

Implications of Alternative Operational Risk Modeling Techniques*

Patrick de Fontnouvelle, Eric Rosengren
Federal Reserve Bank of Boston

John Jordan
FitchRisk

June, 2004

Abstract

Quantification of operational risk has received increased attention with the inclusion of an explicit capital charge for operational risk under the new Basle proposal. The proposal provides significant flexibility for banks to use internal models to estimate their operational risk, and the associated capital needed for unexpected losses. Most banks have used variants of value at risk models that estimate frequency, severity, and loss distributions. This paper examines the empirical regularities in operational loss data. Using loss data from six large internationally active banking institutions, we find that loss data by event types are quite similar across institutions. Furthermore, our results are consistent with economic capital numbers disclosed by some large banks, and also with the results of studies modeling losses using publicly available “external” loss data.

* This paper was prepared for the NBER Project on the Risks of Financial Institutions. It was substantially completed while John Jordan was with the Federal Reserve Bank of Boston. We thank our colleagues in the Federal Reserve System and in the Risk Management Group of the Basel Committee for the many fruitful interactions that have contributed to this work. However, the views expressed in this paper do not necessarily reflect their views, those of the Federal Reserve Bank of Boston, or those of the Federal Reserve System. Address for correspondence: Patrick de Fontnouvelle, Federal Reserve Bank of Boston, Mail Stop T-10, 600 Atlantic Avenue, P.O. Box 2076, Boston, MA 02106-2076, tel: 617-973-3659, fax: 617-973-3219, email: patrick.defontnouvelle@bos.frb.org.

1. Introduction

Large operational losses as a result of accounting scandals, insider fraud, and rogue trading, to name just a few, have received increasing attention from the press, the public, and from policymakers. The frequency of severe losses, with more than 100 instances of losses at financial institutions exceeding \$100 million, has caused many financial institutions to try to explicitly model operational risk to determine their own economic capital. As financial institutions have begun to comprehensively collect loss data and use it to manage operational risk, bank regulators have increased their expectations for measuring and modeling operational risk. Under the current US rules proposal for implementing the Basle Accord, large internationally active banks will be expected to use internal models to estimate capital for unexpected operational losses. A criticism of this proposal has been that the tools for modeling operational risk are in their infancy, making estimating capital problematic.

This paper uses data supplied by six large internationally active banks to determine if the regularities in the loss data will make consistent modeling of operational losses possible. We find that there are similarities in the results of models of operational loss across institutions, and that our results are consistent with publicly reported operational risk capital estimates produced by banks' internal economic capital models.

We begin the analysis by considering tail plots of each bank's loss data by business line and event type. Three findings clearly emerge from this descriptive analysis. First, loss data for most business lines and event types may be well modeled by a Pareto-type distribution, as most of the tail plots are linear when viewed on a log-log scale. Second, the severity ranking of event types is consistent across institutions. Clients, Products and Business Practices is the highest severity event type, while External Fraud and Employment Practices are the lowest severity

event types. Third, the tail plots suggest that losses for certain business lines and event types are very heavy-tailed. This last finding highlights that while basic measurement approaches such as the tail plot are easy to implement and intuitively appealing, overly simplistic approaches may yield implausible estimates of economic capital. A main contribution of this paper is to show how quantitative modeling can result in more reasonable conclusions regarding tail thickness and economic capital.

We next attempt to model the distribution of loss amounts using a “full-data” approach, whereby one fits all of the available loss data with a parametric severity distribution. We consider nine commonly-used distributions, four of which are light-tailed and five of which are heavy-tailed. We fit each of these distributions by business line and event type at each of the six institutions considered. The heavy-tailed distributions provide consistently good fits to the loss data, which confirms our findings based on visual inspection of the tail plots. The light-tailed distributions do not generally provide good fits. However, we find that some parameter estimates for the heavy-tailed distributions can have implausible implications for both tail thickness and economic capital.

Extreme Value Theory (EVT) is an alternative to the full-data approach that is increasingly being explored by researchers, by financial institutions, and by their regulators. However, it is well-known that EVT techniques yield upward-biased tail estimates in small samples. Huisman, Koedijk, Kool and Palm (2001) have proposed a regression-based EVT technique that corrects for small-sample bias in the tail parameter estimate. Applying their technique (hereafter HKKP) to the six banks in our sample, we obtain estimates that are both reasonable and consistent with earlier estimates using purely external data (de Fontnouvelle et. al., 2003).

It is important to stress that the statistical analysis of operational loss data is a new field, and that this paper's results should be viewed as preliminary. This is particularly true since we only have data for one year from each bank. The paper also raises several technical issues that should be addressed in future research as a longer time series becomes available. The most significant such issue is that even though the data appear to be heavy-tailed, we cannot formally reject the hypotheses that they are drawn from a light-tailed distribution such as the lognormal. To investigate this possibility, we propose a threshold analysis of the lognormal distribution that to our knowledge is new to this paper. This technique also provides a reasonable characterization of the tail behavior of operational losses.

We also examine the frequency of operational losses. We consider both the Poisson distribution and the Negative Binomial distribution as potential models for the number of losses that a bank could incur over the course of one year. Using Monte Carlo Simulation to combine the frequency and severity distributions, we obtain an estimate for the distribution of total annual operational losses. The quantiles of this aggregate loss distribution are interpreted as economic capital estimates for operational risk. These estimates should be viewed with several significant cautions. First, we are assuming that the data are complete, however, banks have moved to more comprehensive data collection platforms which may improve the loss capture. Second, we are only using internal data for one year, and banks will be required to have three years of comprehensive data. Third, analysis of internal loss data will not be the sole determinant of capital for operational risk; banks will be also required to demonstrate that their risk estimates

reflect exposures that are not captured in internal loss data.¹ Given these qualifications, the estimates should be viewed as a preliminary indication of the amount of capital needed.

Despite these caveats, the estimates implied by the modeling of the internal loss data are consistent with capital estimates using purely external data (de Fontnouvelle et. al., 2003). The results imply that for a variety of plausible assumptions regarding the frequency and severity of operational losses, the level of capital needed for operational risk for the typical (median) bank in our sample would be equivalent to 5-9 percent of the bank's current minimum regulatory capital requirement. This range also seems consistent with the 12-15 percent of minimum regulatory capital that most banks are currently allocating to operational risk, given that the banks' models tend to have a broader set of model inputs than those used in this analysis, including, external data, scenarios and qualitative risk assessments.

The remainder of the paper is organized as follows. The next section provides a description of the data. Section three reviews related literature on the measurement of operational risk in financial institutions. Section four discusses some commonly used continuous distributions, and discusses their potential relevance to modeling the severity of operational losses. Section five presents visual analyses of the loss data, and draws preliminary conclusions regarding which distributions may be appropriate for modeling loss severity. Section six explores full-data approaches to modeling operational losses, and formally compares the alternative severity distributions. Section seven explores EVT-based approaches to modeling the loss data. Section eight compares alternative frequency distributions. Section nine provides the

¹ The proposed Basel accord requires banks to measure losses to which they are exposed, but that have not actually occurred (via analysis of "scenarios" and "external" data). Banks would also be required to measure exposures that have arisen since the data collection period (via analysis of "business environment and control factors").

implied capital numbers from estimating different loss distributions using Monte Carlo simulations. The final section provides conclusions on using these techniques for quantifying operational risk.

2. Data

The 2002 Operational Risk Loss Data Collection Exercise (LDCE) was initiated by the Risk Management Group (RMG) of the Basel Committee on Banking Supervision in June 2002. The LDCE asked participating banks to provide information on individual operational losses exceeding €10,000 during 2001, among various other data items. Banks were also asked to indicate whether their loss data were complete. The LDCE data include 47,269 operational loss events reported by 89 banks from 19 countries in Europe, North and South America, Asia, and Australasia. For additional information and summary statistics regarding the LDCE, readers can refer to RMG (2003).

Based on the information provided in the LDCE, and on our knowledge of the banks involved, we identified a list of institutions whose data submissions seem relatively complete. Due to practical considerations, we limit our sample to loss data from six of these banks. This paper presents results for these six banks on a bank-by-bank basis (with the exception of the operational risk exposure figures reported in Table 5). However, the results are presented in a way that makes it impossible to identify the individual banks. Focusing on a cross-sectional study of banks enables us to determine whether the same statistical techniques and distributions apply across institutions that may have very different business mixes and risk exposures.

The LDCE categorizes losses into eight business lines and seven event types. To protect the confidentiality of banks participating in the LDCE, we present results only for those Business

Lines and Event Types where three or more banks reported sufficient data to support analysis. The business lines presented are: Trading and Sales; Retail Banking; Payment and Settlement; and Asset Management. The loss types presented are: Internal Fraud; External Fraud; Employment Practices and Workplace Safety; Clients, Products and Business Practices; and Execution, Delivery and Process Management.²

3. Related literature

Moscadelli (2003) also analyzes data from the 2002 LDCE, and performs a thorough comparison of traditional full-data analyses and extreme value methods for estimating the operational loss severity distribution. He finds that extreme value theory outperforms the traditional methods in all eight Basel Business Lines. He also finds that the severity distribution is very heavy-tailed, and that there is a substantial difference in loss severity across Business Lines.

There are several differences between the current paper and Moscadelli (2003). First, Moscadelli (2003) aggregates the data across all banks in the LDCE sample. In this paper, we analyze data at the individual bank level in order to determine whether the same quantitative techniques “work” for a variety of banks with different business mixes, control infrastructures, and geographic exposures. We believe that doing so provides a useful test of the techniques under consideration, and also yields an indication of their ultimate applicability at individual banks. Second, the current paper explores the newly-developed technique of Huisman et. al.

² The following Business Lines were omitted: Corporate Finance; Commercial Banking; Agency Services; and Retail Brokerage. The following Event Types were omitted: Damage to Physical Assets; Business Disruption and System Failure. To preserve confidentiality, we do not report the cutoff that was used for inclusion of Business Lines and Event Types.

(2001) to correct for potential bias in the tail parameter estimate. Third, we explore several models of the loss frequency distribution, which allows us to obtain indicative estimates of economic capital for operational risk.

4. Distributions for operational loss data.

We begin our empirical analysis by exploring which of various empirical approaches best fits the data. In principle, we are willing to consider any distribution with positive support as an acceptable candidate for modeling operational loss severity. To keep the size of our tables within reason, however, we will focus on nine commonly used distributions. This section discusses the salient features of each. In section six, we consider how well these distributions describe the statistical behavior of losses in our database.

TABLE 1 HERE

Table 1 lists each distribution we consider, together with its density function and its maximal moment (discussed at the end of this section). We begin our discussion with the exponential distribution, which is one of the simplest statistical distributions – both analytically and computationally. The exponential distribution is frequently used to analyze duration data (e.g., time to failure of a machine part), and is the only continuous distribution characterized by a “lack of memory.” In the duration context, lack of memory means that the time until the occurrence of an event (failure) does not depend on the length of time that has already elapsed (time since installation). In the operational loss context, lack of memory implies that the distribution of excess losses over a threshold does not depend on the value of the threshold. So if half of all losses exceeding \$1 are less than \$10, then half of all losses exceeding \$1 Million will be less than \$1,00,010 ($\$1,000,000 + \10). Such a result does not seem plausible. However, the

exponential distribution arises in the context of Extreme Value Theory (EVT) as a possible limiting distribution for excess losses above high thresholds. For this reason (and also because it can be transformed into other interesting distributions), we include it in our analysis.

The Weibull distribution is a two-parameter generalization of the exponential that allows the time until event occurrence to depend on the amount of time that has already elapsed. Thus, the Weibull can capture phenomena such as “burn in,” in which the failure rate is high initially but decreases over time. In the context of operational risk, the Weibull may be appropriate for modeling a business line exposed to many small losses but only a few large losses. The Gamma distribution is another two-parameter generalization of the exponential. A Gamma distributed random variable arises as the sum of n exponentially distributed random variables. Thus, a machine’s failure time is Gamma distributed if the machine fails whenever n components fail, and if each component’s failure time is exponentially distributed. Like the Weibull distribution, the Gamma also allows the time until event occurrence to depend on the amount of time that has already elapsed.

Another generalization of the exponential distribution can be obtained by exponentiating an exponentially distributed random variable. The resulting distribution is called a Type I Pareto, and can also be referred to as a log-exponential or power-law distribution. The lack of memory of the exponential distribution manifests itself as scale invariance in the Pareto distribution. Roughly speaking, scale invariance means that data “look the same” no matter what the unit of measure (e.g., hundreds of dollars vs. millions of dollars). So in the earlier example where half of all losses exceeding \$1 were less than \$10, half of all losses over \$1 Million would be less than \$10 Million. Power law behavior has been observed in phenomena as disparate as city sizes, income distributions and insurance claim amounts, and has been an important research

topic for those interested in the behavior of complex systems (i.e., systems consisting of agents linked via a decentralized network rather than via a market or social planner). A variation of the Pareto distribution can be obtained by exponentiating a Gamma distributed random variable instead of an exponentially distributed random variable. The result is referred to as the Loggamma distribution.

The Pareto distribution also arises in Extreme Value Theory as another limiting distribution of excesses over a high threshold. In this case, the limiting distribution is given by a two parameter variant of the Pareto, which is known as the Generalized Pareto Distribution (GPD). One commonly-used transformation of the GPD is obtained by raising a GPD-distributed variable to a power. The result is called the Burr distribution.

Another distribution that we consider is the Lognormal, which is so widely used that little discussion is required here. However, it is worth noting that as the normal distribution is appropriate for modeling variables that arise as the sum of many different components. It is also a worthwhile exercise to consider which types of operational losses may be characterized in this manner. Consider, for example, losses arising from workplace safety lapses. One could argue that the severity of these losses may be approximated by the lognormal distribution, as it is influenced by many factors, including weather, overall health of the injured party, physical layout of the workplace, and the type of activity involved. The final distribution that we consider is the loglogistic, which is obtained by exponentiating a logistic distributed random variable. The Loglogistic is similar to the Lognormal, but may be more appropriate for modeling operational loss data because it has a slightly heavier tail.

We conclude this section by classifying the distributions discussed above according to their tail thickness. This will facilitate interpretation of the estimation results, as the relevance of

a particular distribution to modeling operational losses will be suggestive of the relevance of other distributions with similar tail thickness. There is no commonly agreed-upon definition of what constitutes a heavy-tailed distribution. However, one such definition can be based upon a distribution's maximal moment, which is defined as $\sup\{r : E(x^r) < \infty\}$. Maximal moments for the distributions under consideration are reported in Table 1. In this paper, we will call a distribution light-tailed if it has finite moments of all orders, and heavy-tailed otherwise. Under this definition four of the distributions being considered are light-tailed (Exponential, Weibull, Gamma and Lognormal), and the remaining five distributions are heavy-tailed (Loggamma, Pareto, GPD, Burr and Loglogistic).

5. Descriptive analysis.

This section considers several tools that provide a visual characterization of the loss data. Suppose one has a series of observations $\{x_i\}$ with a cumulative empirical distribution function denoted by $F(x)$. A tail plot is obtained by plotting $\log(1-F(x_i))$ on the vertical axis against $\log(x_i)$ on the horizontal axis. Figures 1 and 2 present tail plots of the six banks' loss data by Basel event type and Basel business line, respectively.

FIGURES 1 and 2 HERE

Many of the tail plots show linear behavior. This is quite interesting, as a linear tail plot implies that the data are drawn from a power-law distribution. Furthermore, the slope of the plot provides a heuristic estimate of the tail parameter, as $\log(1-F(x_i)) = -a \log(x_i) + c$, where c denotes a constant.

Another feature of these plots is that the slopes associated with the seven Basel event types preserve roughly the same ordering across banks. For example, Clients Products and

Business Practices is one of the heaviest-tailed event types for all of the banks where it is plotted separately. Employment practices and workplace safety is always one of the thinnest-tailed event types. While the tail plots by business line also suggest power-law tail behavior, there is no evident consistent cross-bank ordering of business lines. We interpret this as initial evidence that risk may be better ordered by event type, but will revisit this issue later in the paper.

Each of the tail plots also indicates a reference line with slope of -1. Many of the plots lie near or above this line, thus implying heuristic tail parameter estimates of one or higher. These estimates highlight the shortcomings of using an overly simplistic approach to measuring operational risk: tail parameters exceeding one suggest that the expected loss is infinite for many business lines and event types, and that the capital required for operational risk alone could exceed the amount of capital that large banks are currently allocating to all risks.³ We will argue in this paper that the distribution of operational losses is not as heavy tailed as it first appears, and that it is possible to obtain reasonable estimates of regulatory capital for operational risk.

Another useful diagnostic tool is the mean excess plot. The mean excess for a given threshold is defined as the average of all losses exceeding the threshold, minus the threshold value. The mean excess plot reports the mean excess as a function of the threshold value. The shape of the mean excess plot varies according to the type of distribution underlying the data. For example, a Pareto distribution implies a linear, upward-sloping mean excess plot; an

³ The LDCE data suggest that a \$100 Billion bank could experience 500 operational losses (exceeding \$10,000) per year. If these follow a Pareto distribution with a tail parameter equal to one, then Monte Carlo simulation of the aggregate loss distribution indicates capital of \$5 Billion at the 99.9% soundness level. Tail parameters of greater than one would imply capital levels several times larger than this figure.

exponential distribution implies a horizontal linear mean excess plot; and a lognormal distribution implies a concave upward sloping mean excess plot.

FIGURES 3 AND 4 HERE

Figures 3 and 4 present mean excess plots for loss data by event type and business line, respectively. (Each curve has been rescaled in order to display the different business lines and event types together on one plot. Thus, these plots cannot be used to risk-rank business lines or event types.) Nearly all of the plots slope upwards, which indicates tails that are heavier than exponential. Some of the plots are linear (e.g., event type 7 for bank B), which suggests a Pareto-like distribution. Some are concave, which suggests a lognormal or Weibull-like distribution. It is also difficult to establish a consistent pattern across either business line or event type. Potentially, this issue would be less severe with more data.

6. Fitting the distributions.

In this section, we fit each of the distributions listed in Table 1 to the LDCE data via Maximum Likelihood. Results are reported separately for each bank under consideration, and are also broken down by business line and event type.

TABLE 2 HERE

Table 2 reports probability values for Pearson's χ^2 goodness of fit statistic.⁴ In general, the heavy-tailed distributions (Burr, LogGamma, LogLogistic, Pareto) seem to fit the data quite well. The reported probability values exceed 5% for many business lines and event types, which suggests that we cannot reject the null that data are in fact drawn from the distribution under

⁴ We calculated χ^2 goodness of fit tests because EDF-based tests can be sensitive to data rounding, which is prevalent in the LDCE data. One can accommodate rounding within the χ^2 test by choosing bin values appropriately.

consideration. Conversely, most of the light-tailed distributions rarely provide an adequate fit to the data. This is not surprising, as the tail plots suggested that most of the data are heavy-tailed. What is somewhat surprising is the degree to which the Lognormal distribution fits the data. In fact, this light-tailed distribution fits the loss data for roughly as many business lines and event types as many of the heavier-tailed distributions.

TABLE 3 HERE

Table 3 presents parameter estimation results for the GPD and Lognormal distributions. To preserve bank confidentiality, we present only the estimate of the tail parameter ξ for the GPD and only the value of $\mu + \sigma^2/2$ for the Lognormal distribution. While the χ^2 statistics presented in Table 2 suggested that these two distributions provide a reasonable fit to the data, the parameter estimates generally suggest the opposite. Panel A reports estimates of the GPD tail parameter ξ . The parameter estimates are at or above one for many business lines and event types, and also above one when data is pooled across business lines and event types. Note that a tail parameter of one or higher has implausible implications for both expected losses and regulatory capital. Panel b of Table 3 reports the estimated value of $\mu + \sigma^2/2$ for the lognormal distribution, which enables one to calculate the average loss severity via the formula $\exp(\mu + \sigma^2/2)$. While estimates of the average loss vary by business line and event type, one can see that it is less than $\exp(0)$ dollars for multiple business lines and event types. Thus, neither the Pareto nor the Lognormal distribution consistently yields plausible parameter estimates.

Because of space considerations, we do not provide parameter estimates for the other distributions that were estimated. However, the GPD is of special interest because of its role in Extreme Value theory, and the Lognormal is of special interest because it is the only light-tailed distribution that seems to fit the data (according to the χ^2 test.). Parameter estimates for other

heavy-tailed distributions were qualitatively similar to those of the GPD, in that they had unreasonable implications for tail-thickness of the aggregate loss distribution.

a. For which business lines and event types can full data be fit?

In this subsection, we ask whether there seem to be particular event types for which the full-data approach might work. Losses due to Employment Practices and Workplace Safety (Event Type 3) are well fit by most of the heavy-tailed distributions as well as the lognormal. Furthermore, the parameter estimates for both the GPD and Lognormal are reasonable. There are two event types (Internal Fraud; Clients, Products and Business Practices) where several banks' data are well-fit by multiple distributions, but where the resulting parameter estimates are not reasonable. External Fraud losses are not consistently well fit by any distribution on a cross-bank basis. Results for Execution, Delivery and Process Management are less consistent across banks, with two institutions failing the goodness of fit tests, but the others having good fits and (perhaps) reasonable parameter estimates.

The results are broadly similar in the case of estimation by business line. There are two business lines (Agency Services; Asset Management) that pass the goodness of fit tests, and yield reasonable parameter estimates for several banks. Another business line (Retail Banking) fails the fit tests at most banks, and the final business line (Payment and Settlement) yields implausible parameter estimates.

b. What might individual banks do?

Our discussion to this point has searched for features of operational loss data that hold across all of the six banks in our sample. However, the measurement of operational risk will ultimately take place at individual banks, who may not have the luxury of seeing whether their choices and assumptions are also valid at other institutions. We begin our discussion by focusing

on bank F. Bank staff might begin by fitting one statistical distribution across all business lines and event types, but poor goodness of fit statistics would quickly lead them to alternate approaches. They might consider fitting a separate loss severity distribution to each of the seven event types. However, they would find that losses from the most frequent event type (External Fraud) were not well-modeled by any of the distributions. The next most frequent event type (Clients, Products and Business Practices) is modeled quite well by several heavy-tailed distributions. However, they would be quite surprised to find tail parameter estimates exceeding one, and might conclude that this was not a reasonable way to model operational risk. If they next attempted to fit separate loss severity distributions for each business line, they would discover that loss data for the most common business line (Retail Banking) were not well-modeled by any of the distributions considered.

Bank F was chosen at random for discussion. If presented with their bank's results from Tables 2 and 3, risk management staff from the other five institutions might reach similar conclusions. They would discover that for many of the important business lines and event types, none of the statistical distributions considered adequately captured the behavior of operational losses. They would also discover that some Business Lines and Event Types were well-modeled by heavy-tailed distributions, but that the resulting parameter estimates had implausible implications for their overall operational risk exposure.

7. Threshold analysis of loss data.

The previous section's results suggest that it may be difficult to fit parametric loss severity distributions over the entire range of loss amounts, even if separate analyses are conducted for each business line and event type. In this section, we focus on the largest losses,

as these are most relevant for determining a bank’s operational risk exposure. The main theoretical result underlying this “Peaks Over Threshold” (POT) approach is that if the distribution of excess losses converges to a limiting distribution as the threshold increases, then this limiting distribution is either the Exponential distribution or the Generalized Pareto distribution.

Implementation of the peaks over threshold approach begins with choosing an estimator for the tail index parameter ξ , the most common being the Hill estimator. The appeal of this estimator derives from its conceptual and computational simplicity. For a set of losses exceeding a given threshold, the Hill estimator equals the average of the log of the losses minus the log of the threshold. If the underlying loss distribution is a Type I Pareto, then the Hill estimator is the maximum likelihood estimate of the tail thickness parameter. This property is quite useful, as it enables one to conduct likelihood ratio tests of various hypotheses.

FIGURE 5 HERE

Let k denote the number of observations exceeding a given threshold value. The quantity k is often referred to as the number of *exceedances*. Figure 5 presents plots of the Hill estimator for the six banks under consideration. The solid black line represents the Hill estimator calculated across all business lines and all event types for various values of k between 1 and 200. Traditionally, the final estimate of the tail index parameter has depended heavily on the choice of k . However, Huisman, Koedijk, Kool and Palm (2001) [hereafter HKKP] have recently proposed a regression-based enhancement to the Hill estimator that minimizes the role of threshold selection. HKKP note that the Hill estimator is biased in small samples, and that the bias is approximately linear in k , so that

$$E[\gamma(k)] = \xi + c k, \tag{1}$$

where $\gamma(k)$ denotes the Hill estimator calculated using k exceedances, and ξ denotes the true value of the tail index parameter. HKKP use (1) to motivate the following regression

$$\gamma(k) = \beta_0 + \beta_1 k + \varepsilon(k), \quad (2)$$

which is estimated for k in $\{1, \dots, K\}$. The estimate of β_0 is interpreted as a bias-corrected estimate of ξ . This method also requires the researcher to choose the number of exceedances to include in the analysis. However, HKKP conclude that the estimate of β_0 is robust to the choice of K .

We apply the HKKP technique to the six Hill plots presented in Figure 5. The results are presented in Table 4.

TABLE 4 HERE

The second column reports the number of exceedances (K) that were used to estimate the above regression. HKKP suggest setting K equal to half the sample size N , and also note that the function $\gamma(k)$ should be approximately linear over the range $k = \{1, \dots, K\}$. In results not reported, we found that setting $K = N/2$ would not be appropriate, as none of the six Hill plots were linear over such a wide range.⁵ However, each of the plots in Figure 5 does indicate a range of k over which $\gamma(k)$ is approximately linear. We have chosen K accordingly.

The third column of Table 4 reports the estimate of β_0 that was obtained using the “optimal” K . The estimates vary between 0.50 and 0.86, which implies that the maximal moment $\alpha = 1/\xi$ varies between 1.16 and 2.00. These findings confirm the intuition that operational losses have a heavy-tailed severity distribution. The last row of the table reports results obtained for a sample consisting of all six banks. Interestingly, the resulting parameter

⁵ To preserve the banks’ confidentiality, we do not report Hill plots using either N or $N/2$ exceedances, as doing so would reveal the number of losses at each bank.

estimate of 0.68 is consistent with the results of de Fontnouvelle et. al. (2003), who reported tail parameter estimates of about 0.65. This consistency is remarkable, given that de Fontnouvelle et. al. (2003) used external publicly reported loss data (rather than internal data), as well as substantially different empirical techniques than the current paper.

The final three columns of Table 4 report tail index estimates obtained using different numbers of exceedances in the regression procedure of HKKP. For all six banks, the results do not change materially when the number of exceedances is reduced from K to $0.75K$ or $0.5K$. The results change more when $0.25K$ exceedances are used. Overall, Table 4 confirms that the estimation results are not highly sensitive to the choice of K .

Estimation by business line and event type.

Our Hill plot analyses have thus far taken place at the “top of the house” level, where data are aggregated across both business line and event type. It is reasonable to ask whether this approach is appropriate, or whether the tail behavior of the loss severity distribution might vary by business line and event type. To investigate this issue, we calculated for each value of k (the number of exceedances) separate Hill estimators for each business line and event type. For each k , we then calculated likelihood ratio test statistics for the hypotheses that the tail index is constant across business lines, and that it is constant across event types. The probability values for these statistics are reported graphically in Figure 5. The results indicate that both hypotheses can sometimes be rejected at the 10% level when k is near 200. However, neither hypothesis can be rejected at the 10% level for values of k where the Hill estimator is constant (banks A, C and E) or decreasing (banks B, D and F). Because choosing a small k provides a less biased value of the Hill estimator, segregating the analysis by business line or event type does not seem to be called for. This finding does not mean that tail behavior of operational losses is constant across

business lines and event types. Rather, the ability of statistical estimation techniques to meaningfully differentiate tail behavior across business lines is hindered by a lack of data on large losses using only internal data for one year, and by the concentration of these data in one or two business lines and event types.

On the possibility of thin-tailed severity distributions.

The results presented in Table 4 suggest that loss severity distributions at the six banks under consideration have tail indices ranging between 0.5 and 0.86. The reported standard errors also seem to exclude the possibility that $\xi = 0$, which would indicate a thin-tailed loss distribution. However, the Hill estimator is designed for situations where $\xi > 0$. Thus, it cannot be used to reject the hypothesis of a thin-tailed loss distribution. This is an interesting hypothesis, because thin-tailed distributions such as the lognormal could have significantly different implications for capital than fatter-tailed distributions such as the Pareto.

FIGURE 6 HERE

Dekkers, Einmahl and de Haan (1989) show how to extend the Hill estimator so that it is valid for any ξ in \mathfrak{R} . The graph of this estimator as k varies is commonly referred to as a DEdH plot. Figure 6 reports DEdH plots for the six banks under consideration. These plots indicate that for the low values of k for which the Hill estimator was constant or decreasing, we cannot reject the null of a thin-tailed severity distribution at any of the six banks. This is problematic. The choice of fat versus thin tailed loss severity distribution will have significant impact on the capital calculation, yet based on limited data for only one year, available statistical techniques provide little guidance on which choice is more appropriate. We expect that as banks accumulate more data on large losses, the DEdH plots will either be able to reject the null of

$\xi = 0$ or will indicate tail estimates close enough to zero that the choice does not matter so much. For now, we explore the empirical consequences of assuming a thin-tailed loss severity distribution.

Extreme value theory suggests that the exponential distribution is an appropriate choice for modeling loss severity under the thin-tailed assumption. Thus, we wish to construct a threshold plot showing how the exponential parameter varies as the threshold increases (k decreases). Because the maximum likelihood estimate of this parameter is given by the mean excess, these threshold plots would be identical to the mean excess plots already presented in Figures 3 and 4. As was discussed earlier, the mean excess plots suggest that the exponential distribution does not provide an accurate description of the tail behavior of operational losses. All six banks' excess plots are concave and increasing, whereas exponentially distributed data imply a linear and horizontal excess plot.

Since the DEdH plots do suggest that tail behavior of operational losses might be modeled with a light-tailed distribution, we consider whether some other such distribution provides a better fit to the data than the exponential. Because the log-normal was the one light-tailed distribution investigated in section six that provided a good fit across multiple banks, business lines and event types, we investigate whether it might also provide a useful description of the tail behavior of operational losses.

FIGURE 7 HERE

Figure 7 presents threshold plots for the six banks, under the assumption that losses above high thresholds follow a lognormal distribution. For each value of k (the number of exceedances), estimates of the lognormal parameters were obtained via maximum likelihood. (Vertical axis scales have been omitted to protect data confidentiality. However, a reference line

in Figure 7a indicates the location of $\mu=0$.) One can discern a common pattern in the estimates of μ and σ across all six banks. For example, consider the plots for bank B. Both suggest that the Lognormal is not a good fit for more than 100 exceedances, in that the estimates are unstable as k varies and the point estimate for μ is less than zero. We have already argued that this is not a reasonable characterization of operational loss data. However, parameter estimates are both stable and reasonable when 30 to 70 exceedances are used for estimation. The μ estimate lies between 4 and 8, while the σ estimate lies between 0 and 2.⁶ Three of the other banks display a similar pattern, with stable (and reasonable) parameter estimates emerging over high thresholds. The two remaining banks' (C and E) POT plots become unstable for small numbers of exceedances. This finding could indicate a small sample problem, but it could also indicate that the distribution of large losses at these banks does not follow a Lognormal distribution.

8. The operational loss frequency distribution.

We have thus far focused on the loss severity distribution, which describes the potential size of an operational loss, given that the loss has occurred. Operational risk capital will also depend on the loss frequency distribution, which describes how many losses might actually occur over a given time period. The Poisson distribution is a reasonable starting point for modeling loss frequency, because it arises whenever the loss occurrence rate is constant over time. We thus begin by modeling frequency at bank i by the following:

$$n_i \sim \text{Po}(\lambda_i) \tag{3}$$

⁶ The actual range of variation is significantly narrower, but has not been reported to protect data confidentiality.

That the Poisson distribution has only one parameter makes it particularly attractive in the current context. The LDCE does not provide information regarding the date of an event, beyond the knowledge that all losses occurred sometime during the year 2001. Thus, we have enough information to estimate the Poisson parameter, but not enough to estimate multi-parameter frequency distributions. Maximum Likelihood estimates of the parameter λ are given by the annual number of loss events.⁷

An interesting property of a Poisson variable is that the mean and variance are equal. So if a LDCE bank were to report 10,000 loss events for the year 2001, we would expect (with 95% probability) it to report between 9,800 and 10,200 events the following year. On an intuitive level, this seems like a very narrow range, and one might ask whether frequency should be modeled via a distribution permitting more variability than the Poisson. One such distribution is the Negative Binomial, which is a commonly-used generalization of the Poisson.

As was discussed earlier, the LDCE data do not support estimation of two-parameter frequency distributions at the individual bank level. In order to model excess dispersion in the loss frequency distribution, we take a cross-sectional approach. That is, we estimate the following regression:

$$n_i \sim F(X_i, \mathbf{b}), \quad (4)$$

where $F(\cdot)$ is a discrete non-negative valued distribution, X_i is an observable characteristic of bank i (e.g., asset size), and \mathbf{b} is a parameter vector. Because our data set is purely cross-sectional (i.e., there is no time series element), we cannot estimate any fixed effects. Fixed effects represent bank-specific variation in the frequency of operational losses, which could arise from factors such as the quality of an individual bank's risk control environment. However, it is

⁷ To preserve confidentiality, we have not reported the number of loss events.

worth noting that (3) can be interpreted as a fixed effects model. Seen in this light, (3) and (4) are different but complementary ways of treating the fixed effects issue. Under the latter, the expected number of events is purely a function of a bank's observable characteristics; whereas under the former, the expected number of events is purely bank-specific.

We begin by estimating (4) under the assumption that $F(\cdot)$ is the Poisson distribution, so that $n_i \sim \text{Po}(\text{mean} = bX_i)$. Setting each X_i as bank i 's asset size as of year-end 2001, we obtain an estimate of 8.2 for the parameter b . This indicates that banks in our sample reported on average 8.2 operational events for every billion dollars in assets. Next, we estimate (4) under the assumption that $F(\cdot)$ is the Negative Binomial distribution, so that $n_i \sim \text{NB}(\text{mean} = b_1X_i, \text{dispersion} = b_2)$. We obtain an estimate of 7.4 for b_1 and 0.43 for b_2 .

9. The aggregate loss distribution.

In this section, we combine the severity results of section seven with the frequency results of section eight in order to estimate economic capital for operational risk, which is specified as the 99.9th percentile of the aggregate loss distribution. We explore two alternate assumptions regarding loss frequency (Poisson and Negative Binomial), and three different assumptions regarding loss severity (Pareto, Lognormal, and EDF).

We use Monte Carlo simulation to derive an estimate of the aggregate loss distribution as follows. In the case of the empirical severity distribution, the number of loss events in year i is drawn at random from the frequency distribution, and is denoted N_i . Then, N_i individual losses $\{l(1), \dots, l(N_i)\}$ are drawn from the empirical distribution. The N_i losses are summed to obtain the aggregate loss for year i . This process is repeated for one million simulated years in order to obtain the aggregate loss distribution.

Monte Carlo simulation for the Pareto (Lognormal) severity distribution proceeds similarly, except that Losses in $\{l(1), \dots, l(N_i)\}$ greater than or equal to the relevant threshold value are replaced with random draws from the Pareto (Lognormal) distribution estimated in section six.⁸ The N_i losses are then summed to obtain the aggregate loss for year i , and the process is repeated for one million simulated years in order to obtain the aggregate loss distribution. The use of Monte Carlo techniques in the current context has already been extensively documented, and we refer readers interested in further details to Klugman, Panjer and Wilmot (1998) and Embrechts, Kaufmann and Samorodnitsky (2002), and to their references.

a. Simulations based on a Poisson frequency distribution.

In this subsection, we assume that the frequency of operational losses follows a Poisson distribution with a fixed effects specification as in equation (3). We make three different assumptions for loss severity: the Pareto, the Lognormal, and the empirical distribution. Results are presented in panel a of Table 5. To preserve the confidentiality of the banks in the sample, we scaled each percentile for each bank by that bank's assets. The cross-bank median for each percentile is then reported.

TABLE 5 HERE

The Basel Committee has stated that “a reasonable level of the overall operational risk capital charge would be about 12 percent of minimum regulatory capital.”⁹ If one estimates

⁸ For the Lognormal distribution, the relevant threshold is the same as that used for estimation of the tail parameter. For the Pareto distribution, the relevant threshold is the largest observed loss value. This is because by construction, the HKKP tail parameter estimate β_0 corresponds to zero exceedances.

⁹ See Basel Committee on Banking Supervision (2001).

minimum regulatory capital to be five percent of a bank's assets, then a "reasonable" value for operational risk capital would be 0.6% of assets. According to this criterion, the median value of 0.468% reported in Panel a (for the 99.9th percentile) seems reasonable. It is also worth noting that our estimation is based solely on internal loss data for one year, providing limited data to estimate high severity losses. Banks are also using external loss data and scenario analysis to provide additional information on the tail where they have insufficient high severity losses in a particular business line. Thus, we would view the figure of 0.468% as somewhat of a lower bound on the banks' true operational risk exposure.

The next set of simulations is conducted under the assumption that the severity of operational losses follows a lognormal distribution. The results suggest that cross-bank median of the 99.9th percentile is 0.07% of assets. This figure seems small in comparison with both that obtained in the Pareto-based simulations, and with the 0.6% reasonableness criterion discussed above.

We conducted the final set of simulations by drawing the number of loss events from a Poisson distribution, and the loss amounts from the empirical severity distribution. One may think of the resulting 99.9th percentiles as a lower bound on the true capital requirement. Alternatively, one may think of these percentiles as representing the portion of capital that derives from banks' actual loss experience, rather than from their exposure – as measured by a fitted distribution function which would also include information from external data and scenario analysis. Because the lognormal is a thin-tailed distribution, the 99.9th percentile based on the lognormal severity distribution exceeds that based on the empirical distribution by about 20%. Because the Pareto is a heavy-tailed distribution, the 99.9th percentile based on the Pareto severity distribution exceeds that based on the empirical distribution by a factor of eight.

b. Simulations based on a Negative Binomial frequency distribution.

In the previous section, we assumed that the frequency of operational losses followed a Poisson distribution. We found that assuming a Pareto severity distribution yielded capital estimates that were mostly reasonable when judged against the Basel Committee's expectation that operational risk account for 12% of minimum regulatory capital. Assuming a lognormal severity distribution yielded markedly lower capital estimates. In this section, we investigate how these results change under the assumption that the frequency of operational losses follows a Negative Binomial distribution, as was discussed in section eight.

Panels b and c of Table 5 report quantiles of aggregate loss distributions that were simulated using cross-sectional frequency models based on the Poisson and Negative Binomial distributions, respectively. (Note that the cross-sectional Poisson model is included because it is not informative to directly compare the cross-sectional Negative Binomial results with the fixed effects Poisson results, as differences could be due to either differences in the handling of effects or to differences in the assumed frequency distribution.) The Negative Binomial specification implies significantly more variability in the number of operational losses than does the Poisson specification. Thus, intuition suggests that the aggregate loss distribution should have a heavier tail under the Negative Binomial specification. This intuition proves correct in the case of the lognormal severity distribution. The median 99.9th percentile is about twice as large under the Negative Binomial as under the Cross-Sectional Poisson specification. However, intuition proves incorrect in the case of the Pareto distribution, for which the median 99.9th percentile is not materially different under the Negative Binomial than under the Poisson.¹⁰

¹⁰ It has been argued that intuition can be misleading if risks follow very heavy-tailed Pareto-type distributions (e.g., Embrechts et. al. (2002), Rootzen and Kluppelberg (1999)).

Under the Negative Binomial specification of loss frequency, it is difficult to decide whether the Pareto or the Lognormal provides the more useful characterization of the loss severity distribution. The difference between the two sets of results is within an order of magnitude, which may be considered close given the preliminary nature of the data and techniques.

Conclusion

This paper examines operational risk modeling using only internal operational loss data. By focusing on internal data, it captures the potential modeling issues faced by banking organizations that have only recently started to collect comprehensive loss data. The analysis indicates that the data do show statistical regularities, and that the severity ranking of event types is similar across banks. The analysis also shows that the data is reasonably fit by heavy-tailed distributions (such as the Pareto), and illustrates that certain statistical methods yield plausible tail parameter estimates for these heavy-tailed distributions. In fact, the tail parameter estimates for the severity distribution are quite close to the estimates based on publicly available time series of high severity losses (de Fontnouvelle et. al., 2003).

It is important to qualify our results by noting that they are based on only one year of loss data. This limited data makes it difficult to distinguish between different distributional assumptions, though some thin-tailed distributions do appear inconsistent with the data. At this point, we would conclude that a variety of threshold-based techniques seem to yield results that are consistently plausible across banks. However, we may need to await the arrival of better data before making more definitive conclusions. As banks have three or more years of good operational loss data, the ability to differentiate across alternative distributional assumptions should improve.

References

- Basel Committee on Banking Supervision, 1996, Amendment to the Capital Accord to Incorporate Market Risks.
- Basel Committee on Banking Supervision, 2001, Working Paper on the Regulatory Treatment of Operational Risk.
- de Fontnouvelle, Patrick, Virginia Dejesus-Rueff, John Jordan and Eric Rosengren, 2003, Capital and Risk: New Evidence on Implications of Large Operational Losses, working paper, Federal Reserve Bank of Boston.
- Dekkers, A., J. Einmahl and L. de Haan, 1989, A moment estimator for the index of an extreme-value distribution, *Annals of Statistics* 17, 1833-55.
- Embrechts, Paul, Roger Kaufmann and Gennady Samorodnitsky, 2002, Ruin theory revisited: Stochastic models for operational risk, Working paper, ETH-Zurich and Cornell University.
- Embrechts, Paul, Claudia Klüppelberg and Thomas Mikosch, 1997, *Modelling Extremal Events for Insurance and Finance* (Springer-Verlag, New York).
- Embrechts, Paul, Alexander McNeil and Daniel Straumann, 2002, Correlation and dependence in risk management: properties and pitfalls, In: *Risk Management: Value at Risk and Beyond*, ed. M.A.H. Dempster, Cambridge University Press, Cambridge, pp. 176-223
- Gabaix, Xavier, 1999, Zipf's Law for cities: an explanation, *Quarterly Journal of Economics* 114, 739-767.
- Greene, William, 1997, *Econometric Analysis* (Prentice Hall, Upper Saddle River, NJ).
- Huisman, Ronald, Kees Koedijk, Clemens Kool and Franz Palm, 2001, Tail-index estimates in small samples, *Journal of Business and Economic Statistics* 19, 208-216.

Klugman, Stuart, Harry Panjer and Gordon Willmot, 1998, *Loss Models* (Wiley, New York).

Moscadelli, Marco, 2003, *The modeling of operational risk: the experience from the analysis of the data collected by the Risk Management Group of the Basel Committee*, Working Paper, Bank of Italy.

Netter, Jeffry and Annette Poulsen, 2003, *Operational risk in financial service providers and the proposed Basel capital accord: an overview*, Working Paper, Terry College of Business, University of Georgia.

Risk Management Group, 2003, *The 2002 Loss Data Collection Exercise for Operational Risk: Summary of the Data Collected*, Report to Basel Committee on Banking Supervision, Bank for International Settlements.

Rootzen, Holger and Claudia Klüppelberg, 1999, *A single number can't hedge against economic catastrophes*, *Ambio* 28, No 6, 550-555. Royal Swedish Academy of Sciences.

Table 1. Parametric distributions used for modeling operational loss severity.

Distribution Name	Density, $f(x)$	Maximal Moment
Exponential	$(1/b)\exp(-x/b)$	∞
Weibull	$(\beta x^{\beta-1}/\eta^\beta)\exp(-(x/\eta)^\beta)$	∞
Gamma	$(x/b)^{c-1}[\exp(-x/b)]/[b\Gamma(c)]$	∞
Loggamma	$[\log(x)/b]^{c-1}x^{-1/b-1}/[b\Gamma(c)]$	$1/b$
Pareto	$\xi^{-1}x^{-1/\xi-1}$	$1/\xi$
GPD	$\beta^{-1}(1+\xi x/\beta)^{-1/\xi-1}$	$1/\xi$
Burr	$(\tau/\beta)x^{\tau-1}(1+\xi x^\tau/\beta)^{-1/\xi-1}$	τ/ξ
Lognormal	$(2\pi x^2\sigma^2)^{-1/2}\exp[-(\log(x)-\mu)^2/(2\sigma^2)]$	∞
Loglogistic	$\alpha x^{1/b-1}/[b(1+\alpha x^{1/b})^2]$	$1/b$

Table 2. Goodness of fit across Basel Business Lines and Event Types

The following table reports goodness of fit for each of the distributions under consideration. The test was based on a standard Chisquare procedure, except for the rounding adjustment discussed in section six. The reported figures are Probability values, so that a value of 5% or less indicates a poor fit.

All observations

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	0.0%	6.4%	72.0%	23.3%	13.6%	0.1%
Exponential	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Gamma	0.0%	0.0%	0.0%	0.0%	1.5%	0.0%
LogGamma	0.0%	1.5%	64.1%	33.9%	1.4%	0.7%
LogLogistic	0.0%	6.4%	79.4%	23.8%	2.1%	0.7%
Lognormal	0.0%	0.2%	51.8%	0.0%	3.5%	0.2%
GPD	0.0%	4.5%	75.8%	25.7%	1.6%	0.8%
Weibull	0.0%	0.0%	0.0%	0.0%	54.7%	0.0%

Event Type 1 - Internal Fraud

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	31.7%	86.1%		99.1%		13.0%
Exponential	0.0%	0.0%		0.1%		0.0%
Gamma	12.2%	0.0%		73.4%		0.0%
LogGamma	32.7%	85.6%		98.1%		18.9%
LogLogistic	31.8%	87.4%		98.1%		13.6%
Lognormal	35.0%	86.4%		98.4%		13.7%
GPD	33.1%	87.6%		95.1%		13.3%
Weibull	74.9%	0.4%		40.9%		0.0%

Event Type 2 - External Fraud

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	0.0%	10.8%	6.4%	13.2%	1.7%	0.0%
Exponential	0.0%	0.0%	0.1%	0.0%	0.0%	0.0%
Gamma	0.0%	0.0%	0.7%	0.0%	1.8%	0.0%
LogGamma	0.3%	6.5%	7.6%	7.3%	2.8%	0.0%
LogLogistic	0.0%	5.1%	5.8%	9.3%	2.7%	0.0%
Lognormal	0.0%	0.0%	6.6%	4.8%	3.0%	0.0%
GPD	0.0%	10.5%	6.1%	13.9%	1.7%	0.1%
Weibull	0.0%	0.0%	5.9%	0.1%	9.9%	0.0%

Table 2. Goodness of fit across Basel Business Lines and Event Types*Event Type 3 - Employment Practices and Workplace Safety*

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	87.2%	36.7%		85.0%		23.0%
Exponential	1.0%	0.0%		29.0%		0.0%
Gamma	66.5%	0.0%		86.9%		0.2%
LogGamma	88.1%	7.8%		74.7%		0.2%
LogLogistic	91.7%	59.5%		92.0%		24.5%
Lognormal	95.5%	57.1%		86.7%		24.8%
GPD	92.9%	64.4%		82.0%		35.2%
Weibull	87.0%	0.2%		85.8%		7.6%

Event Type 4 - Clients, Products and Business Practices

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	98.9%	58.0%				37.0%
Exponential	0.0%	0.0%				0.0%
Gamma	0.6%	0.0%				0.0%
LogGamma	80.1%	77.2%				42.2%
LogLogistic	80.8%	58.6%				39.0%
Lognormal	81.4%	36.5%				40.3%
GPD	76.7%	57.4%				34.7%
Weibull	50.1%	0.0%				0.0%

Event Type 7 - Execution, Delivery and Process Management

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	3.0%	1.1%	78.9%	24.7%	78.1%	72.6%
Exponential	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Gamma	0.2%	0.0%	0.0%	0.0%	19.8%	0.0%
LogGamma	0.8%	0.4%	54.7%	22.3%	26.6%	61.7%
LogLogistic	0.1%	0.0%	76.8%	47.7%	89.4%	77.3%
Lognormal	0.1%	0.0%	51.7%	52.1%	68.7%	0.0%
GPD	2.6%	0.0%	77.6%	39.8%	83.6%	89.6%
Weibull	12.0%	0.0%	0.0%	0.6%	47.1%	7.1%

Table 2. Goodness of fit across Basel Business Lines and Event Types*Business Line 2 - Trading and Sales*

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr		1.6%	68.6%	88.1%		58.4%
Exponential		0.0%	0.0%	1.7%		12.4%
Gamma		0.0%	0.0%	1.1%		27.4%
LogGamma		0.0%	65.1%	70.6%		42.1%
LogLogistic		0.0%	69.7%	65.3%		91.8%
Lognormal		0.0%	67.0%	18.8%		86.9%
GPD		0.0%	70.6%	25.1%		58.0%
Weibull		0.0%	0.0%	2.3%		18.3%

Business Line 3 - Retail Banking

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr	0.1%	12.5%	32.3%	8.5%	0.9%	1.7%
Exponential	0.0%	0.0%	0.3%	0.0%	0.0%	0.0%
Gamma	0.0%	0.0%	8.1%	0.0%	0.1%	0.0%
LogGamma	0.0%	0.2%	43.0%	1.3%	5.8%	2.4%
LogLogistic	0.0%	0.0%	35.2%	0.2%	5.6%	3.7%
Lognormal	0.0%	0.0%	46.9%	0.0%	5.5%	3.8%
GPD	0.1%	12.5%	32.2%	9.0%	2.4%	4.7%
Weibull	0.0%	0.0%	15.5%	0.0%	14.7%	0.0%

Business Line 5 - Payment and Settlement

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr		48.5%			11.0%	69.2%
Exponential		0.0%			0.0%	0.0%
Gamma		0.0%			7.2%	1.7%
LogGamma		66.7%			40.2%	62.0%
LogLogistic		49.4%			22.7%	
Lognormal		63.0%			38.5%	63.4%
GPD		45.3%			11.1%	66.8%
Weibull		0.3%			13.3%	52.4%

Business Line 7 - Asset Management

Distribution	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
Burr		64.9%	84.4%		30.1%	20.2%
Exponential		6.4%	0.0%		3.6%	0.0%
Gamma		32.3%	0.0%		43.4%	0.0%
LogGamma		31.3%	79.9%		15.9%	17.6%
LogLogistic		63.1%	63.6%		44.9%	17.4%
Lognormal		45.9%	62.6%		69.8%	18.2%
GPD		67.5%	64.2%		61.1%	20.8%
Weibull		25.7%	4.5%		44.8%	2.3%

Table 3. Parameter estimates for the Generalized Pareto and Lognormal distributions.*Panel a. Estimates of the tail parameter ξ for the GPD.*

	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
All BL & ET	1.28 (0.08)	0.87 (0.03)	0.99 (0.08)	0.92 (0.07)	0.97 (0.11)	1.01 (0.03)
ET1 - IntFrd	1.24 (0.36)	1.31 (0.18)		1.10 (0.38)		1.02 (0.14)
ET2 - ExtFrd	1.17 (0.12)	0.79 (0.05)	0.63 (0.19)	0.69 (0.07)	0.86 (0.14)	0.93 (0.03)
ET3 - EP&WS	0.50 (0.16)	0.42 (0.05)		-0.15 (0.22)		0.50 (0.06)
ET4 - CPBP	1.36 (0.21)	1.25 (0.15)				1.46 (0.13)
ET7 - EDPM	1.42 (0.16)	0.71 (0.05)	0.94 (0.08)	1.00 (0.18)	0.96 (0.17)	0.93 (0.09)
BL2 - T&S		0.68 (0.06)	1.18 (0.13)	0.49 (0.18)		0.42 (0.28)
BL3 - RetBnk	1.15 (0.10)	1.09 (0.05)	0.55 (0.17)	0.94 (0.07)	0.99 (0.14)	0.93 (0.03)
BL5 - P&S		1.06 (0.23)			1.07 (0.35)	1.03 (0.29)
BL7 - AsstMgt		0.49 (0.20)	0.96 (0.21)		0.37 (0.18)	1.64 (0.40)

Panel b. Estimates of $\mu + \sigma^2 / 2$ for the Lognormal distribution.

	Bank A	Bank B	Bank C	Bank D	Bank E	Bank F
All BL & ET	-6.08	> 0	> 0	-21.27	> 0	-9.23
ET1 - IntFrd	-9.85	> 0		> 0		-6.49
ET2 - ExtFrd	-8.35	-22.68	> 0	-5.84	> 0	-22.31
ET3 - EP&WS	> 0	> 0		> 0		> 0
ET4 - CPBP	> 0	> 0				-5.85
ET7 - EDPM	-9.64	> 0	> 0	> 0	> 0	-21.51
BL2 - T&S		> 0	-3.73	> 0		> 0
BL3 - RetBnk	-13.32	-7.45	> 0	-5.24	-11.77	-9.61
BL5 - P&S		-12.09			> 0	-2.01
BL7 - AsstMgt		> 0	> 0		> 0	-2.12

Table 4. Tail parameter estimates based on the HKKP method.

The following table reports tail index estimates calculated under the HKKP regression algorithm. The "Optimal" number of exceedances (K) is chosen to correspond to the linear portion of the Hill plot. Standard errors are reported in parentheses.

Bank ID	"Optimal" K	# of exceedances used in estimation			
		K	0.75K	0.5K	0.25K
A	180	0.823 (0.016)	0.794 (0.020)	0.817 (0.030)	0.717 (0.042)
B	80	0.628 (0.020)	0.591 (0.022)	0.565 (0.029)	0.313 (0.016)
C	30	0.859 (0.085)	0.824 (0.097)	0.952 (0.182)	1.032 (0.353)
D	50	0.498 (0.019)	0.405 (0.015)	0.456 (0.028)	0.415 (0.039)
E	200	0.552 (0.003)	0.534 (0.008)	0.558 (0.013)	0.488 (0.018)
F	50	0.633 (0.030)	0.538 (0.019)	0.536 (0.038)	0.342 (0.026)
All	140	0.681 (0.014)	0.554 (0.012)	0.419 (0.008)	0.305 (0.015)

Table 5. Quantiles of the simulated aggregate loss distribution.

The following table reports quantiles of the simulated aggregate loss distribution. To preserve the confidentiality of the banks in the sample, we scale each percentile for each bank by that bank's assets. The cross-bank median for each percentile is then reported. Panel a presents results under the assumption that loss frequency follows a Poisson distribution whose parameter is estimated separately for each bank (fixed effects model). Panel b presents results under the assumption that loss frequency follows a Poisson distribution whose parameter is a linear function of each bank's asset size (cross-sectional model). Panel c presents results under the assumption that loss frequency follows a Negative Binomial distribution whose parameter is a linear function of each bank's asset size. (Cross-sectional model.)

Panel a. Poisson frequency distribution - fixed effects model.

Severity Distribution	Percentiles of the Aggregate Loss Distribution		
	95	99	99.9
Pareto	0.066%	0.117%	0.468%
Lognormal	0.047%	0.056%	0.070%
Empirical	0.047%	0.053%	0.058%

Panel b. Poisson frequency distribution - cross-sectional model.

Severity Distribution	Percentiles of the Aggregate Loss Distribution		
	95	99	99.9
Pareto	0.106%	0.148%	0.362%
Lognormal	0.089%	0.101%	0.121%
Empirical	0.086%	0.093%	0.102%

Panel c. Negative Binomial frequency distribution - cross-sectional model.

Severity Distribution	Percentiles of the Aggregate Loss Distribution		
	95	99	99.9
Pareto	0.166%	0.237%	0.400%
Lognormal	0.143%	0.198%	0.273%
Empirical	0.146%	0.202%	0.273%

Figure 1. Tail plots of loss data by Basel Event Type.

Event Types are labeled as follows. 1 - Internal Fraud. 2 - External Fraud. 3 - Employment Practices and Workplace Safety. 4 - Clients, Products and Business Practices. 7 - Execution, Delivery and Process Management.

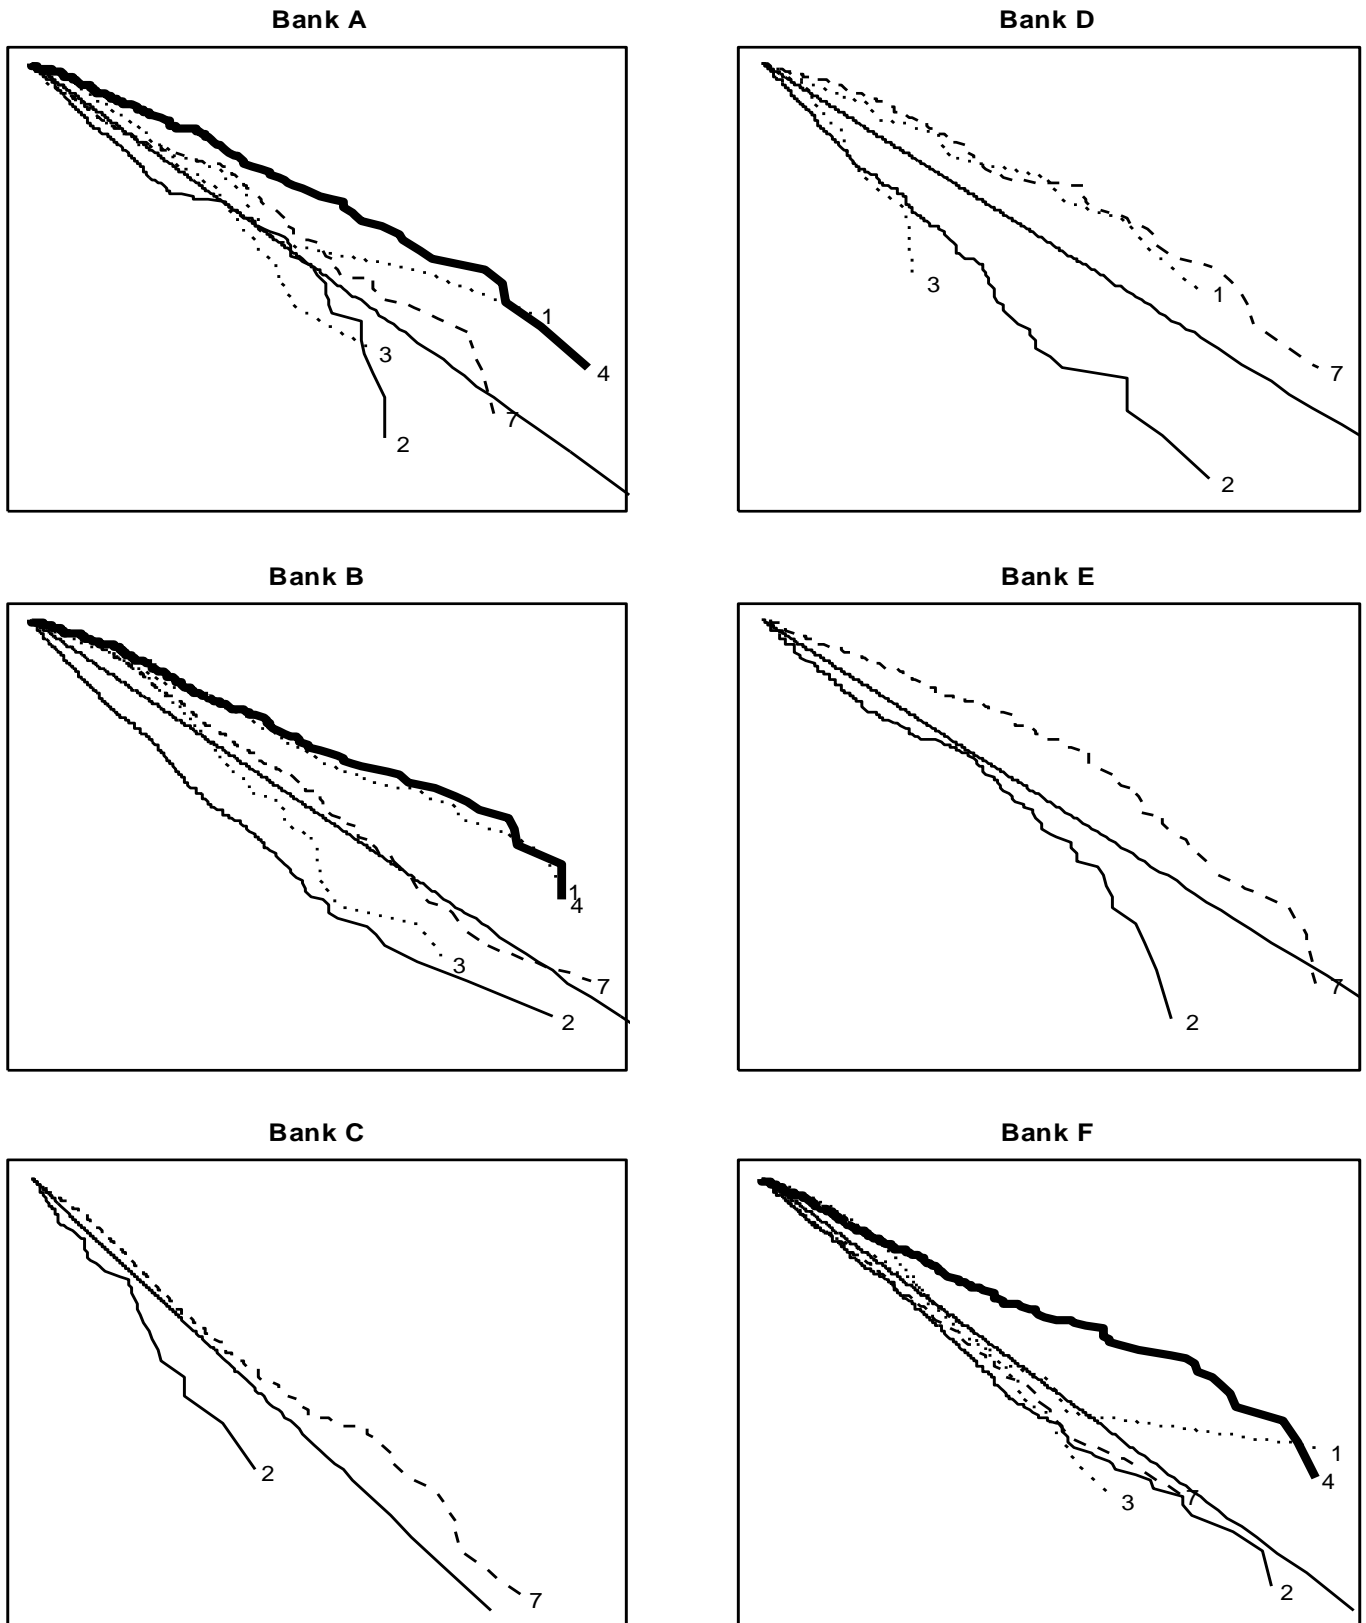


Figure 2. Tail plots of loss data by Basel Business Line.

Business Lines are labeled as follows. 2 - Trading and Sales. 3 - Retail Banking. 5 - Payment and Settlement. 7 - Asset Management.






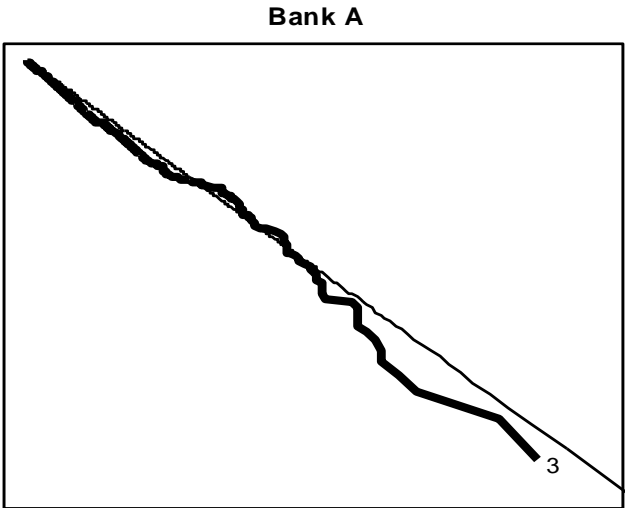


Figure 3. Mean excess plots by Basel Event Type.

Event Types are labeled as follows. 1 - Internal Fraud. 2 - External Fraud. 3 - Employment Practices and Workplace Safety. 4 - Clients, Products and Business Practices. 7 - Execution, Delivery and Process Management.




Figure 4. Mean excess plots by Basel Business Line.

Business Lines are labeled as follows. 2 - Trading and Sales. 3 - Retail Banking. 5 - Payment and Settlement. 7 - Asset Management.









Figure 5. Hill plots of the tail index parameter.

The following are Hill plots of the tail index parameter for the six banks under consideration. The thick dark line indicates the point estimates of the tail parameter as the number of exceedances varies between 1 and 200. The thin dark lines indicate 95% confidence intervals for the point estimates. The thick, medium gray (light gray) line indicates P-values for the Likelihood Ratio test of the hypothesis that the tail parameter is constant across business lines (event types).

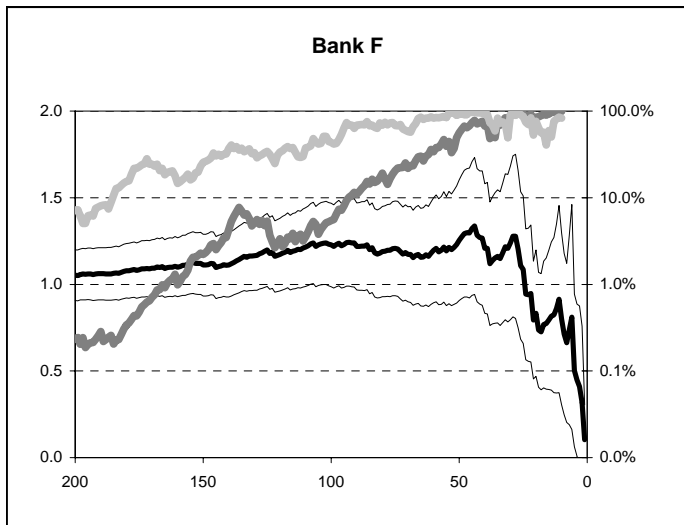







Figure 6. DEdH plots of the tail index parameter.

The following are DEdH plots of the tail index parameter for the six banks under consideration. The solid line indicates the point estimates of the tail parameter as the number of exceedances varies between 1 and 500. The dotted lines indicate 95% confidence intervals for the point estimates.

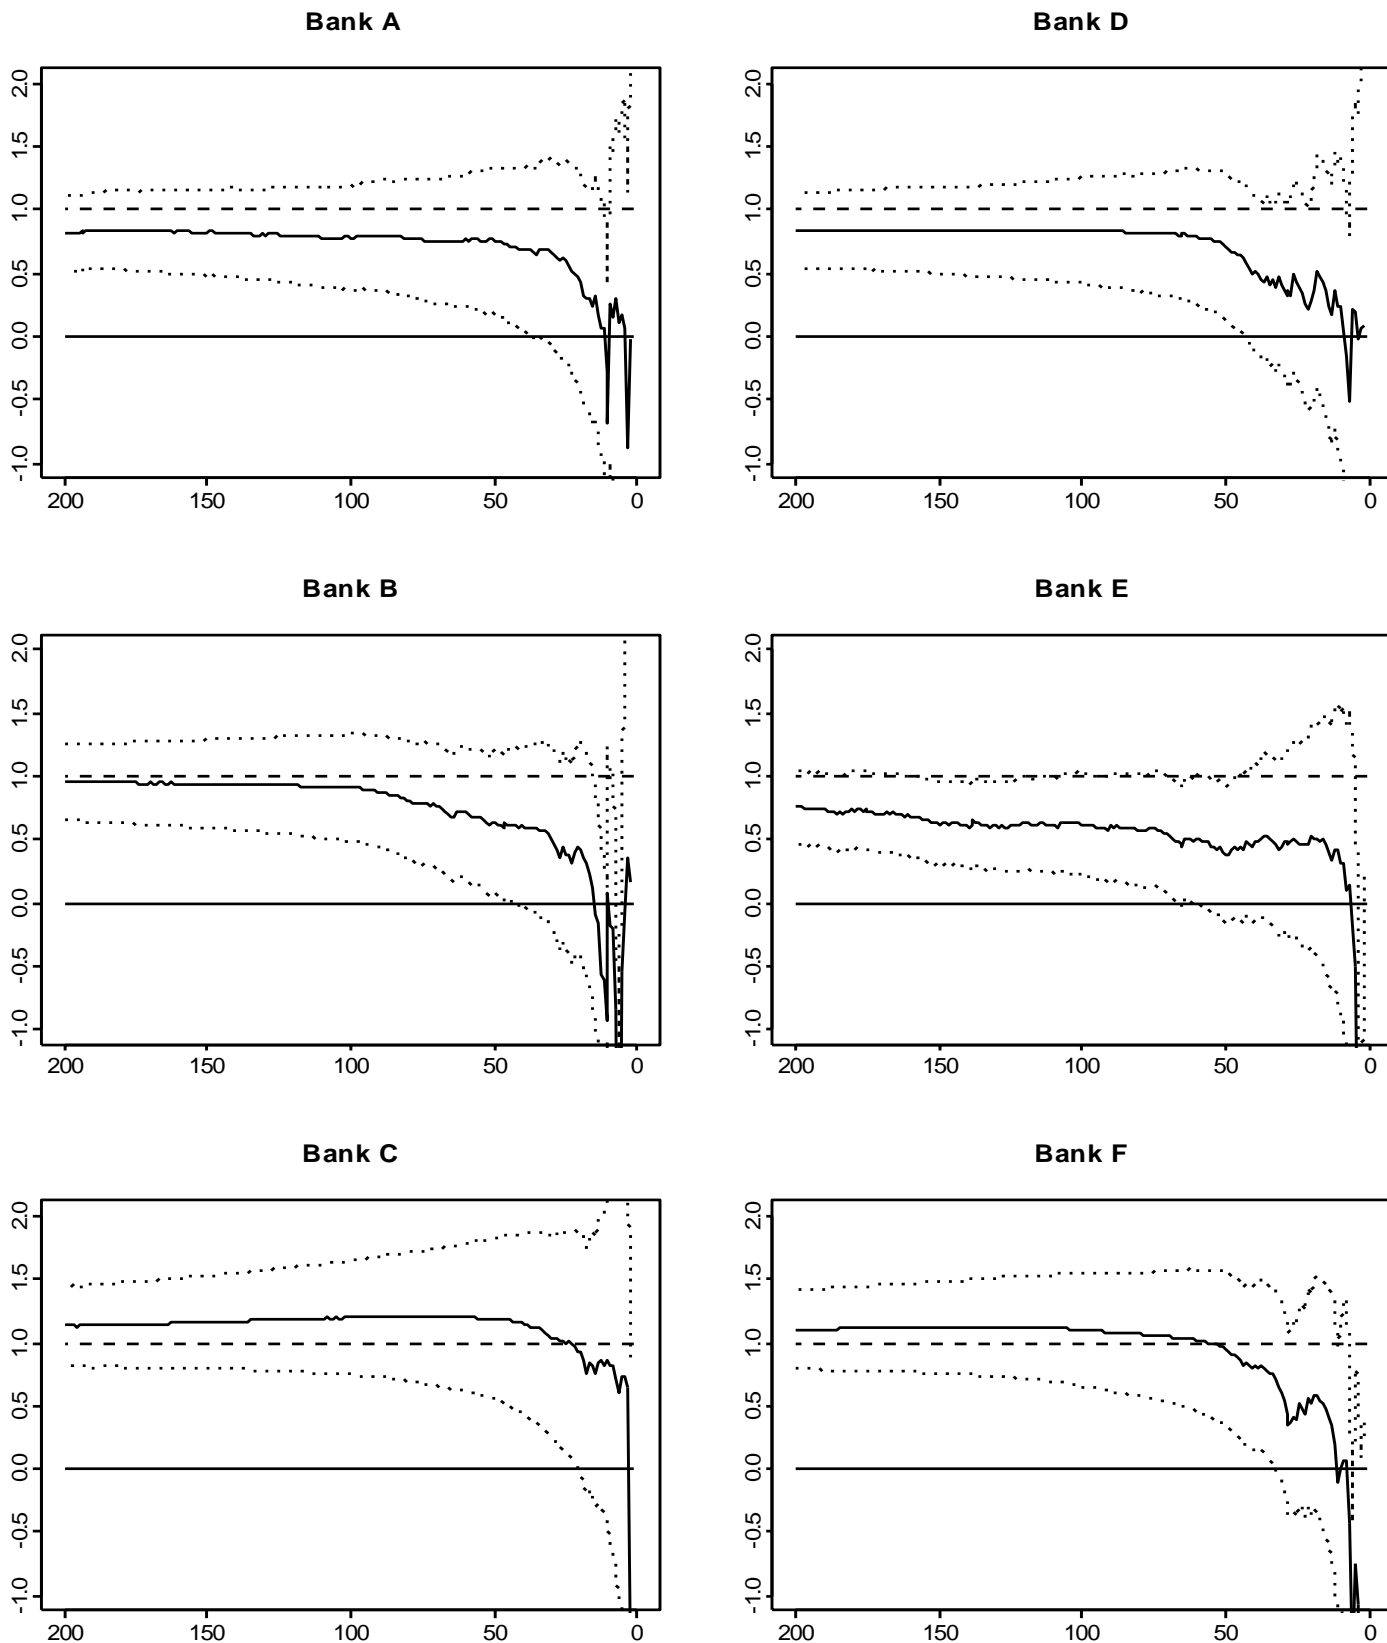


Figure 7a. Threshold plots of the lognormal parameter μ .

The following are threshold plots of the lognormal parameter μ for the six banks under consideration. The solid line indicates the point estimates of μ as the number of exceedances varies between 10 and 200. The dotted lines indicate 95% confidence intervals for the point estimates. Labels are omitted from the vertical axis to preserve confidentiality.






Figure 7b. Threshold plots of the lognormal parameter σ .

The following are threshold plots of the lognormal parameter σ for the six banks under consideration. The solid line indicates the point estimates of σ as the number of exceedances varies between 10 and 200. The dotted lines indicate 95% confidence intervals for the point estimates. Labels are omitted from the vertical axis to preserve confidentiality.


Bank A




Bank D




Bank B



Bank E



Bank C



Bank F

