

BANCA D'ITALIA

Temi di discussione

del Servizio Studi

**The modelling of operational risk: experience with
the analysis of the data collected by the Basel Committee**

by Marco Moscadelli



Number 517 - July 2004

The purpose of the Temi di discussione series is to promote the circulation of working papers prepared within the Bank of Italy or presented in Bank seminars by outside economists with the aim of stimulating comments and suggestions.

The views expressed in the articles are those of the authors and do not involve the responsibility of the Bank.

This paper was submitted to the Risk Management Group of the Basel Committee for evaluation of the confidentiality of individual banks' data. However the views and conclusions expressed in the paper do not necessarily reflect those of the Risk Management Group nor those of its members.

The main purpose of the paper is to increase the level of understanding of operational risk within the financial system, by exploring the statistical behaviour of the data collected by the Basel Committee in the second Loss Data Collection Exercise and going on to derive bottom-up capital charge figures and Gross Income related coefficients for the Business Lines envisaged in the revised framework of the Capital Accord.

To the end of promoting the development of a pragmatic and effective dialogue on the operational risk measurement issues between regulators and institutions in the course of the implementation of the new Capital framework, feedback from academics and practitioners on the topics dealt with in the paper is welcomed. Specifically, the statistical methodologies implemented to model the data and the outcomes of the analysis, in terms of the capital charges and coefficients of the different Business Lines, are the topics on which comments should be focused.

For assistance in the feedback process and in order to submit comments, please refer to the procedures set by your country's competent national supervisory authority or central bank. Comments should also be sent, preferably via e-mail, to the author.

Editorial Board:

STEFANO SIVIERO, EMILIA BONACCORSI DI PATTI, FABIO BusetTI, ANDREA LAMORGESE, MONICA PAIELLA, FRANCESCO PATERNÒ, MARCELLO PERICOLI, ALFONSO ROSOLIA, STEFANIA ZOTTERI, RAFFAELA BISCEGLIA (Editorial Assistant).

THE MODELLING OF OPERATIONAL RISK: EXPERIENCE WITH THE ANALYSIS OF THE DATA COLLECTED BY THE BASEL COMMITTEE

by Marco Moscadelli*

Abstract

The revised Basel Capital Accord requires banks to meet a capital requirement for operational risk as part of an overall risk-based capital framework. Three distinct options for calculating operational risk charges are proposed (Basic Approach, Standardised Approach, Advanced Measurement Approaches), reflecting increasing levels of risk sensitivity. Since 2001, the Risk Management Group of the Basel Committee has been performing specific surveys of banks' operational loss data, with the main purpose of obtaining information on the industry's operational risk experience, to be used for the refinement of the capital framework and for the calibration of the regulatory coefficients. The second loss data collection was launched in the summer of 2002: the 89 banks participating in the exercise provided the Group with more than 47,000 observations, grouped by eight standardised Business Lines and seven Event Types. A summary of the data collected, which focuses on the description of the range of individual gross loss amounts and of the distribution of the banks' losses across the business lines/event types, was returned to the industry in March 2003. The objective of this paper is to move forward with respect to that document, by illustrating the methodologies and the outcomes of the inferential analysis carried out on the data collected through 2002. To this end, after pooling the individual banks' losses according to a Business Line criterion, the operational riskiness of each Business Line data set is explored using empirical and statistical tools. The work aims, first of all, to compare the sensitivity of conventional actuarial distributions and models stemming from the Extreme Value Theory in representing the highest percentiles of the data sets: the exercise shows that the extreme value model, in its Peaks Over Threshold representation, explains the behaviour of the operational risk data in the tail area well. Then, measures of severity and frequency of the large losses are gained and, by a proper combination of these estimates, a bottom-up operational risk capital figure is computed for each Business Line. Finally, for each Business Line and in the eight Business Lines as a whole, the contributions of the expected losses to the capital figures are evaluated and the relationships between the capital charges and the corresponding average level of Gross Incomes are determined and compared with the current coefficients envisaged in the simplified approaches of the regulatory framework.

JEL classification: C11, C13, C14, C19, C29, C81, G21, G28.

Key words: operational risk, heavy tails, conventional inference, Extreme Value Theory, Peaks Over Threshold, median shortfall, Point Process of exceedances, capital charge, Business Line, Gross Income, regulatory coefficients.

* Bank of Italy, Banking Supervision Department.

Contents

1. Introduction and main results	9
2. Data characteristics and assumptions	15
3. Exploratory data analysis	18
4. Conventional inference: Business Lines tail-severity estimate.....	21
5. Extreme Value Theory: theoretical background	26
6. Peaks Over Threshold approach: Business Lines tail-severity estimate	34
7. Advanced tests on the severity results	43
8. Business Lines tail-severity measures and magnitude	48
9. Peaks Over Threshold approach: Business Lines tail-frequency estimate	53
10. Business Lines capital charge.....	59
11. Business Lines capital charge and Gross Income relationship: bottom-up coefficients..	66
12. Conclusions	67
References	70

1. Introduction and main results ¹

Operational risk has become an area of growing concern in banking. The increase in the sophistication and complexity of banking practices has raised both regulatory and industry awareness of the need for an effective operational risk management and measurement system. From the time of the release of the second consultative document on the New Capital Accord in 2001, the Basel Committee on Banking Supervision has established a specific treatment for operational risk: a basic component of the new framework is represented by Pillar 1, which explicitly calls for a minimum capital charge for this category of risk ².

The proposed discipline establishes various schemes for calculating the operational risk charge, ranging from a crude Basic Approach, based on a fixed percentage of Gross Income - the indicator selected by the Committee as a proxy of banks' operational risk exposure - passing through an intermediate Standardised Approach, which extends the Basic method by decomposing banks' activities and, hence, the capital charge computation, into eight underlying business lines³, to the most sophisticated approaches, the Advanced Measurement Approaches (AMA), based on the adoption of banks' internal models. The framework gives banks a great deal of flexibility in the choice of the characteristics of their internal models ⁴, provided they comply with a set of eligible qualitative and quantitative criteria and can demonstrate that their internal measurement systems are able to produce reasonable estimates of unexpected losses.

¹ For the people who have no time to explore the whole paper, the reading of the Sections 1, 2, 10, 11 and 12 may suffice in order to have a quick understanding of the main, significant, topics dealt with in this analysis.

The author would like to thank Giovanni Carosio, Stefano de Polis, Paolo Zaffaroni, Fabrizio Leandri, Giuseppe De Martino and two anonymous referee for comments and fruitful discussions and Giorgio Donato for his careful reading and corrections, which helped the author in improving the final form of this paper. A special thanks to Michele Romanelli who implemented the algorithms used for the bootstrapping and conventional analysis. E-mail: marco.moscadelli@bancaditalia.it

² The new Accord is based on a three Pillar concept, where Pillar 1 corresponds to a Minimal Capital requirement, Pillar 2 stands for a Supervisory Review process and Pillar 3 concerns Market discipline.

³ The eight business lines established by the Accord are: *Corporate Finance, Trading & Sales, Retail Banking, Commercial Banking, Payment & Settlement, Agency Services, Asset Management and Retail Brokerage*.

⁴ The new discipline establishes that the capital calculation must be based on a sound combination of qualitative and quantitative elements: internal data, relevant external data, scenario analysis and bank-specific business environment and internal control factors.

Since the first release of the new Basel proposal, regulators, practitioners and academics have been engaged in discussion on how to define and measure operational risk and, hence, how to determine appropriate capital requirements.

As regards the definition aspects, the Risk Management Group (RMG) of the Basel Committee and industry representatives have agreed on a standardised definition of operational risk, i.e. “the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events”. This definition, which includes legal risk and excludes strategic and reputational risk, relies on the categorisation of operational risks based on the underlying causes. A standardised classification matrix of operational risk into eight Business Lines (BLs) and seven Event Types (ETs) has also been defined, in order to encourage greater consistency of loss data collection within and between banks.

As regards the measurement issue, a growing number of articles, research papers and books have addressed the topic from a theoretical point of view. In practice, this objective is made hard by the relatively short period over which operational risk data have been gathered by banks; obviously, the greatest difficulty is in collecting information on infrequent, but large losses, which, on the other hand, contribute the most to the capital charge. The need to evaluate the exposure to potentially severe tail events is one of the reasons why the new Capital framework requires banks to supplement internal data with further sources (external data, scenario analysis) in order to compute their operational risk capital charge.

Since 2001, the RMG has been performing surveys of banks’ operational loss data, with the main purpose of obtaining information on the industry’s operational risk experience, useful for improving the capital framework and calibrating the regulatory coefficients. In particular, the second Loss Data Collection Exercise (2002 LDCE), serving to collect the operational risk losses borne by banks in the financial year 2001, was an extension and refinement of the previous exercises sponsored by the RMG⁵. Overall, 89 banks participated in the 2002 survey, providing the RMG with more than 47,000 observations, mapped in the standardised matrix BLs/ETs. Feedback on the data collected, which focuses on the description of the range of individual gross loss amounts and of the distribution of these

⁵ A description of the information collected in the previous exercises can be found in the “Working Paper on the Regulatory Treatment of Operational Risk” released in September 2001 and in the paper “The Quantitative

losses across the BLs/ETs categories, was provided to industry in March 2003 (see the document “The 2002 Loss Data Collection Exercise for Operational Risk: Summary of the Data Collected”, published on the BIS website).

The objective of the present paper is to move forward with respect to that document, by illustrating the methodologies and the outcomes of the inferential analysis carried out on the operational risk losses collected through 2002. In order to statistically explore the data, first a pooling exercise of the banks’ losses according to a BL criterion is performed, then the operational riskiness of each BL data set is examined by means of empirical and statistical tools. Several practical and theoretical reasons support this choice rather than exploring any or some individual banks’ database. From a practical point of view, the objective of measuring and comparing the operational riskiness of the BLs in addition to that of providing protection to the confidentiality of the LDCE banks’ data. From a theoretical point of view, the fact that the aggregation of banks’ operational risk data, collected in short time windows (1-year, say), is a viable solution to actually overcome the threats of non-repetitiveness and dependence of the observations, which typically affect any individual banks’ historical database; furthermore, each BL data set, obtained by assembling 1-year period data from n banks having similar size and characteristics (the 2002 LDCE banks), can be thought as referred to a medium-sized (large internationally active) bank and collected over a time window of n -year. In practice, by a cross-section pooling procedure, long time-series of i.i.d. operational risk data are reproduced.

The first purpose of the work is to compare the sensitivity of conventional actuarial distributions and models stemming from the Extreme Value Theory (EVT) in representing the extreme percentiles of the data sets (i.e. the large losses). Then, measures of severity and frequency of the large losses in each data set are gained and, by a proper combination of these estimates, a bottom-up operational risk capital charge is computed. Finally, for each BL and in the eight BLs as a whole, the contributions of the expected losses to the capital figures are evaluated and the relationships between the capital charges and the corresponding average level of Gross Incomes are determined and compared with the current regulatory coefficients.

It is evident that the reliability of the exercise is strictly connected with the (unknown) actual quality of the overall data, which, as the 2002 LDCE summary stresses, have been gathered by banks according to different levels of consistency.

The results indicate a low performance of conventional actuarial severity models in describing the overall data characteristics, summarizable in very high levels of both skewness to the right and kurtosis. Indeed, any traditional distribution applied to all the data in each BL tends to fit central observations, hence not taking the large losses into adequate consideration. On the other hand, the exercise shows that the Extreme Value model, in its severity representation (Peaks Over Threshold-Generalised Pareto Distribution, POT-GPD), provides an accurate estimate of the actual tail of the BLs at the 95th and higher percentiles; this is confirmed by the results of three goodness-of-fit tests and a severity VaR performance analysis.

The POT-GPD model reveals that, while each BL severity riskiness increases substantially at the highest percentiles because of the heaviness of the tail, the ranking of the riskiness of the BLs does not change significantly. In particular *Corporate Finance* and *Commercial Banking* are found to be the riskiest BLs with an estimated severity loss at the 99.9th percentile of € 260 million and € 151 million, respectively. On the other hand, *Retail Banking* and *Retail Brokerage* are the least risky BLs, showing severity loss at the 99.9th percentile of € 17 million and € 27 million respectively.

In light of its supremacy in the estimate of the loss tail-severity distribution, the Extreme Value model, in its Peaks Over Threshold - Point Process representation (POT-PP), is also used to estimate the loss tail-frequency distribution, that is to derive the probability of occurrence of the large losses in each BL.

The results show the significant per-bank variability of the number of large losses in each BL. The reasons for this may be found in the different level of comprehensiveness in the collection of (large) losses between the banks participating in the RMG survey and perhaps also in the short time horizon of the 2002 LDCE (1-year data collection), which might have caused, for some banks, a few gaps in the collection of very rare and large losses. Another likely cause of the variability of the frequency of large losses could lie in the participation, in the 2002 LDCE, of banks having different size and hence potentially in a

position to produce, in some BLs, a lower or higher number of large losses in a given time horizon. These issues are specifically addressed in this paper and their, potential, misleading effects on the estimate of the BLs frequency of large losses mitigated. In particular the possible incompleteness of very large losses is overcome by placing a floor on the, 1-year period, probability of occurrence of the losses with a single-impact magnitude bigger than the 99th percentile of the severity distribution: the floor is represented by the number of large losses occurring at the 99th percentile of the frequency distribution. The potential differences in banks' size is treated by assuming the existence in the panel of two distinct groups of banks – a “lower group”, consisting of banks having smaller size (in fact domestic banks), and an “upper group”, consisting of banks having larger size (in fact internationally active banks) – for which separate analyses are made on the basis of the estimated, distinct, 1-year numbers of large losses. In particular, for a typical international active bank, the model reveals about 60 losses bigger than € 1 million per year; this figure is absolutely comparable with that actually borne by large internationally active banks.

On the basis of the POT tail severity and frequency estimates, an aggregate figure for each BL and for the eight BLs as a whole is computed by means a semiparametric approach. The POT approach appears to be a viable solution to reduce the estimate error and the computational costs related to the not analytical techniques, like the MonteCarlo simulation, usually implemented in the financial industry to reproduce the highest percentiles of the aggregate loss distribution. The findings clearly indicate that operational losses represent a significant source of risk for banks, given a 1-year period capital charge against expected plus unexpected losses at the 99.9th percentile which amounts to € 1,325 million for a typical international active bank and to € 296 million for a domestic bank. Owing to the higher frequency of losses, *Retail Banking* and *Commercial Banking* are the BLs which absorb the majority of the overall capital figure (about 20 per cent each), while *Corporate Finance* and *Trading & Sales* are at an intermediate level (respectively close to 13 per cent and 17 per cent) and the other BLs stay stably under 10 per cent. These figures are comparable with the allocation ratios of economic capital for operational risk reported by banks in the 2002 LDCE (see Table 21 of the cited summary). Moreover, the results show the very small contribution of the expected losses to the total capital charge: on average across the BLs, they amount to less than 3 per cent of the overall capital figure for an international active

bank, with a minimum value of 1.1 per cent in *Corporate Finance* and a maximum of 4.4 per cent in *Retail Banking*. Once again, these outcomes confirm the very tail-driven nature of operational risk.

Finally, for the banks belonging to the “upper group” (the international active banks), the relationships between the BLs overall capital figures and the average level of the Gross Incomes are computed and compared with the current regulatory coefficients envisaged in the Basic and Standardised Approach of the Capital Accord (the so-called Alpha and Betas). For the eight BLs as a whole, the results show a slightly lower ratio than the current regulatory coefficient, hence giving an incentive to move from the Basic to the Standardised Approach and meeting, at the same time, the objective of not increasing the industry overall level of capital requirement for operational risk. Nevertheless, adjustment of the coefficient of some BLs might more effectively capture the actual operational riskiness shown by the data (in particular, a sizable reduction in the *Trading & Sales* and *Retail Banking* and an increase in the *Payment & Settlement* and *Retail Brokerage* Betas).

This paper is organised as follows: Section 2 describes the main characteristics of the raw data used in the work and the choices made in terms of data assumption and treatment; in Section 3 each BLs data set is explored by an empirical analysis which focuses on a bootstrapping procedure and a graphical representation of the density function of the BLs; Section 4 illustrates the main results of the conventional inference on the severity of losses; Section 5 describes the theoretical background of EVT, while Sections 6, 7 and 8 are devoted to estimate, test and measure the tail-severity of the eight BLs by the POT-GPD model. Section 9 is devoted to computing the probability of occurrence of the tail of the BLs, by means of the frequency component of the POT approach (POT-PP). In Section 10 the capital charge of the BLs against expected plus unexpected losses is computed and compared with the contribution pertaining to expected losses alone. Section 11 focuses on the relationship between the estimated capital figures of the BLs and the pertaining average level of Gross Incomes. Section 12 concludes.

2. Data characteristics and assumptions

The data used for the analysis are those collected in the 2002 LDCE by the RMG, which required the 89 banks participating in the survey to provide individual gross operational losses above the threshold of €10,000 for the year 2001, grouped by quarters.

In order to statistically explore these data, all the individual banks' data were pooled according to a BL criterion, leading to eight distinct data sets, each one relative to a different BL. Practical and theoretical reasons determined the decision to pool the individual banks' data by BLs:

- a) the need to protect the confidentiality of the individual banks' losses;
- b) the need to have, for each data set, a number of losses high enough to be modelled;
- c) the objective of assessing and comparing the operational riskiness of the eight BLs.

In doing so, it may be that, due to the particular nature of operational risk, a few inconsistencies in some loss data sets may have arisen (e.g. how should the riskiness stemming from pooling very different ETs as *Internal Fraud* and *Damage to Physical Assets* be interpreted?). Anyway, it should be observed that data inconsistencies could also arise from pooling losses according to an ET criterion (*Internal Fraud* in *Trading & Sales*, for example, seems to appear completely different from *Internal Fraud* in *Retail Banking*) or simply from handling data referred to a specific BL/ET combination, originating from banks with very different operational risk profiles.

Moreover, it should be noted that, since operational risk spreads over the different activities of a bank organisation, any loss analysis is potentially exposed to the threat of inconsistencies of data, when they refer to sources that are not properly categorised: the problem, which could condition the quality and the results of the inference, therefore lies within banks and only later between banks. Sound practices require banks to conduct a rigorous and detailed classification of their products, functions and processes and to adopt a clear and widespread definition of operational risk in their organisational units before any loss event identification and mapping is conducted and a statistical analysis of losses is made (on the issues of operational risk definition and categorisation, see, for example, Samad-Khan, 2003).

Table 1 reports the results of the pooling exercise for each BL, in terms of the number of banks providing at least one loss figure and the total number of observations.

Table 1: BLs data pooling exercise

BUSINESS LINE	n. banks providing loss data	n. total observations
BL 1 (Corporate Finance)	33	423
BL 2 (Trading & Sales)	67	5,132
BL 3 (Retail Banking)	80	28,882
BL 4 (Commercial Banking)	73	3,414
BL 5 (Payment & Settlement)	55	1,852
BL 6 (Agency Services)	40	1,490
BL 7 (Asset Management)	52	1,109
BL 8 (Retail Brokerage)	41	3,267

Overall, the number of observations after the pooling exercise is 45,569, a value slightly lower than that of the 2002 LDCE, the difference being caused by the decision to exclude data not having a business line breakdown (1,699 observations, see Table 3 of the 2002 LDCE summary) as well 1 observation in BL8, which proves to be an outlier. In BL4, 5 out of 3,414 observations tend to be outliers, too: even if kept in the sample, they are not taken into consideration in the fitting exercise, and are considered only in the VaR performance analysis made in Section 7.

Each data set, referred to a distinct BL, can be considered as a sample extracted from the corresponding unknown population, whose properties, characteristics and riskiness are to be detected and compared. It is assumed that, in each BL, the observations are the realisations of independent, identically distributed (i.i.d.) random variables.

The issue of not dependence of the operational risk losses, like the non-stationarity, and its possible effect on the modelling exercise has been more recently addresses by industry and academia. In particular, Embrechts *et al.*, 2003, identify the causes of time-

structural changes of operational losses both in survival bias ⁶ and in changes in banks external/internal environment (e.g. changes in the economic cycle, volume of business, or organisational or internal control systems). As the authors state, the non-stationarity condition can distort the results of the applied statistical models, which are mainly based on the i.i.d. assumption: the authors therefore stress the importance of modelling the non-stationarities of data before a statistical analysis can be made.

In the current exercise, the data independence assumption is mainly based on the idea that any pooling exercise of banks' operational risk losses collected in a reasonably short time (say, 1 or 2 years), can in fact mitigate the threat of dependence of the data, which, on the contrary, might be present in individual banks' historical data-bases. This assumption arises from the consideration that the operational risk losses usually occur independently in each bank, as they are mainly caused by banks internal drivers (process, human resource and internal system). Moreover, if the losses refer to moderately short time horizons (e.g. 1 or 2 years), the risks of non-stationarity of each bank's database - caused, as noted before, by possible survival bias or changes in the bank's external/internal environment as time evolves – should also be reduced ⁷.

The identically distributed data assumption is based on the consideration that banks having characteristics not too dissimilar (as the banks participating in the 2002 LDCE) are not distinguishable by the severity of losses, since they are indifferently exposed to losses of any size ⁸.

⁶ Survival or selection bias is defined as the fact that operational losses that occurred some distance in the past have not survived in current bank databases. As the authors state, the early losses seem not only sparser, but also larger. In practice, one only remembers the largest losses (on this topic, see also Embrechts *et al.*, 2004).

⁷ With regard to event types caused by bank external drives, some form of dependence between data collected from different banks may arise (e.g. an earthquake can contemporaneously damage the physical assets of several banks). However, in this exercise, the data independent assumption for such events was tested and verified on the eight BLs (see Section 3 below). On the data dependence, the interested readers can see the examples discussed by Embrechts and Samorodnitsky, 2002, or the analysis conducted by Ebnother *et al.*, 2001, aimed to assess the impact of specific risk factors (i.e. fraud, system failures, error, external catastrophes) on several production processes.

⁸ In this exercise the issue of the potential non-homogeneity of the operational risk data is addressed only with regard to the frequency component of losses (see Section 9 below, in particular footnote 36). The similarities in the loss severity across the 2002 LDCE banks have been also detected by de Fontnouvelle *et al.* (2004), which have recently examined the empirical regularities of operational risk in six large international active banks participating in the survey launched by the Basel Committee.

As a result of the pool exercise and the i.i.d. assumptions, two practical features originate:

1. the overcome of the threat of the non-repetitiveness of the losses, which indeed represents one of the biggest concerns in the statistical treatment of the operational risk data (see Embrechts *et al.*, 2003);
2. the fact that the collection of 1-year period data from n banks having similar sizes and characteristics can reasonably represent the collection of n -year period data referred to a bank having average size and characteristics. Under this perspective, the current analysis assumes thus that the pool of 1-year period data from the 89 banks participating in the 2002 LDCE (cross-section analysis) is equivalent to the collection of data from a medium-sized LDCE bank during a time window of 89-years (time-series analysis). In other words the whole data set can be thought of as referred to a large internationally active bank and collected over time.

3. Exploratory data analysis

A preliminary exploratory analysis of the operational raw data is conducted in order to gain information on the actual underlying structure of the severity of the eight BL data sets. In particular, owing to the known nature of operational risk and to the ultimate goal of assessing the riskiness of the BLs, the analysis focuses on the evaluation of the levels of asymmetry and tail-heaviness of the data sets (that is measures of skewness and kurtosis) rather than on the location and scale.

In addition, instead of exploring the eight-pooled data sets, new data groups are generated from the original ones on the basis of a bootstrapping procedure. The aim is twofold: to strengthen the informative power of the raw data on the unknown moments of the population and, above all, to provide further protection to the confidentiality of the losses

As a general rule, if substantial differences in terms of the behaviour of the losses were detected for some banks, suitable statistical treatments (so-called “scaling methodologies”) would be required to make data comparable and to ensure that merging all the individual databases leads to unbiased estimates (for a recent scaling proposal, see Frachot and Roncalli, 2002, who address the problem of mixing banks internal and external data).

reported by the individual banks in the 2002 LDCE ⁹. In practice, a resampling technique with replacement is applied to each original BL data set. The steps of the bootstrap are the following:

- a) generating a random number from integers 1, 2,...,n, where n is the BL sample size. Let j be this number;
- b) obtaining the j-th member of the original sample;
- c) repeating the first 2 steps n times (because of replacement, the same value from the original sample may be selected more than once);
- d) computing the parameter estimates from these n new values;
- e) repeating 1,000 times steps 1 through 4.

The large number of bootstrapping estimates can be considered as a random sample from the sampling distribution of each parameter estimator being calculated: the mean of the bootstrap samples is a good indicator of the expected value of the estimator.

In Table 2, the outcomes of the bootstrapping procedure are reported:

Table 2: BLs bootstrapping results

BUSINESS LINE	Mean (euro ,000)	Standard deviation (euro ,000)	Skewness	Kurtosis
BL 1 (Corporate Finance)	646	6,095	16	294
BL 2 (Trading & Sales)	226	1,917	23	674
BL 3 (Retail Banking)	79	877	55	4,091
BL 4 (Commercial Banking)	356	2,642	15	288
BL 5 (Payment & Settlement)	137	1,320	24	650
BL 6 (Agency Services)	222	1,338	13	211
BL 7 (Asset Management)	195	1,473	25	713
BL 8 (Retail Brokerage)	125	1,185	32	1,232

⁹ As regard the first purpose, according to the bootstrap theory, this procedure provides a reliable indication of the properties of the population parameters estimator (see Efron and Tibshirani, 1993). Concerning the confidentiality issue, the replacement of the raw data estimates with the bootstrapped ones definitely removes the risk of an individual bank being identifiable due to its relative importance in the determination of some moments of the BLs.

The estimates of the bootstrapped moments indicate that the empirical distributions of the eight BLs are very skewed to the right and, above all, very heavy in the tail. In order to better appreciate the peculiarities of these data, it should be remembered that the skewness and the kurtosis of a standard LogNormal distribution are equal, respectively, to 6 and 114.

Therefore, despite the short time-window of the 2002 LDCE, the pooled data sets appear to capture the large-impact events, a very important pre-condition for obtaining consistent estimates of the capital figures. One reason may be the circumstance that the reference year for the 2002 LDCE was 2001, the year of the September 11th terroristic attack. As the 2002 LDCE summary paper remarks “.. *the distribution of gross loss amounts, in particular, is likely to be sensitive to the incidence of relatively few very large-impact events. This phenomenon is certainly evident in the 2001 data, which contain some large individual loss amounts associated with events of September 11, for example*”¹⁰. In light of that, it may be that, owing to the September 11th event, the pooled data sets violate the assumption of independence of the observations. In reality, a deeper analysis conducted on the pooled data clearly indicates that the very large losses are spread out across the BLs and do not converge to one or just a few BLs¹¹. In any case, the assumption that, in each BL data set, the observations are independent should be preserved or, at most, moderately weakened.


In order to have a preliminary idea on the graphical behaviour of the losses, a kernel smoothing technique¹² was performed for each original (that is, before the bootstrapping) BL data set. The kernel procedure makes it possible to obtain a nice graphical representation of the BLs density function, by smoothing the histogram figure driven by the original data (see Figure 1). In practice, a probability value is assigned to each observation based on the mass of data close to it: the denser the data close to the observation whose probability is to be evaluated, the higher the probability assigned.

¹⁰ Table 6 of that paper shows that more than three-quarters of the total gross loss arise from only 2.5 per cent of the events.

¹¹ The reason could lie in the fact that the losses, even if caused by just a few common drivers (as, for instance, the September 11th attack) may have affected distinct businesses of the banks (i.e., the September 11th event may have affect the different activities conducted by the banks in the Twin Towers building).

¹² The Epanechnikov kernel was used.

Figure 1: kernel density function for some BLs



The kernel density functions clearly show the skewness of the data, but not the kurtosis.

4. Conventional inference: Business Lines tail-severity estimate

It should be first observed that the cut-off limits (€ 10,000) established in the 2002 LDCE to report losses are not taken into account in modelling the severity of the data. The reason mainly lies in the fact that some banks used different minimum cut-off levels in providing the RMG with their data; therefore each BL pooled data set also contains losses below the threshold of € 10,000. Moreover, as it will become clearer in the remainder of this paper, given the actual nature of the operational risk losses, any statistical model which correctly represents the body of the data (that is the small/medium-sized losses) may have serious drawbacks in fitting the tail area. In the light of the objective of the analysis, i.e. to

gain information on the tail magnitude of the eight BLs, the choice between a ground-up, a truncated or a shifted distribution to estimate the severity of the data becomes immaterial¹³.

The aim of this section is to apply, separately for each BL, conventional inference to the original pooled data, bearing in mind the ultimate goal of detecting the curve that best explains the behaviour of the severity of losses in the tail area. This is done by fitting parametric distributions to the eight overall data sets and obtaining a parameters estimate that optimises the criterion of maximum likelihood.

Several distributions are fitted to the data, according to an increasing level of kurtosis, i.e. starting from light-tail distributions (as Weibull), passing through medium-tail curves (as Gamma, Exponential, Gumbel and LogNormal) to heavy-tail models (as Pareto).

The Kolmogorov-Smirnov (K-S) and Anderson-Darling (A-D) tests are adopted to reject or to accept the null hypothesis that the data originate from the selected distribution with the estimated parameters. In light of the objectives of the exercise, the A-D test seems to be more suitable because it is much more sensitive to the tails of data.¹⁴

The results indicate that Gumbel and LogNormal are the distributions that best fit the data in each BL. Both lighter-tail (as Gamma) and heavier-tail (as Pareto) functions result in much higher test values.

In Figure 2, referred to BL1 (*Corporate Finance*), the plots of the Gumbel and LogNormal cumulative distribution functions can be compared with the empirical distribution functions, the latter being defined, for a sample of size n , as


$$F_n(x) = \frac{1}{n} \sum_{i=1}^n 1_{\{X_i \leq x\}}, \text{ i.e. the number of observations less than or equal to } x \text{ divided by } n.$$

The LogNormal curve seems to provide a reasonable fit to the whole data set, while the Gumbel curve fits poorly in the body of the data but better in the tail.

¹³ Frachot and Roncalli, 2002, and Baud *et al.*, 2002, discuss these and other related issues in the context of pooling internal and external databases which are truncated from below.


¹⁴ The attractive feature of the K-S test is that it is “distribution free”, in the sense that the critical values do not depend on the specific distribution being tested. Despite that, it has the disadvantage of being more sensitive near the center of the distribution than at the tails. On the other hand, the A-D test makes use of any specific distribution in calculating critical values. The A-D advantage of allowing a more sensitive test would thus seem to be counterbalanced by the burden of calculating, different, critical values for each distribution to be investigated. In reality, the differences between the A-D test values are not so important: for example, the tabulated A-D test values for LogNormal, Gumbel and Weibull differ only at the third decimal.

Figure 2: BL1 (Corporate Finance). LogNormal and Gumbel fit




In fact, from this picture it is difficult to investigate if the selected distributions provide a good fit in the region we are most interested in, that is the tail area. When the graphical analysis is limited to the tail, i.e. the last 10 per cent of the distribution, it can immediately be seen that both distributions fit the data very poorly: LogNormal underestimates the empirical, actual, tail from the 90th percentile, Gumbel from the 96th percentile (see Figure 3).

Figure 3: BL1 (Corporate Finance). LogNormal and Gumbel fit (focus on the tail)




If the analysis is performed on another BL (BL3: *Retail Banking*), it is easy to see, by graphical analysis, the poor fit of both the distributions in the tail area: LogNormal underestimates the tail from the 90th percentile, Gumbel from the 96th (see Fig. 4).

Figure 4: BL 3 (Retail Banking). LogNormal and Gumbel fit (focus on the tail)



This phenomenon does not change if the other BLs are considered (Figure 5 shows the tail fit for BL6, BL7 and BL8).

Figure 5: BL6, BL7 and BL8. LogNormal and Gumbel fit (focus on the tail)



The goodness-of-fit test values for the LogNormal and Gumbel distributions, reported in Table 3, confirm the graphical perception: in all the BLs, the test values are much higher than the critical ones, at both the 90 per cent and 99 per cent significance levels (the BLs test values are highlighted in yellow if higher than the critical ones. In Table 3 only the critical values for $\alpha = 90$ per cent are reported). The biggest differences can be observed for the A-D values, owing to the greater weight this test gives to the distance between the estimated and the empirical distribution function in the tail area.

Table 3: Conventional inference results

BUSINESS LINE	n. obs.	LogNormal Distribution				Gumbel Distribution				Critical values ($\alpha = 90^\circ$)	
		Parameters estimate		Fitting tests results		Parameters estimate		Fitting tests results			
		μ	σ	Kolmogorov-Smirnov	Anderson-Darling	μ	σ	Kolmogorov-Smirnov	Anderson-Darling	Kolmogorov-Smirnov	Anderson-Darling
BL 1 (Corporate Finance)	423	3.58	1.71	0.18	22.52	93.96	602.30	0.43	124.62	0.06	0.63
BL 2 (Trading & Sales)	5,132	3.64	1.27	0.14	180.52	51.76	185.25	0.37	1,224.03	0.02	0.63
BL 3 (Retail Banking)	28,882	3.17	0.97	0.18	1,653.03	25.63	58.80	0.34	6,037.35	0.01	0.63
BL 4 (Commercial Banking)	3,414	3.61	1.41	0.16	173.94	48.30	203.53	0.37	830.57	0.02	0.63
BL 5 (Payment & Settlement)	1,852	3.37	1.10	0.15	73.74	35.86	109.93	0.36	436.48	0.03	0.63
BL 6 (Agency Services)	1,490	3.74	1.28	0.12	46.33	54.82	181.19	0.35	332.74	0.03	0.63
BL 7 (Asset Management)	1,109	3.79	1.28	0.11	25.68	56.78	153.72	0.32	203.94	0.04	0.63
BL 8 (Retail Brokerage)	3,267	3.58	1.08	0.12	87.67	41.03	93.51	0.31	576.51	0.02	0.63

The main lesson learnt from modelling the severity of operational risk losses by conventional inference methods is that, even though some selected distributions fit the body of the data well, these distributions would underestimate the severity of the data in the tail area: the extent of this error will be investigated in Section 8, where a severity VaR performance analysis is employed.

In practice, the considerable skewness to the right of the empirical distribution causes each curve parameters estimate to be mainly influenced by the observations located in the left and middle area of the empirical distribution, hence reducing the informative power of the data located in the tail area and providing lower than actual figures for the extreme

quantiles. In such a situation, using all the observations to measure the size of the tail could be therefore misleading.

5. Extreme Value Theory: theoretical background

As seen in the conventional inference, the influence of the small/medium-sized losses in the curve parameters estimate does not permit models that fit the tail data accurately to be obtained. An obvious solution to this problem is not to take into consideration the body of the distribution, focusing the analysis only on the large losses. In practice large and small/medium-sized losses are treated separately. If one is interested in obtaining information on some average values of the distribution, conventional or empirical analysis on the data located in the small/medium-sized region may be used (in the current exercise, for example, the outcomes of the conventional inference will be used to derive the expected losses of the BLs, see Section 10 below).

Concerning the tail area, quite a number of different distributions could be adopted; for example, LogNormal and Pareto curves are commonly accepted in insurance to model large claims. However, in this analysis, extreme distributions, stemming from the Extreme Value Theory (EVT), are utilised. The reason lies in the fact that EVT has solid foundations in the mathematical theory of the behaviour of extremes and, moreover, many applications have indicated that EVT appears to be a satisfactory scientific approach in treating rare, large losses. It has been widely applied in structural engineering, oceanography, hydrology, reliability, total quality control, pollution studies, meteorology, material strength, highway traffic and, more recently, in the financial and insurance fields¹⁵. For a comprehensive source on the application of EVT to finance and insurance, see Embrechts *et al.*, 1997, and Reiss and Thomas, 2001.

In general, operational risk losses undoubtedly present characteristics analogous to data originating from the above-mentioned fields (immediate analogies, for example, can be

¹⁵ In recent years, there have been a number of extreme value studies and applications in finance and insurance: for example McNeil studies the estimation of the tails of loss severity distributions (1997), examines the quantile risk measures for financial time series (1998) and provides an extensive overview of the extreme value theory for risk managers (1999); Embrechts studies the potentials and limitations of the extreme value theory (1999 and 2000); McNeil and Frey study the estimation of tail-related risk measures for heteroschedastic financial time series (2000).

found in insurance, reinsurance, reliability and total quality control). In fact, operational risk data appear to be characterised by two “souls”: the first one, driven by high-frequency low-impact events, constitutes the body of the distribution and refers to expected losses; the second one, driven by low-frequency high-impact events, constitutes the tail of the distribution and refers to unexpected losses. In practice, the body and the tail of data do not necessarily belong to the same, underlying, distribution or even to distributions belonging to the same family. More often their behaviour is so different that it is hard to identify a unique traditional model that can at the same time describe, in an accurate way, the two “souls” of data: the conventional inference on the BLs whole data sets in Section 4 furnishes a clear proof of that ¹⁶.

Consequently, in all the cases in which the tail tends “to speak for itself”, EVT appears to be an useful inferential instrument with which to investigate the large losses, owing to its double property of focusing the analysis only on the tail area (hence reducing the disturbance effect of the small/medium-sized data) and treating the large losses by an approach as scientific as the one driven by the Central Limit Theorem for the analysis of the high-frequency low-impact losses ¹⁷. Clearly, EVT is not a “panacea”, since specific conditions are required for its application and even then it is still open to some criticisms, extensively investigated in the literature (on this topic, see for example Embrechts *et al.*, 1997, Diebold *et al.*, 1998, and Embrechts *et al.*, 2003).

Unlike traditional methods, EVT does not require particular assumptions on the nature of the original underlying distribution of all the observations, which is generally unknown. EVT is applied to real data in two related ways.

The first approach (see Reiss and Thomas, 2001, p. 14 ff) deals with the maximum (or minimum) values the variable takes in successive periods, for example months or years.

¹⁶ Some mixture distributions could be investigated in order to identify a model that provides a reasonable fit to both the body and the tail of data. However, the disadvantage of such distributions is that they are more complex and, hence, less easy to handle. Furthermore a mixture model would be an arbitrary choice, not supported by a robust theory and, because of that, one would have less confidence in extrapolating the outcomes beyond the empirical data.

¹⁷ To cite, respectively, Diebold *et al.*, 1998, and Smith, 1987, “EVT helps the analyst to draw smooth curves through the extreme tails of empirical survival functions in a way that is guided by powerful theory and hence provides a rigorous complement to alternatives such as graphical analysis or empirical survival functions” and “There is always going to be an element of doubt, as one is extrapolating into areas one doesn’t know about. But what EVT is doing is making the best use of whatever data you have about extreme phenomenon”.

These observations constitute the extreme events, also called block (or per-period) maxima. At the heart of this approach is the “three-types theorem” (Fisher and Tippett, 1928), which states that there are only three types of distributions which can arise as limiting distributions of extreme values in random samples: the Weibull type, the Gumbel type and the Frechet type. This result is very important, since the asymptotic distribution of the maxima always belongs to one of these three distributions, regardless of the original one. Therefore the majority of the distributions used in finance and actuarial sciences can be divided into these three classes, according to their tail-heaviness:

- light-tail distributions with finite moments and tails, converging to the Weibull curve (Beta, Weibull);
- medium-tail distributions for which all moments are finite and whose cumulative distribution functions decline exponentially in the tails, like the Gumbel curve (Normal, Gamma, LogNormal);
- heavy-tail distributions, whose cumulative distribution functions decline with a power in the tails, like the Frechet curve (T-Student, Pareto, LogGamma, Cauchy).

The Weibull, Gumbel and Frechet distributions can be represented in a single three parameter model, known as the Generalised Extreme Value distribution (GEV):

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

where: $1 + \xi \frac{x - \mu}{\sigma} > 0$

The parameters μ and σ correspond to location and scale; the third parameter, ξ , called the shape index, indicates the thickness of the tail of the distribution. The larger the shape index, the thicker the tail.

The second approach to EVT (see Reiss and Thomas, 2001, p. 23 ff) is the Peaks Over Threshold (POT) method, tailored for the analysis of data bigger than preset high thresholds.

The severity component of the POT method is based on a distribution (Generalised Pareto Distribution - GPD), whose cumulative function is usually expressed as the following two parameter distribution:

$$\text{GPD}_{\xi, \sigma}(x) = \begin{cases} 1 - \left(1 + \xi \frac{x}{\sigma}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - \exp\left\{-\frac{x}{\sigma}\right\} & \text{if } \xi = 0 \end{cases} \quad (2)$$

where: $x \geq 0$ if $\xi \geq 0$, $0 \leq x \leq -\sigma/\xi$ if $\xi < 0$

and ξ and σ represent respectively the shape and the scale parameter

It is possible to extend the family of the GPD distributions by adding a location parameter μ . In this case the GPD is defined as:

$$\text{GPD}_{\xi, \mu, \sigma}(x) = \begin{cases} 1 - \left(1 + \xi \frac{x - \mu}{\sigma}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - \exp\left\{-\frac{x - \mu}{\sigma}\right\} & \text{if } \xi = 0 \end{cases} \quad (3)$$

The interpretation of ξ in the GPD is the same as in the GEV, since all the relevant information on the tail of the original (unknown) overall distribution is embedded in this parameter¹⁸: when $\xi < 0$ the GPD is known as the Pareto ‘‘Type II’’ distribution, when $\xi = 0$ the GPD corresponds to the Exponential distribution. The case when $\xi > 0$ is probably the most important for operational risk data, because the GPD takes the form of the ordinary

Pareto distribution with tail index $\alpha = 1/\xi$ and indicates the presence of heavy-tail data ¹⁹; in this particular case there is a direct relationship between ξ and the finiteness of the moments of the distribution:

$$E(x^k) = \infty \quad \text{if } k \geq 1/\xi \quad (4)$$

For instance, if $\xi \geq 0.5$ the GPD has an infinite variance, if $\xi \geq 1$ there is no finite moment, not even the mean. This property has a direct consequence for data analysis: in fact the (heavier or lighter) behaviour of data in the tail can be easily directly detected from the estimate of the shape parameter.

Now, let $F_x(x)$ be the (unknown) distribution function of a random variable X (with right-end point x_F) which describes the behaviour of the operational risk data in a certain BL and let $F_u(y)$ be its excess distribution at the threshold u . The excess distribution can be introduced as a conditional distribution function, that is:

$$F_u(y) = P(X - u \leq y | X > u) = \frac{F_x(x) - F_x(u)}{1 - F_x(u)} \quad \text{for } y = x - u > 0 \quad (5)$$

It represents the probability that a loss exceeds the threshold u by at most an amount y , given that it exceeds the threshold.

The theory (Balkema-De Haan, 1974, and Pickands, 1975) maintains that for a large class of underlying distributions, the excess distribution $F_u(y)$ converges asymptotically to a GPD as the threshold is progressively raised to the right endpoint x_F of the distribution ²⁰:

$$\lim_{u \rightarrow x_F} \sup | F_u(y) - GPD_{\xi, \beta}(y) | = 0 \quad (6)$$

¹⁸ The maxima of samples of events from GPD are GEV distributed with shape parameter equal to the shape parameter of the parent GPD. There is a simple relationship between the standard GDP and GEV such that $GPD(x) = 1 + \log GEV(x)$ if $\log GEV(x) > -1$

¹⁹ The ordinary Pareto is the distribution with distribution function $F(x) = 1 - (a/x)^\alpha$ and support $x > a$. This distribution can be rewritten as $F(x) = 1 - (1 + (x - a)/a)^\alpha$ so that it can be seen to be a GPD with shape $\xi = 1/\alpha$, scale $\sigma = a/\alpha$ and location $\mu = a$. In practice it is a GPD where the scale parameter is constrained to be the shape multiplied by the location, hence it is a little less flexible than a GPD, where the scale can be freely chosen.

²⁰ The conditions under which excess losses converge to GPD distributions are very large. For an extensive treatment, see Embrechts *et al.*, 1997.

$$\begin{aligned}
& 1 - \left(1 + \xi \frac{y}{\beta}\right)^{-1/\xi} && \text{if } \xi \neq 0 \\
\text{where } \text{GPD}_{\xi, \beta}(y) = & && \\
& 1 - \exp\left\{-\frac{y}{\beta}\right\} && \text{if } \xi = 0
\end{aligned} \tag{7}$$

with: $y = x - u = \text{excess}$, $\xi = \text{shape}$, $\beta = \text{scale}$;

and support $y \in [0, x_F - u]$ if $\xi \geq 0$

$y \in [0, -\beta/\xi]$ if $\xi < 0$

In this work, the $\text{GPD}_{\xi, \beta}(y)$ will be called the ‘‘excess GPD’’, to stress the fact that the argument y represents the excesses, that is to say the exceedances x (i.e. the data larger than the threshold u) minus the threshold u itself.

Equivalently, the limit condition (6) holds if the exceedances x are used in place of the excesses y : changing the argument, the $F_u(y)$ and $\text{GPD}_{\xi, \beta}(y)$ transform respectively to $F_u(x)$ and $\text{GPD}_{\xi, u, \beta}(x)$, with the threshold u , now, representing the location parameter and $x > u$. Therefore, when the threshold tends to the right endpoint x_F , the exceedance distribution $F_u(x)$ converges asymptotically to a GPD with the same shape ξ , scale β and location $\mu = u$. The $\text{GPD}_{\xi, u, \beta}(x)$ will be called the ‘‘exceedance GPD’’ because it deals with the exceedances x at u .

One of the most important properties of the GPD is its stability under an increase of the threshold.

To show that, let isolate $F_x(x)$ from (5):

$$F_x(x) = [1 - F_x(u)]F_u(y) + F_x(u)$$

Looking at the limit condition (6), both the excess distribution $F_u(y)$ and the exceedance distribution $F_u(x)$ can be approximated well by suitable GPDs. By using the ‘‘exceedance GPD’’, one obtains:

$$F_x(x) \approx [1 - F_x(u)]GPD_{\xi,u,\beta}(x) + F_x(u) \quad (8)$$

Substituting the $GPD_{\xi,u,\beta}$ expression in (8):

$$F_x(x) \approx [1 - F_x(u)] \left[1 - \left(1 + \xi \frac{x-u}{\beta} \right)^{-1/\xi} \right] + F_x(u)$$

The only element now required to identify $F_x(x)$ completely is $F_x(u)$, that is to say the value of the (unknown) distribution function in correspondence with the threshold u . To this end, the empirical estimator of $F_x(x)$ introduced in section 4, computed at u , can be a viable solution:

$$F_n(u) = \frac{1}{n} \sum_{i=1}^n 1_{\{x_i \leq u\}} = \frac{n - n_u}{n} \quad (9)$$

where: n is the total number of observations

n_u the number of observations above the threshold u

The threshold u should be set at a level that let enough observations exceeding u to obtain a reliable empirical estimate of $F_x(u)$.

Consequently, $F_x(x)$ can be completely expressed by the parameters of the $GPD_{\xi,u,\sigma}(x)$ and the number of observations (total and over the threshold):

$$F_x(x) \approx \frac{n_u}{n} \left[1 - \left(1 + \xi \frac{x-u}{\beta} \right)^{-1/\xi} \right] + \left(1 - \frac{n_u}{n} \right)$$

which simplifies to

$$F_x(x) \approx 1 - \frac{n_u}{n} \left[1 + \left(1 + \xi \frac{x-u}{\beta} \right)^{-1/\xi} \right] \quad (10)$$

This quantity is defined as the “tail estimator” of $F_x(x)$, as it is valid only for $x > u$. It is possible to demonstrate that the “tail estimator” is also GPD distributed: it is the semiparametric representation of the $GPD_{\xi,\mu,\sigma}$ referred to all the original data, with the same

shape ξ and location and scale equal to μ and σ respectively. The $\text{GPD}_{\xi,\mu,\sigma}$ will be called the “full GPD” because it is fitted to all the data in the tail area²¹.

Semiparametric estimates for the “full GPD” parameters can be derived from those of the “exceedance GPD”:

$$\begin{aligned}\sigma &= \beta \left(\frac{n_u}{n} \right)^\xi \\ \mu &= u - \frac{\beta}{\xi} \left[1 - \left(\frac{n_u}{n} \right)^\xi \right]\end{aligned}\tag{11}$$

As there is a one-to-one relationship between the “full GPD” ($\text{GPD}_{\xi,\mu,\sigma}$) and the “exceedance GPD” ($\text{GPD}_{\xi,u,\beta}$), it is also possible to express the scale parameter of the latter by the former: $\beta = \sigma + \xi(u - \mu)$. It should be noted that, while the scale (β) of the “exceedance GPD” depends on where the threshold is located, the shape (ξ), the location (μ) and scale (σ) of the “full GPD” are independent of the threshold. Hence a nice practical method to check the robustness of the model for some specific data is to evaluate the degree of stability of these latter parameters over a variety of thresholds²².

By applying the GPD stability property, it is possible to move easily from the excess data ($y = x - u$) to the tail of the original data ($x > u$) and from the excess distribution $F_u(y)$ to the underlying (unknown) distribution $F_x(x)$.

An immediate consequence of the GPD stability is that if the exceedances of a threshold u follow a $\text{GPD}_{\xi,u,\beta}$, the exceedances over a higher threshold $v > u$ are $\text{GPD}_{\xi,v,\beta + \xi(v-u)}$, that is they are also GPD distributed with the same shape ξ , the location equal to v (the new threshold) and the scale equal to $\beta + \xi(v - u)$. This property will be extensively adopted in the current exercise.

²¹ Examining (10) it can be immediately seen that, in contrast with the “excess GPD” (and the “exceedance GPD”), the “full GPD” in its semiparametric representation provides information on the frequency with which the threshold u is pierced by the exceedances: the empirical quantity $F_n(u)$, built on the number of observations (total and above the threshold) takes care of this aspect. Nevertheless, if one wants to move on from the semiparametric “full GPD” to its completely parametric form, all the information on the original data must be available (that is, it is necessary to know the amounts of the data under the threshold, in addition to their number).

6. Peaks Over Threshold approach: Business Lines tail-severity estimate

In light of the EVT features described in the previous Section, the POT method is implemented in each BL data set, by fitting the “excess GPD”, $GPD_{\xi, \beta}(y)$, to the excess losses of a selected threshold.

As seen before, the GPD fitting work depends on three elements:

- a) the threshold (u), to be set by the analyst;
- b) the excess data, i.e. the original data minus the selected threshold and
- c) two parameters (ξ and β) to be estimated from the excess data.

A key modeling aspect with the GPD is the selection of the threshold, that is the point where the tail starts. The choice of u should be large enough to satisfy the limit law condition (theoretical condition: u should tend to the right-end point x_F), while at the same time leaving sufficient observations for the estimation (practical condition). Furthermore, any inference conclusion on the shape parameter – which, as noted, governs the heaviness of the tail – should be insensitive to increases in the threshold above this suitable level.

A number of diagnostic instruments have been proposed in the literature for threshold selection, including a bootstrap method that produces an optimal value under certain criteria (see Danielsson *et al.*, 2000). Owing to its handiness and simplicity, one of the most used techniques is the mean excess plot (see Davison and Smith, 1990), a graphical tool based on the Sample Mean Excess Function (SMEF), defined as:

$$SMEF(u) = \frac{\sum_{i \leq n} x_i - u}{\sum_{i \leq n} 1_{\{x_i > u\}}} \quad (12)$$

²² In practice, for each threshold u , the estimate of the “exceedance GPD” parameters (ξ and β) must be obtained and hence the corresponding values of the “full GPD” parameters (μ and σ) gained. Then the approximate equality of ξ , μ and σ for increasing thresholds must be investigated.

i.e. the sum of the excesses over the threshold u divided by the number of data points that exceed the threshold itself. The SMEF is an estimate of the Mean Excess Function (MEF), defined as:

$$MEF(u) = E(X - u | X > u) \quad (13)$$

which describes the expected overshoot of a threshold once an exceedance occurs.

It can be demonstrated (see Embrechts *et al.*, 1997) that if the plot shows a downward trend (negative slope), this is a sign of short-tailed data. Exponentially distributed data would give an approximately horizontal line while data from a heavy-tail distribution would show an upward trend (positive slope). In particular, if the plot is a positively sloped straight line above a certain threshold u , it is an indication that the data follow a GPD with a positive shape parameter ξ in the tail area above u . This is clear, since for the GPD the MEF is linear:

$$GPD_{MEF}(u) = \frac{\beta + \xi u}{1 - \xi} \quad (14)$$

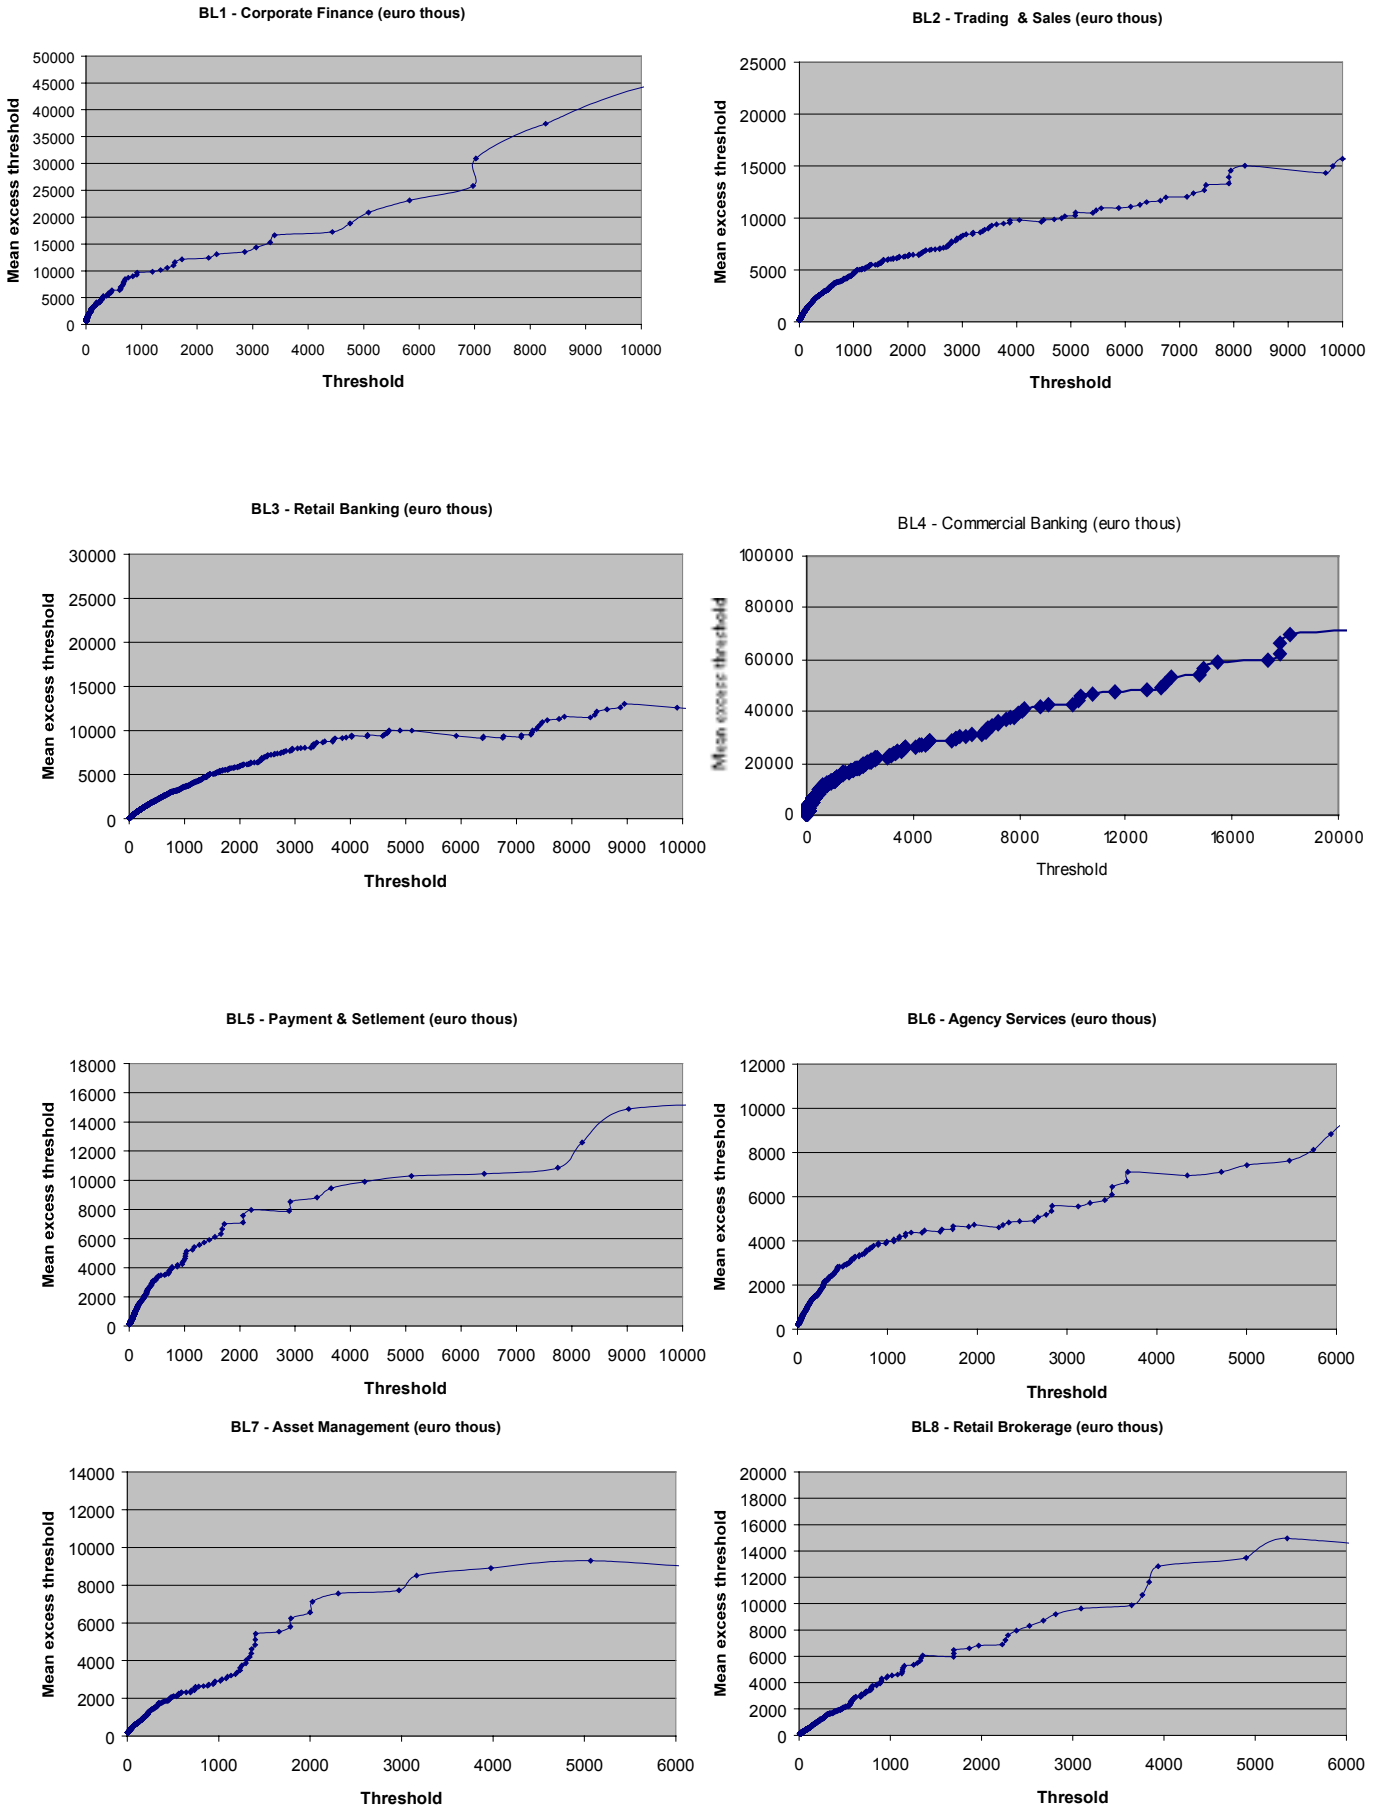
where $(\beta + \xi u) > 0$

In applying the mean excess plot to the eight BLs, the goal is to detect a straightening out or, at least, a change in the slope of the plot above a certain threshold, in order to be able to fix that threshold as the start of the tail and to fit the GPD to the excess data. For each data set, the SMEF against increasing thresholds from the initial value is plotted; since the plot can be hard to interpret for very large thresholds (because there are few exceedances and, hence, high variability in the sample mean), it was decided to end the plot at the fourth order statistic (see Figure 6)²³.

²³ One issue of the mean excess plot is that the MEF does not exist for GPD with $\xi > 1$. In that case, a trimmed version of the MEF could be used, since the last function always exists, regardless of the shape values (on this topic see Reiss and Thomas, 2001, p. 56).

In the current exercise, a comparison between the SMEF and its trimmed version (applying a data truncation percentage of 5 per cent from below and from above) was made. As the differences in the resulting plots were negligible, the whole, untrimmed, version of the SMEF was adopted for the selection of the threshold.

Figure 6: mean excess plot for the BLs



In all the data sets, clear evidence of straightening out of the plot is found from the starting level. However, owing to the need to satisfy the theoretical condition (u should be large enough), the threshold is set close to the 90th empirical percentile for all the BLs except for *Retail Banking*, where it is shifted to the 96.5th percentile because of the greater number of observations.

In Table 4 the threshold selection results are reported for all the BLs.

Table 4: BLs threshold selection

BUSINESS LINE	n. obs.	threshold	related empirical percentile	n. excesses
BL 1 (Corporate Finance)	423	400.28	89.85%	42
BL 2 (Trading & Sales)	5,132	193.00	89.85%	512
BL 3 (Retail Banking)	28,882	247.00	96.50%	1,000
BL 4 (Commercial Banking)	3,414	270.00	90.66%	315
BL 5 (Payment & Settlement)	1,852	110.00	89.85%	187
BL 6 (Agency Services)	1,490	201.66	89.20%	158
BL 7 (Asset Management)	1,109	235.00	90.00%	107
BL 8 (Retail Brokerage)	3,267	149.51	89.99%	326

The set thresholds leave a large enough number of exceedances to apply statistical inference in all the BLs²⁴, except for BL1.

²⁴ This affirmation is supported by the results of a simulation study conducted by McNeil and Saladin, 1997, aimed to detect the minimal number of data and exceedances to work with in order to obtain reliable estimates of high quantiles of given distributions. In particular, the exercise showed that, when the data presented a Pareto heavy tail with shape parameter $\alpha=1/\xi=1$, a minimum number of 1,000 (2,000) data and 100 (200) exceedances was required to have a reliable GPD estimate of the 99th (99.9th) percentile. As it can be seen in Table 4, apart from BL1, the number of excesses appears to be large enough to obtain reliable estimates of the 99.9th percentile for the majority of the BLs.

For each BL, the maximum likelihood method is adopted to estimate the shape and scale parameters of the GPD. For $\xi > -0.5$, Hosking and Wallis, 1987, present evidence that maximum likelihood regularity conditions are fulfilled (consistency and asymptotic efficiency) and the maximum likelihood estimates are asymptotically normally distributed²⁵. As regards the BL1, a Bayesian correction of the (prior) maximum likelihood estimates is conducted to obtain more consistent values for the parameters of the GPD²⁶.

To reinforce the judgements on the shapes estimate, a shape plot, based on the comparison of the ξ estimates across a variety of thresholds, is used (see Embrechts *et al.*, 1997, p. 339). In practice, different GPD models are fitted to the excesses, with thresholds increasing from the starting limit (see Table 4) to the values located approximately at the 99.9th empirical percentile. The degree of stability of ξ is then evaluated in the, right to medium-left, range of the estimates interval: if the ξ values do not vary somewhat, the inference is not too sensitive to the choice of the threshold in that range and a final estimate of the shape parameter can be obtained as an average value of the ξ estimates²⁷.



Figure 7 shows, for some BLs (BL3, BL5, BL7 and BL8) the number of exceedances identified for the corresponding thresholds (horizontal axis) and the estimates of the shape parameter (vertical axis).

²⁵ The main competitors of the maximum likelihood estimator in the GPD model are the methods of moments (simple matching and probability weighted) and, limited to ξ , the Hill estimator. However, the moments methods assume the existence of the second moment, hence they perform poorly when the second moment does not exist (that is when $\xi > 0.5$), while the Hill estimator may be inaccurate if the shape parameter ξ estimate is large.

²⁶ If there are few data (less than 40, say) the estimation errors become very large and resort has to be made to credibility or Bayesian methods. For the application of Bayesian methods, see for example Smith and Goodman, 1996, Medova, 2000, and Reiss and Thomas, 2001.

²⁷ To evaluate the finite sample properties of the maximum likelihood estimator (i.e. sensitivity to threshold and sample size) McNeil and Frey, 2000, and Nystrom and Skoglund, 2002, conducted MonteCarlo experiments for various distributions and sample sizes. The results were encouraging in all the cases, because

Figure 7: Shape plot for some BLs



the estimator hardly varied with the choice of threshold within reasonable limits of the number of excesses (5-13 per cent of the whole data set).

In each BL plot, a satisfactory stability of ξ can be found in the range from the starting threshold (right-side of the plot) to high thresholds (medium-left side of the plot)²⁸. In light of that, each BL shape finale figure is set as the median of the ξ values in that range, pending confirmation of the estimate in a while by suitable tests. The final value for the scale parameter β is the maximum likelihood estimate calculated at the starting threshold.

Table 5 reports, for each BL, the GPD shape and scale estimates, together with the values of the K-S and A-D tests. For the shape parameters, confidence intervals at the significance level of $\alpha = 95$ per cent are also computed using a bootstrap procedure.

Table 5: GPD fitting results

BUSINESS LINE	n. excesses	Parameters estimate		ξ confidence interval ($\alpha=95^\circ$)		Fitting tests results				
		β	ξ	lower limit	upper limit	Kolmogorov-Smirnov		Anderson-Darling		
						test results	critical values ($\alpha = 90^\circ$)	test results	critical values ($\alpha = 90^\circ$)	critical values ($\alpha = 99^\circ$)
BL 1 (Corporate Finance)	42	774	1.19	1.06	1.58	0.099	0.189	0.486	0.630	1.030
BL 2 (Trading & Sales)	512	254	1.17	0.98	1.35	0.027	0.054	0.508	0.630	1.030
BL 3 (Retail Banking)	1,000	233	1.01	0.88	1.14	0.020	0.023	0.675	0.630	1.030
BL 4 (Commercial Banking)	315	412	1.39	1.20	1.62	0.058	0.070	1.541	0.630	1.030
BL 5 (Payment & Settlement)	187	107	1.23	0.96	1.37	0.028	0.090	0.247	0.630	1.030
BL 6 (Agency Services)	158	243	1.22	1.03	1.42	0.064	0.097	0.892	0.630	1.030
BL 7 (Asset Management)	107	314	0.85	0.57	1.18	0.060	0.118	0.217	0.630	1.030
BL 8 (Retail Brokerage)	326	124	0.98	0.76	1.20	0.033	0.068	0.291	0.630	1.030

It can be observed that:

- the K-S test values are lower than the critical ones in all the BLs, while the A-D test values show a slight rejection of the null hypothesis for BL3 and BL6 (at the significance level of $\alpha=90$ per cent) and BL4 (even at the higher significance level of $\alpha=99$ per cent). Even if, in general, the hypothesis BLs tail originates from a GPD seems reasonable, the


²⁸ In the above plots, the stability is evident, for BL3, from a number of exceedances equal to 1,000 (in correspondence with the starting threshold) to about 100; for BL5, from 187 to about 70; for BL7, from 107 to about 40; for BL8, from 326 to about 100.







accuracy of the model needs to be confirmed by the outcomes of testing procedures, which are more appropriate for large losses. This exercise will be conducted in the next section;

- the shape parameters estimate (ξ) confirms the previous bootstrapping results (see Section 3) as regards the high kurtosis of the loss data for the BLs. Owing to the ξ estimates, the GPD models have infinite variance in all the BLs (the ξ values are always greater than 0.5) and, furthermore, have infinite mean in six BLs ($\xi > 1$ in BL1, BL2, BL3, BL4, BL5, BL6) and almost infinite mean in BL8 ($\xi = 0.98$);
- although there is some uncertainty in each shape estimate, reflected in the extent of its confidence interval, the range of the ξ values provides evidence of the different riskiness of the BLs. This impression will be confirmed in Section 8 by directly exploring and measuring the tail of the BLs.

Figure 8 below shows, for each BL, the GPD curve together with the empirical distribution function; the graphical analysis is limited to the last 50 per cent of the data or to the tails. Unlike the conventional inference, closeness of the GPD and the empirical distribution is now found even at the highest percentiles.

Figure 8: All the BLs. Generalised Pareto Distribution fit





7. Advanced tests on the severity results

As seen in the previous Section, the K-S and A-D tests provide some evidence of the closeness of the GPD to the data. However, these tests are usually employed in actuarial applications to measure the goodness of distributions fitted to a data set as a whole, not being tailored for distributions of excesses over some thresholds. A specific, more tailored, test for the GPD assumption is the W-statistics proposed by Davison, 1984. The test is based on the residuals, defined as:

$$W_i = \frac{1}{\xi} \log \left[1 - \xi \frac{x_i - u}{\sigma + \xi(u - \mu)} \right] \quad (15)$$

where u is the threshold and σ and μ are the parameters of the “full GPD”.

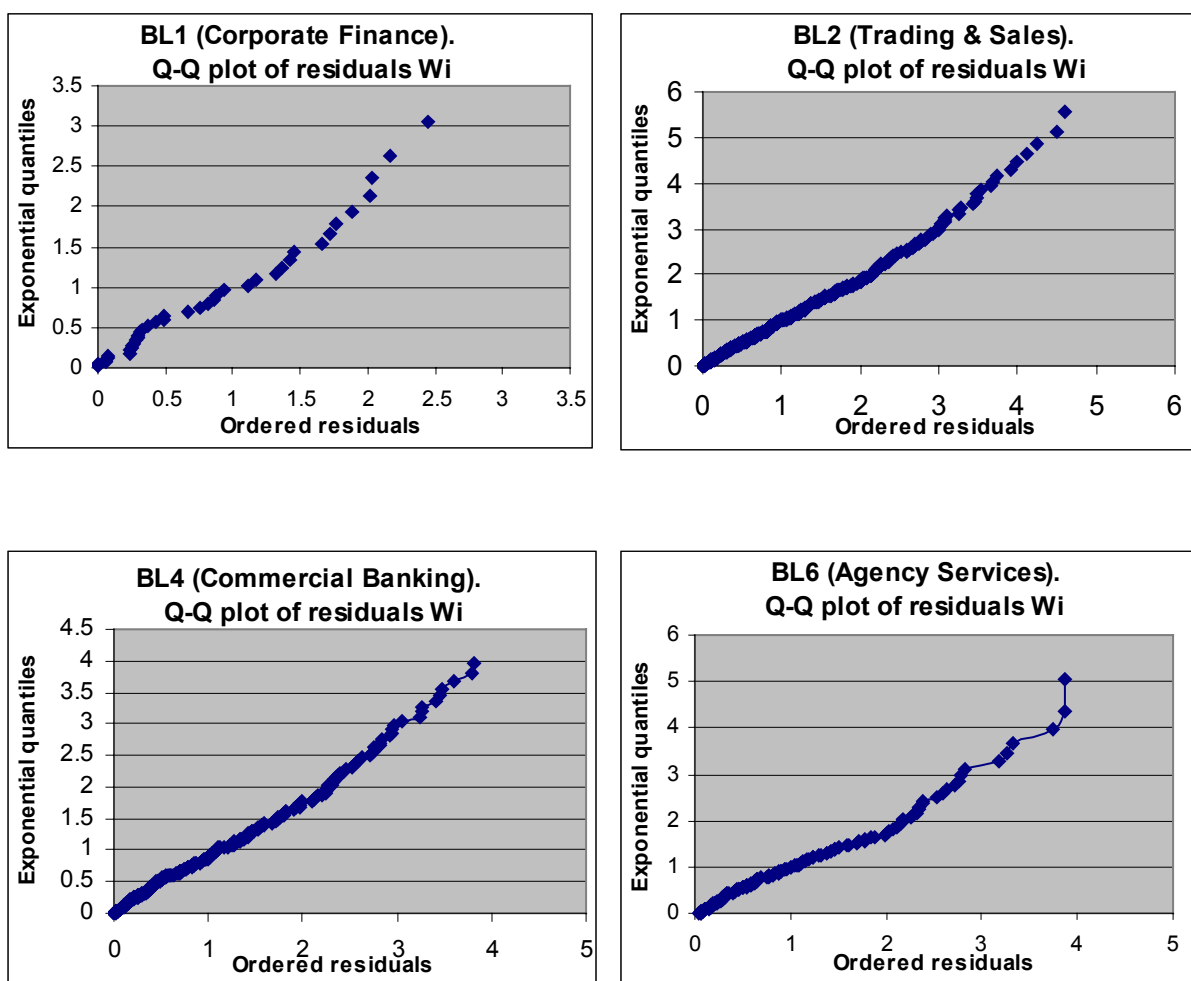
If the excesses $(x_i - u)$ are i.i.d. from a $\text{GPD}_{\xi, \mu, \sigma}$, the residuals W_i should be i.i.d. Exponentially distributed with mean $\gamma=1$. As a result, two types of plots can be used to show whether these assumptions are in fact supported by the data:

- a scatter plot of residuals against their (time) order of occurrences. Systematic variation of the residuals with time would indicate the presence of a trend in the model. Consequently, the assumption of data stationarity may not be correct, given that the losses become smaller or larger on average as time evolves;
- a Q-Q plot of residuals against the expected order statistics under the Exponential distribution. If the plot stays close to a straight line, the Exponential assumption for the residuals and hence the GPD assumption for the excesses, may be tenable.

In light of the pooling exercise performed in this exercise, only the latter assumption (the GPD behaviour of the excesses) may be tested. Nevertheless, as stated in Section 3, the short temporal window on which the losses were collected provides a significant guarantee for the absence of a trend in the data gathered in each individual bank’s database.

After obtaining the residuals W_i ²⁹, the Q-Q plots are plotted (in Figure 9, the plots for the BLs not, graphically, addressed in the previous shape plots analysis are shown). The Q-Q plots appear to be quite close to a straight line of unit slope, hence, indicating an acceptable fit of the GPD to the excesses.

Figure 9: Q-Q plot of residuals W_i for some BLs



²⁹ In order to obtain the residuals W_i , the “full GPD” parameters μ and σ are needed. To this end, semiparametric estimates are derived by substituting in (11) the estimates of the “exceedance GPD” parameters (ξ and β), in addition to the threshold value and the number of observations (total and over the threshold).

In order to check in a more accurately way the GPD assumption for all the BLs, the Exponential hypothesis for residuals W_i is analytically tested within the Generalised Pareto model ($H_0 : \xi = 0$ ³⁰ against $H_1: \xi \neq 0$) by a suitable Likelihood Ratio test.

The statistic is T_{LR} (see Reiss and Thomas, 2001, p. 154), defined as:

$$T_{LR} = 2 \log \left[\prod_{i \leq n} GPD_{\xi, \mu, \sigma}(W_i) \prod_{i \leq n} GPD_{0, \mu, \sigma}(W_i) \right] \quad (16)$$

Since the parameter sets have dimension 3 and 2, the T_{LR} -test is asymptotically distributed as a χ^2 with 1 degree of freedom under the null hypothesis. Consequently, the p-value is:

$$p_{LR} = 1 - \chi_1^2(T_{LR})$$

For each BL, the T_{LR} statistic is applied to the W_i residuals and the corresponding p-value derived (see Table 6).

Table 6: Test of residuals W_i for the GPD model: mean estimate (γ) and p-value

BUSINESS LINE	γ estimate	p-value of T_{LR}
BL 1 (Corporate Finance)	1.001	0.81
BL 2 (Trading & Sales)	1.001	0.18
BL 3 (Retail Banking)	0.999	0.52
BL 4 (Commercial Banking)	1.070	0.29
BL 5 (Payment & Settlement)	1.002	0.74
BL 6 (Agency Services)	0.960	0.77
BL 7 (Asset Management)	0.900	0.52
BL 8 (Retail Brokerage)	1.003	0.92

The results confirm the hypothesis that the GPD is an appropriate model to represent, even at high confidence levels, the tail of all the BLs.

³⁰ For ξ close to 0, the GPD tends to an Exponential distribution (see Section 5).

In the last part of this Section, a severity Value at Risk (VaR_{Sev}) performance analysis is carried out for each BL in order to compare the different levels of accuracy of the GPD and the conventional Gumbel and LogNormal distributions in representing the highest percentiles of data.

The relative VaR_{Sev} performance of each model (see the next section for the GPD_{VaR} calculation) is backtested by comparing the estimated and the expected number of violations: a violation occurs when the actual loss exceeds the VaR_{Sev} value. A number of violations higher than the expected one indicates that the model consistently underestimates the risk at the tail ³¹.

In practice, the expected number of violations in each BL is obtained by comparing the total number of observations with the desired percentile. For instance, if a BL contains 1,000 data overall, the expected number of violations at the 99th percentile is equal to $0.01 * 1,000 = 10$. Therefore, if the parametric model were correct, one would expect only 10 observations to be greater than the 99th percentile singled out by the model. If the violations are more than 10, the 99th parametric percentile lies at a lower level and hence underestimates the actual tail of data.

In Table 7, the theoretical number of violations, calculated at the 95th, 97.5th, 99th, 99.5th, 99.9th percentiles are compared with the estimated one, drawn from, respectively, the GPD, LogNormal and Gumbel distributions.

³¹ This test is equivalent to that usually adopted in market risk to evaluate the VaR sensitivity of the model.

Table 7: VaR_{Sev} performance analysis

	Percentile	N. VIOLATIONS			
		Theoretical	GPD	LogNormal	Gumbel
BL 1 (Corporate Finance) n. obs = 423	0.950	21.15	21	36	16
	0.975	10.58	12	22	15
	0.990	4.23	3	16	13
	0.995	2.12	2	13	12
	0.999	0.42	0	5	10
BL 2 (Trading & Sales) n. obs = 5,132	0.950	256.60	259	351	211
	0.975	128.30	129	261	184
	0.990	51.32	56	185	160
	0.995	25.66	26	144	137
	0.999	5.13	2	89	113
BL 3 (Retail Banking) n. obs = 28,882	0.950	1,444.10	1,386	2,062	1,234
	0.975	722.05	722	1,551	1,023
	0.990	288.82	294	1,094	812
	0.995	144.41	139	837	719
	0.999	28.88	31	514	560
BL 4 (Commercial Banking) n. obs = 3,414	0.950	170.70	173	241	168
	0.975	85.35	102	175	155
	0.990	34.14	48	137	137
	0.995	17.07	20	106	130
	0.999	3.41	5	71	106
BL 5 (Payment & Settlement) n. obs = 1,852	0.950	92.60	95	115	60
	0.975	46.30	44	89	49
	0.990	18.52	20	58	43
	0.995	9.26	10	46	42
	0.999	1.85	2	33	34
BL 6 (Agency Services) n. obs = 1,490	0.950	74.50	73	98	66
	0.975	37.25	41	73	58
	0.990	14.90	16	51	50
	0.995	7.45	7	41	46
	0.999	1.49	0	30	39
BL 7 (Asset Management) n. obs = 1,109	0.950	55.45	55	72	55
	0.975	27.73	32	53	48
	0.990	11.09	9	38	40
	0.995	5.55	6	27	38
	0.999	1.11	1	10	30
BL 8 (Retail Brokerage) n. obs = 3,267	0.950	163.35	166	220	134
	0.975	81.68	88	149	117
	0.990	32.67	30	105	99
	0.995	16.34	16	73	87
	0.999	3.27	6	37	58

It can immediately be seen that, while the number of violations for the GPD model are very close to the theoretical ones – this occurs both at the highest (99th, 99.5th and 99.9th) and lowest (95th and 97.5th) percentiles – the LogNormal and Gumbel numbers of violations are

always larger than the expected ones. This outcome once again confirms the excess of optimism of the conventional inference in representing the operational riskiness of the BLs.

8. Business Lines tail-severity measures and magnitude

The results from the graphical inspection, the goodness-of-fit tests and the VaR_{Sev} performance analysis clearly indicate that the GPD appears to be a consistent and accurate model to with which represent the extreme quantiles of each BL's operational risk data set. In light of that, in this section the GPD will be used to get information on the size of the tail of the eight BLs, that is on their operational severity riskiness.

The first risk measure that is computed is the VaR introduced in the previous Section for backtesting purposes (GPD_{VaR}). In standard statistical language, the VaR at confidence level p is the smallest loss that is greater than the p^{th} percentile of the underlying loss distribution:

$$VaR_p(X) = \inf\{x : F_x(X) \geq p\} \quad (17)$$

As the GPD only deals with the severity component of the losses and does not address the matter of the frequency of their occurrence, the GPD_{VaR} identifies merely the time-unconditional loss percentile. In this exercise, the conditional losses referred to a 1-year time horizon (and hence the conditional 1-year loss percentiles) will be computed after completely implementing the POT approach in order to take into account its frequency component (see Sections 9 and 10).

In the GPD model, it is possible to obtain a formula for the VaR by the semiparametric representation of the “full GPD”.

In fact, for a given confidence level $p > F_x(u)$, the VaR expression can be obtained by inverting the tail estimator (10) to get:

$$VaR_p(x) = u + \frac{\hat{\beta}}{\hat{\xi}} \left\{ \left[\frac{n}{n_u} (1-p) \right]^{-\hat{\xi}} - 1 \right\} \quad (18)$$

For each BL, Table 8 reports the GPD_{VaR} at different confidence levels (95th, 99th and 99.9th), computed on the basis of the estimates of ξ and β gained in the previous Section.

Table 8: BLs GPD_{VaR} (Euro ,000)

BUSINESS LINE	$p = 95^\circ$	$p = 99^\circ$	$p = 99.9^\circ$
BL 1 (Corporate Finance)	1,222	9,743	154,523
BL 2 (Trading & Sales)	463	3,178	47,341
BL 3 (Retail Banking)	176	826	8,356
BL 4 (Commercial Banking)	668	6,479	159,671
BL 5 (Payment & Settlement)	230	1,518	25,412
BL 6 (Agency Services)	501	3,553	58,930
BL 7 (Asset Management)	511	2,402	17,825
BL 8 (Retail Brokerage)	272	1,229	11,539

By looking at Table 8, one can get the desired information on the size of the tail of the BLs, on the basis of the figure revealed by the GPD_{VaR} at the various percentiles.

However it should be observed that one of the most serious problems in using the VaR in practical applications is that, when losses do not have a like-Normal behaviour, the VaR is unstable and difficult to work with and, further, it fails to be a coherent measure of risk (in the sense introduced by Artzner *et al.* in a well-known article in 1999)³². Moreover, the VaR provides no handle on the extent of the losses that might be suffered beyond the amount indicated by the VaR measure itself. It merely gives a lowest bound for the losses in the tail and, in doing so, has a bias toward optimism instead of conservatism in the measure of the riskiness of the businesses.

³² The concept of “coherent” measures of risk was introduced in a famous article by Artzner *et al.*, 1999, in which the authors identified the specific properties a measure of risk had to respect in order to be classified as coherent. Except for elliptical distributions, VaR was proved not to satisfy the property of subadditivity and hence not to take into consideration the principle of risk diversification. Despite that, VaR was proved to satisfy the other axioms: monotonicity, positive homogeneity and translation invariance.

In light of that and given the actual characteristics of the 2002 LDCE operational risk losses –a high level of kurtosis and a distributional behaviour very far from the Normal one - different measures of risk are called for in order to have a consistent and reliable view of the actual riskiness of the eight BLs. Quantities which estimate the shortfall risk appear to be the proper tool, since they provide information on the magnitude of the whole tail and, moreover, they were proved to be coherent measures of risk (see the abovementioned article by Artzner *et al.*).

The most popular of these measures is the Expected Shortfall (ES), which estimates the potential size of the loss exceeding a selected level L of the distribution (the level L may be associated, for example, to a preset threshold u or to the VaR_p itself). The expression for the ES is:

$$ES(L) = L + E(X - L | X > L) = L + MEF(L) \quad (19)$$

where the second term in the formula is simply the Mean Excess Function (13), introduced in Section 6.

In the GPD model with threshold u and parameters ξ and β , the expression for the ES is the following³³:

$$GPD_{ES}(u) = \frac{\beta + \xi u}{1 - \xi} \quad \text{in the case of } L = u$$

$$GPD_{ES}(VaR_p) = \frac{VaR_p}{1 - \xi} + \frac{\beta - \xi u}{1 - \xi} \quad \text{in the case of } L = VaR_p > u$$

which is defined only for values of the shape $\xi < 1$.

In light of the BLs shape estimate obtained in the current exercise (greater than 1 or very close to 1, see Table 5), it is evident that the GPD_{ES} cannot be consistently used. Accordingly, alternative measures of shortfall risk are called for.

³³ Given the POT stability propriety, the conditional distribution over the level L , $F_L(y) = P(X-L \leq y | X > L)$, is also GPD with the same ξ , location = L and scale = $\beta + \xi(L-u)$. The expected value of $F_L(y)$ is the $MEF(L)$ and assumes the form $[\beta + \xi(L-u)]/(1-\xi)$. Substituting in (19), it is possible to calculate the expression for the $GPD_{ES}(L)$, in both the cases of $L = u$ and $L = VaR_p$.

A solution lies in a quantity which resorts to the Median Excess Function [MEDEF (u)], that is to the median of excesses over a threshold u: $MEDEF(u) = [F_u(1/2)]^{-1}$.

In particular, the $MEDEF(u)$ for the GPD model can be derived, firstly, by inverting the expression (7) – that is the “excess GPD” at u, $GPD_{\xi,\beta}(y)$ – for a generic probability p, to get:

$$[GPD_{\xi,\beta}(p)]^{-1} = \frac{\beta}{\xi} [(1-p)^{-\xi} - 1]$$

and then imposing $p = 1/2$

$$GPD_{MEDEF}(u) = \frac{\beta}{\xi} [2^\xi - 1] \quad (20)$$

Making use of the $GPD_{MEDEF}(u)$, an appropriate measure of shortfall risk, strictly connected to the $GPD_{ES}(u)$, can be easily derived. The expression is:

$$GPD_{MS}(u) = u + GPD_{MEDEF}(u) = u + \frac{\beta}{\xi} [2^\xi - 1] \quad (21)$$

which may be identified as the (GPD) Median Shortfall at level u.

The nice feature of the $GPD_{MS}(u)$ is that, unlike the GPD_{ES} , it is defined regardless of the values of the shape (see Reiss and Thomas, 2001, p. 56), while preserving, similarly to the GPD_{ES} , the property of the POT stability (see Rootzen and Tajvidi, 1997). In particular, given the quantity $GPD_{MS}(u)$ computed at the threshold u, the $GPD_{MS}(v)$ at the higher level $v > u$ can be expressed as:

$$GPD_{MS}(v) = v + \frac{\beta + \xi(v-u)}{\xi} [2^\xi - 1] \quad (22)$$

In light of its features, the GPD_{MS} represents the suitable and reliable risk measure used in this exercise to compute the tail-severity riskiness of the eight BLs.

In Table 9, for each BL, the GPD_{MS} values computed at the starting threshold u – as previously noted, close to the 90th empirical percentile for all the BLs except for BL3 (96.5th) – and at the higher thresholds v (set close to significant empirical percentiles in the range 90th - 99.9th) are reported.

Table 9: BLs tail-severity magnitude (Euro ,000)

BUSINESS LINE	90°	91°	92°	93°	94°	95°	96°	97°	98°	99°	99.5°	99.9°
BL 1 (Corporate Finance)	1,234	1,383	1,740	1,861	2,245	3,383	4,256	7,311	11,927	19,030	111,577	260,415
BL 2 (Trading & Sales)	464	535	599	705	878	1,121	1,436	2,180	3,493	7,998	17,824	70,612
BL 3 (Retail Banking) ^(*)	481	481	481	481	481	481	481	550	831	1,694	3,232	17,411
BL 4 (Commercial Banking) ^(**)	750	764	896	1,070	1,322	1,694	2,776	4,490	8,453	20,063	39,246	151,553
BL 5 (Payment & Settlement)	227	262	302	342	415	551	697	914	1,714	3,910	9,950	80,518
BL 6 (Agency Services) ^(**)	466	579	661	779	950	1,208	1,669	2,619	5,324	10,107	22,636	51,805
BL 7 (Asset Management)	531	580	656	716	887	1,076	1,431	1,972	2,457	4,264	13,630	79,423
BL 8 (Retail Brokerage)	273	295	329	376	427	520	643	909	1,200	2,222	4,954	27,628
TOTAL	4,428	4,879	5,664	6,331	7,604	10,034	13,390	20,946	35,400	69,287	223,050	739,364

(*) As the starting threshold for *Retail Banking* is close to the 96.5th percentile, the GPD_{MS} is constant under this level.

(**) As the GPD_{MS} expression (22) makes use of the thresholds v and u as non-parametric components, it is not directly comparable with the GPD_{VaR} (18), whose non-parametric components are represented, other than the threshold u itself, by the number (n) of total observations and the number (n_u) of exceedances with respect to the threshold u . It follows that at some percentiles, the GPD_{MS} figure may be lower than the GPD_{VaR} (in the current exercise, this is the case for *Commercial Banking* and *Agency Services* at the percentile of 99.9th).

Table 9 clearly indicates that:

- owing to the tail heaviness of the data, each BL's severity riskiness increases remarkably at the highest percentiles. On average over the eight BLs, the ratio between the 99.9th and the 99th percentiles is about 10, with a peak of about 20 in BL5 (*Payment & Settlement*) and BL7 (*Asset Management*);
- the ranking of the riskiness of the BLs does not change significantly as the threshold is progressively raised up to the 99.9th percentile. In particular *Corporate Finance* and *Commercial Banking* are found to be the riskiest BLs with an estimated 99.9th percentile of severity loss of € 260 million and € 151 million, respectively. On the other hand, *Retail Banking* and *Retail Brokerage* prove to be the least risky BLs, showing a 99.9th severity loss of € 17 million and € 27 million, respectively.

9. Peaks Over Threshold approach: Business Lines tail-frequency estimate

The exercises carried out in the previous Sections focused on the estimate and the measurement of the (tail) severity component of the distribution of losses. As noted, the severity figure provides information only on the potential size of the losses that a (large internationally active) bank can bear, regardless of the holding period in which the losses occur. In practice such as unconditional figure does not take into consideration the frequency of the losses incurred in a given time horizon and therefore it is not yet adequate to represent the BLs capital charge.

In fact, supposing that the operational risk capital charge should be determined for 1-year holding period at the 99th percentile, if the final estimate were based only on the unconditional $GPD_{MS}(99^{th})$ amount, this figure would underestimate or overestimate the actual risk, depending on whether the probability of a 1-year occurrence of the losses with a single-impact magnitude bigger than the 99th percentile of the severity distribution is respectively higher or lower than 1 per cent.

The purpose of this Section is thus to supplement the GPD analysis carried out in the previous Sections, by estimating and measuring the frequency of the large losses for each BL.

To achieve this, the POT approach, so far limited to its severity component (POT-GPD), is now totally implemented by means of the Point Process representation of the exceedances (POT-PP).

The basic assumption of this method - developed as a probabilistic technique by Leadbetter *et al.*, 1983, and Resnick, 1987, and as a statistical tool by Smith, 1989 - is to view the number of exceedances and the excesses as a marked point process with a proper intensity, that, in its basic representation, converges to a two-dimensional Poisson process. In practice:

- a) the exceedances (x) over a threshold u occur at the times of a Poisson process with intensity λ ;
- b) the corresponding excesses ($y=x-u$) are independent and have a GPD distribution;
- c) the number of exceedances and the excesses are independent of each other.

The parameter λ measures the intensity of the exceedances at u per unit of time, that is if the number of large losses is stable over time or if it becomes more or less frequent.

In the basic case of stationarity of the process ³⁴, the number of exceedances occurs as a homogeneous Poisson process with a constant intensity, which can be written as:

$$\lambda_u = \left(1 + \xi \frac{u - \mu}{\sigma} \right) \quad (23)$$

where ξ , μ and σ , as usual, represent the shape, the location and the scale parameters of the “full GPD” and λ_u is supplied with the subscript to stress the dependence on the threshold u (it is assumed in fact that this expression is valid only for $x \geq u$).

Since λ_u should be measured in the same time units as used for the collection of the data, an estimate of the time-adjusted number of exceedances in a certain period T can be simply obtained by $\lambda_u T$. For instance, if the per-bank number of exceedances in a 1-year period is to be determined and the data collection refers indistinctly to working days and holidays, the annualised intensity of exceedance will be:

$$N_{1-year} = 365 \lambda_u \quad (24)$$

³⁴ It should be noted, however, that the Point Process is robust against the non stationarity of data. In fact if evidence of time-dependence of large losses is detected, some, or all, the parameters can be a function of time and, in that case, the model would be a non-homogeneous Poisson process (for example, if the intensity rate was not constant over time, a smoothing of $\lambda=\hat{\lambda}(t)$ over t may be appropriate). Furthermore, specific techniques exist to handle other kinds of dependence of data, such as clustering of excesses. The only assumption that has to be preserved to guarantee the robustness of the model is that the distribution of excesses is approximated by a GPD. For examples of the application of the POT method in presence of trends of exceedances, see the case studies of “tropospheric ozone” and “wind-storm claims”, carried out respectively by Smith and Shively, 1995, and Rootzen and Tajvidi, 1997. For a smoothing technique to incorporate trends in the model, see Chavez-Demoulin and Embrechts, 2003. For declustering methods, see Smith, 1989, Rootzen *et al.*, 1992, and Ferro and Segers, 2003.

Equation (24) is tailored to the operational risk data, since the time of occurrence of a single operational loss (either small or large) does not appear to be dependent on the working days³⁵.

One of the nicest properties of the POT-PP, similarly to the POT-GPD, is its stability under an increase of the threshold: if a Point Process at the threshold u converges to a Poisson process with intensity λ_u , the Point Process at level $v > u$ also converges to a Poisson process. The new intensity is:

$$\lambda_v = \lambda_u \left(1 + \xi \frac{v-u}{\beta} \right) \quad (25)$$

where β is the scale parameter of the “exceedance GPD” $\text{GPD}_{\xi,u,\beta}$ (see Rootzen and Tajvidi, 1997). The same relationship holds if the intensity for unit of time λ_u is substituted by a certain time-adjusted intensity, which, as noted, identifies the mean number of exceedances in a given holding period (see Reiss and Thomas, 2001, p. 286):

$$N_{T,v} = N_{T,u} \left(1 + \xi \frac{v-u}{\beta} \right) \quad (26)$$

where $N_{T,\cdot}$ represents the per-bank mean number of exceedances (at v or at u) in a period of length T .

By expression (25) or (26), it is thus possible to analytically derive frequency figures for large losses in correspondence with higher thresholds than the initial level.

In order to employ the POT-PP approach in the current exercise, an estimate of $N_{T,u}$ in each BL at the starting threshold u is needed; subsequently the (26) may be used to get the mean number of exceedances $N_{T,v}$ at higher thresholds v . The holding period is, obviously, the time window employed in the 2002 LDCE, i.e. 1-year.

In general, an empirical estimate of the average annualised number of exceedances ($N_{1\text{-year},u}$) occurred in a bank can simply be obtained from the total number of exceedances occurred over the years of the data collection divided by the number of years itself.

³⁵ Operational loss events such as business disruption, fraud, external events, etc. may occur either on working days or on holidays. On the contrary, the annualised intensity for financial losses tends to be connected to the actual trading days, i.e. $N_{1\text{-year}} = 250 \lambda_u$.

Therefore, if the observations pooled in each BL data set were referred to only one bank, $N_{1\text{-year},u}$ would have been merely the count of exceedances over the threshold identified in the GPD analysis (see the last column of Table 4).

However, it is to remind that each BL data set is the result of pooling the observations of a, distinct, number n of banks (see Table 1) and, as stated in Section 2, it can reasonably be assumed equivalent to the collection of i.i.d. losses from a medium-sized LDCE bank during a time window of n -years. In light of that, a consequent, easy, way to identify an empirical average annualised intensity N_u^i pertaining to the i -th BL (omitting the 1-year notation for simplicity) is the ratio between the total number of exceedances that occurred in that BL and the number of banks providing data to the i -th BL itself (column 5 of Table 4 divided by column 2 of Table 1).

Nevertheless, a per-bank analysis of the actual number of exceedances at the GPD starting threshold occurring in each BL reveals that this number is rather widespread over the panel of banks: on average across the BLs, it results equal to 0 in 30 per cent of the cases (to say, no exceedances occur) and assumes values much larger than the mean in other, not negligible, cases. The reasons for this may lie in the different levels of comprehensiveness in the collection of (large) losses among the banks participating in the RMG survey and, perhaps, also in the short time horizon of the 2002 LDCE (1-year data collection), which might have caused, for some banks, a few gaps in the collection of very rare and large losses, especially if driven by banks' external sources (see footnote 8). Another likely cause of the variability of the frequency of large losses across the panel may lie in the actual presence of banks having different size and hence potentially in a position to produce, in some BLs, a lower or higher number of large losses in a given time horizon ³⁶.

These frequency issues are specifically addressed and mitigated in this Section in order to reduce the threat of obtaining biased estimates of the BLs frequency of large losses and hence of the capital figures.

³⁶ Although it is realistic to assume that banks having similar characteristics are not distinguishable by the severity of the losses (see Section 2, in particular footnote 8), the repetitiveness of losses, both small and large, appears to be actually linked to the volume of business carried out, that is to the size of the bank. A "frequency scaling methodology" is therefore called for.

In particular, to explicitly take into account the potential differences in banks' size – which, as noted, may affect the frequency of large losses – it is assumed the existence in the 2002 LDCE of two distinct groups of banks: a “lower group”, consisting of banks with smaller size (in fact **domestic banks**), and an “upper group”, consisting of banks with larger size (in fact **internationally active banks**). For both a typical domestic bank and a typical international active bank, suitable estimates of the annualised intensity of exceedances are gained.

In order to get such estimates, in each BL the number of per-bank exceedances is fitted by a Poisson and Binomial Negative model. Owing to the high skewness to the right, the Binomial Negative distributions result the best-fitting ones in all the datasets: accordingly, in each BL, the “mean” and the “mean plus two standard deviations” of that distribution are assumed to be respectively the average annualised intensity of exceedances for a typical domestic bank (N^{low}) and the average annualised intensity of exceedances for a typical international active bank (N^{high}).

Table 10 shows, for each BL, the number of banks providing data and the number of banks with at least one exceedance over the relevant GPD threshold, as well the parameters estimate of the Binomial Negative distribution (r and p). On the basis of such estimates, the mean number of exceedances at u for typical domestic and international active banks are computed (the last two columns of the Table).

Table 10: BLs average annualised frequency of large losses at the initial threshold

BUSINESS LINE	n. banks providing data	n. banks with at least one exceedance	Binomial Negative parameters estimate		Annualised intensity at threshold u	
			r	p	N^{low}	N^{high}
BL 1 (Corporate Finance)	33	15	0.45	0.25	1.30	5.83
BL 2 (Trading & Sales)	67	48	0.37	0.05	7.78	33.92
BL 3 (Retail Banking)	80	66	0.26	0.02	12.36	61.75
BL 4 (Commercial Banking)	73	54	0.47	0.10	4.36	17.68
BL 5 (Payment & Settlement)	55	36	0.51	0.13	3.42	13.66
BL 6 (Agency Services)	40	23	0.30	0.07	3.97	19.10
BL 7 (Asset Management)	52	29	0.52	0.20	2.08	8.53
BL 8 (Retail Brokerage)	41	26	0.24	0.03	7.98	41.16

After computing the average annualised frequency of large losses at the threshold u , the figure for the mean number of exceedances at higher thresholds (set in correspondence with the same significant empirical percentiles as those of the GPD_{MS} ; see Table 9) may be derived from (26). In Table 11, the BLs tail-frequency magnitudes for typical international active and domestic banks are reported.

Table 11: BLs tail-frequency magnitude

Typical international active bank : N^{high}

BUSINESS LINE	90°	91°	92°	93°	94°	95°	96°	97°	98°	99°	99.5°	99.9°
BL 1 (Corporate Finance)	5.83	5.38	4.55	4.33	3.77	2.75	2.29	1.48	0.99	0.68	0.15	0.08
BL 2 (Trading & Sales)	33.92	30.24	27.53	24.08	20.08	16.37	13.30	9.36	6.28	3.10	1.57	0.48
BL 3 (Retail Banking) ^(*)	61.75	61.75	61.75	61.75	61.75	61.75	61.75	53.88	35.45	17.37	9.14	1.73
BL 4 (Commercial Banking)	17.68	17.47	15.63	13.80	11.89	9.98	7.02	4.98	3.17	1.70	1.05	0.40
BL 5 (Payment & Settlement)	13.66	12.00	10.59	9.49	8.04	6.31	5.17	4.12	2.45	1.24	0.58	0.11
BL 6 (Agency Services)	19.10	15.98	14.34	12.52	10.64	8.73	6.70	4.63	2.59	1.53	0.79	0.40
BL 7 (Asset Management)	8.53	7.85	6.97	6.40	5.16	4.22	3.12	2.20	1.72	0.92	0.24	0.03
BL 8 (Retail Brokerage)	41.16	37.69	33.48	28.90	25.17	20.39	16.27	11.30	8.46	4.47	1.96	0.34
TOTAL	202	188	175	161	146	130	116	92	61	31	15	4

Typical domestic bank: N^{low}

BUSINESS LINE	90°	91°	92°	93°	94°	95°	96°	97°	98°	99°	99.5°	99.9°
BL 1 (Corporate Finance)	1.30	1.20	1.02	0.97	0.84	0.61	0.51	0.33	0.22	0.15	0.03	0.02
BL 2 (Trading & Sales)	7.78	6.94	6.31	5.52	4.61	3.75	3.05	2.15	1.44	0.71	0.36	0.11
BL 3 (Retail Banking) ^(*)	12.36	12.36	12.36	12.36	12.36	12.36	12.36	10.79	7.10	3.48	1.83	0.35
BL 4 (Commercial Banking)	4.36	4.30	3.85	3.40	2.93	2.46	1.73	1.23	0.78	0.42	0.26	0.10
BL 5 (Payment & Settlement)	3.42	3.00	2.65	2.37	2.01	1.58	1.29	1.03	0.61	0.31	0.15	0.03
BL 6 (Agency Services)	3.97	3.33	2.98	2.60	2.21	1.82	1.39	0.96	0.54	0.32	0.16	0.08
BL 7 (Asset Management)	2.08	1.91	1.70	1.56	1.26	1.03	0.76	0.54	0.42	0.23	0.06	0.01
BL 8 (Retail Brokerage)	7.98	7.30	6.49	5.60	4.88	3.95	3.15	2.19	1.64	0.87	0.38	0.07
TOTAL	43	40	37	34	31	28	24	19	13	6	3	1

(*) As the starting threshold for *Retail Banking* is close to the 96.5th percentile, N is constant below this level.

For an international active bank, the number of exceedances N^{high} at the starting thresholds is 202. In correspondence with the threshold of € 1 million ³⁷, the POT-PP model reveals a total number of exceedances in the region of 60. This value is absolutely comparable with that of large internationally active banks, which provide evidence of an average number of losses above \$1 million, ranging from 50 to 80 per year (see De Fontnouvelle, Deleuss-Rueff, Jordan and Rosengren, 2003).

However, in both the “lower group” and the “upper group” of banks, the number of exceedances identified by the model at the highest percentiles (99.5th and 99.9th) declines remarkably, giving force to the above-mentioned hypothesis of some gaps in the collection of the very large losses, which the pooling exercise has not been able to recover sufficiently.

In the next section, devoted to computing the capital figure of the BLs, this possible incompleteness in the number of extreme losses will be mitigated by placing a floor on the, 1-year period, probability of occurrence of the losses with a single-impact magnitude bigger than the $GPD_{MS}(99^{\text{th}})$ amount: the floor will be the number of exceedances identified by the POT-PP model just at the 99th percentile ³⁸.

Finally, it should be noted that the number of exceedances across the range of percentiles is, for a typical international active bank, up to 4-5 times that of a typical domestic bank; as it will be shown in the next Section, this will have a significant effect on the magnitude of the capital charge pertinent to this category of banks.

10. Business Lines capital charge

On the basis of the POT tail-frequency and tail-severity values gained in the previous Sections, it is now possible to compute, for each BL and at any desired percentile, an estimate of the aggregated figure, which represents the operational risk capital charge required to cover expected plus unexpected losses in a 1-year holding period.

³⁷ Obviously, the threshold of € 1 million is placed in each BL at a specific empirical percentile, which is not necessarily identical across the BLs.

³⁸ As it will be seen in the next Section, this assumption implies that the frequency figures used to compute the 99.5th and the 99.9th percentiles of the aggregated losses are stable and equal to the (higher) frequency value identified by the POT-PP model at the 99th percentile.

An easy and sound way to compute such an estimate is by means of a similar measure to that adopted in the Risk Theory for the calculation of the “excess claims net premium”, i.e. the “average frequency of exceedances” times the “average severity of excess ” (see Reiss and Thomas, 2001, p. 279-87). In the current exercise it is sufficient to substitute the “average severity of excesses” with the “average severity of exceedances” (the excesses plus the threshold) to obtain outcomes that, at any percentile of the aggregated losses, move on from the “unexpected” to the “expected plus unexpected” side of the overall distribution.

However, before such an exercise may be performed, it is worthwhile to draw attention to the fact that the extent of the capital figures for rising percentiles depends on the increase of the severity of the exceedances compared with the reduction of their frequency of occurrence. Under this perspective, the nice feature of the POT approach, in its POT-GPD and POT-PP representations, is that it binds analytically the severity of large losses to their frequency as the corresponding thresholds (say, percentiles) are progressively raised, thus allowing the two components of the aggregated losses to be jointly addressed. The intensity of exceedances, λ_u , identified by (23), or even its time-adjusted figure, $N_{T,u}$, is the “bridge” which naturally connects the severity of large losses to their frequency at the initial threshold, u , while expression (22) for the severity ³⁹, and equations (25) or (26) for the frequency, constitute the tools that take care, from the starting threshold to the highest percentiles, of both the components in a mathematical, mutually coherent, manner. So, through a suitable combination of these outcomes (for instance simply by multiplying them), a reliable analytical figure for the aggregated losses at any desired (high) percentile can be derived.

This approach differs sharply from the conventional actuarial approach, where - except for the rare case in which the expression for the compound distribution of the aggregated losses is analytically derivable from the distributions of its components of frequency and severity - the computation of a (high) percentile of the aggregated losses is obtained by treating the estimate of the severity and the frequency components as a separate, disjointed, problem and, afterwards, aggregating the corresponding outcomes using numerical,

³⁹ Obviously, appropriate expressions, alternative to (26), would be needed if different measures of risk (for instance GPD_{ES} or GPD_{VaR}) were adopted in place of the GPD_{MS} .

approximation or simulation methods (i.e. the MonteCarlo procedure)⁴⁰. Owing to the absence of an analytical basis, these methods require many steps to be generated to calculate the highest percentiles of the aggregated distribution of losses⁴¹.

Moreover, even supposing that only the large losses constitute the input of the analysis (for example the data exceeding the 90th percentile of the empirical distribution) and that the GPD is the selected distribution for the severity of the losses, whenever the frequency is modelled by conventional actuarial models and an approximation or simulation method is then implemented to derive the aggregated losses, the figure for the highest percentiles of the aggregated losses (for example the 99th or 99.9th ones) would be overestimated⁴². This occurs because the frequency of large losses stemming from the procedure would be, on average, anchored to the (higher) values observable in correspondence with the start of the tail area (in this case, the 90th percentile) wherein the data are more abundant. Not effectively addressing the reduction of the frequency of large losses that occurs as their size increase depends exclusively on the, disjointed and not analytical, method employed to compute the aggregated losses; this may produce a significant bias in the estimate of the highest percentiles.

The advantage of the POT approach in the estimate of the tail of the aggregated losses therefore appears directly connected to the two following properties:

Property 1: the POT method takes into consideration the (unknown) relationship between the frequency and the severity of large losses up to the end of the distribution;

⁴⁰ Under the condition of homogeneity of the actuarial risk model – i.e. with the loss severity variables (X_i) i.i.d. and independent from the loss frequency variable (N) – the only quantities of the aggregated loss variable (S) that are always analytically derivable from the distributions of frequency, $P(N)$, and severity, $F(X)$, are its moments. In particular the expectation of total losses, $E(S)$ – the so called net premium – can be obtained simply as the product of the expectations of the frequency, $E(N)$, and the severity, $E(X)$. On the other hand, if some specific ordinal statistic of S is to be computed (as VaR_p or even ES_p), it is necessary to identify the complete form of the compound distribution of the total losses, $S = \sum_k P(k) F^{k*}(x)$, where $F^{k*}(x)$ is the k -fold convolution of $F(X)$. Since the compound distribution, S , can be represented in an analytical way only in very rare cases, a MonteCarlo simulation is usually used. At each step, the MonteCarlo procedure generates a number, n , of losses from $P(n)$ and n amounts of losses, x_i , from $F(x)$ and then computes the additive quantity, $S_n = \sum_{i=1..n} x_i$. Finally, from the empirical distribution of S_n it is possible to identify the desired percentile.

⁴¹ A MonteCarlo simulation performed by David Lawrence (Citigroup) and presented to the 10th annual ICBI conference in Geneva, 2003, showed that about 1,000,000 data points are required to calculate the 99th percentile of the aggregated losses, stemming from a compound Poisson (4000)-LogNormal (-4,1.8,3.5) distribution. More data would be necessary to estimate equivalent or higher percentiles of compound distributions originated from severity and frequency components with greater values of skewness and kurtosis.

⁴² This approach is largely adopted by practitioners. For theoretical references, see, *inter alia*, King, 2001.

Property 2: the POT method makes it possible to employ a semiparametric approach to compute the highest percentiles of the aggregated losses, hence reducing the computational cost and the estimate error related to a not analytical representation of the aggregated losses themselves. In the POT model, it suffices to select a suitable (high) threshold, on which basis the model can be built and the relevant parameters estimated. Once the model is correctly calibrated⁴³, the total losses (and their percentiles) are easily obtainable by proper analytical expressions.

Having said that, all the information required to get the capital charge for the BLs are in Tables 9 and 11, which report, respectively, the tail-severity and the tail-frequency magnitudes identified by the POT approach. For each BL, the adoption of a related measure of risk to the “excess claims net premium” produces, as previously noted, the following expressions:

$$\text{CaR}_i = N_i^{\text{low}} \text{GPD}_{MS}(i) \quad \text{for a typical domestic bank}$$

$$\text{CaR}_i = N_i^{\text{high}} \text{GPD}_{MS}(i) \quad \text{for a typical international active bank}$$

where CaR is the Capital at Risk and i ranges from the starting threshold, u (as observed, close to the 90th percentile for all the BLs, except for BL3, around the 96.5th), to that close to the 99.9th percentile.

The BLs capital charges for typical international active and domestic banks are reported in Table 12, where, as anticipated in Section 9, the aggregated losses at the highest levels of 99.5th and 99.9th are forced to make use of the frequency floor (i.e. the number of exceedances identified by the POT-PP model at the 99th percentile) in place of the (lower) pertaining annualised intensities⁴⁴. In the last column of Table 12, the contribution of each BL to the total capital charge is also reported.

⁴³ Depending on the data, it might be necessary, for instance, to relax the hypothesis of homogeneity of the parameters in order to incorporate a trend in the model (see footnote 34).

⁴⁴ Although this assumption violates the Property 1 of the POT model, it seems a reasonable compromise between the objectives of circumventing the problem of the, likely, incompleteness in the frequency of the largest losses in the 2002 LDCE (as noted in Section 9) and not introducing a too high level of discretionality in the determination of the extreme percentiles of the total losses.

Generally speaking, as the data collection is reputed less satisfactory, a proportionally higher floor for the frequency to be associated to the highest percentiles of the severity should be adopted. However, in order to not

Table 12: BLs capital charge**Typical international active bank (Euro ,000)**

BUSINESS LINE	90°	95°	96°	97°	98°	99°	99.5°	99.9°	99.5° with floor = N ₉₉	99.9° with floor = N ₉₉	99.9° with floor = N ₉₉ - % on Total
BL 1 (Corporate Finance)	4,859	8,192	8,837	10,253	11,458	12,587	17,147	19,693	75,116	175,676	13.3%
BL 2 (Trading & Sales)	9,205	15,199	16,539	18,593	20,710	24,209	27,608	34,027	54,687	218,423	16.5%
BL 3 (Retail Banking) ^(*)	14,459	14,459	14,459	16,320	20,715	25,142	27,280	29,629	51,860	298,218	22.5%
BL 4 (Commercial Banking)	8,494	14,209	17,605	21,030	25,919	33,707	40,980	60,198	66,376	257,633	19.4%
BL 5 (Payment & Settlement)	1,599	2,781	3,035	3,314	3,927	4,728	5,713	8,515	12,244	100,052	7.6%
BL 6 (Agency Services)	5,058	8,789	9,827	11,185	13,245	15,142	17,707	20,661	34,296	78,886	6.0%
BL 7 (Asset Management)	2,527	3,549	3,730	3,820	3,830	3,727	3,237	2,429	12,389	73,237	5.5%
BL 8 (Retail Brokerage)	5,065	7,551	8,026	8,581	8,883	9,262	9,419	9,290	21,469	122,788	9.3%
TOTAL	51,267	74,729	82,059	93,097	108,687	128,503	149,090	184,442	328,437	1,324,912	100%

Typical domestic bank (Euro ,000)

BUSINESS LINE	90°	95°	96°	97°	98°	99°	99.5°	99.9°	99.5° with floor = N ₉₉	99.9° with floor = N ₉₉	99.9° with floor = N ₉₉ - % on Total
BL 1 (Corporate Finance)	1,087	1,832	1,976	2,293	2,562	2,815	3,834	4,404	16,798	39,286	13.3%
BL 2 (Trading & Sales)	2,111	3,486	3,793	4,264	4,750	5,552	6,332	7,804	12,543	50,097	16.9%
BL 3 (Retail Banking) ^(*)	2,895	2,895	2,895	3,267	4,147	5,033	5,461	5,932	10,382	59,702	20.2%
BL 4 (Commercial Banking)	2,093	3,501	4,337	5,181	6,386	8,305	10,096	14,831	16,353	63,474	21.5%
BL 5 (Payment & Settlement)	400	696	760	829	983	1,183	1,430	2,131	3,064	25,036	8.5%
BL 6 (Agency Services)	1,053	1,829	2,045	2,328	2,756	3,151	3,685	4,300	7,137	16,417	5.6%
BL 7 (Asset Management)	616	865	909	931	933	908	789	592	3,018	17,841	6.0%
BL 8 (Retail Brokerage)	981	1,463	1,555	1,663	1,721	1,795	1,825	1,800	4,160	23,791	8.0%
TOTAL	11,235	16,566	18,270	20,756	24,238	28,742	33,452	41,793	73,455	295,644	100%

(*) As the starting threshold for *Retail Banking* is close to the 96.5th percentile, the CaR is constant below this level.

The findings clearly indicate that operational losses are a significant source of risk for banks: overall, the 99.9th, 1-year period, CaR amounts to € 1,325 million for a typical international active bank and to € 296 million for a typical domestic bank. Owing to the

enlarge the final aggregated figure excessively, this “rule of thumb” should also establish a minimum, preset, level (for example the 95th percentile) at which the frequency floor can be derived (as an example, interested readers can verify the increase of the 99.9th capital charges that would derive from setting the floor equivalent to the number of exceedances pertaining to the 90th percentile (N_{90°}) instead of the current (N_{99°})).

higher frequency of occurrence of losses, *Retail Banking* and *Commercial Banking* are the BLs which absorb the majority of the overall capital figure (about 20 per cent each in both the groups), while *Corporate Finance* and *Trading & Sales* come in at an intermediate level (respectively close to 13 per cent and 17 per cent). The other BLs stay stably below 10 per cent in both the groups, with *Asset Management* and *Agency Services* showing the smallest capital charges (around 6 per cent each). These figures are comparable with the allocation ratios of economic capital for operational risk reported by the banks participating in the 2002 LDCE (see Table 21 of the 2002 LDCE summary).

The last part of this Section is devoted to measuring, for an international active bank, the contribution of the expected losses to the overall capital figures.

In fact, as stated in Section 5, the results of the conventional inference support the hypothesis that, in each BL, the small/medium-sized operational risk data (i.e. the body of the loss distribution) have a different statistical “soul” than the tail-data. This means that the outcomes of the GPD analysis - in particular the implausible implications for any mean value of the distribution, owing to an estimate of the shape parameter (ξ) higher or close to 1 - are not applicable to the body-data. Therefore, by focusing the analysis on the body-data only and making use of the conventional inference, an estimate of the operational risk expected severity is tenable in each BL: in particular, the results of the LogNormal model are used (see Table 3), since this distribution proved to have satisfactory properties in representing the small/medium-sized area of the data.

Besides, in order to gain information on the (1-year) frequency of occurrence of the small/medium-sized losses, for each BL, the Poisson and the Binomial Negative distributions are fitted to the per-bank number of losses having a single-impact magnitude lower than the threshold identified in the GPD analysis. For all the BLs with the exception of *Retail Banking*, the Binomial Negative curve shows a more accurate fit, in consideration of the substantial skewness to the right of the frequency of data: accordingly, BLs related parameters estimate for the Binomial Negative models are gained.

The average values of the LogNormal and Binomial Negative models, arisen from the estimated parameters, are consequently assumed to represent, respectively, the expected

severity and the expected frequency of the operational risk datasets⁴⁵. The expected losses are obtained by the simple product of these values.

For each BL, Table 13 shows the LogNormal parameters already computed in Section 4 and the Binomial Negative parameters just now estimated. In the last two columns, the CaR at the 99.9th percentile (with frequency floor at $N_{99.9}$) and the ratio between the expected losses and the CaR is reported.

Table 13: BLs expected losses (absolute and relative to CaR_{99.9}) (Euro, 000)

BUSINESS LINE	LogNormal parameters estimate		Binomial Negative parameters estimate		Expected severity	Expected frequency	Expected Losses	CaR _{99.9%} with floor = $N_{99.9}$	Expected Losses/ CaR _{99.9%}
	μ	σ	r	p					
BL 1 (Corporate Finance)	3.58	1.71	0.59	0.04	154	12.67	1,953	175,676	1.1%
BL 2 (Trading & Sales)	3.64	1.27	0.45	0.01	85	74.45	6,359	218,423	2.9%
BL 3 (Retail Banking)	3.17	0.97	n.a.	n.a.	38	347.45	13,172	298,218	4.4%
BL 4 (Commercial Banking)	3.61	1.41	0.52	0.01	100	43.90	4,405	257,633	1.7%
BL 5 (Payment & Settlement)	3.37	1.10	0.61	0.02	53	32.00	1,711	100,052	1.7%
BL 6 (Agency Services)	3.74	1.28	0.47	0.01	96	35.03	3,375	78,886	4.3%
BL 7 (Asset Management)	3.79	1.28	0.60	0.03	100	20.02	2,011	73,237	2.7%
BL 8 (Retail Brokerage)	3.58	1.08	0.34	0.00	64	75.55	4,811	122,788	3.9%
TOTAL							37,797	1,324,912	2.9%

⁴⁵ For *Retail Banking*, both the Poisson and the Binomial Negative models show a low performance; consequently the figure for the expected frequency in each BL is gained by an empirical estimate of the average number of small/medium-sized losses borne by the banks providing data to that BL.

The results reveal the very small contribution of the expected losses to the total charge: on average over the BLs it is less than 3 per cent ($\text{CaR}_{99,9}/\text{EL} = 36$), with a minimum of 1.1 per cent in *Corporate Finance* ($\text{CaR}_{99,9}/\text{EL} = 89$) and a maximum of 4.4 per cent in *Retail Banking* ($\text{CaR}_{99,9}/\text{EL} = 22$)⁴⁶.

Once again, these outcomes confirm the very tail-driven nature of operational risk.

11. Business Lines capital charge and Gross Income relationship: bottom-up coefficients

The exercise performed in this Section is limited to the “upper group” of banks, which, as previously noted, is assumed to contain the internationally active banks. Not enough information is available in the database to carry out a similar analysis on the “lower group” of banks, which is supposed to contain the domestic banks.

The objective of the analysis is to determine for each BL (and in the eight BLs as a whole) the relationship between the capital charge required to an international active bank to cover (expected plus unexpected) operational risk losses and the value of an indicator measuring the average business produced in that BL. The indicator of volume of business chosen is the Gross Income (GI), that is the variable identified by the Basel Committee to derive the regulatory coefficients in the simpler approaches (Basic and Standardised) of the new Capital Accord framework.

Therefore, for each BL, an average GI is obtained from the individual GIs of the banks which furnished operational risk data to that BL; the total GI across the eight BLs is obtained as the simple sum of these average figures. Then the ratios between the BLs capital charges ($\text{CaR}_{99,9}$) and the average GIs are computed and compared with the current regulatory coefficients envisaged in the Basic and Standardised Approach of the Capital Accord (the so-called Alpha and Betas; see Table 14).

⁴⁶ The same analysis carried out on the “lower group” of banks (the domestic banks) reveals a slightly larger contribution of the expected losses to the total capital charge: on average they measure about 12 per cent of the $\text{CaR}_{99,9}$ (with frequency floor $N_{99,9}$), with a minimum of 5 per cent in *Corporate Finance* and a maximum of 22 per cent in *Retail Banking*.

Table 14: BLs capital charge and Gross Income relationship (Euro ,000).**Bottom-up coefficients vs. current regulatory coefficients**

BUSINESS LINE	n. banks providing data	Capital charge for op. risk (1-year at 99.9°) A	Average Gross Income B	Bottom-up coefficients A/B	Current regulatory coefficients
BL 1 (Corporate Finance)	33	175,676	1,056,568	16.6%	18%
BL 2 (Trading & Sales)	67	218,423	1,723,483	12.7%	18%
BL 3 (Retail Banking)	80	298,218	3,580,369	8.3%	12%
BL 4 (Commercial Banking)	73	257,633	1,829,454	14.1%	15%
BL 5 (Payment & Settlement)	55	100,052	300,153	33.3%	18%
BL 6 (Agency Services)	40	78,886	375,638	21.0%	15%
BL 7 (Asset Management)	52	73,237	454,130	16.1%	12%
BL 8 (Retail Brokerage)	41	122,788	632,445	19.4%	12%
	TOTAL	1,324,912	9,952,241	13.3%	15%

The results indicate that, for an international active bank, the BLs overall capital charge amounts to 13.3 per cent of the GI, a figure slightly lower than the value of 15 per cent envisaged in the Basic Approach. On the other hand, there is a more marked variety of ratios across the BLs. In particular *Retail Banking* shows the lowest ratio (below 10 per cent) and *Payments & Settlements* the highest (above 30 per cent), *Trading & Sales*, *Commercial Banking*, *Asset Management* and *Corporate Finance* are in a medium-low area (around 15 per cent) while *Retail Brokerage* and *Agency Services* are in a medium-high region (about 20 per cent).

12. Conclusions

The aim of this paper was to illustrate the methodology and the outcomes of the inferential analysis carried out on the operational risk data collected by the RMG through 2002, pooled by BLs. The exercise aimed, first of all, to compare the sensitivity of

conventional actuarial distributions and models stemming from the Extreme Value Theory (EVT) in representing the extreme percentiles of the data sets (i.e. the large losses). Then, measures of severity and frequency of the large losses in each data set were gained and, by a proper combination of these estimates, a bottom-up operational risk capital charge was computed. Finally, for each BL and in the eight BLs as a whole, the contributions of the expected losses to the capital figures were evaluated and the relationships between the capital charges and the corresponding average level of Gross Incomes were determined.

From a statistical point of view, the results indicate a low performance of conventional severity models in describing the overall data characteristics, summarizable in very high levels of both skewness to the right and kurtosis. In fact any traditional distribution applied to all the data in each BL tends to fit the central observations, hence failing to take the extreme percentiles into adequate consideration. On the other hand, the exercise shows that the Extreme Value model, in its Peaks Over Thresholds representation, explains the behaviour of the operational risk data in the tail area well.

Moreover, in the conventional actuarial approach, the distribution of the aggregated losses is usually obtained by treating the estimate of its severity and frequency components as a separate, disjointed, problem and, afterwards, aggregating the corresponding outcomes using not analytical methods. One of the main remarks coming out of this paper is that, if the aim of the analysis is to estimate the extreme percentiles of the aggregated losses, the treatment of these two components within a single overall estimation problem may reduce the estimate error and the computational costs ⁴⁷. The Peaks Over Thresholds approach appears to be a suitable and consistent statistical tool to tackle this issue, since it takes into account the (unknown) relationship between the frequency and the severity of large losses up to the end of the distribution and hence makes it possible to employ a semiparametric approach to compute any high percentile of the aggregated losses.

As the exercise makes clear, the EVT analysis requires that specific conditions be fulfilled in order to be worked out, the most important of which are the i.i.d. assumptions for the data and, as concerns the GPD model, a satisfactory stability of the inference to an

⁴⁷ The extent of the estimate error which derives from employing approximation or simulation methods, as the MonteCarlo procedure, for the detection of high percentiles of the operational risk aggregated losses will be matter of a next paper.

increase of the pre-set (high) threshold. Actually, this seems to be the case for each BL data set that originates from pooling the operational risk losses collected through the 2002 LDCE. In any case, it should borne in mind that, even if suitable tools are available in the Extreme Value literature for handling data with specific characteristics (such as trend, seasonality and clustering), a high sensitivity of the model remains for the largest observed losses and the very extreme quantile estimates. As Embrechts states (see Embrechts *et al.*, 1997): “*The statistical reliability of these estimates becomes very difficult to judge in general. Though we can work out approximate confidence intervals for these estimators, such constructions strongly rely on mathematical assumptions which are unverifiable in practice*”.

Concerning the outcomes of the analysis, there is clear evidence of the considerable magnitude of operational risk in the businesses carried out by the 2002 LDCE banks as well as of the differences in the riskiness of the BLs (in terms of both the time-unconditional severity and the 1-year aggregated capital figure). These differences persist after comparing, for a typical international active bank, the BLs capital figures with the average Gross Incomes and obtaining ratios as the coefficients set in the revised framework of the Capital Accord. In practice, the bottom-up analysis of the 2002 LDCE data suggests that the actual operational riskiness of the BLs may be captured in a more effectively way by setting, for the regulatory coefficients of the Standardised Approach, a wider range than the current one; besides, for the eight BLs as a whole, the implied capital ratio results to be a slightly lower figure than the coefficient envisaged in the Basic Approach.

Anyway, as frequently noted along this paper, the soundness of the exercise relies on the (unknown) actual quality of the data provided by the banks participating in the 2002 LDCE. In light of that, it would be extremely valuable that, in the course of the implementation of the new Capital Accord and when more data will be available to individual banks, consortia, academic circles, etc., the people involved in the quantification of operational risk make use of the statistical model implemented in this exercise to get pertinent figures of BLs capital charges and implied coefficients and hence test the robustness of the model itself and the consistency of its outcomes.

References

- Artzner, P., Delbaen, F., Eber, J.M. and Heath, D., (1999), "Coherent measures of risk", *Mathematical Finance*, 9, pp. 203-28.
- Basel Committee on Banking Supervision, (2003), *The 2002 Loss Data Collection Exercise for Operational Risk: Summary of the Data Collected*, Basel, BIS.
- Basel Committee on Banking Supervision, (2003), *Overview of the New Basel Capital Accord*, Consultative document, Basel, BIS.
- Basel Committee on Banking Supervision, (2004), *International Convergence of Capital Measurement and Capital Standards*, Basel, BIS.
- Baud, N., Frachot, A. and Roncalli, T., (2002), "Internal data, external data and consortium data for operational risk measurement: How to pool data properly?", Groupe de Recherche Opérationnelle, Credit Lyonnais, Working Paper.
- Balkema, A.A. and de Haan, L., (1974), "Residual life time at great age", *Annual Probability*, 2, pp. 792-804.
- Chavez-Demoulin, V. and Embrechts, P., (2003), "Smooth extremal models in finance and insurance", to appear in *Journal of Risk and Insurance*.
- Danielsson, J. and de Vries, C.G., (1997), "Tail index and quantile estimation with very high frequency data", *Journal of Empirical Finance*, 4, pp. 241-57.
- Danielsson, J., de Haan, L., Pend, L. and de Vries, C.G., (2000), "Using a bootstrap method to choose the sample fraction in tail index estimation", *Journal of Multivariate Analysis*, 76, pp. 226-48.
- Davison, A.C., (1984), "Modelling excesses over high thresholds, with an Application", in J. de Oliveira (ed.), *Statistical extremes and applications*, D. Reidel, pp. 461-82.
- Davison, A.C. and Smith, R.L., (1990), "Models for exceedances over high thresholds (with discussion)", *Journal of the Royal Statistical Society*, B 52, pp. 393-442.
- Diebold, F., Schuermann, T. and Stroughair, J., (1998), "Pitfalls and Opportunities in the Use of Extreme Value Theory in Risk Management", *Advances in Computational Finance*, Boston, Kluwer Academic Press.
- de Fontnouvelle, P., Deleus-Rueff, V., Jordan, J. and Rosengren, E., (2003), "Using loss data to quantify operational risk", Federal Reserve Bank of Boston, Working Paper.
- de Fontnouvelle, P., Rosengren, E. and Jordan, J., (2004), "Implications of alternative operational risk modelling techniques", Federal Reserve Bank of Boston, Working Paper.
- Ebnother, S., Vanini, P., McNeil, A.J. and Antolinez-Fehr, P., (2001), "Modelling Operational Risk", Zurich Cantonal Bank and Swiss Federal Institute of Technology Zurich (Department of Mathematics), Working Paper.

- Efron, B. and Tibshirani, R.J., (1993), *An Introduction to the Bootstrap*, London, Chapman & Hall.
- Embrechts, P., (1999), "Extreme value theory in finance and insurance", Manuscript, Department of Mathematics, ETH, Swiss Federal Technical University.
- Embrechts, P., (2000), "Extreme value theory: Potentials and limitations as an integrated risk management tool", *Derivatives Use, Trading & Regulation*, 6, pp. 449-56.
- Embrechts, P. (ed.), (2000), *Extremes and Integrated Risk Management*, London, Risk Books and UBS Warburg.
- Embrechts, P., Furrer, H. and Kaufmann, R., (2003), "Quantifying regulatory capital for operational risk", *Derivatives Use, Trading & Regulation*, 9, pp. 217-33.
- Embrechts, P., Kaufmann, R. and Samorodnitsky, G., (2004), "Ruin theory revisited: stochastic models for operational risk", to appear in *ECB, Volume on Foreign Reserves Risk Management*.
- Embrechts, P., Kluppelberg, C. and Mikosch, C., (1997), *Modelling Extremal Events for Insurance and Finance*, New York, Springer.
- Embrechts, P., Resnick, S. and Samorodnitsky, G., (1999), "Extreme value theory as a risk management tool", *North American Actuarial Journal*, 3, pp. 30-41.
- Embrechts, P. and Samorodnitsky, G., (2003), "Ruin problem and how fast stochastic processes mix", *The Annals of Applied Probability*, 13, pp. 1-36.
- Ferroy, C. and Segers, J., (2003), "Inference for clusters of extreme values", *Journal of the Royal Statistical Society*, B 65, pp. 545-56.
- Fisher, R.A. and Tippett, L.H.C., (1928), "Limiting forms of the frequency distribution of the largest or smallest member of a sample", *Proceeding of the Cambridge Philosophical Society*, 24, pp. 180-90.
- Frachot, A. and Roncalli, T., (2002), "Mixing internal and external data for managing operational risk", Groupe de Recherche Opérationnelle, Credit Lyonnais, Working Paper.
- Gencay, R. and Selcuk, F., (2004), "Extreme value theory and Value-at-Risk: Relative performance in emerging markets", *International Journal of Forecasting*, 20, pp. 287-303.
- Hauksson, H.A., Dacorogna, M., Domenig, T., Muller, U. and Samorodnitsky, G., (2000), "Multivariate extremes, aggregation and risk estimation", Manuscript, School of Operations Research and Industrial Engineering, Cornell University.
- Hosking, J.R.M. and Wallis, J.R., (1987), "Parameter and quantile estimation for the generalized pareto distribution", *Technometrics*, 29, pp. 339-49.
- King, J.L., (2001), *Measurement and modelling operational risk*, Wiley Finance.
- Lawrence, D., (2003), "Tackling the problem of fitting empirical data to estimate risk capital at high confidence levels", Presentation to the 10th ICBI Annual Meeting, Geneva.

- Leadbetter, M.R., Lidgren, G. and Rootzen, H., (1983), *Extremes and Related Properties of Random Sequences and Processes*, New York, Springer.
- McNeil, A.J., (1997), “Estimating the tails of loss severity distributions using extreme value theory”, *ASTIN Bulletin*, 27, pp. 1117-37.
- McNeil, A.J., (1998), “Calculating quantile risk measures for financial time series using extreme value theory”, ETH Zurich, Working Paper.
- McNeil, A.J., (1999), “Extreme value theory for risk managers”, *Internal Modeling and CAD II, RISK Books*, 93-113.
- McNeil, A.J. and Frey, R., (2000), “Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach”, *Journal of Empirical Finance*, 7, pp. 271-300.
- McNeil, A.J. and Saladin, T., (1997), “The peaks over thresholds method for estimating high quantiles of loss distributions”, *Proceedings of XXVIIth International ASTIN Colloquium, Cairns Australia*, pp. 23-43.
- Medova, E., (2000), “Measuring risk by extreme values”, *RISK*, pp. 20-6.
- Nyström, K. and Skoglund, J., (2002), “Univariate Extreme Value Theory, GARCH and Measures of Risk”, Swedbank, Group of Financial Risk Control, Working Paper.
- Pickands, J., (1975), “Statistical inference using extreme order statistics”, *Annals of Statistics*, 3, pp. 119-31.
- Reiss, R. and Thomas, M., (2001), *Statistical Analysis of Extreme Values*, 2^a edition, Basel, Birkhauser.
- Resnick, S.I., (1987), *Extreme Values, Regular Variation and Point Process*, New York, Springer.
- Rootzen, H., Leadbetter, M.R. and de Haan, L., (1992), “Tail and quantile estimators for stationary processes”, Report, Department of Statistics, University of North Carolina.
- Rootzen, H. and Tajvidi, N., (1997), “Extreme value statistics and wind storm losses: a case study”, *Scandinavian Actuarial Journal*, pp. 70-94.
- Samad-Khan, A., (2003), “An integrated approach to monitoring operational risk”, Presentation to the 10th ICBI Annual Meeting, Geneva.
- Smith, R.L., (1987), “Estimating tails of probability distributions”, *Annals of Statistics*, 15, pp. 1174-207.
- Smith, R.L., (1989), “Extreme Value Analysis of environmental time series: an application to trend detection in ground-level ozone”, *Statistical Science*, 4, pp. 367–93.
- Smith, R.L. and Goodman, D., (1996), “Bayesian risk analysis”, in Embrechts, P. (ed.), *Extremes and Integrated Risk Management*, London, Risk Books and UBS Warburg.
- Smith, R.L. and Shively, T., (1995), “Point process approach to modelling trends in tropospheric ozone based on exceedances of a high threshold”, *Atmospheric Environment*, 29, pp. 3489-99.

RECENTLY PUBLISHED “TEMI” (*)

- N. 489 – *Information technology and productivity changes in the banking industry*, by L. CASOLARO and G. GOBBI (March 2004).
- N. 490 – *La bilancia dei pagamenti di parte corrente Nord-Sud (1998-2000)*, by L. CANNARI and S. CHIRI (March 2004).
- N. 491 – *Investire in Italia? Risultati di una recente indagine empirica*, by M. COMMITTERI (March 2004).
- N. 492 – *Centralization of wage bargaining and the unemployment rate: revisiting the hump-shape hypothesis*, by L. FORNI (June 2004).
- N. 493 – *Endogenous monetary policy with unobserved potential output*, by A. CUKIERMAN and F. LIPPI (June 2004).
- N. 494 – *Il credito commerciale: problemi e teorie*, by M. OMICCIOLI (June 2004).
- N. 495 – *Condizioni di credito commerciale e differenziazione della clientela*, by L. CANNARI, S. CHIRI and M. OMICCIOLI (June 2004).
- N. 496 – *Il debito commerciale in Italia: quanto contano le motivazioni finanziarie?*, by P. FINALDI RUSSO and L. LEVA (June 2004).
- N. 497 – *Funzionamento della giustizia civile e struttura finanziaria delle imprese: il ruolo del credito commerciale*, by A. CARMIGNANI (June 2004).
- N. 498 – *Does trade credit substitute for bank credit?*, by G. DE BLASIO (June 2004).
- N. 499 – *Monetary policy and the transition to rational expectations*, by G. FERRERO (June 2004).
- N. 500 – *Turning-point indicators from business surveys: real-time detection for the euro area and its major member countries*, by A. BAFFIGI and A. BASSANETTI (June 2004).
- N. 501 – *La ricchezza delle famiglie italiane e americane*, by I. FAIELLA and A. NERI (June 2004).
- N. 502 – *Optimal duplication of effort in advocacy systems*, by G. PALUMBO (June 2004).
- N. 503 – *Il pilastro privato del sistema previdenziale. Il caso del Regno Unito*, by F. SPADAFORA (June 2004).
- N. 504 – *Firm size distribution and employment protection legislation in Italy*, by F. SCHIVARDI and R. TORRINI (June 2004).
- N. 505 – *Social mobility and endogenous cycles in redistribution*, by F. ZOLLINO (July 2004).
- N. 506 – *Estimating expectations of shocks using option prices*, by A. DI CESARE (July 2004).
- N. 507 – *Estimating state price densities by Hermite polynomials: theory and application to the Italian derivatives market*, by P. GUASONI (July 2004).
- N. 508 – *The interaction between face-to-face and electronic delivery: the case of the Italian banking industry*, by E. BONACCORSI DI PATTI, G. GOBBI and P. E. MISTRULLI (July 2004).
- N. 509 – *Bad loans and entry into local credit markets*, by M. BOFONDI and G. GOBBI (July 2004).
- N. 510 – *Does wealth affect consumption? Evidence for Italy*, by M. PAIELLA (July 2004).
- N. 511 – *Information variables for monetary policy in a small structural model of the euro area*, by F. LIPPI and S. NERI (July 2004).
- N. 512 – *Monetary union with voluntary participation*, by W. FUCHS and F. LIPPI (July 2004).
- N. 513 – *Monetary policy and stock prices: theory and evidence*, by S. NERI (July 2004).
- N. 514 – *Financial structure and the transmission of monetary shocks: preliminary evidence for the Czech Republic, Hungary and Poland*, by A. ANZUINI and A. LEVY (July 2004).
- N. 515 – *The pricing behaviour of Italian firms: new survey evidence on price stickiness*, by S. FABIANI, A. GATTULLI and R. SABBATINI (July 2004).
- N. 516 – *Business cycle non-linearities and productivity shocks*, by P. PISELLI (July 2004).

(*) Requests for copies should be sent to:

Banca d'Italia – Servizio Studi – Divisione Biblioteca e pubblicazioni – Via Nazionale, 91 – 00184 Rome (fax 0039 06 47922059). They are available on the Internet www.bancaditalia.it.

"TEMI" LATER PUBLISHED ELSEWHERE

1999

- L. GUISO and G. PARIGI, *Investment and demand uncertainty*, Quarterly Journal of Economics, Vol. 114 (1), pp. 185-228, **TD No. 289 (November 1996)**.
- A. F. POZZOLO, *Gli effetti della liberalizzazione valutaria sulle transazioni finanziarie dell'Italia con l'estero*, Rivista di Politica Economica, Vol. 89 (3), pp. 45-76, **TD No. 296 (February 1997)**.
- A. CUKIERMAN and F. LIPPI, *Central bank independence, centralization of wage bargaining, inflation and unemployment: theory and evidence*, European Economic Review, Vol. 43 (7), pp. 1395-1434, **TD No. 332 (April 1998)**.
- P. CASELLI and R. RINALDI, *La politica fiscale nei paesi dell'Unione europea negli anni novanta*, Studi e note di economia, (1), pp. 71-109, **TD No. 334 (July 1998)**.
- A. BRANDOLINI, *The distribution of personal income in post-war Italy: Source description, data quality, and the time pattern of income inequality*, Giornale degli economisti e Annali di economia, Vol. 58 (2), pp. 183-239, **TD No. 350 (April 1999)**.
- L. GUISO, A. K. KASHYAP, F. PANETTA and D. TERLIZZESE, *Will a common European monetary policy have asymmetric effects?*, Economic Perspectives, Federal Reserve Bank of Chicago, Vol. 23 (4), pp. 56-75, **TD No. 384 (October 2000)**.

2000

- P. ANGELINI, *Are banks risk-averse? Timing of the operations in the interbank market*, Journal of Money, Credit and Banking, Vol. 32 (1), pp. 54-73, **TD No. 266 (April 1996)**.
- F. DRUDI and R. GIORDANO, *Default Risk and optimal debt management*, Journal of Banking and Finance, Vol. 24 (6), pp. 861-892, **TD No. 278 (September 1996)**.
- F. DRUDI and R. GIORDANO, *Wage indexation, employment and inflation*, Scandinavian Journal of Economics, Vol. 102 (4), pp. 645-668, **TD No. 292 (December 1996)**.
- F. DRUDI and A. PRATI, *Signaling fiscal regime sustainability*, European Economic Review, Vol. 44 (10), pp. 1897-1930, **TD No. 335 (September 1998)**.
- F. FORNARI and R. VIOLI, *The probability density function of interest rates implied in the price of options*, in: R. Violi, (ed.) , *Mercati dei derivati, controllo monetario e stabilità finanziaria*, Il Mulino, Bologna, **TD No. 339 (October 1998)**.
- D. J. MARCHETTI and G. PARIGI, *Energy consumption, survey data and the prediction of industrial production in Italy*, Journal of Forecasting, Vol. 19 (5), pp. 419-440, **TD No. 342 (December 1998)**.
- A. BAFFIGI, M. PAGNINI and F. QUINTILIANI, *Localismo bancario e distretti industriali: assetto dei mercati del credito e finanziamento degli investimenti*, in: L.F. Signorini (ed.), *Lo sviluppo locale: un'indagine della Banca d'Italia sui distretti industriali*, Donzelli, **TD No. 347 (March 1999)**.
- A. SCALIA and V. VACCA, *Does market transparency matter? A case study*, in: *Market Liquidity: Research Findings and Selected Policy Implications*, Basel, Bank for International Settlements, **TD No. 359 (October 1999)**.
- F. SCHIVARDI, *Rigidità nel mercato del lavoro, disoccupazione e crescita*, Giornale degli economisti e Annali di economia, Vol. 59 (1), pp. 117-143, **TD No. 364 (December 1999)**.
- G. BODO, R. GOLINELLI and G. PARIGI, *Forecasting industrial production in the euro area*, Empirical Economics, Vol. 25 (4), pp. 541-561, **TD No. 370 (March 2000)**.
- F. ALTISSIMO, D. J. MARCHETTI and G. P. ONETO, *The Italian business cycle: Coincident and leading indicators and some stylized facts*, Giornale degli economisti e Annali di economia, Vol. 60 (2), pp. 147-220, **TD No. 377 (October 2000)**.
- C. MICHELACCI and P. ZAFFARONI, *(Fractional) Beta convergence*, Journal of Monetary Economics, Vol. 45, pp. 129-153, **TD No. 383 (October 2000)**.

- R. DE BONIS and A. FERRANDO, *The Italian banking structure in the nineties: testing the multimarket contact hypothesis*, Economic Notes, Vol. 29 (2), pp. 215-241, **TD No. 387 (October 2000)**.
- 2001
- M. CARUSO, *Stock prices and money velocity: A multi-country analysis*, Empirical Economics, Vol. 26 (4), pp. 651-72, **TD No. 264 (February 1996)**.
- P. CIPOLLONE and D. J. MARCHETTI, *Bottlenecks and limits to growth: A multisectoral analysis of Italian industry*, Journal of Policy Modeling, Vol. 23 (6), pp. 601-620, **TD No. 314 (August 1997)**.
- P. CASELLI, *Fiscal consolidations under fixed exchange rates*, European Economic Review, Vol. 45 (3), pp. 425-450, **TD No. 336 (October 1998)**.
- F. ALTISSIMO and G. L. VIOLANTE, *Nonlinear VAR: Some theory and an application to US GNP and unemployment*, Journal of Applied Econometrics, Vol. 16 (4), pp. 461-486, **TD No. 338 (October 1998)**.
- F. NUCCI and A. F. POZZOLO, *Investment and the exchange rate*, European Economic Review, Vol. 45 (2), pp. 259-283, **TD No. 344 (December 1998)**.
- L. GAMBACORTA, *On the institutional design of the European monetary union: Conservatism, stability pact and economic shocks*, Economic Notes, Vol. 30 (1), pp. 109-143, **TD No. 356 (June 1999)**.
- P. FINALDI RUSSO and P. ROSSI, *Credit constraints in Italian industrial districts*, Applied Economics, Vol. 33 (11), pp. 1469-1477, **TD No. 360 (December 1999)**.
- A. CUKIERMAN and F. LIPPI, *Labor markets and monetary union: A strategic analysis*, Economic Journal, Vol. 111 (473), pp. 541-565, **TD No. 365 (February 2000)**.
- G. PARIGI and S. SIVIERO, *An investment-function-based measure of capacity utilisation, potential output and utilised capacity in the Bank of Italy's quarterly model*, Economic Modelling, Vol. 18 (4), pp. 525-550, **TD No. 367 (February 2000)**.
- F. BALASSONE and D. MONACELLI, *Emu fiscal rules: Is there a gap?*, in: M. Bordignon and D. Da Empoli (eds.), *Politica fiscale, flessibilità dei mercati e crescita*, Milano, Franco Angeli, **TD No. 375 (July 2000)**.
- A. B. ATKINSON and A. BRANDOLINI, *Promise and pitfalls in the use of "secondary" data-sets: Income inequality in OECD countries*, Journal of Economic Literature, Vol. 39 (3), pp. 771-799, **TD No. 379 (October 2000)**.
- D. FOCARELLI and A. F. POZZOLO, *The determinants of cross-border bank shareholdings: An analysis with bank-level data from OECD countries*, Journal of Banking and Finance, Vol. 25 (12), pp. 2305-2337, **TD No. 381 (October 2000)**.
- M. SBRACIA and A. ZAGHINI, *Expectations and information in second generation currency crises models*, Economic Modelling, Vol. 18 (2), pp. 203-222, **TD No. 391 (December 2000)**.
- F. FORNARI and A. MELE, *Recovering the probability density function of asset prices using GARCH as diffusion approximations*, Journal of Empirical Finance, Vol. 8 (1), pp. 83-110, **TD No. 396 (February 2001)**.
- P. CIPOLLONE, *La convergenza dei salari manifatturieri in Europa*, Politica economica, Vol. 17 (1), pp. 97-125, **TD No. 398 (February 2001)**.
- E. BONACCORSI DI PATTI and G. GOBBI, *The changing structure of local credit markets: Are small businesses special?*, Journal of Banking and Finance, Vol. 25 (12), pp. 2209-2237, **TD No. 404 (June 2001)**.
- G. MESSINA, *Decentramento fiscale e perequazione regionale. Efficienza e redistribuzione nel nuovo sistema di finanziamento delle regioni a statuto ordinario*, Studi economici, Vol. 56 (73), pp. 131-148, **TD No. 416 (August 2001)**.

2002

- R. CESARI and F. PANETTA, *Style, fees and performance of Italian equity funds*, Journal of Banking and Finance, Vol. 26 (1), **TD No. 325 (January 1998)**.
- L. GAMBACORTA, *Asymmetric bank lending channels and ECB monetary policy*, Economic Modelling, Vol. 20 (1), pp. 25-46, **TD No. 340 (October 1998)**.
- C. GIANNINI, *"Enemy of none but a common friend of all"? An international perspective on the lender-of-last-resort function*, Essay in International Finance, Vol. 214, Princeton, N. J., Princeton University Press, **TD No. 341 (December 1998)**.
- A. ZAGHINI, *Fiscal adjustments and economic performing: A comparative study*, Applied Economics, Vol. 33 (5), pp. 613-624, **TD No. 355 (June 1999)**.
- F. ALTISSIMO, S. SIVIERO and D. TERLIZZESE, *How deep are the deep parameters?*, Annales d'Economie et de Statistique, (67/68), pp. 207-226, **TD No. 354 (June 1999)**.
- F. FORNARI, C. MONTICELLI, M. PERICOLI and M. TIVEGNA, *The impact of news on the exchange rate of the lira and long-term interest rates*, Economic Modelling, Vol. 19 (4), pp. 611-639, **TD No. 358 (October 1999)**.
- D. FOCARELLI, F. PANETTA and C. SALLO, *Why do banks merge?*, Journal of Money, Credit and Banking, Vol. 34 (4), pp. 1047-1066, **TD No. 361 (December 1999)**.
- D. J. MARCHETTI, *Markup and the business cycle: Evidence from Italian manufacturing branches*, Open Economies Review, Vol. 13 (1), pp. 87-103, **TD No. 362 (December 1999)**.
- F. Busetti, *Testing for stochastic trends in series with structural breaks*, Journal of Forecasting, Vol. 21 (2), pp. 81-105, **TD No. 385 (October 2000)**.
- F. LIPPI, *Revisiting the Case for a Populist Central Banker*, European Economic Review, Vol. 46 (3), pp. 601-612, **TD No. 386 (October 2000)**.
- F. PANETTA, *The stability of the relation between the stock market and macroeconomic forces*, Economic Notes, Vol. 31 (3), **TD No. 393 (February 2001)**.
- G. GRANDE and L. VENTURA, *Labor income and risky assets under market incompleteness: Evidence from Italian data*, Journal of Banking and Finance, Vol. 26 (2-3), pp. 597-620, **TD No. 399 (March 2001)**.
- A. BRANDOLINI, P. CIPOLLONE and P. SESTITO, *Earnings dispersion, low pay and household poverty in Italy, 1977-1998*, in D. Cohen, T. Piketty and G. Saint-Paul (eds.), *The Economics of Rising Inequalities*, pp. 225-264, Oxford, Oxford University Press, **TD No. 427 (November 2001)**.
- L. CANNARI and G. D'ALESSIO, *La distribuzione del reddito e della ricchezza nelle regioni italiane*, Rivista Economica del Mezzogiorno (Trimestrale della SVIMEZ), Vol. XVI(4), pp. 809-847, Il Mulino, **TD No. 482 (June 2003)**.

2003

- F. SCHIVARDI, *Reallocation and learning over the business cycle*, European Economic Review, , Vol. 47 (1), pp. 95-111, **TD No. 345 (December 1998)**.
- P. CASELLI, P. PAGANO and F. SCHIVARDI, *Uncertainty and slowdown of capital accumulation in Europe*, Applied Economics, Vol. 35 (1), pp. 79-89, **TD No. 372 (March 2000)**.
- P. PAGANO and G. FERRAGUTO, *Endogenous growth with intertemporally dependent preferences*, Contribution to Macroeconomics, Vol. 3 (1), pp. 1-38, **TD No. 382 (October 2000)**.
- P. PAGANO and F. SCHIVARDI, *Firm size distribution and growth*, Scandinavian Journal of Economics, Vol. 105(2), pp. 255-274, **TD No. 394 (February 2001)**.
- M. PERICOLI and M. SBRACIA, *A Primer on Financial Contagion*, Journal of Economic Surveys, **TD No. 407 (June 2001)**.
- M. SBRACIA and A. ZAGHINI, *The role of the banking system in the international transmission of shocks*, World Economy, **TD No. 409 (June 2001)**.

- E. GAIOTTI and A. GENERALE, *Does monetary policy have asymmetric effects? A look at the investment decisions of Italian firms*, *Giornale degli Economisti e Annali di Economia*, Vol. 61 (1), pp. 29-59, **TD No. 429 (December 2001)**.
- L. GAMBACORTA, *The Italian banking system and monetary policy transmission: evidence from bank level data*, in: I. Angeloni, A. Kashyap and B. Mojon (eds.), *Monetary Policy Transmission in the Euro Area*, Cambridge, Cambridge University Press, **TD No. 430 (December 2001)**.
- M. EHRMANN, L. GAMBACORTA, J. MARTÍNEZ PAGÉS, P. SEVESTRE and A. WORMS, *Financial systems and the role of banks in monetary policy transmission in the euro area*, in: I. Angeloni, A. Kashyap and B. Mojon (eds.), *Monetary Policy Transmission in the Euro Area*, Cambridge, Cambridge University Press, **TD No. 432 (December 2001)**.
- F. SPADAFORA, *Financial crises, moral hazard and the speciality of the international market: further evidence from the pricing of syndicated bank loans to emerging markets*, *Emerging Markets Review*, Vol. 4 (2), pp. 167-198, **TD No. 438 (March 2002)**.
- D. FOCARELLI and F. PANETTA, *Are mergers beneficial to consumers? Evidence from the market for bank deposits*, *American Economic Review*, Vol. 93 (4), pp. 1152-1172, **TD No. 448 (July 2002)**.
- E. VIVIANO, *Un'analisi critica delle definizioni di disoccupazione e partecipazione in Italia*, *Politica Economica*, Vol. 19 (1), pp. 161-190, **TD No. 450 (July 2002)**.
- F. BUSETTI and A. M. ROBERT TAYLOR, *Testing against stochastic trend and seasonality in the presence of unattended breaks and unit roots*, *Journal of Econometrics*, Vol. 117 (1), pp. 21-53, **TD No. 470 (February 2003)**.

2004

- P. CHIADES and L. GAMBACORTA, *The Bernanke and Blinder model in an open economy: The Italian case*, *German Economic Review*, Vol. 5 (1), pp. 1-34, **TD No. 388 (December 2000)**.
- M. PAIELLA, *Heterogeneity in financial market participation: appraising its implications for the C-CAPM*, *Review of Finance*, Vol. 8, pp. 1-36, **TD No. 473 (June 2003)**.
- E. BONACCORSI DI PATTI and G. DELL'ARICCIA, *Bank competition and firm creation*, *Journal of Money Credit and Banking*, Vol. 36 (2), pp. 225-251, **TD No. 481 (June 2003)**.
- R. GOLINELLI and G. PARIGI, *Consumer sentiment and economic activity: a cross country comparison*, *Journal of Business Cycle Measurement and Analysis*, Vol. 1 (2), **TD No. 484 (September 2003)**.

FORTHCOMING

- A. F. POZZOLO, *Research and development regional spillovers, and the localisation of economic activities*, *The Manchester School*, **TD No. 331 (March 1998)**.
- F. LIPPI, *Strategic monetary policy with non-atomistic wage-setters*, *Review of Economic Studies*, **TD No. 374 (June 2000)**.
- P. ANGELINI and N. CETORELLI, *Bank competition and regulatory reform: The case of the Italian banking industry*, *Journal of Money, Credit and Banking*, **TD No. 380 (October 2000)**.
- L. DEDOLA AND F. LIPPI, *The Monetary Transmission Mechanism: Evidence from the industry Data of Five OECD Countries*, *European Economic Review*, **TD No. 389 (December 2000)**.
- M. BUGAMELLI and P. PAGANO, *Barriers to Investment in ICT*, *Applied Economics*, **TD No. 420 (October 2001)**.
- D. J. MARCHETTI and F. NUCCI, *Price Stickiness and the Contractionary Effects of Technology Shocks*, *European Economic Review*, **TD No. 392 (February 2001)**.
- G. CORSETTI, M. PERICOLI and M. SBRACIA, *Correlation analysis of financial contagion: what one should know before running a test*, *Journal of International Money and Finance*, **TD No. 408 (June 2001)**.
- D. FOCARELLI, *Bootstrap bias-correction procedure in estimating long-run relationships from dynamic*

panels, with an application to money demand in the euro area, Economic Modelling, **TD No. 440 (March 2002)**.

- A. BAFFIGI, R. GOLINELLI and G. PARIGI, *Bridge models to forecast the euro area GDP*, International Journal of Forecasting, **TD No. 456 (December 2002)**.
- F. CINGANO and F. SCHIVARDI, *Identifying the sources of local productivity growth*, Journal of the European Economic Association, **TD No. 474 (June 2003)**.
- E. BARUCCI, C. IMPENNA and R. RENÒ, *Monetary integration, markets and regulation*, Research in Banking and Finance, **TD No. 475 (June 2003)**.
- G. ARDIZZI, *Cost efficiency in the retail payment networks: first evidence from the Italian credit card system*, Rivista di Politica Economica, **TD No. 480 (June 2003)**.
- L. GAMBACORTA and P. E. MISTRULLI, *Does bank capital affect lending behavior?*, Journal of Financial Intermediation, **TD No. 486 (September 2003)**.