Economic Analysis Letters

Homepage: https://anser.press/index.php/EAL

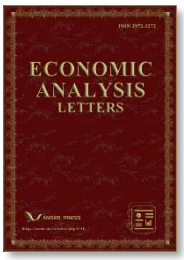

Reasonableness and Correctness for Operational Value-at-Risk

Peter Mitic [a, b, *]

[a] *Department of Computer Science, University College London, London, UK*

[b] *Risk Analytics, Santander Bank, London, UK*

ABSTRACT

Calculating the amount of regulatory capital to cover unexpected losses due to operational events in the upcoming year has caused problems because of difficulties in fitting probability distributions to data. It is consequently difficult to judge an appropriate level of capital that reflects the risk profile of a financial institution. We provide theoretical and empirical analyses to link the calculated capital to the sum of losses using appropriate statistical approximations. We conclude that, in order to reasonably reflect the associated risk, the capital should be approximately half the sum of losses, with a wide bound for the ratio of capital to sum.

KEYWORDS

Value at Risk; Central Limit Theorem; Loss Distribution; Loss Sum; Operational Risk; Regulatory Capital

* Corresponding author: Peter Mitic
E-mail address: p.mitic@ucl.ac.uk

ISSN 2972-3272
doi: 10.58567/eal02030005
This is an open-access article distributed under a CC BY license
(Creative Commons Attribution 4.0 International License)

Received 17 May 2023; Accepted 12 June 2023; Available online 24 June 2023

1. Introduction

Every year, banks and other financial institutions are required to assess their reserves for unexpected adverse financial events for the following year. That assessment covers financial risk (due to credit and market events), and also non-financial risk. Operational risk (*OpRisk*) constitutes the major part of non-financial risk, and is the subject of this paper. *OpRisk* can be loosely defined as the risk of financial systems failure, and typically might constitute approximately 10% of total capital reserves. Reserved monies for potential upcoming adverse events cannot be used for lending. Hence, it is important that the amount of the reserve should be consistent with the risk profile of the bank. If it is too high, lending will be impaired, and the knock-on effect on supporting economic activity may be significant. If it is too low, a bank may not be able to fund unexpected losses, and could, in extreme circumstances, fail. The recent collapse of Credit Suisse is a prime example, and has been widely reported in the financial press. It has been attributed to lax internal risk controls coupled with scandals involving its senior management, which resulted in a loss of confidence from customers, and withdrawal of funds (Stuart, 2023). *OpRisk* is regulated by national supervisory banks using *Basel* regulations (BIS, 2006) as a basis.

Particular problems are associated with calculating operational risk reserves. They are discussed below. The purpose of this paper is to provide a simple decision rule to determine whether or not a calculated operational risk reserve is consistent with the data used to calculate it. Informally, "Is the *OpRisk* value *correct*?". The proposal is to find an approximate theoretical relationship between *OpRisk* reserve value and the sum of *OpRisk* losses.

1.1. Glossary of banking technical terms used in this paper

In this paper, the technical terms can be divided into those referring to banking practice, and common statistical terminology. The banking terms are listed below. The statistical terms are listed in section (3.1).

- *Operational Risk* (*OpRisk*) is defined (BIS, 2006) as the risk of financial loss due to flawed or failed processes, policies, systems or events that disrupt business operations. It covers losses due to fraud, customer compensation, fines, physical damage and system outages.
- *Value-at-Risk* (*VaR*) is a common financial risk metric. *VaR* is the maximum acceptable loss over a given time horizon within a given probability bound. The regulations (BIS, 2006) specify that a fixed probability 0.999 should be used, and we state this as 99.9% *VaR*.
- *Capital* is the name given to the reserves that are held from year to year in order to be able to pay for *OpRisk* events that materialize. It can amount to billions of $US
- The most widely applicable metric for measuring *OpRisk VaR* is the *Loss Distribution Approach* (*LDA*), originally formulated by Frachot et al. (2001).

1.2. Problems associated with calculating OpRisk capital

The *LDA* is a Monte Carlo simulation, incorporating a convolution of frequency and severity distributions, both of which have to be fitted to the data. The problem lies with the severity distribution. In most cases it is possible to find, for any one data set, multiple statistical distributions that fit the data. The resulting capitals can differ widely. Consequently, which distribution to choose is not a simple matter. One choice is the distribution with the optimal fit. Another is a distribution that has fewer outliers. Very often the selected capital is one that corresponds closely to capitals calculated in previous years, thereby stabilizing capital year-on-year.

1.2.1. Illustration of wide variation on VaR

The following example illustrates the problem. A sample of size 500 was generated from a *LogNormal* distribution, resulting in data that has the 'fat-tailed' characteristics that are typical of operational risk data. Figure

(1) shows a histogram of the sample. A range of distributions were fitted to the data, and *VaR* was calculated for each. The results are shown in table (1), in decreasing order of fit suitability, The huge variation in *VaR* for the distributions shown is clear. It is a phenomenon that is very common in this context, and arises for two main reasons. First, distribution fitting can be imprecise. Iterative processes sometimes converge to non-optimal states, and small changes in parameter values can produce large changes in *VaR*. Second, sampling in the *LDA* Monte Carlo process can generate excessive sample values for some distributions, which are unrepresentative of the original data. Section (6) has a discussion of the *VaR* values for this case in the light of the model proposed in section (4). Some entries in table (1) are extraordinarily high (many billions – marked as "Reject: not credible"). They exceed the capitalization of most banks, and would be rejected immediately. However, there is, as yet, no objective criterion for doing so.

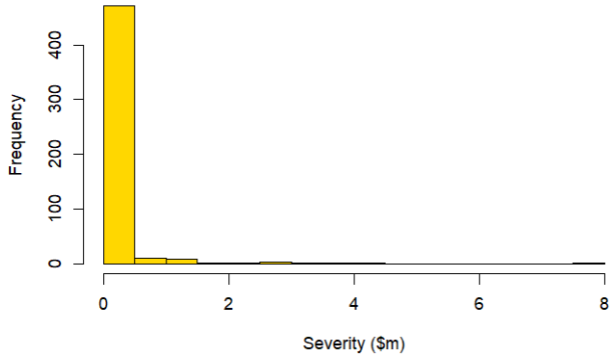

Figure 1. Histogram: typical fat-tailed operational risk data, generated from a *LogNormal*(10,2) distribution.

Table 1. Significance and *VaR* for candidate distributions, illustrating a wide range of *VaR* values.

Distribution	VaR ($m)	% Significance	Comment
LogNormal	133.9	0.63	
LogNormal Gamma Mix	134.5	0.63	
LogLogistic	1.72×10^{4}	0.71	Reject: not credible
LogGamma	686.2	0.94	
LogNormal Mix	30.4	1.07	
Generalized Pareto	2.34×10^{4}	1.21	Reject: not credible
Tukey G&H	8.2	1.77	
Weibull	34.3	2.68	
Gumbel	2.99×10^{4}	2.74	Reject: not credible
Gamma	3.04×10^{6}	2.96	Reject: not credible
Gamma Mix	27.3	3.18	
LogCauchy	1.21×10^{5}	3.19	Reject: not credible
Burr Type VII	7.33×10^{9}	5.18	Reject: not credible. *GoF* fail
Frechet	7.54×10^{6}	16.08	Reject: not credible. *GoF* fail

2. Literature Review

Basing *VaR* calculations directly on properties of the data has received little attention to date. Authors have concentrated, instead, on attempts to predict *OpRisk* losses using economic variates. Examples include Abdymomunov and Curti (2020), and Abdymomunov, Curti and Mihov (2020). In contrast, Berkowitz and O'Brian (2002) link *OpRisk* losses with corporate P&L. Chavez-Demoulin et al. (2016) have implemented an enhanced Generalized Additive Model (*GAM*) using economic co-variates as distribution hyper-parameters. Their model does use simple statistical properties of the data, such as *Mean*, *Standard Deviation, Median, Third Quartile, Maximum*, and *Skewness*. It is hard, with complex models such as GAM, to identify specific input features as drivers of output.

Specific properties of a set of largest losses, which are known to be principal determinants of *VaR*, are not considered (see Abdymomunov, Curti and Mihov, 2020). Two papers by Curti and Migueis (2016 and 2023) incorporate historic *loss frequen*cy and *mean total loss* into a quantile regression model. Their intention is to test whether future *VaR* can be predicted using those historic features values. They report that those features, with balance sheet co-variates such as *Market Capitalization*, are significant predictors of future losses.

Machine Learning has been applied in the context of *OpRisk*, but rarely using statistical properties of the data directly. Pena et al. (2021) use a convolutional neural network enhanced by fuzzy cognitive maps, which makes it extremely difficult to link *VaR* to any particular feature. Total loss severity and annual frequency are used. Chen and Wen (2010) investigated prediction of *OpRisk* losses using operational control variables, such as fraud prevention measures.

3. Theoretical Background

3.1. Statistical terminology

The following is a list of the principal statistical terms, expressed informally, used for the analysis in this paper. Full statistical definitions may be found in, for example, Weisstein (2023).

- *Distribution* (or *Statistical Distribution*) is a description of the relative numbers of times each possible outcome will occur in a number of trials. In particular, the Normal distribution is the "bell curve".
- *Quantile*: specific ranges within a distribution. For example, the 99th quantile is the value that exceeds 99% of all data values. In particular, the *Median* is the middle data value (i.e. the 50% quantile).
- *Order Statistic*: If data are ordered by size, the n^{th} *Order Statistic* is the n^{th} in the list.
- *Goodness-of-Fit* (*GoF*): A test to find the probability that a given result could not have occurred by chance.
- The *Central Limit Theorem* (*CLT*) states that if samples are drawn from a data set, the distribution of the mean of those samples approximates a *Normal* distribution as the sample size increases.
- *Maximum Likelihood* is a method for estimating distribution parameters by progressively improving parameter estimates, so as to maximize the probability that the data fits the statistical model.

3.2. The Loss Distribution Approach - Usage

Consider a set of M losses $L=\{x_1, x_2, \ldots, x_M\}$, assumed to be independently and identically distributed, and drawn from a fat-tailed severity distribution. Each loss is date-stamped. If the date range is t years, it is useful to use the annual loss frequency, $n = M/t$. A frequency distribution, usually *Poisson* or *Negative Binomial*, can also be assigned to L. The base *LDA* formulation is shown in Appendix A1. In its convolution stage, the single draw from the frequency distribution (step b(i)) is a random variable. That draw determines how many samples to draw from the severity distribution. The sum of draws from the severity distribution constitutes a single estimate of annual loss. After repeating multiple times (typically 1 million), the annual loss estimates can be ordered for easy extraction of the regulatory 99.9% quantile. In the analysis below, the variable single draw from the frequency distribution is replaced by a constant.

4. Central Limit Theorem *VaR* Approximation

The basis of the following analysis is the *CLT*. Its impact is to model sums of random variables by a single Normal distribution, from which a quantile can be extracted. In this analysis, it is assumed that the losses L are modelled by a *LogNormal* distribution with parameters μ and σ, denoted here by $LN(\mu, \sigma^2)$.

In the *Fixed Frequency LDA* variant (Mitic, Cooper and Bloxham, 2020) samples of fixed size and are drawn, and

the 99.9% quantile is extracted. using the "base" *LogNormal* distribution, it is possible to calculate, separately, expressions for the expected value of the sum of all *M* losses, and for the expected value for the extracted quantile, which is used as an approximation for *VaR*. For calculations using the data only, the resulting *VaR* value is the *Bootstrap VaR* (Efron and Tibshirani, 1994), and is expected to be less than the corresponding value calculated from the distribution fitted to the data (the *Distribution VaR*). This is because sampling from the data excludes sample points that are greater than the largest empirical datum.

The mean and variance, μ_X and σ_X^2 respectively, of the $LN(\mu, \sigma^2)$ distribution are:

$$\mu_X = e^{\mu+\frac{\sigma^2}{2}};\ \sigma_X^2 = \left(e^{2\mu+\sigma^2}\right)\left(e^{\sigma^2}-1\right)^{\frac{1}{2}} \tag{1}$$

The analysis in Mitic, Cooper and Bloxham (2020) proved that minimal error is introduced by replacing that random variable by a constant, the annual frequency, *n*. Step (b) in the LDA then reduces to drawing samples of size *n* from the severity distribution. Since the sample size is constant, the *CLT* applies. The result of applying the CLT for samples of size *n* drawn with replacement is a Normal distribution, S_n, representing the sum of *n* losses.

$$S_n \sim N(n\mu_X, n\sigma_X^2) \tag{2}$$

The 99.9% quantile of S_n now acts as an estimator of *VaR*, *R*, at 99.9% over a *t*-year horizon. From equations (1) and (2), with $z = 3.09$ for 99.9% confidence:

$$R = ne^{\mu+\frac{\sigma^2}{2}}\left(1 + z\sqrt{\frac{e^{\sigma^2}-1}{n}}\right) \tag{3}$$

Now consider the expected value of the *M* losses. That sum, *S*, can be modelled by a sum of *t* random variables, each of which is a sum of *n* $LN(\mu, \sigma^2)$ random variables, S_n, from equation (2). The distribution and expected value of *S* are given in equation (4).

$$S \sim N(tn\mu_X, tn\sigma_X^2);\ \mathbb{E}(S) = tne^{\mu+\frac{\sigma^2}{2}} \tag{4}$$

The ratio $\frac{R}{\mathbb{E}(S)}$ follows immediately (equation (5)).

$$\frac{R}{\mathbb{E}(S)} = \left(\frac{1}{t}\right)\left(1 + z\sqrt{\frac{e^{\sigma^2}-1}{n}}\right) \tag{5}$$

It is also easy to show that, theoretically, the sum of losses exceeds *VaR*. Equation (6) shows the difference between $\mathbb{E}(S)$ and R. The expression on the right-hand side is positive provided that $t > 1$. Typically, *t* would be at least 5, as most banks have an *OpRisk* loss history of at least 5 years. A further result, equation (7), gives an expression for a lower bound for *VaR* by considering the condition $R - \lambda\mathbb{E}(S) > 0$ for some number λ. The condition leads to an expression which shows that λ depends on the parameters *t, n, z* and σ. The dependency is examined in section (5).

$$\mathbb{E}(S) - R = ne^{\mu+\frac{\sigma^2}{2}}\left(t - 1 - z\sqrt{\frac{e^{\sigma^2}-1}{n}}\right) > 0 \tag{6}$$

$$R - \lambda \mathbb{E}(S) > 0 \Rightarrow \lambda < \left(\frac{1}{t}\right)\left(1 + z\sqrt{\frac{e^{\sigma^2} - 1}{n}}\right) < \left(\frac{1}{t}\right)\left(1 + z\frac{e^{\sigma}}{\sqrt{n}}\right) \quad (7)$$

4.1. Scaling

Numerical results for the preceding analysis were based on a 5-year time window. If the time window is not 5 years, a scaling adjustment is needed. A simple scaling formula can be derived for large *n*, which is the general case. With that assumption, equation (5) reduces to the expression in equation (8). If *t* is then replaced by $1/(c\ t)$, where *c* is a positive scale factor, the right-hand mapping applies. The implications are discussed in section (6).

$$\frac{R}{\mathbb{E}(S)} \sim \left(\frac{1}{t}\right) \rightarrow \left(\frac{1}{ct}\right) \quad (8)$$

5. Results

The main thrust of this paper is to develop the relationship between the expected loss sum and *VaR* in equation (5). The numerical results presented in this section serve as a sense check for that relationship.

5.1. Data and software

The loss data comprised approximately 500 data sets drawn from 'fat-tailed' distributions (listed in table 1), with a nominal time span t = 5 years. The overriding principle for data generation was random selection of distribution, distribution parameters and sample size, and then to randomly mix generated data. For each data set, a distribution was selected at random, with a bias towards distributions that we encounter most often: *LogNormal, LogNormal Mixtures, LogNormal-Gamma Mixtures, Weibull, Tukey G&H* and *LogGamma*. Distribution parameters were selected at random using uniform distributions based on appropriate parameter ranges. For example, the ranges for the *LogNormal* "μ" and "σ" parameters were [6, 15] and [0.5, 2.5] respectively. The sample sizes for all data sets were selected using uniform distributions with lower and upper bounds 100 and 500 respectively. In order to make each data set more realistic, pairs of data sets were randomly selected and merged, such that it would not be possible to definitively assign a distribution *a priori* to any given data set.

Each candidate distribution was fitted to each data set using *Maximum Likelihood*, and optimal *GoF* was determined using the *TNA GoF* test (Mitic, 2015). The *TNA* test is a formalization of a *QQ*-plot (a plot of predicted against actual empirical quantiles , and is robust with respect to sample sizes. For each data set, the loss sum was recorded, and 99.9% *VaR* was calculated using the *LDA*. All calculations were done using the *R* statistical language (https://www.r-project.org/).

5.2. Comparison of loss Sum and VaR

Figure (2) shows a scatter plot of the loss sum against 99.9% *VaR*. The plot shows a distinctive 'fan' shape with a concentration in the region *VaR* < 200. The gradient of the best linear fit (the blue dotted line) is approximately 0.48, which can be reasonably rounded to 0.5. The observed gradient is consistent with the theoretical value from equation (5) provided that parameter values are within typical ranges. Table (2) shows values of $\lambda = \frac{R}{\mathbb{E}(S)}$ for a typical range of *n* and σ. The highlighted values show regions that are broadly consistent with a range (0.3, 0.7), centered on 0.5. This 'consistent' region can be characterized by "small σ with small *n*, or medium σ with medium *n*, or large σ with large *n*". The condition $R < \mathbb{E}(S)$ applies in most other parts of table (2). When *s* is large, *R*

exceeds $\mathbb{E}(S)$. In those cases, the value of R should be questioned.

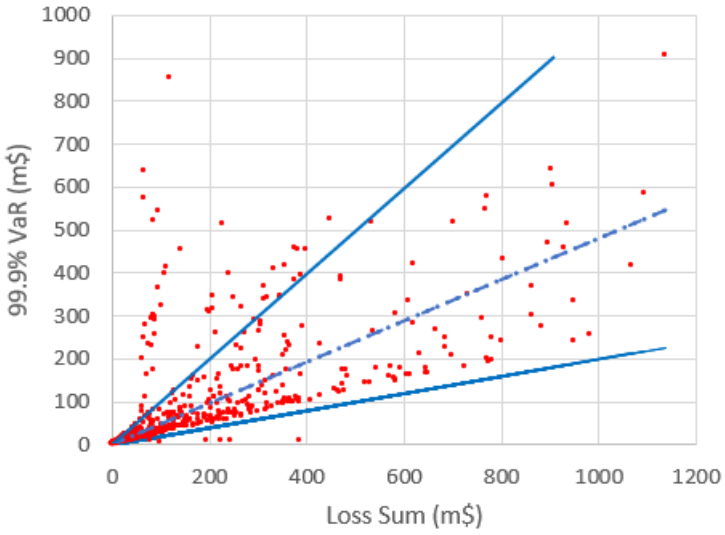

Figure 2. Loss Sum/VaR scatter with best linear fit, upper bound and lower bound.

Table 2. Variation of λ (equation 7) with *LogNormal* σ parameter, and Mean Annual Loss frequency, n.

	n	20	40	60	80	100	120	140	160	180	200
	0.5	0.27	0.25	0.24	0.24	0.23	0.23	0.23	0.23	0.22	0.22
	0.75	0.32	0.28	0.27	0.26	0.25	0.25	0.24	0.24	0.24	0.24
	1	0.38	0.32	0.30	0.29	0.28	0.27	0.27	0.26	0.26	0.26
	1.25	0.46	0.38	0.35	0.33	0.32	0.31	0.30	0.29	0.29	0.28
σ	1.5	0.59	0.48	0.43	0.40	0.37	0.36	0.35	0.34	0.33	0.32
	1.75	0.81	0.63	0.55	0.50	0.47	0.45	0.43	0.41	0.40	0.39
	2	1.18	0.89	0.77	0.69	0.64	0.60	0.57	0.55	0.53	0.51
	2.25	1.88	1.39	1.17	1.04	0.95	0.89	0.84	0.79	0.76	0.73
	2.5	3.25	2.36	1.96	1.73	1.56	1.45	1.35	1.28	1.22	1.16

Three features that are apparent in Figure (2) can be rationalized in terms of equations (5) and (7). They apply for typical annual frequencies, n~(50-250).

- The lower *VaR* bound (blue, solid) corresponds approximately to σ<1 with t = 1/5. It constitutes an absolute lower bound, λ=0.2 in this case, and applies for very large sample size (i.e. $n \to \infty$). Typically, λ~(0.2, 0.3) near this line. Very rarely, $\lambda < 1/t$ (in this case λ < 0.2). Effectively, $R = (1/t)\mathbb{E}(S)$ marks *minimum VaR*.
- The upper *VaR* bound (blue, solid) corresponds approximately to σ>2, with λ = 1. Ordinates with λ>1 are common, but indicate unreliable *VaR* estimations as λ increases.
- The mid-line (blue, dashed) is the best linear fit, and corresponds approximately to σ~(1,2), with a major concentration for λ~(0.3, 0.7)

5.3. Scale factor for other data windows

The 'consistency' range (0.3, 0.7), centered on 0.5 and highlighted in table (2), depends on having used a 5-year data window. For windows of other lengths, the table (2) entries would have to be scaled to account for an increased data sum. A scale factor is introduced in section (6.1) for this purpose.

6. Discussion

This work is an attempt to resolve the problem of which distribution and *VaR* value to select if many choices are available. The solution proposed is to relate *OpRisk VaR* to a simple function of the original data: the *Sum* of losses. The underlying intuition comes from known properties of the *LDA VaR* calculation: *VaR* increases as loss frequency and severity increase. Further, combinations such as *Large Var/Small Sum*, or *Small Var/ Large Sum* do not make sense in practice sense they would be inconsistent with the risk profile of the organization concerned. The theoretical treatment in section (4) confirms the empirical view.

6.1. Decision Rule

Given the empirical results of section (5), the following decision rule can be formulated. The results presented in section (5) rely on a 5-year data window. Therefore, a scale factor, *c*, has to be incorporated to account for other data windows. Using the approximation in equation (8), we can define a scaled ratio $\lambda' = c\lambda$ The 'base' case is *c* = 1, which corresponds to the 5-year window. The value of the scale factor is 2 for a 10-year time window, 3 for a 15-year time window etc. With empirical data, the expected loss sum, $\mathbb{E}(S)$, is the observed loss sum, *S,* so that $\lambda = \frac{R}{S}$ and $\lambda' = \frac{c\,R}{S}$. The decision rule takes the form of the choices in equation (9). It incorporates a *de minimis* value for cases where the calculated VaR is close to zero, based on the minimum VaR $R = (1/ct)\mathbb{E}(S)$ (case 1 in section (5.2) and equation (5) for large *n*). In practice, *n* needs to be approximately 100.

$$\begin{cases} Scaled\ Lambda & Applicable\ Range & Decision \\ \lambda' \sim 0.5 & (0.2-1) & \text{Accept: } R \text{ nearest to } \frac{S}{2} \\ \lambda' < 0.2 & 0-0.2 & \text{Reject: } VaR \text{ is too low} \\ \lambda' > 1 & 1-2 & \text{Reject: } VaR \text{ is too high} \\ \lambda' \sim 0 & 0 < \lambda < 0.01 & \text{Use minimum } R = \frac{S}{t} \end{cases} \quad (9)$$

Referring to the 5-year time window, the 'ideal' value of *VaR* is approximately 0.5 (i.e. $R \sim$ ½S). The value ½ should be relaxed within a wide range between 0.2 and 1. Values less than 0.2 can be rejected immediately. Values greater than 1 are harder to assess. In theory, VaR can be unbounded above, so some leeway in this case is permissible. Overall, the decision rule in equation (9) should be regarded as guidelines.

Looking again at the example in table (1), the decision rule does provide a clear guide. The sum of losses for the sample was \$75.59m, and the time horizon was 5 years. Setting $\lambda = 1/2$ sets a guide *VaR* value of \$37.80m. The nearest *VaR* with a *GoF* pass to that value is \$34.3m, using a *Weibull* distribution. The *LogNormal Mix* or *Gamma Mix* distributions also resulted in acceptable distributions. One of them might be selected if, for example, it produced fewer outliers. If $\lambda = \frac{1}{t} = 0.2$, a minimum *VaR* value is \$15.19m. The low *VaR* from the *Tukey G&H* distribution would therefore be rejected. All others would be rejected on the grounds that they are too high. Of those, the *LogNormal* and *LogNormal Gamma Mix* distributions are high, but credible. The *LogGamma* distribution is borderline credible (just below 10 times the loss sum – a somewhat arbitrary limit, obtained from observation in practice). The others are not credible, and can now be rejected on an objective basis. It should be noted that although the *LogNormal* distribution is the best fit, its resulting *VaR* is not optimal.

Two remaining problems are what to do if no distributions pass the *GoF* test, or if no distributions with a *GoF* pass satisfy the criteria specified in the equation (9) decision rule. In both cases, a 'solution' is to select the distribution that expresses the minimum *GoF* fail, and is nearest to the nominal "*Sum*/2" value. Such cases are rare.

It is nearly always possible to amend data thresholds such that the decision rule can be applied successfully.

6.2. Advice for Practitioners

Using the method in practice is very simple using the scaled ratio $\lambda' = \frac{c\,R}{S}$. *VaR* is calculated, and is multiplied by the data sum and the scale factor *c*. The result can then be compared with the numerical values in the decision rule in equation (9). For example, a recent *fraud* capital calculation for a medium-to-large European bank resulted in *VaR* ~ €28.2m, and a data sum of €88.4m, with a 10-year data window. In this case $c = 10/5 = 2$. The scaled ratio $2\times 28.2/88.4 = 0.64$ falls within the range (0.2, 1), and the measured *VaR* is deemed to be acceptable. Any measured *VaR* in excess of €44.2m would be rejected as excessive. The minimum acceptable *VaR* would be €8.84m.

6.3. Limitations of the method

The decision rule in equation (9) should be treated as a guide, rather than an absolute rule. It arises from approximations for the *LDA VaR*, and is formulated with a view to assessing whether or not a measured *VaR* is 'sensible'. If a measured *VaR* value is orders of magnitude greater than the data sum, it is clear that the severity distribution used was unrepresentative of the data. The decision rule is most useful if an empirical *VaR* value is 'large', but not 'too large'. Considering, again the example in section (6.2), the proposed upper 'reasonable' *VaR* value was stated as €44.2m. Had the measured *VaR* been €50m (i.e. approximately 10% above the proposed limit), it would have been accepted. Had the measured *VaR* been €100m (i.e. nearly four times the proposed limit), we would seek an alternative distribution and recalculate *VaR*.

A further limitation of the proposed decision rule is that it says nothing about an upper ceiling for *VaR*. In other words, how 'large', is 'too large'? This is a long-standing problem in *OpRisk*, and a paper on the subject is currently under at the review stage (Mitic, 2024). The approach for that study was similar to the approach taken for this paper, and the form of the result is also similar: "Ceiling ~ $7\frac{1}{3}\times$ Annual Sum of Losses". The multiplier $7\frac{1}{3}$ is fluid to a limited extent.

Funding Statement

This research received no external funding.

Conflict of interest

The author claims that the manuscript is completely original. The author declares no conflict of interest.

Appendix

A1. The Loss Distribution Approach (LDA) algorithm

The *LDA* calculates *VaR* (at 99.9%) using a Monte Carlo simulation with a convolution of frequency and severity distributions. Inputs: *M* losses spanning *t* years, and a severity distribution *D*. Output: 99.9% *VaR*.

a) Calculate an annual loss frequency, $n = M/t$.
b) Repeat *n* times.
    i) Obtain a sample size *Z* by drawing a random sample of size 1 from a frequency distribution (usually Poisson or Negative).

ii) Draw a sample of size Z, D_Z, from D.

iii) Sum the losses in D_Z to obtain $\sum D_Z$, the estimate of annual loss.

c) Calculate the 99.9th percentile of the retained sums $\sum D_Z$.

References

Abdymomunov, A. and Curti, F. (2020). Quantifying and stress testing operational risk with peer banks' data. *Journal of Financial Services Research*, 57, 287-313. https://doi.org/10.1007/s10693-019-00320-w

Abdymomunov, A., Curti, F., and Mihov, A. (2020). U.S. banking sector operational losses and the macroeconomic environment. *Journal of Money, Credit, and Banking*, 52(1), 115–144. https://doi.org/10.1111/jmcb.12661

Berkowitz, J. and O'Brian, J. (2002). How accurate are value-at-risk models at commercial banks? *Journal of Finance*, 57(3), 1093–1111. https://doi.org/10.1111/1540-6261.00455

BIS (2006). *Bank for International Settlements BCBS128 - International Convergence of Capital Measurement and Capital Standards.* https://www.bis.org/publ/bcbs128.htm

Chen,Q. and Wen,Y. (2010). A BP-Neural Network Predictor Model for Operational Risk Losses of Commercial Bank. *3rd Int. Symposium on Information Processing, Qingdao*, 291-295. https://doi.org/10.1109/ISIP.2010.43

Efron, B and Tibshirani, R.J. (1994) An Introduction to the Bootstrap, Chapman Hall

Frachot, A., Georges, P., and Roncalli, T. (2001). Loss distribution approach for operational risk. *Working paper, Groupe de Recherche Operationnelle, Credit Lyonnais, France.* http://ssrn.com/abstract=1032523

Mitic, P., Cooper, J. and Bloxham, N. (2020). Incremental Value-at-Risk. *Journal of Model Risk Validation*, 14(1), 1-37. https://doi.org/10.21314/JRMV.2020.216

Mitic, P. (2015). Improved Goodness-of-Fit tests for Operational Risk. *Journal of Operational Risk*, 15(1), 77-126. https://doi.org/10.21314/JOP.2015.159

Mitic, P. (2023). Credible Value-at-Risk. Submitted to the *Journal of Operational Risk*

Pena,A., Patino, A., Chiclana, F., Caraffini, F. and Congora, M. (2021). Fuzzy convolutional deep-learning model to estimate the operational risk capital using multi-source risk events. *Applied Soft Computing* 107(107381). https://www.sciencedirect.com/science/article/pii/S1568494621003045

Stuart, R. (2023) Why did Credit Suisse fail? *Economics Observatory* https://www.economicsobservatory.com/

Weisstein, E. W. (2023) Wolfram MathWorld. https://mathworld.wolfram.com