MONASH University

Australia

Department of Econometrics and Business Statistics

Non-parametric Estimation of Operational Risk and Expected Shortfall

Ainura Tursunalieva and Param Silvapulle

October 2013

Working Paper 23/13

Nonparametric Estimation of Operational Risk and Expected Shortfall*

Ainura Tursunalieva and Param Silvapulle[†]

Department of Econometrics and Business Statistics, Monash University, Australia

Abstract

This paper proposes improvements to advanced measurement approach (AMA) to estimating operational risks, and applies the improved methods to US business losses categorised into five business lines and three event types operational losses. The AMA involves, among others, modelling a loss severity distribution and estimating the Expected Loss and the 99.9% operational value-at-risk (OpVaR). These measures form a basis for calculating the levels of regulatory and economic capitals required to cover risks arising from operational losses. In this paper, Expected Loss and OpVaR are estimated consistently and efficiently by nonparametric methods, which use the large (tail) losses as primary inputs. In addition, the 95% intervals for the underlying true OpVaR are estimated by the weighted empirical likelihood method. As an alternate measure to OpVaR, the Expected Shortfall - a coherent risk- is also estimated. The empirical findings show that the interval estimates are asymmetric, with very large upper bounds, highlighting the extent of uncertainties associated with the 99.9% OpVaR point estimates. The Expected Shortfalls are invariably greater than the corresponding OpVaRs. The heavier the loss severity distribution the greater the difference between OpVaR and Expected Shortfall, from which we infer that the latter would provide the right level of capital to cover risks than would the former, particularly during crises.

Keywords: Heavy-tailed distribution, Loss severity distribution, Data tilting method, OpVaR, Expected shortfall
JEL Classification: C13, C14, C46

1 Introduction

The last two decades witnessed that globalization of international financial markets and the developments of e-commerce, out-sourcing, automated technology, large-scale mergers, and the like, increased the risk of running a business and encouraged fraud. As financial scandals have regularly surfaced with huge losses producing devastating economic and financial impacts on investors and consumers, *operational risk* has been receiving increasingly significant attention from the academics, practitioners and regulators. The operational risk is defined as the difference between the 99.9% quantile - operational-value-at-risk (OpVaR) - of a loss distribution and the Expected Loss (i.e. mean of the loss distribution). The Basel II Accord (Basel Committee on Banking Supervision 2004) on capital adequacy requires each international bank to develop statistical models of operational risk and use them to determine the capital to be held against operational risk. It is now well-accepted by the banking industry that this requirement is ongoing, as banks strive to improve their operational risk models.

The Basel II Accord requires banks to hold regulatory capital against operational risk, offering three approaches for this purpose: (i) the basic indicators approach (BIA), (ii) the standardised approach (STA), and (iii) the advanced measurement approach (AMA)[1]. While the first two are clearly simple approaches, the AMA requires explicit modelling of operational risk. It is also believed that the AMA should produce lower regulatory capital than that produced by either the BIA or the STA. Under AMA, loss distribution approach is employed for modelling operational losses from which operational risk is computed. It is evident that despite the widespread interest in and

---

[1](i) The BIA sets regulatory capital to 15% of average gross income, (ii) the STA sets regulatory capital to the sum over a number of business lines of the products of gross income and a multiplier (between 12% and 18%) per business line, and (iii) the AMA allows banks to base regulatory capital on bank-internal models.

the mushrooming of research on operational risk by academics and practitioners, Basel Committee for Banking Supervision (Basel Committee on Banking Supervision 2011b), is concerned that the required capitals calculated during crises are still inadequate, and has provided some guidelines for improvements to estimating operational risk, to which this paper pays considerable attention.

The primary objective of this paper is to propose some improvements to AMA for efficiently and consistently estimating both the Expected Loss and the 99.9% OpVaR of business operational losses. These two measures form a basis for computing the levels of economic and regulatory capitals to cover various risks arising from operational losses. The data used in our empirical investigation include five business lines and three event types of US business operational losses. The specific objectives are: (i) to estimate the Expected Loss and 99.9% OpVaR of a heavy-tailed loss distribution; (ii) to construct 95% interval estimates for the underlying true OpVaR and to evaluate the properties of these interval estimates in terms of their coverage probabilities and lengths; and (iii) to estimate the Expected Shortfall - a coherent risk measure - which is an alternative measure to OpVaR.

This paper makes principal contributions to AMA by proposing some improvements that are based on the recent developments in nonparametric statistics. The main contributors to this literature include Peng (2001), Hall & Yao (2003) and Peng & Qi (2006) who developed nonparametric methods for estimating the mean and high quantiles of heavy-tailed distributions. An attractive nature of these nonparametric methods is that they were developed without having to estimate either the overall loss distribution or the loss severity (tail) distribution. Only the large (tail) losses are used as primary inputs for computing the required statistics. Furthermore, this paper proposes a simple extension to the method proposed by Huisman et al. (2001) for consistently and efficiently

estimating the tail index. We extend this method to estimate simultaneously both the tail index and the optimum threshold loss, which determines the tail losses. Both of these quantities are critical for correct estimation the Expected Loss as well as OpVaR measures.

In addition, the three different methods that we use for constructing the 95% confidence interval estimates for the 99.9% OpVaR include the normal approximation based method, the empirical likelihood method and the data tilting method which is also known as the weighted empirical likelihood method. The last two are nonparametric methods that were developed specifically for estimating confidence intervals for high-quantiles of heavy-tailed distributions. The data tilting method appears to be well-suited for our applications involving heavy-tailed loss distributions, because it allocates zero weights to non-tail losses and positive and non-decreasing weights to losses in the tail region - the most important part of the loss distribution in estimating the 99.9% OpVaR and the other associated measures. These confidence interval estimates indicate the level of uncertainty associated with the point estimates of the OpVaR[2]. To date, to our knowledge, such interval estimates have not been provided for the true 99.9% OpVaR. In this paper, not only do we introduce these recent advances in nonparametric methods to the area of operational risk, but also provide detailed algorithms to facilitate their implementation; see Section 2.3 for details.

According to Basel III guidelines published in the latest consultative document (Basel Committee on Banking Supervision 2011b), identifying the practical challenges associated with the successful development, implementation and maintenance of an AMA framework is mandatory. Furthermore, it has been argued that a bank must be able to demonstrate that its approach captures potentially severe 'tail' loss events. Following

---

[2]In a limited simulation study conducted in this paper, the weighted empirical likelihood method turned out to be better than the other two methods; see Appendix A for details.

the global financial crisis, Basel Committee on Banking Supervision (2011a) has reported that the regulatory capital requirements that the banks provided under AMA were not sufficient enough to withstand losses during the crisis, due to inadequate modelling of large tail losses. However, for the purpose of calculating the required capital level, the 99.9% OpVaR is primarily used, and one problem associated with this measure is that the large losses over and above this OpVaR are not included in the calculation of the amount of regulatory capital. As a consequence, the regulatory and economic capitals can be vastly underestimated, particularly during crises which instigate large business losses. As said, this paper computes the Expected Shortfall, which will provide much greater regulatory and economic capitals to cover risks when mostly needed.

In quantifying operational risk measures, the AMA takes two dimensions into account. In dimension one, business environments and internal control factors are taken into account and conditional losses are calculated. Then, in dimension two, the loss severity distribution is fitted to conditional operational losses. We say at the outset that the unconditional operational losses were modelled in this paper. Addressing the issues related to dimension one is beyond the scope of this paper, because the dataset that we use contains only the irregularly spaced operational losses, and does not contain any firm specific details, business environments or internal control factors; see Section 3 for the description of the data. However, the approaches/methods that we use for estimating the Expected Loss and OpVaR of a severity loss distribution can be applied to conditional losses as well as unconditional losses. Moreover, Brown et al. (2008) addressed the mandatory disclosure as an important regulatory tool that allows market participants to assess manager risks without unnecessarily constraining manager actions. These and the similar issues raised by Brown et al. (2008) are not considered in this paper.

The distribution of operational losses is traditionally estimated by fitting empirical or

parametric extreme-value distributions. Several studies have employed semi-parametric methods, known as the peak-over-threshold (POT) method (Chavez-Demoulin et al. (2006)), in that a nonparametric distribution is fitted to the small to medium losses and a parametric distribution such as generalised Pareto distribution (GPD) to the large losses. However, these approaches have several limitations.

First, the Expected Loss (i.e., the mean) is estimated as the sample average of the losses or as the modified mean as in Johansson (2003) via the POT method. For a heavy-tailed loss distribution, these methods can lead to biased and inconsistent estimators of the Expected Loss. Furthermore, the GPD may not adequately fit the very large (tail) losses, and consequently the high quantile (OpVaR) can be underestimated. The Expected Loss can be correctly estimated by an improved method that exploits the tail properties of heavy-tailed loss distributions; see Peng (2001) for details.

Second, a serious problem which is associated with the parametric and semi-parametric approaches used widely in the literature is that if the underlying true loss severity distribution is unwittingly misspecified, then statistical inferences about the true Expected Loss as well as OpVaR are unreliable. On the other hand, in the nonparametric methods employed in this paper, neither the loss distribution nor the tail region needs to be estimated. The nonparametric methods that we study in this paper are sophisticated and not straightforward to implement them in empirical applications. We provide a step-by-step description of an application of the nonparametric methodologies for estimating the Expected Loss, OpVaR and its intervals; see Section 4.1 for details. This detailed description is designed to assist applied researchers and practitioners to overcome the complexities associated with the implementation of these nonparametric methodologies.

The remainder of the paper is organised as follows. Section 2 briefly discusses the methodologies and our improvements to AMA. Section 3 describes the operational loss

data. Section 4 reports and analyses the results of the empirical applications. Section 5 concludes the paper. Appendix A reports the results of a limited simulation study, which evaluates the performances of the three methods for constructing confidence interval estimates for the true 99.9% OpVaR of a heavy-tailed loss distribution. Appendix B provides the relevant graphs and tables.

2   Methodologies

As discussed in the introduction, the key variables that form a basis for calculating economic and regulatory capital requirements include the Expected Loss and the 99.9% OpVaR. However, these quantities are very sensitive to the threshold loss that separates the loss severity (tail) distribution with huge losses from the rest, as well as to the tail index that measures the heaviness of the loss distribution. Therefore, we pay considerable attention to estimate consistently and efficiently both the tail index and the threshold loss of a heavy-tailed loss distribution. In this section, we explain the methods for estimating both the tail index and the threshold loss, followed by the Expected Loss, OpVaR and Expected Shortfall. In addition, we explain three method for constructing 95% confidence interval estimates for the underlying high quantile - OpVaR.

2.1   Estimation of the tail index and the optimal threshold loss of a heavy-tailed distribution

We will first describe some basic concepts of the tail index, followed by the method proposed by Huisman et al. (2001) for estimating the tail index consistently and efficiently. Then, we will propose a simple extension to estimating the optimum threshold loss that determines the threshold length k - the number of losses that lie in the farthest right

tail region of the loss severity distribution.

The well-known and widely used nonparametric tail index estimator is the Hill's index (Hill 1975). This estimator is unstable for small values of k due to its large variance, and as it is biased for large values of k. Therefore, finding an optimal value of k is crucial to produce tail index estimator that is both unbiased and efficient. As a result, the Expected Loss (the mean) and the OpVaR (high quantile) of a heavy-tailed distribution can also be reliably estimated.

To define the Hill index, consider $X_{n,1} \leq X_{n,2} \leq \ldots \leq X_{n,n}$, which are the order statistics of an $i.i.d.$ random variables $X_1, ..., X_n$. For a given value of k, a selected threshold $X_{n,n-k}$ that splits the distribution into the non-tail and tail regions. Then, using k tail losses, define the Hill index estimator as:

$$\hat{\alpha}_k = \left(\frac{1}{k} \sum_{j=1}^{k} ln(X_{n,n-j+1}) - ln(X_{n,n-k}) \right)^{-1}. \qquad (1)$$

For the majority of heavy-tailed distributions, the Hill index is below two. When the tail index below two, the distribution of interest is in the domain of attraction of a stable law. On the other hand, when the tail index greater than two, the distribution of interest is in the domain of attraction of a normal distribution. The popularity of this Hill estimator is due to its computational simplicity and asymptotic unbiasedness[3]. In the context of a heavy-tailed distribution, an overestimated tail index leads to an underestimation of OR, which is an undesirable outcome for regulators.

To explain the tail index estimator that is both unbiased and efficient in small sam-

---

[3]The asymptotic normality of the Hill estimator has been established by many authors, such as Beirlant et al. (2004), and consistency was established by Mason (1982). However, its main disadvantage is that it requires a priori selection of k. In addition, it tends to overestimate the tail thickness when the true distribution is stable (McCulloch 1997).

ples, consider the model:

$$\alpha(k) = \beta_0 + \beta_1 k + \epsilon(k), \qquad (2)$$

where $\alpha(k)$ is the tail index, $k = 1, ..., \kappa$, $\kappa \approx n/2$, and n is the sample size. See Huisman et al. (2001) for details. The property of the Hill index estimator $\alpha(k)$ is exploited in the development of the new estimator, as the former is unbiased only when $k \to 0$. Based on this property, the estimator of β_0 is found to be an unbiased estimator of the tail index. However, the ordinary least squares estimate of (2) is inefficient, for two reasons: (i) the variance of $\alpha(k)$ is not constant for different values of k, and thus, the error term is heteroscedastic; and (ii) an overlapping data problem exists, due to the way in which $\alpha(k)$ is constructed. Huisman et al. (2001) suggest using a weighted least squares method to overcome these problems.

Multiplying equation (2) by the weight $w(k)$, we obtain

$$\alpha(k)w(k) = \beta_0^{wls}w(k) + \beta_1^{wls}kw(k) + u(k), \qquad (3)$$

where $w(k)$ are the weights in a vector form, defined as $W = (\sqrt{1}, \sqrt{2}, \ldots, \sqrt{\kappa})$, $\kappa = n/2$, and $u(k) = \epsilon(k)w(k)$. The asymptotic properties of the modified tail index estimator $\hat{\beta}_0^{wls}$ are discussed by Huisman et al. (2001). In small samples, this tail index estimator, obtained as the weighted average of Hill estimators, was found to be robust with respect to the choice of k. The tail index estimator $\hat{\beta}_0^{wls}$ can be obtained by OLS of equation (3).

We propose to select k by minimising the distance between the value of $\hat{\beta}_0^{wls}$ and the values of $\hat{\alpha}(k)$ for $k = 1, ..., n/2$. That is, for the optimal value of k, $|\hat{\alpha}(k) - \hat{\beta}_0^{wls}|$ for $k = 1, ..., n/2$ is minimised.

When the k and the corresponding $\hat{\alpha}(k)$ are estimated, one can ensure that these are

9

indeed estimated efficiently and consistently, by employing two goodness-of-fit methods, one is known as the mean excess function (MEF) proposed by Davison & Smith (1990). The MEF plot presents the sample mean excess values over the optimal threshold. An upward trend in the plot means that the data have a heavy-tailed distribution. For the optimum threshold k, the MEF plot is nearly linear.

The other popular tool for selecting the optimal threshold level is the Hill plot, $(k, \hat{\alpha}_k) : k = 2, ..., n$. This Hill plot is constructed using the Hill estimators defined in (1). The Hill estimators for a given range of values of k are plotted together with the 95% confidence intervals for the Hill indices. The optimal k is found, at which the Hill index plot is stable. For very small values of k, the Hill estimates are highly unstable with much wider confidence intervals.

We want to say at the outset that in the nonparametric methods that we discuss in the following sections, and we employ them for estimating the Expected Loss and the 99.9% OPVaR and its confidence intervals are invariably functions of the tail index and the tail length k. In the literature, the Hill estimator of the tail index and the resulting k are used in the applications of these methods. In this paper, they will be replaced with the WLS estimator $\hat{\beta}_0^{wls}$ and the optimal value of k estimated in this section.[4]

2.2   Estimation of the mean of a heavy right tailed distribution

For a heavy-tailed distribution, the sample average is not a consistent estimator of the population mean (see for example Peng 2001). To define his consistent estimator of the mean of a heavy-tailed distribution, let $X_1, ..., X_n$ be *i.i.d.* observations, with a common

---

[4]The double bootstrap method of Danielsson et al. (2001) was initially used in this paper to estimate the tail lengths for the simulated data from various heavy-tailed distributions. However, this method produced unstable estimates for the tail lengths, even in large samples, for example, a sample size of 1000. The results are not reported to save space, but available from the authors upon request.

distribution function F. We assume that F has regularly varying tails with tail index $-\alpha_k < -1$ and satisfies the following tail balance conditions:

$$\lim_{t\to\infty} \frac{1 - F(tx) + F(-tx)}{1 - F(t) + F(-t)} = x^{-\alpha_k} \quad \text{for} \quad x > 0, and$$

$$\lim_{t\to\infty} \frac{1 - F(t)}{1 - F(t) + F(-t)} = p \in [0, 1], \tag{4}$$

If $1 < \alpha_k < 2$, then equation (4) is equivalent to the statement that F is in the domain of attraction of a stable law. On the other hand, if $\alpha_k \geq 2$, then equation (4) implies that F is in the domain of attraction of a normal distribution. In order to estimate the mean, the distribution is partitioned into tail and non-tail regions. Peng (2001) proposed an estimator of the mean of an α_k-stable heavy-tailed distribution with $\alpha_k > 1$, which is the sum of the two mean estimators, of which one is for the tail region and the other is for the non-tail region. The population mean of a heavy-tailed distribution $E(X)$ is defined as follows:

$$E(X) = \int_0^1 F^-(u)du = \int_0^{1-k/n} F^-(u)du + \int_{1-k/n}^1 F^-(u)du, \tag{5}$$

where $F^-(s) := \inf\{x : F(x) \geq s\}$, $0 < s < 1$ denotes the inverse function of F. This method is employed to obtain a consistent and unbiased estimator of the mean of a heavy-tailed loss distribution, which in this paper is called the adjusted mean, μ_{adj}.

To implement this method, the data are decomposed into the tail and non-tail regions. Then, the adjusted mean is defined as:

$$\hat{\mu}_{adj} = \hat{\mu}_n^{(1)} + \hat{\mu}_n^{(2)}, \tag{6}$$

11

where $\hat{\mu}_n^{(1)}$ and $\hat{\mu}_n^{(2)}$ are the estimated means of the non-tail and tail regions, respectively. The component $\mu_n^{(1)}$ is the simple sample average of observations in the non-tail region of the distribution, while $\hat{\mu}_n^{(2)}$ is estimated using the tail observations as follows:

$$\hat{\mu}_n^{(2)} = \frac{k}{n} X_{n,n-k+1} \frac{\hat{\alpha}_k}{(\hat{\alpha}_k - 1)}, \tag{7}$$

where $\hat{\alpha}_k$ is the tail index estimator.

This way of computing the mean of a heavy-tailed distribution will provide a consistent and unbiased estimator of the Expected Loss, which is one of the key variables that form a basis for calculating the economic capital required to cover risks arising from operational losses.

2.3 Point and interval estimation of high quantiles of heavy-tailed distributions

In this section, both the 99.9% quantile (OpVaR) of a heavy-tailed distribution and the confidence intervals for this quantile are estimated using empirical likelihood method and data tilting method, which is in fact the weighted empirical likelihood method. For comparison purpose, the interval estimates based on normal approximation were reported. Furthermore, proposed in this section are algorithms which will facilitate the implementation of the first two methods.

The maximum likelihood method is one of the most frequently used methods for estimating quantiles of heavy-tailed distributions. Smith (1985) showed that, for $\alpha_k > -0.5$, the asymptotic properties of the maximum likelihood estimates, such as consistency, efficiency and normality, hold. In this paper, a high quantile (OpVaR) of a heavy-tailed distribution is estimated using the method proposed by Peng & Qi (2006). The confi-

dence intervals for OpVaR are constructed using the following methods: (i) the normal approximation method; (ii) the empirical likelihood method developed by Peng & Qi (2006); and (iii) the data tilting method developed by Hall & Yao (2003). The last two methods are largely nonparametric and therefore have the advantage that they allow the data to determine both the skewness of the loss severity distribution and its tail thickness.

To explain these methodologies, let the observations $X_1, X_2, \ldots, X_n$ be independent random variables with a distribution function F which satisfies:

$$1 - F(x) = L(x)x^{-\alpha_k} \quad \text{for} \quad x > 0, \tag{8}$$

where α_k is an unknown tail index and $L(x)$ is a slowly varying function. That is, $lim_{t\to\infty} L(tx)/L(t) = 1$ for all $x > 0$. Let $X_{n,1} \leq X_{n,2} \leq \ldots \leq X_{n,n}$ denote the order statistics of $X_1, \ldots, X_n$. Then, for a given small value of p_n, the $(100 - p_n)\%$ quantile x_p for the distribution F is defined as $x_p = (1 - F)^-(p_n)$, where $(\cdot)^-$ denotes the inverse function of $(\cdot)$. Assuming $L(x) = c$, let the equation (8) to have a simpler form,

$$1 - F(x) = cx^{-\alpha_k} \quad \text{for} \quad x > T. \tag{9}$$

Then, the likelihood function can be defined as:

$$L(\alpha_k, c_k) = \prod_{i=1}^{n} (c_k \alpha_k X_i^{-\alpha_k - 1})^{\delta_i} (1 - c_k X_{n,n-k}^{-\alpha_k})^{1-\delta_i}, \tag{10}$$

where $\delta_i = \mathbf{I}(X_i > X_{n,n-k})$, and $\mathbf{I}$ is an indicator function which defines the tail and non-tail losses, taking the values 0 and 1 respectively.

The estimators of the parameters c_k and α_k can be obtained by maximising the

likelihood function in equation (10).

Let $(\hat{\alpha}_k, \hat{c}_k)$ denote the maximum likelihood estimator of (α_k, c_k), and it can be derived as follows:

$$\hat{c}_k = \frac{k}{n} X_{n,n-k}^{\hat{\alpha}_k} \tag{11}$$

where $\hat{\alpha}_k$ is the WLS estimator $\hat{\beta}_0^{wls}$ obtained from (1). Using (9), a natural estimator for x_p is:

$$\hat{x}_p = \left(\frac{p_n}{\hat{c}_k}\right)^{-1/\hat{\alpha}_k}. \tag{12}$$

In the following subsections, three methods for constructing confidence intervals estimators for high quantiles of heavy-tailed distributions are briefly discussed.

2.3.1 Normal approximation method

For a given small value of p_n, the $100(1 - p_n)\%$ normal approximation based confidence interval estimator for x_p is defined as follows:

$$I_{x_p}^n = \left[\hat{x}_p exp\left(-z_\alpha log\left(\frac{k}{np_n}\right)\left(\hat{\alpha}_k\sqrt{k}\right)\right), \ \hat{x}_p exp\left(z_\alpha log\left(\frac{k}{np_n}\right)\left(\hat{\alpha}_k\sqrt{k}\right)\right)\right], \tag{13}$$

where $P(|N(0,1)| \leq z_\alpha) = \alpha$, k is the optimum tail length and $\hat{x}_p$ is the quantile estimator. This confidence interval estimator has asymptotically the correct coverage probability as $n \to \infty$. In a limited simulation study, Peng & Qi (2006) showed that the normal approximation based confidence interval estimator has an under-coverage problem in small to moderate sample sizes.

14

2.3.2 Empirical likelihood method

The empirical likelihood is a flexible nonparametric method for constructing a confidence interval estimator for high quantiles when the distribution of the quantile is unknown, as is often the case in practice. The empirical likelihood method that was introduced by (Owen 1988, 1990) combines the efficiency of the likelihood approach with the flexibility to incorporate the skewness of a heavy-tailed distribution, the shape of which is assumed to be unknown.

Let $\hat{\alpha}_k$ and $\hat{c}_k$ be defined as the WLS estimator $\hat{\beta}_0^{wls}$ obtained from (3) and $\hat{c}_k$ be defined as in equation (11). The empirical log-likelihood function as:

$$l_1 = max_{\alpha_k>0, c_k>0} logL(\alpha_k, \ c_k) = logL(\hat{\alpha}_k, \ \hat{c}_k), \tag{14}$$

where $\hat{c}_k$ and $\hat{\alpha}_k$ depend only on the data.

Then, maximise $logL(\alpha_k, c_k)$ subject to the constraints:

$$\alpha_k > 0, \quad c_k > 0, \quad \alpha_k logx_p + log\left(\frac{p_n}{c_k}\right) = 0, \tag{15}$$

where the value of p_n is small, such as 0.001, and x_p is the corresponding 99.9% quantile of the loss distribution. The explicit expression for $L(\alpha_k, \ c_k)$ is given as follows:

$$logL(\alpha_k, \ c_k) = \left(\frac{\alpha_k\sqrt{k}log(\hat{x}_p)x_p}{log(k/(np_n))}\right)^2. \tag{16}$$

Denote the maximised log-likelihood function by $l_2(x_p)$. It can then be shown that

$$l_2(x_p) = logL(\bar{\alpha}_k(\lambda), \bar{c}_k(\lambda)), \tag{17}$$

where

$$\bar{\alpha}_k(\lambda) = k \left(\sum_{i=1}^{k} (logX_{n,n-i+1} - logX_{n,n-k}) + \lambda logX_{n,n-k} - \lambda logx_p \right)^{-1} \qquad (18)$$

and

$$\bar{c}_k(\lambda) = X_{n,n-k}^{\bar{\alpha}(\lambda)} \frac{k - \lambda}{n - \lambda}, \qquad (19)$$

and λ satisfies the following two conditions:

$$\bar{\alpha}_k(\lambda)logx_p + log\left(\frac{p_n}{c_k(\lambda)}\right) = 0, \qquad (20)$$

$$\bar{\alpha}_k(\lambda) > 0 \quad \text{and} \quad \lambda < k. \qquad (21)$$

The log likelihood ratio multiplied by -2 takes the form:

$$l(x_p) = -2(l_2(x_p) - l_1). \qquad (22)$$

At the true quantile $x_{p,0}$, $l(x_p)$ has a χ^2 distribution with one degree of freedom, $l(x_{p,0}) \overset{d}{\to} \chi^2_{(1)}$, and the $100(1-\alpha)\%$ confidence interval estimator for x_p is $I_\alpha^{lr} = \{x_p : l(x_p) \le u_\alpha\}$, where u_α is the α-level critical point of a $\chi^2_{(1)}$ distribution. See Peng & Qi (2006) for the derivations of these results.

2.3.2a: An algorithm for constructing an interval estimator for a high quantile using the empirical likelihood method

Steps:

1. Estimate the tail length k and the corresponding threshold loss $X_{n,n-k}$, and then estimate the tail index as $\hat{\beta}_0^{wls}$ using the methods described in Section 2.1. Then

estimate the $100(1 - p_n)\%$ quantile as $\hat{x}_p$ using equation (12).

2. Find the range for λ from:

$$|\lambda| < \frac{k}{(\hat{\alpha}_k log\,(\hat{x}_p/X_{n,n-k}))^{-2}} . \qquad (23)$$

Denote this range by $\lambda_l < \lambda < \lambda_u$, where λ_l and λ_u are the lower and upper limits of λ respectively. If $\lambda_u \geq k$, then set the range to $\lambda_l < \lambda < k$.

3. Choose values for λ_i in this range for $i = 1, ..., N$, where N is large, say $5,000$.

4. For each value of λ_i, compute $\bar{\alpha}_k(\lambda_i)$ and $\bar{c}_k(\lambda_i)$ using equations (18) and (19), respectively. Discard values of $\bar{\alpha}_k(\lambda_i)$ and $\bar{c}_k(\lambda_i)$ which do not satisfy the positivity constraints in equation (15), for $i = 1, ..., N$.

5. Using equation (12), compute $\bar{x}_p$ for each pair of values of $\bar{\alpha}_k(\lambda_i)$ and $\bar{c}_k(\lambda_i)$, for $i = 1, ..., N$.

6. For $i = 1, ..., N$, calculate the likelihood function as follows:

$$l(\bar{x}_{pi}) = -2k\left(log(\psi) - (\psi - 1) - 2klog\left(1 - \frac{\lambda_i}{n}\right) + 2nlog\left(1 - \frac{\lambda_i}{n}\right)\right), \qquad (24)$$

where $\psi = \bar{\alpha}_k(\lambda_i)/\hat{\alpha}_k$.

7. The upper and lower bounds of the 95% confidence interval for the 99.9% high quantile x_p can be found as follows: (i) plot the profile likelihood function $l(\bar{x}_{pi})$ against $\bar{x}_{pi}$, for $i = 1, ..., N$; and (ii) the lower and upper bounds are found by minimizing the distance between the 5% critical value of the $\chi^2_{(1)}$ distribution and the profile likelihood function $l(\bar{x}_p)$. That is, the values of $\bar{x}_p$ that correspond to

17

this minimum distance constitute the empirical likelihood interval estimator for the true 99.9% high quantile x_p. The details of an empirical application of this method are provided in Section 4.1.

2.3.3 Data tilting method

The data tilting method is a weighted empirical likelihood method. One desirable property of the data tilting method is that it allocates small constant weights to losses in the non-tail region of the loss distribution, but large weights to the most important part of the distribution, namely the tail region, which contains only a small number of losses. Using the notations introduced by Peng & Qi (2006), we explain the data tilting method for confidence interval estimation of a high quantile.

Consider weights $q = (q_1, ..., q_n)$, such that $q_i \geq 0$ and $\sum_{i=1}^{n} q_i = 1$, and solve:

$$(\tilde{\alpha}_k(q), \tilde{c}_k(q)) = argmax_{(\alpha_k, c_k)} \sum_{i=1}^{n} log((c_k \alpha_k X_i^{-\alpha_k - 1})^{\delta_i} (1 - c_k X_{n,n-k}^{-\alpha_k})^{(1-\delta_i)}), \qquad (25)$$

where δ_i is defined in equation (10).

Solving equation (25) results in:

$$\tilde{\alpha}_k(q) = \frac{\sum_{i=1}^{n} q_i \delta_i}{\sum_{i=1}^{n} q_i \delta_i (log X_i - log X_{n,n-k})}, \qquad (26)$$

and

$$\tilde{c}_k(q) = X_{n,n-k}^{\tilde{\alpha}_k(q)} \sum_{i=1}^{n} q_i \delta_i. \qquad (27)$$

Now, define a function $D_p(q)$ which is the distance between q and a uniform distribution $q_i = n^{-1}$ as follows

$$D_p(q) = \sum_{i=1}^{n} q_i log(n q_i), \qquad (28)$$

18

where the weights $q_i(i = 1, ..., n)$ are subject to constraints: $q_i \geq 0$, $\sum_{i=1}^{n} q_i = 1$ and

$$\tilde{\alpha}_k(q)log\frac{x_p}{X_{n,n-k}} = log\frac{\sum_{i=1}^{n} q_i\delta_i}{p_n}.$$ (29)

This minimization procedure results in solving $(2n)^{-1}L(x_p) = min_q D_p(q)$, which gives the following two measures:

$$A_1(\lambda_1) = 1 - \frac{n-k}{n}e^{-1-\lambda_1}$$ (30)

and

$$A_2(\lambda_1) = A_{(\lambda_1)}\frac{log(x_p/X_{n,n-k})}{log(A_1(\lambda_1)/(p_n)}.$$ (31)

By the standard method of Lagrange multipliers, we have $q_i = q_i(\lambda_1, \lambda_2) =$

$$\begin{cases} n^{-1}exp(-1 - \lambda_1), & \text{if } \delta_i = 0 \\ n^{-1}exp\{-1 - \lambda_1 + \lambda_2(\varphi A_2^{-1}(\lambda_1) - A_1^{-1}(\lambda_1) - A_1(\lambda_1)\phi\varphi A_2^{-2}(\lambda_1))\} & \text{if } \delta_i = 1, \end{cases}$$ (32)

where $\varphi = log(\hat{x}_p X_{n,n-k}^{-1})$ and $\phi = log(X_j X_{n,n-k}^{-1})$ for $j = n-k, ..., n$, and both λ_1 and λ_2 satisfy

$$\sum_{i=1}^{n} q_i = 1, \quad \text{and} \quad \tilde{\alpha}_k(q)log\frac{x_p}{X_{n,n-k}} = log\frac{\sum_{i=1}^{n} q_i\delta_i}{p_n}.$$ (33)

According to Theorem 1 of Peng & Qi (2006), there exists a solution $(\tilde{\lambda}_1(x_p), \tilde{\lambda}_2(x_p))$, subject to the following inequalities:

$$-log\left(1 + \frac{\sqrt{k}\sqrt{\omega}}{n-k}\right) \leq 1 + \lambda_1 \leq -log\left(1 - \frac{\sqrt{k}\sqrt{\omega}}{n-k}\right),$$ (34)

19

and

$$|\lambda_2| \leq k^{-1/4} \frac{k}{n\omega}, \tag{35}$$

where $\omega = log(k(np_n)^{-1})$. $L(x_{p,0}) \xrightarrow{d} \chi^2_{(1)}$, with $(\lambda_1, \lambda_2) = (\tilde{\lambda}_1(x_{p,0}), \tilde{\lambda}_2(x_{p,0}))$ in the expression of $L(x_{p,0})$ defined in (16). Therefore, the $100(1 - \alpha)\%$ data tilting confidence interval estimator for x_p is $I_\alpha^{dt} = \{x_p : l(x_p) \leq u_\alpha\}$, where u_α is the α-level critical point of a $\chi^2_{(1)}$ distribution.

2.3.3a: An algorithm for constructing the confidence interval estimator for the quantile using thr data tilting

Steps:

1. Estimate the tail length k and the corresponding threshold loss $X_{n,n-k}$, then estimate the tail index as $\hat{\beta}_0^{wls}$ using the methods described in Section 4.1. Next, estimate the $100(1 - p_n)\%$ quantile as $\hat{x}_p$ using equation (12).

2. Using equations (34) and (35), compute the upper and lower bounds for λ_1 and λ_2, respectively. Choose the values of λ_{1i} and λ_{2i}, for $i = 1, ..., N$ (for a large value of N), within their respective ranges.

3. Compute $A_1(\lambda_{1i})$ and $A_2(\lambda_{1i})$ for $i = 1, ..., N$ using equations (30) and (31) respectively.

4. Using equation (32), calculate the weights q_i.

5. Compute $\tilde{\alpha}_k(q)$ and $\tilde{c}_k(q)$ using equations (26) and (27), respectively.

6. Using equation (12), compute $\tilde{x}_{pi}$ for each pair of $\tilde{\alpha}_k(\lambda_i)$ and $\tilde{c}_k(\lambda_i)$, for $i = 1, ..., N$.

7. Calculate the likelihood function $L(\tilde{x}_{pi})$ as follows:

$$L(\tilde{x}_{pi}) = \tilde{\alpha}_k^2(q_i)k \left[\frac{log(\tilde{x}_{pi}/x_p)}{log(k/np_n)} \right]^2. \tag{36}$$

8. The 95% data tilting method based confidence interval estimator for the true 99.9% x_p can be found by following Step 5 of Section 2.3.2a. An empirical application of this method is given in Section 4.1.

2.4   Estimation of Expected Shortfall

As discussed in the introduction, the OpVaR is a standard risk measure used for calculating capital requirements. In the context of operational risk, the OpVaR is commonly expressed as a high-quantile of a loss distribution at the small probability level, say α. For a given probability density function $p(X)$, and the observations $X_1, X_2, \ldots, X_n$, the $OpVaR_q$ is defined as follows:

$$OpVaR_\alpha(X) = inf\{x \in \mathbb{R} : P(X > x) \leq 1 - \alpha\}.$$

At the same probability level, the economic capital requirement is calculated based on the difference between the $(1 - \alpha)\%$ OpVaR and the Expected Loss (i.e. the mean of the loss distribution) which are estimated as $\hat{x}_p$ and $\hat{\mu}_{adj}$ defined in equations (12) and (6) respectively.

It is well known that quantile-based risk measures, such as OpVaR, may underestimate the true risk of operational losses, particularly when the losses are huge. The main reason for this is that such measures disregard losses beyond OpVaR. On the other hand, an alternative risk measure, the Expected Shortfall (ES), does consider losses beyond

21

OpVaR:

$$ES_\alpha = E\{X_i | X_i > OpVaR\}, \tag{37}$$

where OpVaR is the 99.9% quantile. The Expected Shortfall (ES) measure is therefore more sensitive to the shape of the tail region of the loss distribution than OpVaR, and thus, ES is a more appropriate measure for computing the level of capital required to cover risks.

The main contributor to the development of ES as a coherent risk measure is Acerbi & Tasche (2001). To briefly discuss ES, consider a set V of real-valued random variables. In order to be called a coherent risk measure, a function $\rho : V \to \mathbb{R}$ must satisfy the following four conditions:

1. monotonous: $X \in V, X \geq 0 \implies \rho(X) \leq 0$,

2. sub-additive: $X, Y, X + Y \in V \implies \rho(X + Y) \leq \rho(X) + \rho(Y)$,

3. positively homogeneous: $X \in V, hX \in V, \implies \rho(h, X) = h\rho(X)$,

4. translation invariant: $X \in V, a \in \mathbb{R}, \implies \rho(X + a) = \rho(X) - a$.

The sub-additivity axiom (#2) does not hold for OpVaR, and therefore it is not a coherent risk measure (Acerbi & Tasche 2001). On the other hand, ES does satisfy the above four conditions, including that of sub-additivity, and therefore it is a coherent risk measure. Sub-additivity is a desirable property when calculating the capital adequacy requirements, as it promotes diversification. This means that if a risk measure is computed for a bank that has a number of branches, the sum of the individual ES values will constitute an upper bound for ES of the bank. Since OpVaR is calculated merely as a measure of the quantile of a loss distribution, it does not have this property.

2.5 Bootstrap confidence intervals for the tail index of heavy-tailed distribution

Here, the bootstrap methodology proposed by Caers & Maes (1998) for constructing confidence intervals for the parameters of heavy-tailed distributions is briefly explained. The main idea of this bootstrap method is that re-sampling is performed from the semi-parametric distribution $\hat{F}_s$, which consists of the empirical distribution below the threshold value, and Generalised Pareto Distribution (GPD) above the threshold value. That is, $\hat{F}_s$ is defined as:

$$\hat{F}_s(x|u) = \begin{cases} (1 - \hat{F}(u))\ (1 - (1 + \frac{\alpha_k}{\sigma}(x-u))^{-1/\alpha_k}) + \hat{F}(u) & x > u \\ \\ \hat{F}(x) & x \leq u. \end{cases} \tag{38}$$

The sizes of the sub-samples for heavy-tailed distributions should be smaller than the total sample size. An extreme value statistic of interest is the tail index estimator $\hat{\alpha}_{k,b}$ which is computed for each sub-sample, where $b = 1, ..., B$ and B is the number of sub-samples. The set of estimates $\hat{\alpha}_{k,b}$ constitutes the sampling distribution for the tail index estimator. A simple percentile method (Efron & Tibshirani 1993) is then used to obtain confidence intervals for the tail index parameter from the following ordered B tail index parameter estimates:

$$\hat{\alpha}_{k,1}(u_n) \leq \hat{\alpha}_{k,2}(u_n) \leq ... \leq \hat{\alpha}_{k,B}(u_n), \tag{39}$$

where u_n is the optimal threshold. This bootstrapping method is employed for computing confidence intervals for the tail index estimates for eight types of operational losses in Section 4..

3   Description of operational loss data

The methods discussed earlier in this paper are now applied to operational losses in excess of $1M incurred by some large bank/business losses in the United States. That is, $1M is the minimum loss - the losses above which were collected for the purposes of empirical modelling operational risks. The data consists of irregularly-spaced losses covering the period May 1984 to May 2007. Due to these operational losses that we use in this study occurred over the period 1983 to 2007, we computed the real losses at the January 1997 consumer price index level in order to avoid the influence of inflation on the results. As a consequence, the real losses incurred prior to 1997 are higher than the nominal losses, whereas the losses incurred post 1997 are smaller than the nominal losses[5]. Furthermore, for comparison purpose, the methodologies discussed in Section 2 were applied to both nominal and real operational losses.

The losses were pooled across the various sectors and categorised according to the Basel business lines and event types. The losses are assumed to be *i.i.d.*[6]. The pie charts in Figure 1 and Figure 2 show the breakdown of losses by business lines and event types respectively. Two business lines—the *Trading and Sales* and *Commercial Banking* losses—account for approximately the same proportions of losses, 29% and 28% respectively. *Internal Fraud* losses constitute by far the highest proportion (more than 80%) of total value of losses, classified by event type.

To highlight the nature of the various sectors which are included in the operational loss data, a breakdown of *Employment PWS* losses by various business sectors is pro-

---

[5]As said in the introduction, the business environment variables are not available, and as a consequence loss severity distributions were fitted to unconditional losses in this paper. However, the same approaches/methods can also be applied to conditional losses, should these environmental variables become available.

[6]As more data become available, operational risks can be estimated separately for each business sector. Moreover, when a sufficient amount of data is available, one can construct quarterly or even monthly loss series and study the time series structure of these losses before estimating OR.

vided in the pie chart in Figure 3. The *Employment PWS* losses constitute the losses incurred by the following sectors: "Automotive", "Corporation", "Energy Corporation", "Insurance Company", "Agency", "Full Service Bank", "Retail", "Corporate Conglomerate", "Brokerage", "Software Corporation", and other sectors. It appears that the largest proportion of *Employment PWS* losses occurred in the "Automotive" sector.

Table 4 reports some descriptive statistics and the sizes of the losses at the 50th, 75th, and 95th percentiles of empirical distributions of five business lines and three event types operational losses. The percentiles of the real losses are slightly lower than (or close to) those of the nominal losses, with an exception that the 95th percentile of real *External Fraud* event-type losses is notably higher than that of nominal losses. It appears that some of the large losses occurred prior 1997. For all types of losses, the mean is much larger than the median, indicating that the loss distributions are heavily skewed to the right. The skewness for all types of losses is above 5.0. The standard deviation ranges from 66 for the *External Fraud* losses to 443 for the *Trading and Sales* losses.

A comparison of the sizes of the losses at different percentiles demonstrates the following: (i) *External Fraud* losses tend to be much smaller than other losses, as is evident from their lower standard deviation and smaller maximum loss size; and (ii) the *Trading and Sales* losses tend to be larger than other losses, as is evident from their larger standard deviation and larger maximum loss size. The descriptive statistics indicate the varying nature of the heavy-tailed distributions that would fit across the five business lines and three event types operational losses that we study in this paper.

4   Empirical applications and results

In this section, the methodologies discussed in Section 2 are applied to five business lines
and three event type operational losses in order to estimate the Expected Loss, 99.9%
OpVaR and Expected Shortfall. In addition, three methods are used to construct interval
estimates for the true OpVaR. As said, the application of some of these methods is not
straightforward and involves many stages. A step-by-step description of the estimation
stages is provided for the business line *Employment PWS* losses. Then, the overall
results are briefly discussed.

4.1   An application to *Employment PWS* losses

The Expected Loss (i.e. the adjusted mean denoted by μ_{adj}) and the 99.9% OpVaR (high-
quantile) of the loss distribution for *Employment PWS*, and the Expected Shortfall are
estimated. Furthermore, 95% confidence intervals are constructed for the true 99.9%
OpVaR using three methods.

   The time series plot of the *Employment PWS* nominal losses is presented in Figure
4. It is evident that this business line incurred many large losses during the period
1984-2007, with the largest loss occurring prior to 1997. This largest loss makes up
about 16% of the total loss. The number of losses seems to have increased since 1998.
The probability distribution function for the *Employment PWS* losses plotted in Figure
5 indicates that this distribution is heavily skewed to the right.

   In what follows, a step-by-step approach to application of methodologies explained
in Section 2 is provided.

   Steps:

   1. Estimate the tail index of the distribution of *Employment PWS* nominal losses as

follows: (i) compute the Hill estimate $\hat{\alpha}_k$ defined in (1) for $k = 1, ..., n/2$ $(n = 237)$; and (ii) estimate the tail index as $\hat{\beta}_0^{wls} = 1.3$ using (3).

2. Find the optimal tail length k that minimizes the distance between the $\hat{\beta}_0^{wls} = 1.3$ and $\hat{\alpha}_k$ for $k = 1, ..., n/2$. Such a k is found to be 21 and the corresponding threshold loss of \$70.4M. This means there are 21 losses that exceed this threshold loss.

3. To ensure that the optimum value for k has been found, MEFs for $k = 0$, 10 and 21 are plotted in Figure 6. Clearly, the first two are far from linear. The third MEF of the loss data excluding the optimum number of largest losses, which is 21, is approximately linear, indicating that the k is correctly estimated[7].

4. In order to assess the accuracy of the tail index estimate, construct the 95% bootstrap confidence interval for the true tail index as the percentiles of the bootstrapped sampling distribution discussed in Figure 7. The 95% interval estimate for the tail index is [1.1 1.9], which is asymmetric and it indicates some degree of uncertainty associated with the point estimate of the tail index, as the upper bound is relatively large.

5. Estimate the adjusted mean, defined in (6): $\hat{\mu}_{adj} = \$41.8M$, which is the sum of the sample average of non-tail losses, $\hat{\mu}_n^{(1)} = \$11.1M$, and the mean of the right tail losses computed using (7), $\hat{\mu}_n^{(2)} = \$30.7M$.

6. Estimate the 99.9% OpVaR of the loss distribution, which is defined in (36):

---

[7]The double bootstrap method of Danielsson et al. (2001) was initially used in this paper to estimate the tail lengths for the simulated data from various heavy-tailed distributions. However, this method produced unstable estimates for the tail lengths, even in large samples, for example, a sample size of 1000. The results are not reported to save space, but available from the authors upon request. See also footnote 7 to our paper.

$\hat{x}_{p_n=0.001}$ =\$2,475M. From this amount, the regulatory capital is calculated.

7. Find the difference between the 99.9% OpVaR and the $\hat{\mu}_{adj}$ as 2,433M. From this amount the required economic capital is calculated.

8. Estimate the Expected Shortfall defined in (37), $\widehat{ES}$ =\$3,571M.

9. For the true 99.9% OpVaR, the 95% interval estimate (based on normal approximation) defined in (13), is [\$538M   \$11,040M].

10. Using the algorithm provided in Section 6.2.3a, the 95% empirical likelihood confidence interval for the 99.9% OpVaR is estimated as follows:

    The estimated empirical likelihood function $l(x_p)$, which is defined in (24), is first estimated, then plotted in Figure 8 against a possible range of values of quantiles. To construct a 95% confidence interval estimate: (i) draw a horizontal line $\chi^2_{(1)} =$ 3.84 (5% critical value) through the profile log-likelihood curve; (ii) draw vertical lines through the two points of intersection, A and B; and (iii) take the points at which these two vertical lines intersect the x_p axis as the lower and upper limits of the confidence interval estimate for x_p. Now, it is easy to deduce that the empirical likelihood confidence interval estimate for the quantile (OpVaR) is [\$1,615M \$4,567M].

11. The 95% data tilting confidence interval estimate for the 99.9% OpVaR is estimated using the algorithm provided in Section 6.3, as follows. The weights for the likelihood function are computed using (32). It is worth noting that the weights allocated to non-tail losses are very small and constant, while losses in the important part of the loss distribution, the tail, are allocated higher weights. The larger the losses in the tail, the higher the weights allocated to them. The tail weights

are plotted in Figure 9, and the weighted empirical likelihood function is plotted in Figure 10. Adopting the previous step, the data tilting confidence interval estimate for the true 99.9% OpVaR is found to be [$1,583M $3,876M], which is much narrower than the empirical likelihood interval estimate.

In summary, in the above analysis, for the business line *Employment PWS* nominal losses, the point estimate of 99.9% OpVaR is $2,475M and its 95% data tilting interval estimate is [$1,583M $3,876M]. The estimate of Expected Shortfall is $3,571M, which also lie within this confidence interval estimate. The reason for preferring the data tilting method for estimating confidence intervals is that, in a simulation study, this method was found to have the correct coverage probabilities and narrower interval estimates than its counterparts. Although the normal approximation confidence interval estimate is easy to compute, its coverage probabilities were found to be lower than the nominal probabilities. However, for small samples such as 30 data points or fewer, the normal approximation confidence interval is the only interval estimate that can be constructed; see Appendix A for the results of this simulation study.

For the real operational losses for the business line *Employment PWS*: (i) the estimate of the tail index is 1.2, which is smaller than for the nominal losses, indicating that severity loss distribution of real losses are heavier than that of nominal losses; (ii) the Expected Loss is almost the same as that for nominal losses; and (iii) the 99.9% OpVaR and the Expected Shortfall are $2,504M$ and $3,641M$ are respectively, larger than those correspond to nominal losses. The data-tilting 95% CI estimate of OpVaR is [$1,681M $3,874M] which is much narrower for the real losses than for the nominal losses. A comparison of the results for both nominal and real operational losses for the business line *Employment PWS* indicates that this business line incurred heavy losses prior to 1997 at which the real losses are calculated.

4.2 Discussion of the empirical results

Tables 5 and 6 present the nonparametric estimates of the 99.9% OpVaR of loss severity distributions, and their 95% confidence intervals, as well as the estimates of the Expected Shortfall for all five business lines and three event types losses. In all cases, Expected Shortfall estimates lie within the 95% interval estimates of OpVaR. The estimates of the tail length k for all losses are reported in column 2 of Table 5. The point and interval estimates for the tail index are reported in columns 3 and 4 of Table 5, respectively. All of the tail index estimates are below two, with many being 1.2, indicating that all of the loss distributions have extremely heavy right tails. The bootstrap confidence interval estimates for the true tail indices of all losses indicate that even the upper bounds are below two.

Column 5 of Table 5 reports the estimates of the adjusted mean. This column and Figure 11 show that the adjusted mean estimates exceed the sample means, except for *Retail Banking* losses, for which the sample mean is greater than the adjusted mean. The difference between these two means is the smallest (4%) for the *Retail Brokerage* losses, and the largest (30%) for the *External Fraud* losses. The heavier the tails the greater the difference between the sample mean and the nonparametric estimator of the mean - the adjusted mean.

The estimates of the Operational Risk (defined as the difference between OpVaR and the Expected Loss), and Expected Shortfall are reported in Table 6. The OpVaR estimates (Table 5) vary from \$2,128M for the *Retail Brokerage* losses to \$7,149M for the *Internal Fraud* losses. The Operational Risk estimates vary from \$2,196M for the *External Fraud* losses to \$7,032M for the *Internal Fraud* losses. The Expected Shortfall estimates are invariably greater than their OpVaR counterparts. We found that the heavier the loss severity distribution the greater this difference, and this difference ranges

from a minimum of 22% for *Asset Management* losses to a maximum of 37% for *Trading and Sales* losses.

Three sets of 95% confidence interval estimates for each OpVaR are reported in Table 6. The normal approximation based confidence interval estimates are very wide, with huge upper bounds. The empirical likelihood confidence interval estimates and data tilting confidence interval estimates are much narrower than those of the normal approximation confidence intervals. However, the data tilting method provides the most reliable confidence interval estimates in terms of their width and coverage probability; see the Appendix A for details.

On the analysis of real operational losses, the estimated quantities, discussed in the previous two paragraphs, are in general smaller (or the same) than those for the nominal losses. However, for the *Commercial Banking* and *Asset Management* real losses, tail indexes are much smaller than those for nominal losses, indicating heavier loss severity distributions for the real losses. As a result, the OpVaR of these two type business-lines are much larger for real losses than for nominal losses.

5   Conclusion

This paper proposes improvements to advanced measurement approach (AMA) to estimating operational risks, and applies the improved methods to US business losses. The AMA involves, among others, modelling a loss severity distribution and estimating the Expected Loss and the 99.9% operational value-at-risk (OpVaR). These measures form a basis for calculating the levels of regulatory and economic capitals required to cover risks arising from operational losses. In this paper, Expected Losses and 99.9% OpVaRs of five business lines and three event types operational losses were estimated

consistently and efficiently by nonparametric methods. An attractive feature of these methods is that there is no need to estimate neither the overall loss distribution nor the tail distribution. These methods use only the (large) tail losses as primary inputs. Additionally, this paper also adapts the approach proposed by Huisman et al. (2001) for estimating the tail index efficiently and consistently, and extended this approach to estimating the optimum tail loss that determines the number of tail losses - the tail length. These quantities are crucial for correctly estimating the Expected Loss and the 99.9% OpVaR based on which capitals required to cover risks are calculated.

In a limited simulation study, the weighted empirical likelihood (data tilting) 95% interval estimates were found to have better properties than the traditional empirical likelihood method and normal approximation interval estimates. In particular, the data tilting method provides the narrowest interval in comparison to other two methods. Furthermore, an attractiveness of this method in this specific context is that it allocates zero weights to the central losses and positive and non-decreasing weights to tail losses - the most important part of losses in estimating OpVaR.

The estimates of the 99.9% OpVaR of the five business lines and three event types of losses range from \$2,092M to \$7,032M, while those of Expected Shortfall range from \$3,234M to \$10,920M. The *Internal Fraud* losses have the largest 99.9% OpVaR, while the *Trading and Sales* losses have the largest Expected Shortfall estimates. Moreover, for all eight business line/event types of losses, the estimates of Expected Shortfall lie within the 95% interval estimates of 99.9% OpVaR. The weighted empirical likelihood (data tilting) 95% interval estimates for the underlying true 99.9% OpVaR were the narrowest in comparison to empirical likelihood and normal approximation interval estimates. A noteworthy observation is that the simulation evidence presented in the Appendix A is corroborated with the empirical evidence presented in this paper. However, the

95% confidence interval estimates for the true 99.9% OpVaR are asymmetric, with very large upper limits, indicating that there is a high degree of uncertainty associated with the point estimates of the OpVaR. We also estimated Expected Shortfall, which are invariably greater than the OpVaR.

The results of this investigation indicate that the regulatory capital charges computed based on the point estimates of the 99.9% OpVaR may not be adequate to cover risks arising from extremely large operational losses. The empirical findings show that the heavier the loss severity distribution the greater the difference between OpVaR and Expected Shortfall, from which we infer that the latter would provide the right level of capitals required to cover risks than would the former, particularly during crises.

Appendix A

Coverage probabilities of interval estimates of the 99.9% quantile of heavy-tailed distributions: Simulation evidence

Frequently, operational losses are assumed to follow a normal distribution. However, for positively skewed loss data, nonparametric methods have a clear advantage in obtaining reliable point and interval estimates for high-quantiles of heavy-tailed loss severity distributions; see (Peng & Qi 2006). Sections 2.3 discussed the point estimator as well as the three interval estimators of the high quantiles. The purpose of this section is to study the coverage properties of the two nonparametric interval estimators, namely empirical likelihood ratio method and the data tilting method. The simple normal approximation method is included as a benchmark for comparison purpose.

A1: The design of the simulation study

The samples are generated from two heavy-tailed distributions using a range of sample sizes and tail indices. First, 10,000 random samples of size n are generated from the following two distributions:

1. Fréchet(α_k): $F(x) = exp\left\{-x^{-\alpha_k}\right\}$,

2. Burr$(\alpha_k, \alpha_k + \beta)$: $F(x) = 1 - (1 - x^{\beta - \alpha_k})^{\frac{-\alpha_k}{\beta - \alpha_k}}$ for $x > 0$, where $\beta = \alpha_k$.

The following sample sizes and tail indices were used for both distributions:

1. $n = \{200, 300, 500\}$.

   The tail length was selected using the method described in Section 2.1 for each sample size as the following:

   (a) for $n = 200$, the tail length $k = 16$

   (b) for $n = 300$, the tail length $k = 20$

(c) for $n = 500$, the tail length $k = 26$.

2. The tail index α_k is selected from the range of $[1.2, 1.5]$. To ensure that the simulated samples have the correct tail index, the tail index was computed for each of the simulated samples for a given value of k. All of the results reported are for a small tail probability $p_n = 0.001$.

The estimates of the confidence intervals for the 99.9% quantile, which is known as the operational-value-at-risk (OpVaR), are evaluated in terms of their coverage probabilities. The observed coverage probability is calculated as the proportion of confidence interval estimates that cover the true quantile. The true quantile was obtained by taking the 99.9^{th} percentile of one million observations for both distributions and for each tail index. The expected mean length is the sum of the lengths of 95% confidence interval estimates which cover the true quantile divided by the total number of such interval estimates. Furthermore, in order to asses how well the interval estimates capture the positive skewness of the underlying heavy-tailed distribution, the right non-coverage rate is also computed. It shows the proportion of intervals that lie to the left of the true quantile. Using these outcomes of this limited simulation experiment, recommendations are made as to the most appropriate method for constructing interval estimates for the high quantiles of a heavy-tailed distribution. Such an interval estimate would indicate the uncertainty associated with the point estimate of the true 99.9% quantile.

A2: The simulation results

Here, the results of the simulation study conducted to assess the accuracy of the confidence interval estimator of the true (high) 99.9% quantile are reported in Table 1. This table presents the values of the 99.9% population quantile (true quantile). Clearly, the closer the tail index to 1.0, the heavier the loss distribution. Thus, it is noticeable

that the true (high) quantile increases as the tail index decreases. The values of the true quantiles of the Burr distribution are much lower than those corresponding to the Fréchet distribution when the tail index is 1.5. On the other hand, when the tail index is 1.2, the quantiles are very high, and they appear to be close to each other.

Table 1: True 99.9% quantiles for the generated distributions

Distribution and Tail index	1.5	1.2
$Fréchet(\alpha_k)$ $distribution$	149.5	379.1
$Burr(\alpha_k, \alpha_k + \beta)$ $distribution$	133.8	369.8

Tables 2 and 3 present the coverage probabilities and expected lengths of the confidence interval estimates for the two generated distributions. While the normal approximation method does provide the best coverage probabilities, the average interval length is very long. Overall, the data tilting method yields higher coverage probabilities and shorter interval lengths than the empirical likelihood method. Based on the evidence from our simulation study, we recommend the use of the data tilting method for estimating confidence intervals for the 99.9% quantile of a heavy-tailed distribution.

Table 2: Coverage probabilities of 95% confidence intervals for the 99.9% quantile

Methods	Normal approximation		Empirical likelihood		Data tilting	
α_k	1.5	1.2	1.5	1.2	1.5	1.2
n			$Fréchet(\alpha_k)$ distribution			
500	90.1	89.2	63.7	62.4	70.3	69.1
300	89.9	88.1	57.7	53.1	60.1	59.1
200	88.1	87.7	47.1	46.2	54.9	53.1
n			$Burr(\alpha_k, \alpha_k + \beta)$ distribution			
500	90.1	88.8	64.3	61.7	70.6	68.6
300	89.3	88.7	54.7	52.0	62.3	59.8
200	87.8	87.2	47.4	45.7	55.8	53.6

Table 3: Expected lengths of 95% confidence intervals for the true 99.9% quantile

Methods	Normal approximation		Empirical likelihood		Data tilting	
α_k	1.5	1.2	1.5	1.2	1.5	1.2
n			$Fréchet(\alpha_k)$ distribution			
500	561	2,392	434	2,103	397	1,898
300	1,039	4,681	725	2,907	663	2,624
200	1,773	8,784	1,124	5,168	979	4,667
n			$Burr(\alpha_k, \alpha_k + \beta)$ distribution			
500	471	3,229	363	2,695	325	2,455
300	822	5,877	679	4,306	585	3,916
200	1,396	10,148	905	7,502	822	6,717

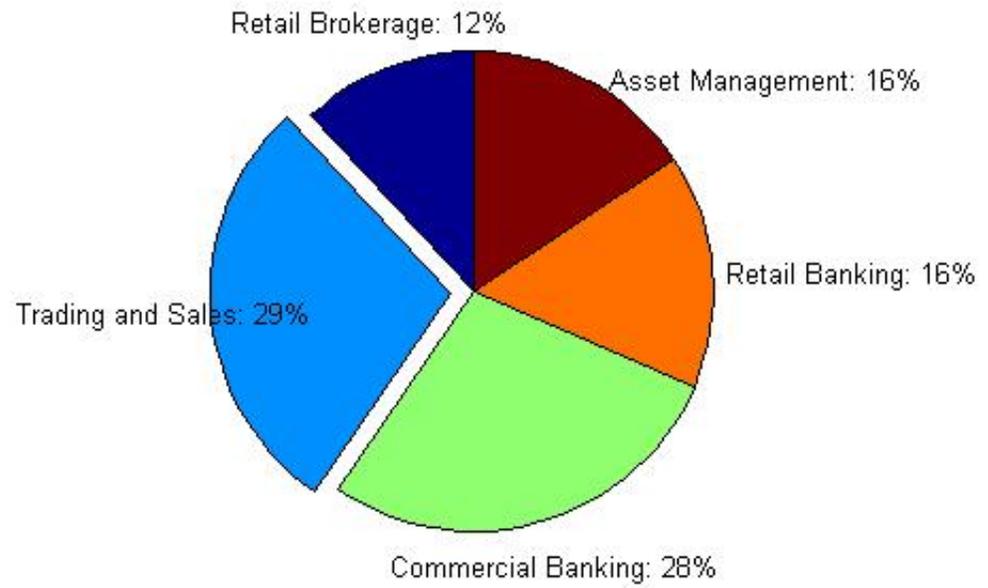**Appendix B Figures and tables**

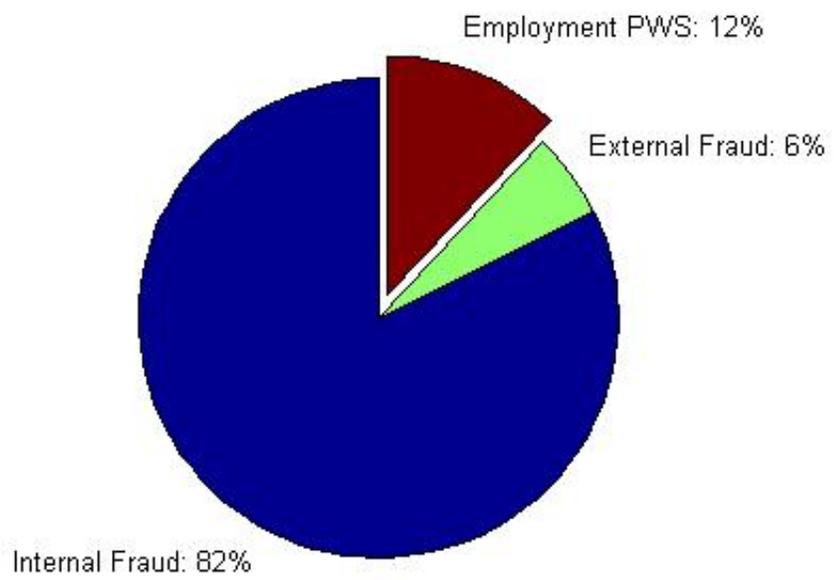Figure 1: Pie chart of losses by five business lines

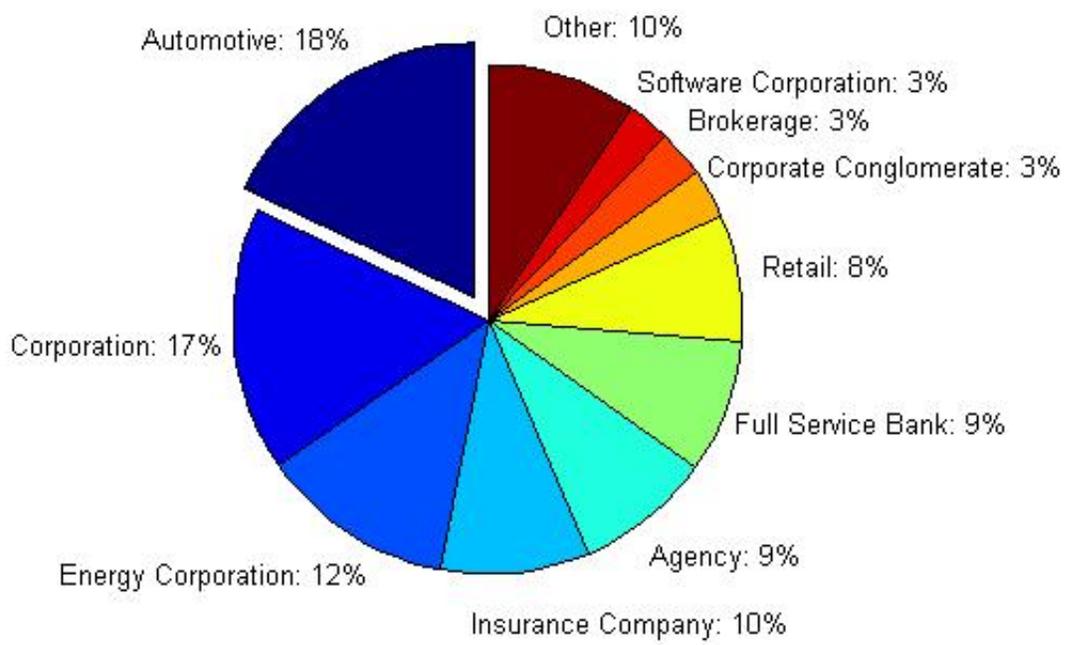Figure 2: Pie chart of losses by three event types

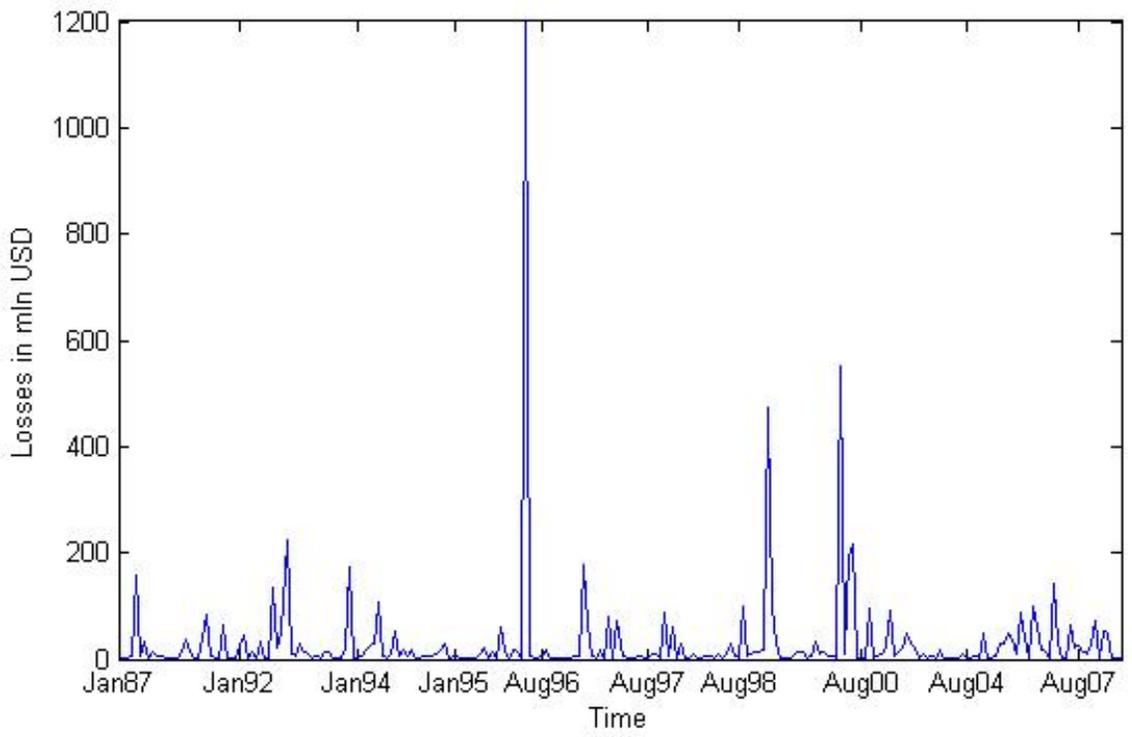Figure 3: Pie chart for the *Employment PWS* losses by various business sectors

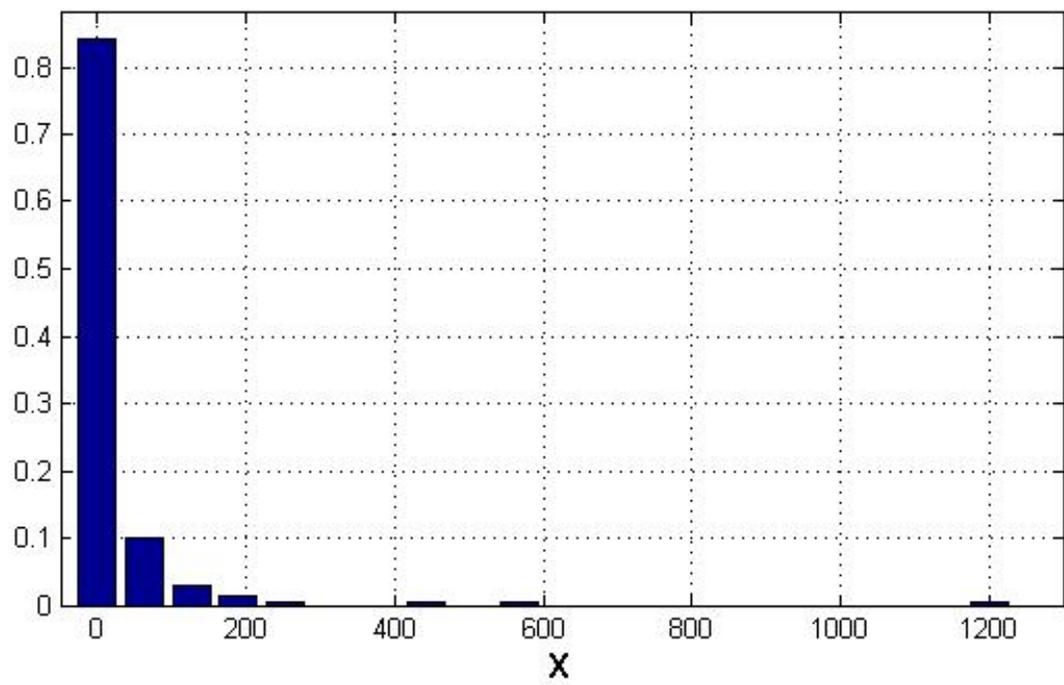Figure 4: The time series plot of the *Employment PWS* losses

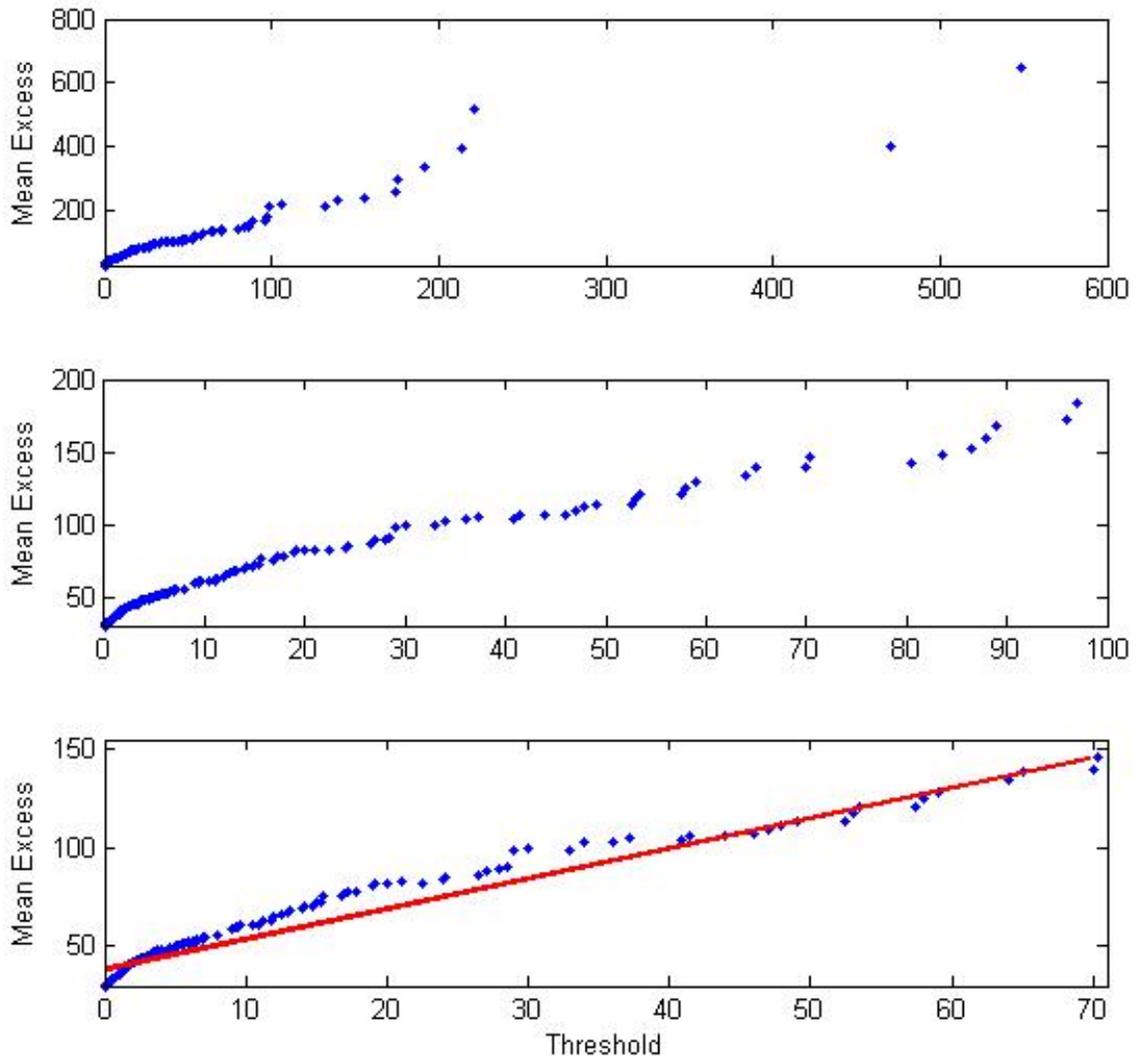Figure 5: Histogram of *Employment PWS* losses

Figure 6: The mean excess function plot of the *Employment PWS* losses, including all losses (top), excluding the ten largest losses (middle), and excluding the 21 largest losses (bottom)
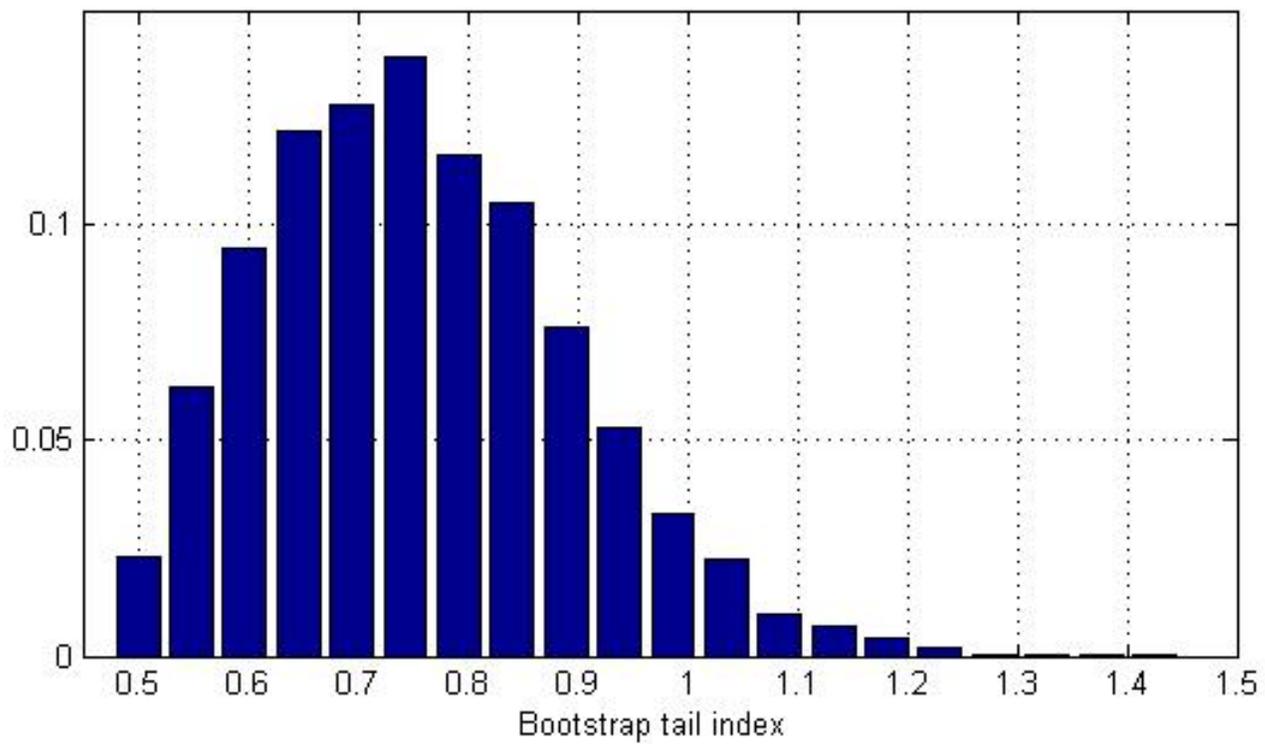
Figure 7: *Employment PWS* losses: bootstrap sampling distribution of the tail index estimate values (with 10,000 replications)
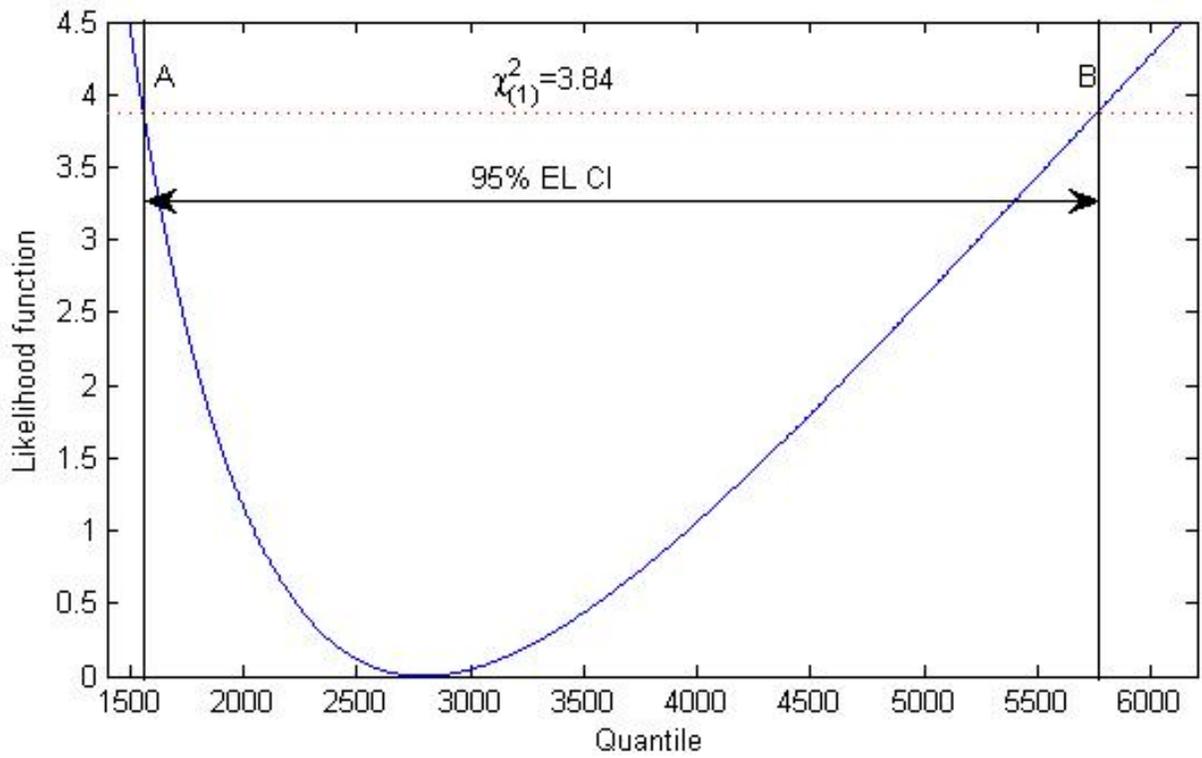
Figure 8: The likelihood function for the empirical likelihood method (see Section 6.2) for the *Employment PWS* losses. The 95% empirical likelihood confidence bound estimate is [\$1,615M, \$4,567M]
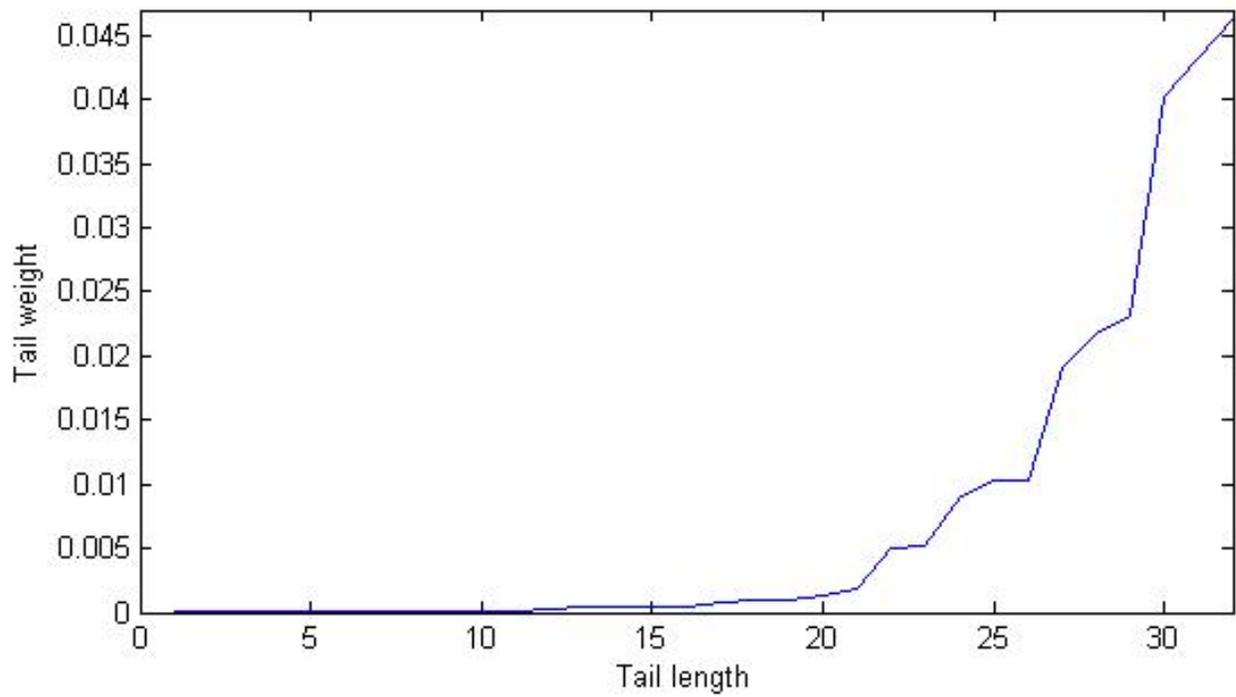
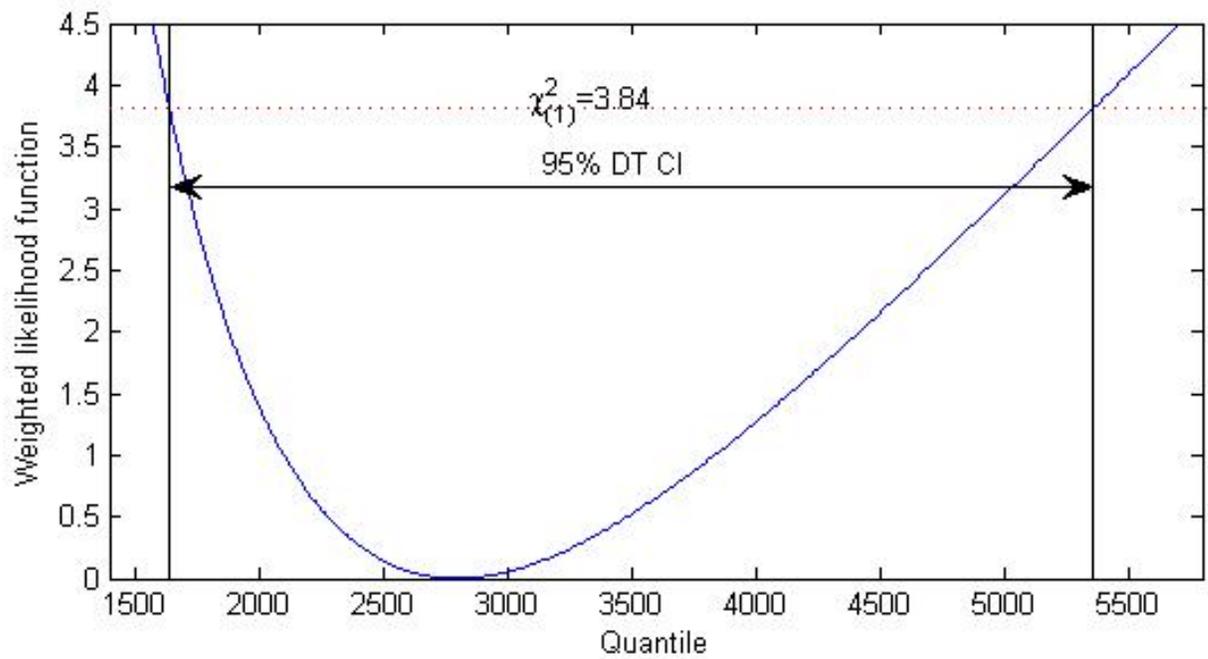Figure 9: The tail weights for the data tilting method for the *Employment PWS* losses, computed using equation (32)

Figure 10: The likelihood function of the data tilting method for the *Employment PWS* losses. The 95% data tilting confidence bound estimate is [$1,583M, $3,876M]
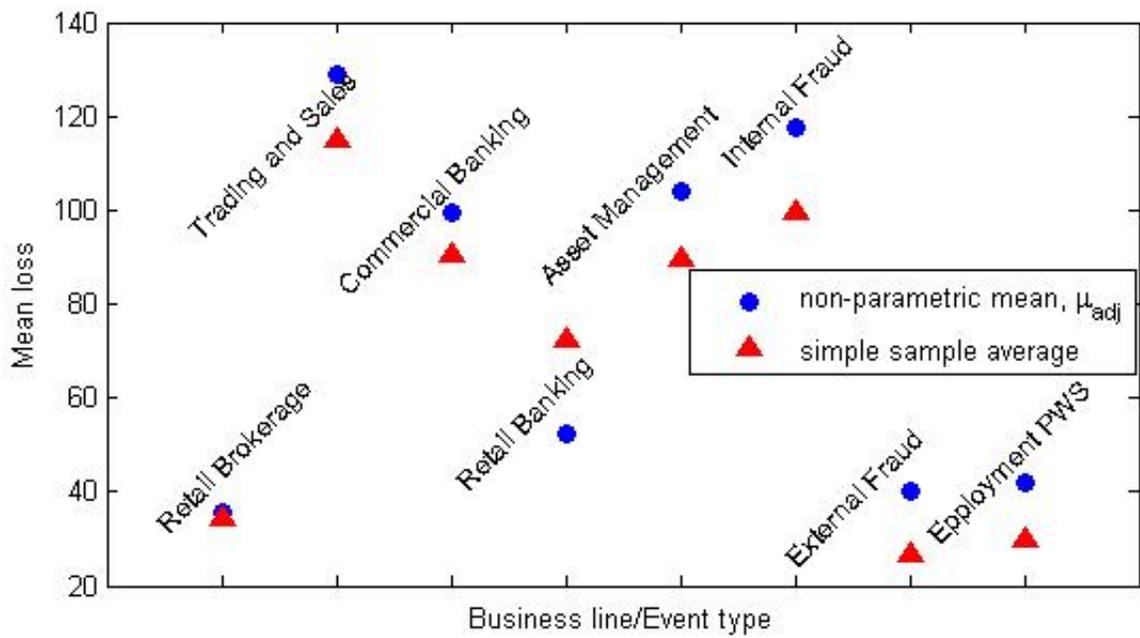
Figure 11: Sample mean and non-parametric adjusted mean estimates for business line/event type losses

Table 4: Descriptive statistics for operational losses by Basel business lines and event types

Loss type	Sample size	% of all losses	Sample mean	Median	75th perc.	95th perc.	99.9th perc.	St. dev.	Skewness	Max
Operational losses: Basel business line										
Retail Brokerage	198	12	34.1	3.6	14.0	98.2	1,500	154.3	7.7	1,499.0
Trading and Sales	141	29	114.9	14.0	50.3	399.0	4,400	443.1	7.4	4,399.0
Commercial Banking	176	28	90.4	9.4	40.9	223.7	3,600	376.2	7.3	3,599.0
Retail Banking	123	16	72.3	5.0	23.0	432.8	2,086	246.4	5.8	2,085.0
Asset Management	100	16	89.4	24.8	78.0	424.8	1,600	196.3	5.1	1,599.0
Operational losses: Basel event type										
Internal Fraud	490	82	99.6	10.7	42.0	504.0	3,367	330.1	6.1	3399.0
External Fraud	130	6	26.5	7.0	16.0	118.3	675	73.7	6.2	674.0
Employment PWS	237	12	29.7	5.9	21.4	103.7	1,200	96.0	8.9	1,199.0
Operational losses: Basel business line (in real terms)										
Retail Brokerage	191	12	33.6	3.0	12.3	86.3	1,571	162.6	8.1	1,570.1
Trading and Sales	141	30	111.1	15.7	49.7	363.0	4,279	430.4	7.5	4,278.0
Commercial Banking	177	28	84.0	9.1	37.2	200.1	3,166	342.2	7.0	3,165.2
Retail Banking	120	14	63.1	3.9	23.4	434.8	1,861	212.7	6.0	1,860.4
Asset Management	100	16	82.4	21.7	72.3	361.3	1,701	194.8	6.1	1,699.5
Operational losses: Basel event type (in real terms)										
Internal Fraud	501	82	91.0	9.2	41.0	450.8	2,663	297.3	6.0	2,665.0
External Fraud	131	6	25.3	5.8	14.9	140.1	529	66.2	5.1	528.3
Employment PWS	241	12	28.9	5.5	20.0	103.8	1,226	95.8	9.2	1,225.2

Note: The denomination of losses is in million.

Table 5: Estimates of the threshold, the tail index and its confidence interval, the adjusted mean and the quantile

Loss type	Threshold	$\hat{\alpha}_k$	CI for $\hat{\alpha}_k$	μ_{adj}	OpVaR
Operational losses: Basel business line					
Retail Brokerage	36.0	1.2	[0.5 1.3]	35.5	2,128
Trading and Sales	64.0	1.1	[0.5 1.3]	128.7	7,042
Commercial Banking	149.0	1.2	[0.4 1.4]	99.6	5,891
Retail Banking	99.0	1.4	[0.6 1.8]	52.2	3,092
Asset Management	203.0	1.7	[0.9 1.9]	103.9	4,023
Operational losses: Basel event type					
Internal Fraud	94.0	1.2	[0.9 1.5]	117.7	7,149
External Fraud	26.0	1.2	[0.8 1.8]	40.0	2,236
Employment PWS	70.4	1.3	[1.1 1.9]	41.8	2,475
Operational losses: Basel business line (in real terms)					
Retail Brokerage	32.8	1.2	[0.5 1.3]	30.1	1,826
Trading and Sales	63.9	1.2	[0.5 1.3]	110.3	6,260
Commercial Banking	156.6	1.1	[0.3 1.3]	138.1	7,484
Retail Banking	87.9	1.4	[0.6 1.6]	42.4	2,344
Asset Management	164.7	1.5	[0.9 1.9]	101.6	4,665
Operational losses: Basel event type (in real terms)					
Internal Fraud	102.0	1.2	[1.0 1.6]	92.9	5,690
External Fraud	23.9	1.2	[0.8 1.7]	27.1	1,506
Employment PWS	64.2	1.2	[1.1 1.9]	41.7	2,504

Notes: $\hat{\alpha}_k$ is the tail index; bootstrap confidence interval estimate for $\hat{\alpha}_k$; and μ_{adj} is the non-parametric adjusted mean.

Table 6: Non-parametric estimates of the operational risk, Expected Shortfall, and confidence intervals for operational risk

Loss type	OR[1]	ES[2]	NA CI[3]	EL CI[4]	DT CI[5]
Operational losses: Basel business line					
Retail Brokerage	2,092	3,234	[426 12,478]	[1,334 4,113]	[1,242 3,657]
Trading and Sales	6,913	10,920	[1,317 41,237]	[4,363 13,640]	[4,421 11,241]
Commercial Banking	5,791	8,929	[807 44,473]	[3,373 13,341]	[3,702 9,369]
Retail Banking	3,040	4,240	[519 19,444]	[1,872 6,358]	[1,856 5,157]
Asset Management	3,919	5,037	[814 19,047]	[2,603 7,498]	[2,467 6,512]
Operational losses: Basel event type					
Internal Fraud	7,032	10,798	[2,843 19,148]	[5,518 10,464]	[5,591 9,149]
External Fraud	2,196	3,425	[380 16,875]	[1,328 4,622]	[1,157 4,313]
Employment PWS	2,433	3,571	[538 11,040]	[1,615 4,567]	[1,583 3,876]
Operational losses: Basel business line (in real terms)					
Retail Brokerage	1,796	2,750	[363 10,911]	[1,141 3,559]	[1,014 3,297]
Trading and Sales	6,150	9,533	[1,223 35,187]	[3,925 11,929]	[3,980 9,867]
Commercial Banking	7,346	11,783	[925 62,708]	[4,163 17,661]	[4,759 11,687]
Retail Banking	2,302	3,151	[401 14,531]	[1,427 4,801]	[1,578 4,325]
Asset Management	4,563	6,148	[809 26,531]	[2,865 9,366]	[2,790 7,788]
Operational losses: Basel event type (in real terms)					
Internal Fraud	5,598	8,330	[2,289 14,992]	[4,404 8,269]	[4,467 7,251]
External Fraud	1,479	2,197	[291 10,010]	[928 2,968]	[707 3,214]
Employment PWS	2,462	3,641	[558 10,861]	[1,645 4,556]	[1,618 3,874]

[1] Operational Risk
[2] Expected Shortfall
[3] Normal approximation based confidence interval estimate
[4] Empirical likelihood confidence interval estimate
[5] Data tilting confidence interval estimate

References

Acerbi, C. & Tasche, D. (2001), Expected shortfall: a natural coherent alternative to value at risk.

Basel Committee on Banking Supervision (2004), Basel II: Revised international capital framework, Technical report, Bank for International Settlements, Basel, Switzerland.

Basel Committee on Banking Supervision (2011*a*), Basel III: A global regulatory framework for more resilient bank and banking systems, Technical report, Bank for International Settlements, Basel, Switzerland.

Basel Committee on Banking Supervision (2011*b*), Operational risk: Supervisory guidelines for the Advanced Measurement Approaches, Technical report, Bank for International Settlements, Basel, Switzerland.

Beirlant, J., Goegebeur, Y., Segers, J. & Teugels, J. L. (2004), *Statistics of Extremes*, Wiley.

Brown, S., Goetzman, W., Liang, B. & Schwartz, C. (2008), 'Mandatory disclosure and operational risk: evidence from hedge fund registration', *The journal of finance* **LXIII**, 2785–2815.

Caers, J. & Maes, M. A. (1998), 'Identifying tails, bounds and end-points of random variables', *Structural Safety* **20**, 1–23.

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

Danielsson, J., de Haan, L., Peng, L. & de Vries, C. G. (2001), 'Using a bootstrap method to choose the sample fraction in tail index estimation', *Journal of Multivariate Analysis* **76**, 226–248.

Davison, A. C. & Smith, R. L. (1990), 'Models for exceedances over high thresholds', *Journal of the Royal Statistical Society* **52(3)**, 393–442.

Efron, B. & Tibshirani, R. J. (1993), *An Introduction to the Bootstrap*, Chapman and Hall.

Hall, P. & Yao, Q. (2003), 'Data tilting for time series', *Journal of the Royal Statistical Society* **Series B 65**, 425–442.

Hill, B. M. (1975), 'A simple general approach to inference about the tail of a distribution', *The Annals of Statistics* **3**, 1163–1174.

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

Johansson, J. (2003), 'Estimating the mean of heavy-tailed distributions', *Extremes* **6**, 91–109.

Mason, D. M. (1982), 'Laws of large numbers for sums of extreme values', *Annals of Probability* **10**(3), 754–764.

McCulloch, J. H. (1997), 'Measuring tail thickness to estimate the stable index: A critique', *Journal of Business and Economic Statistics* **15**, 74–81.

Owen, A. B. (1988), 'Empirical likelihood ratio confidence intervals for a single functional', *Biometrika* **75**, 237–249.

Owen, A. B. (1990), 'Empirical likelihood ratio confidence regions', *The Annals of Statistics* **18**, 90–120.

Peng, L. (2001), 'Estimating the mean of a heavy-tailed distribution', *Statistics and Probability Letters* **52**, 255–264.

Peng, L. & Qi, Y. (2006), 'Confidence regions for high quantiles of a heavy-tailed distribution', *The Annals of Statistics* **34**, 1964–1986.

Smith, R. L. (1985), 'Maximum likelihood estimation in a class of nonregular cases', *Biometrika* **72**, 67–90.