OpVaRs for both cells, that were higher than the lower bound of Alexander.

5 Conclusions

This paper introduces a structural operational risk model, named OpRisk+, that has been inspired from the well known credit risk model, CreditRisk+, which has characteristics transposable to the operational risk modelling.

In a simulation setup, we work out aggregate loss distributions and operational Value-at-Risks (OpVaR) for various confidence levels, including the one required by Basel II. The performance of our model is assessed by comparing our results to theoretical OpVaRs, to an OpVaR issued from a simpler approach, that is, the IMA approach of Alexander (2003), and to a more sophisticated approach proposed in Chapelle *et al.* (2008), or “CCHP” approach which uses a mixture of two distributions to model the body and the tail of the severity distribution separately.

The results show that OpRisk+ produces OpVaRs closer to theoretical ones than the approach of Alexander (2003), but that it is not receptive enough to extreme events. On the other hand, our goal is not to produce a complete compliant AMA model to compute regulatory capital requirements, but rather to propose a first solution to the lack of low frequency operational risk models. Besides, whereas the CCHP approach has better sensitivity to very extreme losses, the simplicity of OpRisk+ gives the model the advantage of requiring no large database in order to be implemented.

Specifically, we view the value-added of the OpRisk+ procedure as twofold. Firstly, it produces average estimates of operational risk exposures that are very satisfactory at the

95% level, which makes it a very useful complement to approaches that specifically target the extreme tail of the loss distribution. Indeed, even though the performance of OpRisk+ is clearly not sufficient for the measurement of unexpected operational losses as defined by the Basel II Accord (the VaR should be measured with a 99.9% confidence level), it could be thought of as a sound basis for the measurement of the body of losses; another more appropriate method must relay OpRisk+ for the measurement of the far end of the distribution. Moreover, it appeared to us, that the 99.9% level of confidence required by Basel II might be quite utopian when we observe that even an LDA approach with 300 losses do not even get close to the theoretical level when the distribution is characterized with a Pareto(100;0.7).

Secondly, despite the fact that we cannot conclude that OpRisk+ is an adequate model to quantify the economic capital associated to the bank's operational risk, its applicability to approximate the loss distribution with small databases is proven. Even for such a small database as one comprising 20 observations, the estimation could make it attractive as a complement to more sophisticated approaches requiring large numbers of data per period. The fit is almost perfect when the Pareto-distribution has a thin tail, and the $OpVaR_{95}$ is the closest among the three specifications tested when the tail gets fatter.

Of course, this approach is still subject to refinements, and could be improved in many ways. Indeed, internal data rarely includes very extreme events (banks suffering those losses probably would no more be there to tell us), whereas the last percentiles are very sensitive to the presence of those events. The problem would therefore be to determine which weight to place on the internal data and on the external ones. From our study, we could imagine that fitting a distribution calibrated with external data, EVT or relying on scenario analysis beyond the 95% percentile would justify the use of OpRisk+ preferably to other models. This advantage can prove to be crucial for business lines or event types where very few internal observations are available, and thus where more data intensive approaches such as the CCHP would be powerless.

6 Appendix - CreditRisk+: The distribution of default losses - Calculation procedure[19]

CreditRisk+ mathematically describes the random effect of the severity distribution through its probability generating function *G(Z)*:

$$G(z) = \sum_{n=0}^{\infty} P(\text{aggregated loss} = \text{n x L}) z^n$$

Comparing this definition with the Taylor series expansion for *G(z)*, the probability of a loss of *n x L*, A_n, is given by:

$$P(\text{loss of nL}) = \frac{1}{n!} \left. \frac{d^n G(z)}{dz^n} \right|_{z=0} = \mathrm{A}_\mathrm{n}$$

In CreditRisk+, *G(Z)* is given in closed form by :

$$G(z) = \prod_{j=1}^{m} e^{-\mu_j + \mu_j z^{\nu_j}} = e^{-\sum_{j=1}^{m} \mu_j + \sum_{j=1}^{m} \mu_j z^{\nu_j}}$$

Therefore, using Leibniz formula we have:

$$\begin{aligned} \frac{1}{n!} \left. \frac{d^n G(z)}{dz^n} \right|_{z=0} &= \frac{1}{n!} \frac{d^{n-1}}{dz^{n-1}} \left. \left(G(z). \frac{d}{dz} \sum_{j=1}^{m} \mu_j z^{\nu_j} \right) \right|_{z=0} \\ &= \frac{1}{n!} \sum_{k=0}^{n-1} \binom{n-1}{k} \frac{d^{n-k-1}}{dz^{n-k-1}} G(z). \frac{d^{k+1}}{dz^{k+1}} \left. \left(\sum_{j=1}^{m} \mu_j z^{\nu_j} \right) \right|_{z=0} \end{aligned}$$

However

$$\cdot \frac{d^{k+1}}{dz^{k+1}} \left. \left(\sum_{j=1}^{m} \mu_j z^{\nu_j} \right) \right|_{z=0} = \begin{cases} \mu_j (k+1)! & \text{if } k = \nu_j - 1 \text{ for some } j \\ 0 & \text{otherwise} \end{cases}$$

and by definition,

$$\cdot \frac{d^{n-k-1}}{dz^{n-k-1}} G(z) \bigg|_{z=0} = (n-k-1)! \mathrm{A}_{\mathrm{n\text{-}k\text{-}1}}$$

Therefore

$$A_n = \sum_{\substack{k \le n-1 \\ k=\nu_j - 1 \text{ for some j}}} \frac{1}{n!}\binom{n-1}{k}(k+1)!(n-k-1)!\,\mu_j\, A_{n-k-1} = \sum_{j:\nu_j \le n}^{m} \frac{\mu_j \nu_j}{n} A_{n-\upsilon}$$

Using the relation $\varepsilon_j = \nu_j \cdot \mu_j$, the following recursive equation is obtained:

$$A_n = \sum_{j:\nu_j \le n} \frac{\varepsilon_j}{n} A_{n-\nu_j}$$

7 REFERENCES

Alexander, C. (2003) Statistical models of operational loss, in *Operational Risk. Regulation, Analysis and Management*, Ft Prentice Hall Financial Times, Upper Saddle River, NJ., pp. 129-170.

Basel Committee on Banking Supervision (2009) Observed range of practice in key elements of advanced measurement approaches, BIS Report.

Carrillo-Menéndez, S. and Suárez, A. (2012) Robust quantification of the exposure to operational risk: Bringing economic sense to economic capital, *Computers & Operations Research*, **39**, 792-804.

Chapelle, A., Crama, Y., Hübner, G. and Peters, J. (2008) Practical methods for measuring and managing operational risk in the financial sector: A clinical study, *Journal of Banking and Finance*, **32**, 1049-1061.

Chavez-Demoulin, V., Embrechts and Neslehova, J. (2006) Quantitative models for operational risk: Extremes, dependence and aggregation, *Journal of Banking and Finance*, **30**, 2635-2658.

Crouhy, M., Galai, D. and Mark, R. (2000) A comparative analysis of current credit risk models, *Journal of Banking & Finance*, **24**, 59-117.

Cruz, M. (2004) *Operational Risk Modelling and Analysis: Theory and Practice*, Risk Waters Group, London.

Cruz, M. G. (2002) Frequency models, in *Modeling, Measuring and Hedging Operational Risk.*, Wiley Finance, New York.

de Fontnouvelle, P., Dejesus-Rueff, V., Jordan, J. and Rosengren, E. (2003) Capital and risk: New evidence on implications of large operational losses, Working Paper No 03-5, Federal Reserve Bank of Boston.

de Fontnouvelle, P. and Rosengren, E. (2004) Implications of alternative operational risk modeling techniques, Federal Reserve Bank of Boston.

Degen, M., Embrechts, P. and Lambrigger, D. (2007) The quantitative modelling of operational risk: between g-and-h and EVT, *Astin Bulletin*, **37**, 265-291.

Dunne, T. and Helliar, C. (2002) The Ludwig report: Implications for corporate governance, *Corporate Governance*, **2**, 26-31.

Embrechts, P., Kluppelberg, C. and Mikosch, T. (1997) *Modelling Extremal Events for Insurance and Finance*, Springer-Verlag, Berlin.

Frachot, A., Georges, P. and Roncalli, T. (2001) Loss distribution approach for operational risk, Crédit Lyonnais.

Gordy, M. (2000) A comparative anatomy of credit risk models, *Journal of Banking and Finance*, **24**, 119-149.

King, J. (2001) *Operational Risk: Measurement and Modelling*, Wiley, New York.

McNeil, A. J. and Saladin, T. (1997) The Peak Over Thresholds method for estimating high quantiles of loss distributions, ETH Preprint.

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

Moscadelli, M. (2004) The modelling of operational risk: Experience with the analysis of the data collected by the Basel Committee, No 517, Banca d'Italia.

Peters, G. W., Shevchenko, P. V., Young, M. and Yip, W. (2011) Analytic loss distributional approach models for operational risk from the -stable doubly stochastic compound processes and implications for capital allocation, *Insurance: Mathematics and Economics*, **49**, 565-579.

Ross, J. (1997) Rogue trader: How I brought down Barings Bank and shook the financial world by Nick Leeson, *Academy of Management Review*, **22**, 1006-1010.

Sahay, A., Wan, Z. and B., K. (2007) Operational Risk Capital: Asymptotics in the case of Heavy-Tailed Severity., *Journal of Operational Risk*, **2**.

Sheaffer, Z., Richardson, B. and Rosenblatt, Z. (1998) Early-warning-signals management: A lesson from the Barings crisis, *Journal of Contingencies and Crisis Management*, **6**, 1-22.

Stonham, P. (1996) Whatever Happened at Barings? Part Two: - Unauthorised Trading and the Failure of Controls, *European Management Journal*, **14**, 269-278.

Walker, P. L., Shenkir, W. G. and Hunn, C. S. (2001) Developing risk skills: An investigation of business risks and controls at Prudential Insurance Company of America, *Issues in Accounting Education*, **16**, 291-313.

Table 1

Allocating losses to bands.

Loss Amount (A)	Loss in L (B)	round-off loss *vj (C)*	band j (D)
1 500	1.5	2.00	2
2 508	2.51	3.00	3
3 639	3.64	4.00	4
1 000	1.00	1.00	1
1 835	1.84	2.00	2
2 446	2.45	3.00	3
7 260	7.26	8.00	8

Illustration of the first three steps of the OpRisk+ approach: 1. Choose a unit amount of loss L. (1000 in the example) 2. Divide the losses of the available database (column A) by L (column B)and round up these numbers (column C). 3. Allocate the losses of different sizes to their band (column D)

Table 2

Exposure, number of events and expected loss.

v_j	μ_j	ε_j
1	9	9
2	121	242
3	78	234
4	27	108
5	17	85
6	15	90
7	8	56
8	4	32
⋮	⋮	⋮

Illustration of step 5 of the OpRisk+ approach: "Compute the expected loss per band, ε_j, equal to the expected number of losses per band μ_j , multiplied by the average loss amount per band, v_j , equal to *j*."

Table 3

Gamma and phi values (no loss severity variability)

μ	100	50	40	30	20	10	8	6
$VaR_{99.9}$	131.81	72.75	60.45	47.81	34.71	20.66	17.63	14.45
Φ	3.18	3.22	3.23	3.25	3.29	3.37	3.41	3.45
γ	0.32	0.46	0.51	0.59	0.74	1.07	1.21	1.41

μ	5	4	3	2	1	0.9	0.8	0.7
$VaR_{99.9}$	12.77	10.96	9.13	7.11	4.87	4.55	4.23	3.91
Φ	3.48	3.48	3.54	3.62	3.87	3.85	3.84	3.84
γ	1.55	1.74	2.04	2.56	3.87	4.06	4.29	4.59

μ	0.6	0.5	0.4	0.3	0.2	0.1	0.05	0.01
$VaR_{99.9}$	3.58	3.26	2.91	2.49	2.07	1.42	1.07	0.90
Φ	3.85	3.90	3.97	4.00	4.19	4.17	4.54	8.94
γ	4.97	5.51	6.27	7.30	9.36	13.21	20.31	89.40

(source: Alexander, C. (2003), p151).

Illustratation of the wide range for the gammas by opposition to the narrow range of the phi's values in the computation of the unexpected loss (UL=. $VaR_{99.9}$ – EL) determining the operational risk requirement

The basic formula of the Internal Measurement Approach (IMA) of Basel II is UL = γ EL, where γ is a multiplier, and EL is the expected loss. As Gamma factors are not easy to evaluate, Alexander, C. (2003) suggests to write unexpected loss as a multiple (Φ) of the loss standard deviation (σ).

Table 4

Characteristics of the twelve databases created for testing the different models. For three Pareto severity distribution (thin-tailed, medium tailed and fat-tailed), we simulate two sets of 1000 years of 20 and 50 operational losses respectively and two sets of 100 series of 200 and 300 operational losses respectively. For each of the 6600 simulated years (3 x 2 x 1100), the aggregate loss distribution is computed with the algorithm described in Section 2.2.

Panel A : Thin-tailed-Pareto distribution (shape parameter = 0.3)				
Poisson parameter μ	20	50	200	300
Theoretical Mean	2857	7143	28571	42857
Mean	2845	7134	28381	42886
Standard deviation	287	472	847	1118
Median	2796	7078	28172	42763
Maximum	4683	9026	30766	45582
Minimum	2268	6071	26713	40383
Number of simulated years	1000	1000	100	100

Panel B : Medium-tailed-Pareto distribution (shape parameter = 0.5)				
Poisson parameter μ	20	50	200	300
Theoretical Mean	4000	10000	40000	60000
Mean	3924	9913	39871	59431
Standard deviation	1093	1827	3585	5504
Median	3676	9594	39777	57947
Maximum	15680	29029	54242	91182
Minimum	2567	7097	33428	52436
Number of simulated years	1000	1000	100	100

Panel C : Fat-tailed-Pareto distribution (shape parameter = 0.7)				
Poisson parameter μ	20	50	200	300
Theoretical Mean	6667	16667	66667	100000
Mean	6264	16165	61711	93724
Standard deviation	5940	13018	13899	24514
Median	5180	13721	57713	87646
Maximum	157134	265621	137699	248526
Minimum	2646	8304	45315	69991
Number of simulated years	1000	1000	100	100

Table 5

Values-at-Risk generated by OpRisk+ for small databases, with 20 and 50 loss events. The OpVaRs are calculated separately for each year of data, and we report their average (Mean), the average value of the spread between the "true" value, *Target*, and the OpVaRs, as percents of the latest (Bias), and the root mean square error as percents of the "true" OpVaRs (RMSE). The "true" value or target is approximated through a Monte Carlo simulation of 100,000 years of data, characterized by a frequency equal to a random variable following a Poisson(*N*) and a severity characterized by the selected Pareto-distribution. The unit amount chosen for the OpRisk+ implementation and the average number of corresponding bands is reported in brackets.

	N = 20				N = 50			
	Target	OpRisk+ (L=10, bands = 9)			Target	OpRisk+ (L=10, bands = 13)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	3770	3880	3%	13%	8573	8882	3%	8%
$OpVaR_{95}$	4073	4173	3%	13%	9030	9334	3%	9%
$OpVaR_{99}$	4712	4744	1%	13%	9942	10209	3%	9%
$OpVaR_{99,9}$	5596	5410	-3%	13%	11141	11250	1%	10%

Panel B : Medium-tailed-Pareto distribution (shape parameter = 0.5)

	N = 20				N = 50			
	Target	OpRisk+ (L=10, bands = 11)			Target	OpRisk+ (L=10, bands = 19)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	5579	5672	2%	40%	12630	12855	-6%	29%
$OpVaR_{95}$	6364	6247	-2%	46%	13862	13734	-7%	32%
$OpVaR_{99}$	8966	7329	-18%	48%	18051	15410	-20%	36%
$OpVaR_{99,9}$	18567	8626	-54%	60%	33554	17338	-52%	55%

Panel C : Fat-tailed-Pareto distribution (shape parameter = 0.7)

	N = 20				N = 50			
	Target	OpRisk+ (L=50, bands = 7)			Target	OpRisk+ (L=50, bands = 13)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	9700	11410	18%	107%	22495	23992	7%	116%
$OpVaR_{95}$	12640	12931	3%	99%	28103	27089	-3%	134%
$OpVaR_{99}$	27261	15583	-43%	72%	55994	32020	-43%	99%
$OpVaR_{99,9}$	114563	18726	-84%	85%	220650	38761	-83%	88%

Table 6

OpVaRs generated by OpRisk+ for databases with 200 and 300 loss events. The OpVaRs are calculated separately for each year of data, and we report their average (Mean), the average value of the spread between the "true" value and the OpVaRs, as percents of the latter (Bias), and the root mean square error as percents of the "true" OpVaRs (RMSE). The "true" value is approximated through a Monte Carlo simulation of 100,000 years of data, characterized by a frequency equal to a random variable following a Poisson(*N*) and a severity characterized by the selected Pareto-distribution. The unit amount chosen for the OpRisk+ implementation and the average number of corresponding bands is reported in brackets.

Panel A : Thin-tailed-Pareto distribution (shape parameter = 0.3)								
	N = 200				N = 300			
	Target	OpRisk+ (L=20, bands = 13)			Target	OpRisk+ (L=50, bands = 8)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	31448	33853	7%	8%	46355	56470	22%	22%
$OpVaR_{95}$	32309	34728	7%	8%	47403	57683	22%	22%
$OpVaR_{99}$	33995	36397	7%	7%	49420	59992	21%	22%
$OpVaR_{99,9}$	36063	38310	6%	7%	51750	62628	21%	21%
Panel B : Medium-tailed-Pareto distribution (shape parameter = 0.5)								
	N = 200				N = 300			
	Target	OpRisk+ (L=50, bands = 14)			Target	OpRisk+ (L=50, bands = 11)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	45757	51836	13%	18%	67104	75723	13%	19%
$OpVaR_{95}$	48259	53816	12%	18%	70264	78161	11%	20%
$OpVaR_{99}$	55919	57668	3%	16%	79718	82817	4%	19%
$OpVaR_{99,9}$	83292	62237	-25%	29%	113560	88309	-22%	27%
Panel C : Fat-tailed-Pareto distribution (shape parameter = 0.7)								
	N = 200				N = 300			
	Target	OpRisk+ (L=50, bands = 21)			Target	OpRisk+ (L=50, bands = 17)		
		Mean	Bias	RMSE		Mean	Bias	RMSE
$OpVaR_{90}$	82381	82539	0%	30%	120654	119943	-1%	29%
$OpVaR_{95}$	96971	88248	-9%	32%	139470	127037	-9%	32%
$OpVaR_{99}$	166962	98972	-41%	47%	234442	140665	-40%	47%
$OpVaR_{99,9}$	543597	111875	-79%	80%	733862	156642	-79%	79%

Table 7

Comparison of the average of the yearly OpVaRs computed with OpRisk+ using resp. an upper bound limit value (rounded up) and an average value (rounded) for the allocations into bands (see step 2 of the OpRisk+ procedure described in section 2.1), for large databases. The average value of the spread between the “true” value and the mean of the yearly OpVaRs, as percents of the latter, is reported under the “Bias” column’s titles. The “true” value is approximated through a Monte Carlo simulation of 100,000 years of data, characterized by a frequency equal to a random variable following a Poisson(N) and a severity characterized by the selected Pareto distribution.

Panel A : Thin-tailed-Pareto distribution (shape parameter = 0.3)										
	N = 200					N = 300				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	31448	33853	8%	30576	-3%	46355	56470	22%	43558	-6%
$OpVaR_{95}$	32309	34728	7%	31404	-3%	47403	57683	22%	44563	-6%
$OpVaR_{99}$	33995	36397	7%	32991	-3%	49420	59992	21%	46486	-6%
$OpVaR_{99,9}$	36063	38310	6%	34813	-3%	51750	62628	21%	48687	-6%

Panel B : Medium-tailed-Pareto distribution (shape parameter = 0.5)										
	N = 200					N = 300				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	45757	51836	13%	44338	-3%	67104	75723	13%	64523	-4%
$OpVaR_{95}$	48259	53816	12%	46222	-4%	70264	78161	11%	66849	-5%
$OpVaR_{99}$	55919	57668	3%	49885	-11%	79718	82817	4%	71296	-11%
$OpVaR_{99,9}$	83292	62237	-25%	54257	-35%	113560	88309	-22%	76544	-33%

Panel C : Fat-tailed-Pareto distribution (shape parameter = 0.7)										
	N = 200					N = 300				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	82381	82539	0%	75696	-8%	120654	119943	-1%	112596	-7%
$OpVaR_{95}$	96971	88248	-9%	81375	-16%	139470	127037	-9%	120850	-13%
$OpVaR_{99}$	166962	98972	-41%	91991	-45%	234442	140665	-76%	135481	-42%
$OpVaR_{99,9}$	543597	111875	-79%	104699	-81%	733862	156642	-79%	152904	-79%

Table 8

Comparison of the average of the yearly OpVaRs computed using OpRisk+ with an upper bound limit value (round up) and an average value (rounded) for the allocations into bands (see step 2 of the OpRisk+ procedure described in section 2.1), for small databases. The average value of the spread between the "true" value and the mean of the yearly OpVaRs, as percents of the latter, is reported under the "Bias" column's titles. The "true" value is approximated through a Monte Carlo simulation of 100,000 years of data, characterized by a frequency equal to a random variable following a Poisson(*N*) and a severity characterized by the selected Pareto distribution.

Panel A : Thin-tailed-Pareto distribution (shape parameter = 0.3)										
	N = 20					N = 50				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	3770	3880	3%	3535	-6%	8573	8882	4%	8074	-6%
$OpVaR_{95}$	4073	4173	2%	3815	-6%	9030	9334	3%	8501	-6%
$OpVaR_{99}$	4712	4744	1%	4363	-7%	9942	10209	3%	9332	-6%
$OpVaR_{99,9}$	5596	5410	-3%	5010	-10%	11141	11250	1%	10311	-7%
Panel B : Medium-tailed-Pareto distribution (shape parameter = 0.5)										
	N = 20					N = 50				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	5579	5672	2%	5332	-4%	12630	12855	2%	11323	-10%
$OpVaR_{95}$	6364	6247	-2%	5901	-7%	13862	13734	-1%	12152	-12%
$OpVaR_{99}$	8966	7329	-18%	6945	-23%	18051	15410	-15%	13668	-24%
$OpVaR_{99,9}$	18567	8626	-54%	7904	-57%	33554	17338	-48%	14377	-57%
Panel C : Fat-tailed-Pareto distribution (shape parameter = 0.7)										
	N = 20					N = 50				
	Target	OpRisk+				Target	OpRisk+			
		Roundup	Bias	Round	Bias		Roundup	Bias	Round	Bias
$OpVaR_{90}$	9700	11410	18%	9413	-3%	22495	23992	7%	25235	12%
$OpVaR_{95}$	12640	12931	2%	10914	-14%	28103	27089	-4%	28537	2%
$OpVaR_{99}$	27261	15583	-43%	13353	-51%	55994	32020	-43%	33837	-40%
$OpVaR_{99,9}$	114563	18726	-84%	16290	-86%	220650	38761	-82%	40024	-82%

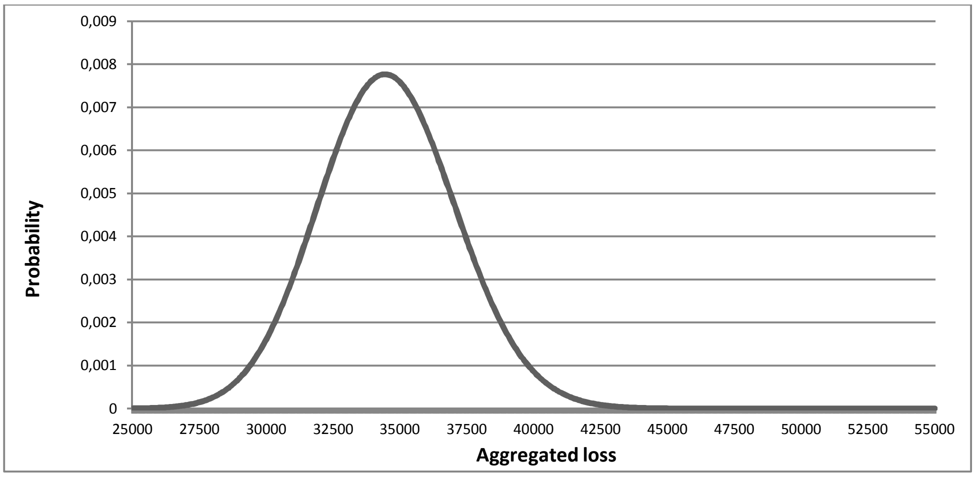

Figure 1. Aggregate loss distribution derived from the application of OpRisk+ for a series of 200 loss events characterized by a Pareto(100;0.3).

Panel A : Thin-tailed Pareto distribution (tail index = 0,3)

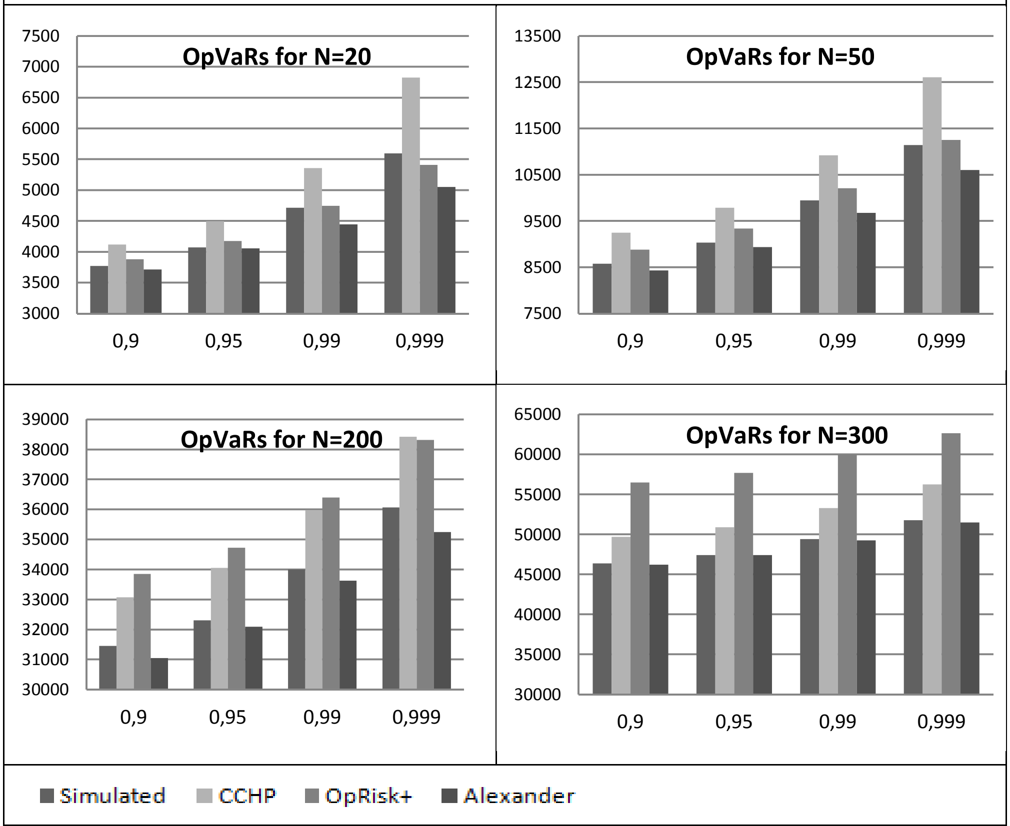

Panel B : medium-tailed Pareto distribution (tail index = 0,5)

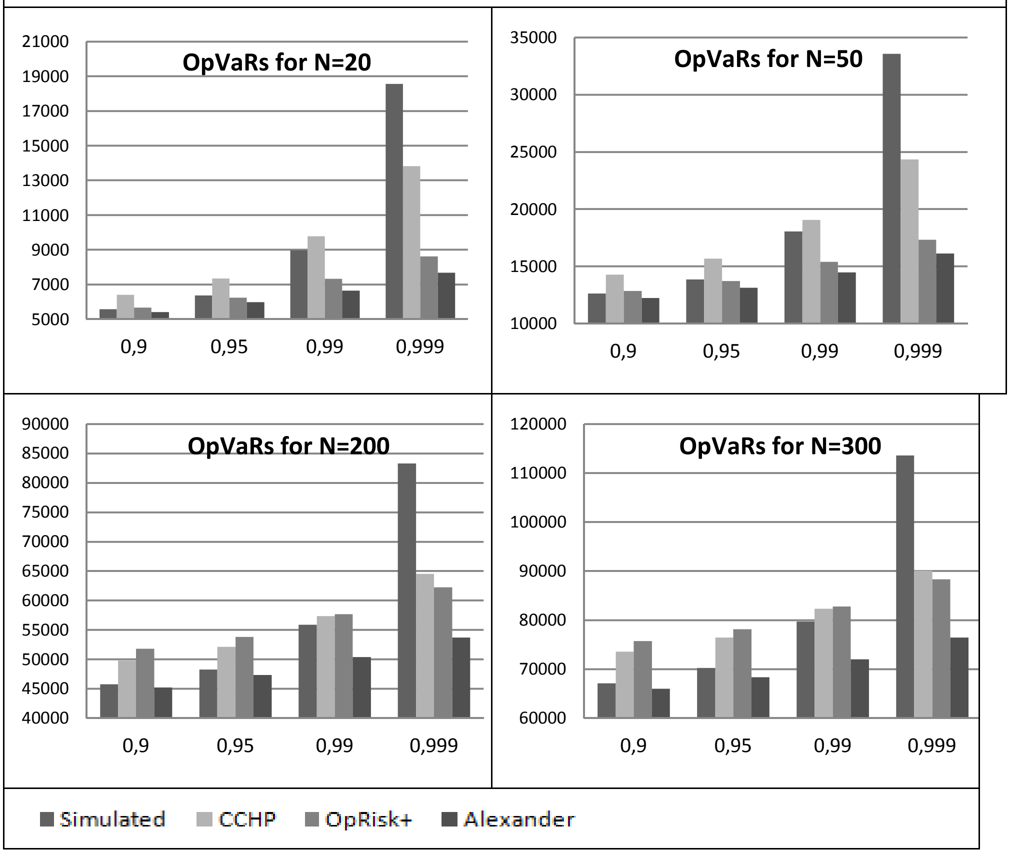

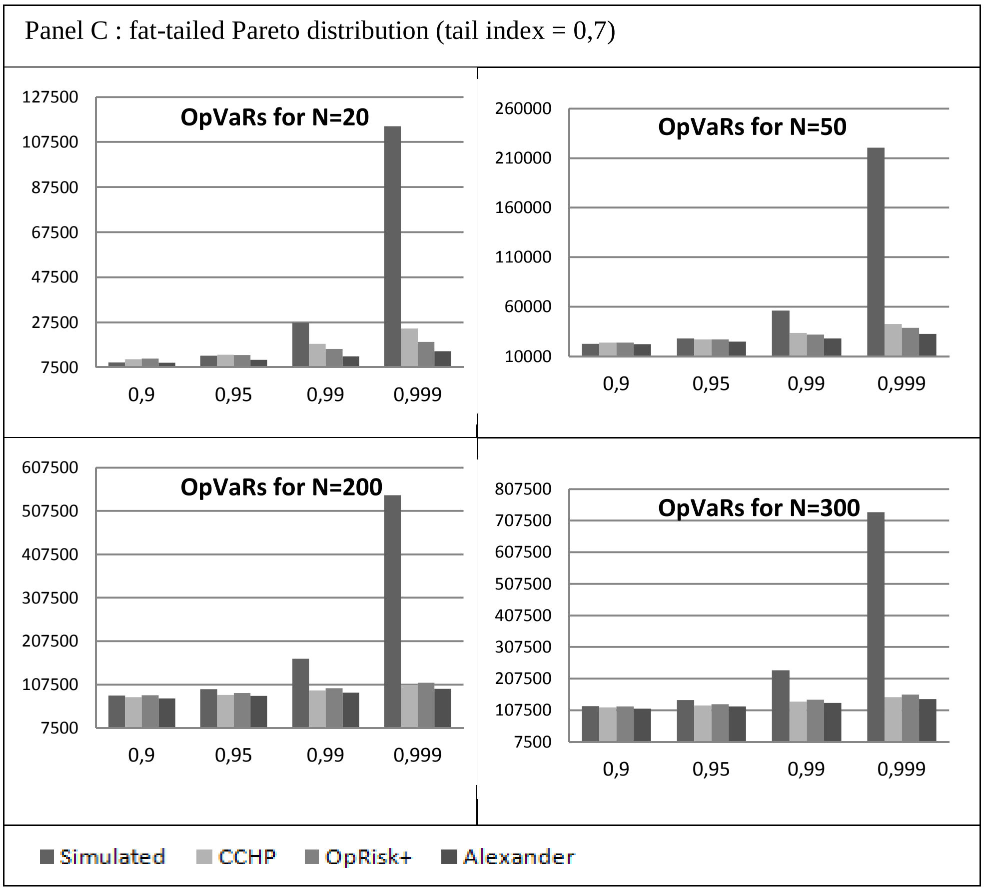

Figure 2. Comparison of CCHP, OpRisk+ and Alexander's IMA Approach.

On the basis of N simulated losses, characterized by a thin, medium or fat-tailed Pareto distribution, we computed OpVaR with level of confidence at 90, 95, 99 and 99.9 percents using three different approaches. The "Simulated" value corresponds to the true value to estimate.

[1] See Ross, J. (1997) Rogue trader: How I brought down Barings Bank and shook the financial world by Nick Leeson, *Academy of Management Review,* **22**, 1006-1010.; Stonham, P. (1996) Whatever Happened at Barings? Part Two: - Unauthorised Trading and the Failure of Controls, *European Management Journal,* **14**, 269-278.; Sheaffer, Z., Richardson, B. and Rosenblatt, Z. (1998) Early-warning-signals management: A lesson from the Barings crisis, *Journal of Contingencies and Crisis Management,* **6**, 1-22..

[2] See Dunne, T. and Helliar, C. (2002) The Ludwig report: Implications for corporate governance, *Corporate Governance,* **2**, 26-31..

[3] See Walker, P. L., Shenkir, W. G. and Hunn, C. S. (2001) Developing risk skills: An investigation of business risks and controls at Prudential Insurance Company of America, *Issues in Accounting Education,* **16**, 291-313..

[4] See Sahay, A., Wan, Z. and B., K. (2007) Operational Risk Capital: Asymptotics in the case of Heavy-Tailed Severity., *Journal of Operational Risk,* **2**. or Degen, M., Embrechts, P. and Lambrigger, D. (2007) The quantitative modelling of operational risk: between g-and-h and EVT, *Astin Bulletin,* **37**, 265-291..

[5] We named our model OpRisk+ to keep its source to mind, that is, the CreditRisk+ model developed by Credit Suisse First Boston. Our model is not a new model but an adaptation of their model to make it useful in our specific situation: that is small samples of operational loss data.

[6] The Value-at-Risk (VaR) is the amount that losses will likely not exceed, within a predefined confidence level and over a given time-period.

[7] Cruz, M. G. (2002) Frequency models, in *Modeling, Measuring and Hedging Operational Risk.,* Wiley Finance, New York. argues that this is due to its simplicity and to the fact that it fits most of the databases very well.

[8] CreditRisk+'s authors argue that the exact amount of each loss cannot be critical in the determination of the global risk.

[9] The purpose of the model is to be applied to real loss data.

[10] See Appendix.

[11] More precisely the ALD is obtained through the *n*-fold convolution of the severity distribution with itself, *n* being a random variable following the frequency density function.

[12] While frequency could also be modelled with other discrete distributions such as the Negative Binomial for instance, many authors use the Poisson assumption (see de Fontnouvelle, P., Dejesus-Rueff, V., Jordan, J. and Rosengren, E. (2003) Capital and risk: New evidence on implications of large operational losses, Working Paper No 03-5, Federal Reserve Bank of Boston. for instance).

[13] This solution has been advocated by many other authors; see for instance King, J. (2001) *Operational Risk: Measurement and Modelling,* Wiley, New York., Cruz, M. (2004) *Operational Risk Modelling and*

Analysis: Theory and Practice, Risk Waters Group, London., Moscadelli, M. (2004) The modelling of operational risk: Experience with the analysis of the data collected by the Basel Committee, No 517, Banca d'Italia., de Fontnouvelle, P. and Rosengren, E. (2004) Implications of alternative operational risk modeling techniques, Federal Reserve Bank of Boston. or Chavez-Demoulin, V., Embrechts and Neslehova, J. (2006) Quantitative models for operational risk: Extremes, dependence and aggregation, *Journal of Banking and Finance,* **30**, 2635-2658..

[14] See Embrechts, P., Kluppelberg, C. and Mikosch, T. (1997) *Modelling Extremal Events for Insurance and Finance,* Springer-Verlag, Berlin. for a comprehensive overview of EVT.

[15] The unexpected loss is defined as the difference between the value-at-risk at the 99.9% confidence level ($VaR_{99.9}$) and the expected loss.

[16] Only 100 years of data were simulated for high-frequency databases as the computation becomes too heavy for a too large number of data. However, we tested our model with 200 years of data for the sample of 200 events characterized by a Pareto (100, 0.7), and did not obtain significantly different OpVaRs. Detailed results are available upon request.

[17] Note that this figure represents the distribution built from one year of data (200 losses), whereas Table 4 displays the average mean of the 100 years of 200 losses.

[18] That is, every loss between 15000 and 25000 would be in band 20, instead of every loss between 10,000 and 20,000 being in band 20.

[19] Source : Credit Suisse (1997); "CreditRisk+ : A Credit Risk Management Framework", Credit Suisse Financial Products, Appendix A4, p36.