ISSN: 2281-1346

Department of Economics and Management

DEM Working Paper Series

Bayesian operational risk models

Silvia Figini
(Department of Political and Social Sciences, University of Pavia)

Lijun Gao
(Management School, Shandong University of Finance, Jinan, China)

Paolo Giudici
(Department of Ecomomics and Management, University of Pavia)

 # 47 (07-13)

Via San Felice, 5
I-27100 Pavia
http://epmq.unipv.eu/site/home.html

July 2013

Bayesian operational risk models

Silvia Figini*, Lijun Gao† and Paolo Giudici‡

July 12, 2013

Abstract: Operational risk is hard to quantify, for the presence of heavy tailed loss distributions. Extreme value distributions, used in this context, are very sensitive to the data, and this is a problem in the presence of rare loss data. Self risk assessment questionnaires, if properly modelled, may provide the missing piece of information that is necessary to adequately estimate operational risks. In this paper we propose to embody self risk assessment data into suitable prior distributions, and to follow a Bayesian approach to merge self assessment with loss data. We derive operational loss posterior distributions, from which appropriate measures of risk, such as the Value at Risk, or the Expected Shortfall, can be derived. We test our proposed models on a real database, made up of internal loss data and self risk assessment questionnaires of an anonymous commercial bank. Our results show that the proposed Bayesian models performs better with respect to classical extreme value models, leading to a smaller quantification of the Value at Risk required to cover unexpected losses.

Keywords: Extreme value distributions, Operational risk management, Self-assessment questionnaires, Value at Risk

1 Introduction

In this paper we improve the state of the art on statistical models for operational risk measurement. Operational risk is the risk that a company occurs into financial losses, resulting from inadequate or failed internal processes, people and systems, or from external events (see e.g. Cruz, 2002 or Panjer, 2006). Operational risks of financial institutions must be measured according to international standards, described in the so-called Basel accords (see www.bis.org). In particular, here we focus our attention on the advanced measurement approach (AMA). This approach gives greater flexibility, in comparison with simpler, more standardised approaches, as it takes into account the particular

*Department of Political and Social Sciences, University of Pavia, Italy
†Management School, Shandong University of Finance, Jinan, China
‡Department of Ecomomics and Management, University of Pavia, giudici@unipv.it

characteristics of financial institutions (see, for example, Cornalba and Giudici, 2004).
Operational risks are hard to quantify, for the presence of heavy tailed loss distributions. Extreme value distributions, used in this context, are very sensitive to the data, and this may be a problem when losses are rare. Self risk assessment questionnaires, if properly modelled, may provide the missing piece of information that is necessary to estimate operational risks.
In this paper we propose to embody self risk assessment data into a suitable parametric prior distribution, and follow a Bayesian approach to merge self assessment with loss data. To achieve this aim, we extend the Bayesian approach of Behrens et al. (2006) to operational risk modelling, and consider a convolution between two distributions: one representing the loss frequency, and the other the loss severity. Prior distributions will be elicited for the parameters of both distributions, under two different settings: an uninformative prior setting, where expert opinions are not available (as in Dalla Valle and Giudici, 2008), and an informative prior setting, where expert opinions are extracted from a risk self assessment process (described, for example, in Bilotta and Giudici, 2004 and Bonafede and Giudici, 2007).
By means of Bayes' theorem we can update each prior distribution with the observed loss data and determine the posterior distribution for the frequency and the severity distributions and, by convolution, the predictive loss distribution. As a closed form analytic solution is not possible, the predictive loss distribution can be approximated by means of Markov Chain Monte Carlo methods (see e.g. Gamerman, 1997). On the basis of the estimated predictive loss distribution we can thus compute summary risk measures, such as the Value at Risk (VaR) and the Expected Shortfall (ES) (see e.g. Artzner et al. 1999).
Our proposed model will be tested on a real database, made up of internal loss data and self risk assessment questionnaires of an anonymous commercial bank. We shall compare the results achieved with our Bayesian model with those obtained under a classical extreme value model, based on a Poisson distribution for the frequency and a generalized Pareto distribution for the severity (see e.g. Chavez-Demoulin et al., 2006).
The main outcome of such comparison is that the application of our Bayesian model reduces the Value at Risk and, therefore, the required capital coverage, an important financial advantage for the adopters of our proposal. A further advantage of what we propose is the possibility to combine observed loss data, that are necessarily "backward- looking", with self assessment data that, being human opinions, may include "forward-looking" aspects.
The paper is organized as follows: Section 2 introduces our proposed Bayesian model; Section 3 describes the real data at hand and underlines the empirical evidences obtained using both classical and Bayesian approaches. Section 4 ends with concluding remarks and further ideas of research.

2 Proposal

There are various reasons for preferring a Bayesian analysis of extremes over the more traditional likelihood approach. Since extreme data are (by their very nature) quite scarce, the ability to incorporate other sources of information through a prior distribution has an obvious appeal (see e.g. Figini and Giudici, 2011). Recently, some Bayesian extreme value models have been proposed in the financial risk management literature (see e.g. Behrens et al. 2006, Diebolt et al. 2005, Panjer 2006). Such models consider the case of developing risk measures for a loss distribution that is assumed to be a Generalized Pareto Distribution (GPD).

In this paper we extend these models to operational risk management, whose modelling is more complex than financial risk management (see e.g. Cruz 2002 or Alexander, 2003). To achieve this aim, we consider a convolution between a GPD distribution for the mean loss (severity), with a Poisson distribution for the number of loss events (frequency), the latter being unknown in the operational risk setting.

More precisely, our Bayesian Loss Distribution Approach (BLDA) starts by expressing each operational loss in terms of Frequency (the number of loss events during a certain time period) and Severity (the mean impact of the event, in terms of financial losses, in the same period). Such expression will hold for a collection of different loss events, classified in a matrix M of Business Lines (BL) and Event types (ET), as prescribed by the AMA approach in the Basel accord (see e.g. Cornalba and Giudici, 2004, or Dalla Valle and Giudici, 2008). Formally, for each business line/event type intersection i (where $i = 1, ..., I$) in the matrix M and for a given time period t, the total operational losses could be defined as the sum of a random number n_{it} (frequency) of losses:

$$L_{it} = X_{i1} + ... + X_{in_t}$$

where, for the business line/event type intersection i and for $t = 1, ..., T$, (T representing the number of time periods available), L_{it} denotes the total operational loss, $X_{i1}, ..., X_{in_t}$ denote individual loss severities and n_{it} denotes the unknown frequency.

Note that, for each intersection and for each time period, the total loss can be expressed as $L_{it} = s_{it} \times n_{it}$, where n_{it} is the frequency, defined as before, and s_{it} (commonly referred to as the severity) is the mean loss for that period.

Once operational loss variables are defined and classified as previously described, our BLDA model proceeds with some independence assumptions.

We assume that:

(1) within each intersection i, and each time period t, the distribution of the frequency n_{it} is independent of the distribution of the severity s_{it};

(2) for any given time period t, the I losses $L_{it} = s_{it} \times n_{it}$, for $i = 1, \ldots, I$, occurring in different intersections, are independent of each other;

(3) for any given intersection i the T losses losses $L_{it} = s_{it} \times n_{it}$, for $t = 1, \ldots, T$ occurring in different time periods, are independent of each other.

The above assumptions are common in the operational risk modelling literature (see e.g. Cruz, 2002). Some authors have considered a copula-based approach to remove one or more of the previous assumptions (see e.g. Fantazzini et al., 2008; Giudici and Politou, 2009), within a simplified modelling context. Our aim here is to evaluate whether a Bayesian model can improve classical models, within a complete extreme value modelling formulation and, therefore, will not consider such extensions.

For the sake of simplicity, in the rest of the paper, and unless otherwise specified, we shall consider a generic intersection event and drop the corresponding index i from the notation. This without loss of generality, following assumption (2). We can thus arrive at the core formulation of our approach.

For a given intersection, we assume a discrete probability density for the frequencies $(n_t, t = 1, \ldots, T)$ and a continuous probability density for the severities $(s_t, t = 1, \ldots, T)$. Following Dalla Valle and Giudici (2008), we can express the likelihood function for each intersection in a form that depends on some unknown parameters. Indicating the severity distributions with $f(s_t|\theta)$ and the frequency distributions with $f(n_t|\eta)$, where θ denotes the parameter vector of the severity distribution and η denotes the parameter vector of the frequency distribution, we have that, according to assumptions (1)-(3):

$$L(s, n|\theta, \eta) = \prod_{t=1}^{T} f(n_t|\theta) f(s_t|\eta), \tag{1}$$

where n and s indicate the data vectors $n = (n_t, t = 1, \ldots, T)$)and $s = (s_t, t = 1, \ldots, T)$.

Within the AMA approach in the Basel framework, each financial institution may choose to use different functional forms for the frequency and severity distributions for each ET/BL intersection. In order evaluate the relative performance of our BLDA model, here we consider a Poisson distribution for the frequency and a Generalised Pareto distribution for the severity, as in Chavez-Dumoulin et al. (2006). Our additional contribution is a Bayesian Poisson-GPD model that allows the combination of quantitative data, coming from the time series of operational losses collected by financial institutions, and categorical data, extracted from risk self-assessment questionnaires. representing expert opinions, as in a proper Bayesian analysis (see e.g. Bernardo and Smith, 1994). Differently from Dalla Valle and Giudici (2008) who also consider a Bayesian model, here our prior distributions are real expert opinions and not "convenient" uninformative priors, whose actual significance and interpretation is difficult.

We now describe how we specify such prior distribution, first of the frequency and then of the severity. As the loss frequency distribution is a Poisson with parameter η, we take a Gamma conjugate prior for the parameter η: $\eta = Gamma(\alpha, \beta)$, where the hyperparameters α and β will be set using expert opinions. Were such opinions not available, we could use an "uninformative" prior distribution, characterised by a very large variance (see e.g. Behrens et al. 2006, or Dalla Valle and Giudici, 2008).

In addition, it is a good idea to consider alternative prior specifications, so to evaluate how the results change and, possibly, select one of them, on the basis

of predictive performance.
According to Bayes Theorem, the posterior distribution of the parameter η can be obtained by multiplying the likelihood function with the prior distribution and normalizing such product. As the likelihood is a prodcut of Poisson distributions, conjugate to the Gamma prior, it can be easily shown that the frequency parameter posterior distribution is again a Gamma distribution:

$$\pi(\eta|n) \sim Gamma(\sum_{t=1}^{T} n_t + \alpha, T + \beta). \tag{2}$$

On the other hand, the distribution of the is a Generalised Pareto Distribution, characterised by the parameter vector $\theta = (\mu, \sigma, \xi)$, a vector with three parameters, respresenting the location, the scale and the shape of the distribution. In a Bayesian framework, we thus need to consider three different prior distributions, one for each parameter. Here we follow Coles and Tawn (1996), who suggest to elicit such priors with respect to parameters that can easily interpreted by the experts, such as the quantiles of the distribution. Let q indicate the $(1-p)\%$ quantile of the severity distribution. It can be shown that, by inverting the cumulative distribution function of the GPD, the expression of such quantile is the following:

$$q = u + \frac{\sigma}{\xi}\left\{\left[\frac{T}{T_u}(1-p)\right]^{-\xi} - 1\right\}, \tag{3}$$

where T_u indicate the number of periods in which the severity exceeds the threshold μ, among the available T.

According to Coles and Tawn (1996), prior specification for the GPD can be done in terms of a triple of quantiles, ordered as $(q_1 < q_2 < q_3)$, that correspond to three different (ordered) probability levels. Prior specification can be done in terms of the three chosen quantiles and this implies a corresponding prior specification for the three unknown parameters.

Having established a correspondence between a triple of quantiles $(q_c, \quad c = 1, 2, 3)$ and the parameter vector θ, we can specify a prior distribution directly over the former. We assume the three priors to be independent and let $q_c \sim Ga(a_c, b_c)$, for $c = 1, 2, 3$.
The joint prior density of the severity parameters can thus be obtained by a change of variables from:

$$\pi(\theta) \propto J \prod_{c=1}^{3} q_c^{a_c - 1} exp(-q_c/b_c)$$

where J is the Jacobian of the transformation from (q_1, q_2, q_3) to $\theta = (\mu, \sigma, \xi)$.
In order to derive the posterior distribution of the severity parameters, the above prior, expressed in terms of θ, should be compound with the likelihood, by means of Bayes' theorem.

The likelihood of the GPD model can be shown to be:

$$f(s|\theta) = \sigma^{-T} \prod_{t=1}^{T} (1 + \xi(s_t - \mu)/\sigma)^{-(1+1/\xi)}, \tag{4}$$

provided that $1 + \xi(s_t - \mu)/\sigma$ is positive for each $t = 1, ..., T$. Given the prior density $\pi(\theta)$ and the likelihood $f(s|\theta)$, we can obtain the posterior density $\pi(\theta|s)$ by means of Bayes' theorem. An exact, analytic, solution cannot be determined. However, Markov Chain Monte Carlo (MCMC) methods (see e.g. Gamerman, 1997 or Robert and Casella, 1999) can approximate the calculation producing stationary sequences of simulated values with marginal density $\pi(\theta|s)$ (see e.g. Bernardo and Smith, 1994).

In operational risk measurement, the aim is to obtain, on the basis of the available data, the distribution of the total loss, for the next period of time. This can be estimated obtaining the predictive distributions for the frequency and for the severity and taking their convolution.
Let y denote a frequency observation, in a future time period, with density function $f(y|\eta)$, with $\eta \in H$. The posterior predictive density of such future frequency y, given the observed data, $n = (n_t, t = 1, ..., T)$ is equal to $f(y|n) = \int_H f(y|\eta)\pi(\eta|n)d\eta$. Let z denote a future observation of the severity with density function $f(z|\theta)$, where $\theta \in \Theta$. The posterior predictive density of z , given the observed data s ,is $f(z|s) = \int_\Theta f(z|\theta)\pi(\theta|s)d\theta$.

The two predictive distributions just described can be used to estimate the predictive loss distribution, embedding the well-known convolution process of operational risk into a full Bayesian paradigm. The actual simulation-based convolution mechanism can be summarised in four main steps, as follows.
Step1. For each time period to be predicted, generate n random observations from the predictive frequency distribution;
Step 2. For the same time periods, generate a number of losses from the predictive severity distribution equal to the corresponding frequency observation drawn in step 1 (that is, if the simulated frequency of events for period k is n_k, we simulate n_k severity losses from the predictive severity marginal distribution);
Step 3. For each period, sum the losses obtained in step 2, obtaining a loss observation for the period, drawn from the convoluted marginal distribution as described, thereby obtaining one loss observations for each period;
Step 4. Using the loss observations obtained in step 3, estimate the predictive loss distribution and obtain the VaR and ES, the summary risk measures that establish how much capital is at risk.

3 Application

In this section we describe how our proposal works on a real operational loss database. Coherently with the rest of the paper, the data we use in our analysis concerns one specific business line (retail banking) and one event type (external fraud) of the internal database of operational losses of an anonymous commercial bank, for a period of time that ranges from October 2006 to December 2010. The number of loss data collected is equal to 1855. As self-assessment data is available only for the year 2009, in this analysis we concentrate the analysis on the monthly losses of the year 2009, composed of 396 loss events. In the following all loss data and summary results, unless otherwise indicated, will be expressed in thousands of Euro. Furthermore, all MCMC algorithm simulations were run for a number of iterations equal to 45000, after 5000 of burn-in.

A summary data analysis reveals that the average monthly loss is equal to 38.77, while the minimum and the maximum monthly losses are equal to 17.29 and 115.66, respectively. We now apply a classical GPD model, as in Chavez-Demoulin et al., 2006. On the basis of parameters estimated by means of the Peak Over the Threshold method (see e.g. Pickands, 1975) , we can compute the combined distribution of frequency and severity via a Monte Carlo simulation. Table 1 reports the VaR and the ES estimated for data at hand.

Table 1 about here

From Table 1 note that the summary risk measures, and especially the Value at Risk, do change sensibly along the extreme tail of the distribution, due to the presence of extreme events, in a small dataset: the results are clearly data-dependent.

In order to build our Bayesian model, we now construct a prior distribution for the frequency using the risk self-assessment data provided by the bank. The information contained in such data is based on opinions expressed by different process owners of the bank. Each owner expresses its opinions on the perceived frequency, severity and adequacy of controls for all event types and business lines he is responsible or involved, through the managed banking process. This opinion is mandatory, as the whole self-assessment result is used for periodic internal audit control of the bank operations.

In order to evaluate the impact of actual prior opinions, we first specify an uninformative prior without self-assessment characterised by a very high variance. Using such prior, the posterior mean frequency is equal to 32.98 and the posterior median frequency is equal to 33 thousand Euro. This result seems coherent with the observed sample mean, which is equal to 38.77.

We now consider an informative prior that uses self assessment opinions. Self assessment opinions on the frequency are given by each expert by choosing one between three ordinal categories: low, medium, high frequency, which can be grouped into relevant=(high) and not relevant= (low, medium). In order to use expert opinions we can thus calculate the observed proportion of relevant frequencies, and build an approximate 95% confidence interval around it. Under this framework we get two reference prior distributions: one that will be called lower informative, with hyperparameters set so that the prior expectation matches the lower bound of the confidence interval, and one, conversely, that will be called upper informative, with hyperparameters set so that the prior expectation matches the upper bound of the confidence interval. Specifically, the lower informative prior turns out to be Gamma (7.37, 0.29) and the upper informative prior Gamma (19.53, 0.48).

We now move to the task of specifying a prior distribution for the severity. In analogy with what done for the frequency, we first consider uninformative Gamma prior characterised by a large variance. We then build an informative prior, based on self-assessment data. Opinions on severity are expressed on a ordinal scale with three levels, similar to that of the frequency. We can thus calculate three binomial proportions, one for each level, and follow the procedure employed for the frequency to derive, for each of the corresponding three ordered quantiles, two moment matching gamma priors, corresponding to the lower and upper approximate binomial confidence interval. Specifically we obtain, as lower informative priors: for the first quantile a Gamma (13.55, 91.76), for the second quantile a Gamma (0.15, 14396.12) and for the third quartile a Gamma (0.19, 26486.18). As upper informative priors for the first quantile we obtain a Gamma (35.89, 56.38), for the second quantile a Gamma (0.40, 8845.81) and for the third quantile a Gamma (0.52, 16274.64).

Using the previous prior specifications, we can obtain the predictive distribution for the severity and combine it with the predictive frequency distribution via the MCMC convolution, explained in Section 2. The combined distribution allows us to derive the Value at Risk and the Expected Shortfall, as summary measures of risk. They are shown in Table 2, for each combination of the three alternative frquency priors that we have chosen: Lower Informative (LIF), Upper Informative (UIF) and Uninformative (UNF) and the three alternative triples of severity priors: Lower Informative (LIS), Upper Informative (UIS) and Uninformative (UNS).

Table 2 about here

Comparing Table 2 with Table 1 note that the VaR calculated with the BLDA approach is always lower than the classical one, with the exception of the 99.99% VaR, which is higher for most BLDA models. This percentile may be, however, too extreme to be estimated accurately, with the data at hand. We can thus conclude that, for the percentiles of usual interest in operational risk management, using the BLDA model, we can obtain a remarkable reduction in terms of VaR (always lower, especially with self-assessment priors, than the classical VaR) and, therefore, a lower capital charge.

On the other hand, the comparison in terms of Expected Shortfall between the classical and the Bayesian model are not so clearcut, and do depend on the chosen model priors. This, again, may depend on the limited data available.

Indeed, the most important test of an operational risk model is a predictive one. For this purpose, we can compare the VaRs estimated in Table1 and Table 2, based on the 2009 monthly data, with the actual losses occurred in the year 2010. Concerning the 99.9% VaR, which is the most used reference in operational risk management, note that all VaRs, except that calculated with the LIFUNS model, cover all actual 2010 losses. Of course so does the classical

VaR, but with the disadvantage of a much higher allocation of capital. We can thus conclude that our proposed BLDA model is a promising alternative to classical extreme value models.

It remains the problem of choosing which prior setting is most suitable for the problem at hand. The specification of a collection of alternative priors allows the evaluation of the proposed model, in terms of stability of its results, under different prior specifications, as well as the selection of the most appropriate prior specification, in terms of predictive purposes. In order to choose among the different combinations of prior distributions we have computed in Table 3 the sum of the differences between each VaR and all 2010 observed losses.

Table 3 about here

From Table 3 the best models, in terms of minimisation of the differences between the VaR and the observed losses are : LIFUNS, UIFLIS, UNFUNS, UIFUNS.

However, LIFUNS can be excluded as the corresponding VaR does not cover all losses. Among the remaining three models, UIFLIS has a high Expected Shortfall and, therefore, will also be excluded, being incoherent. In conclusion, we support choosing the UNFUNS or the UIFUNS model, the former being totally uninformative and the second partly informative.

4 Concluding remarks

The main purpose of this paper is to introduce a new Bayesian methodology for estimating loss distributions in operational risk management, making use of both loss data ("backward-looking") and expert opinions ("forward-looking"). We have presented a methodology able to extract expert opinion from self assessment questionnaire, that are typically collected in financial institutions, for audit and control purpose.

Our main outcome is that the application of our proposed Bayesian methodology causes a reduction of the Value at Risk and, therefore, of the capital charge, in comparison with the classical extreme value analysis method. This is a very important result, in terms of capital saved by the financial institution adopting this approach. In addition, the approach may prove to be extremely useful to estimate the operational loss distribution of events with few observed data, as prior opinions could supplement the scarsity of data.

We remark that the case study we have examined in this paper is based on data which considers a short time horizon, as only one year of self-assessment data is available. Bayesian models, typically more parameterised, are usually characterised by a high model variability, when data is limited. Very likely, the relative superiority of our proposed model, with respect to classical ones, will emerge more clearly on a case study with repeated self-assessment exercises.

In any case, even for a limited period of time, our BLDA shows a loss reduction compared with the classical extreme model and therefore a reduction in terms

of regulatory capital.
Our results also show that, using a self-assessment data to specify an informative prior leads to results that are comparable with uninformative prior models, but more stable, so it is helpful that banks analyse self-assessment expert opinions, which are in any case available, as they ought to be collected for audit and control purposes.
An interesting development of our research could be to remove one or more of the independence assumptions (1)-(2)-(3). While removal of (1) could be tackled with the addition of copula models, the removal of (2) would involve a multivariate analysis of all event type/business lines combinations, although this latter would involve multivariate self-assesment (as in Bonafede and Giudici 2007), a rather difficult task. Finally, removal of (3) would involve a time series modelling of oeprational loss data, which may become possible as longer series of data are being collected.

5 Tables

Percentile	VaR	ES
99%	124374.4	162184.3
99.90%	210456.5	304489.8
99.97%	304337.2	420995.3
99.99%	392317.6	568834.4

Table 1: VaR and ES under the classical model

Model	VaR 99%	VaR 99.9%	VaR 99.97%	VaR 99.99%
LIFUNS	67519	111208	192350	413386
UIFUNS	68244	125752	203349	426953
LIFLIS	65986	139424	248751	643641
LIFUIS	81073	141538	219601	495748
UIFLIS	67830	119599	183164	257015
UIFUIS	82424	154822	264339	354807
UNFLIS	66048	139779	242925	490587
UNFUIS	79916	153766	266226	626557
UNFUNS	67568	122634	236121	359226
Model	ES 99%	ES 99.9%	ES 99.97%	ES 99.99%
LIFUNS	96630	257281	513272	961494
UIFUNS	99756	268337	524409	930257
LIFLIS	133452	606196	1545844	3547730
LIFUIS	113138	274391	511727	901157
UIFLIS	112641	408922	992471	2378685
UIFUIS	113669	244735	354755	460954
UNFLIS	101528	292657	564637	994957
UNFUIS	118552	318824	595871	946322
UNFUNS	110495	384555	875528	1859656

Table 2: VaR and ES under the BLDA models

Model	Sum of difference with estimated VaR (99.9%)
LIFUNS	72430
UIFUNS	86974
LIFLIS	100647
LIFUIS	102761
UIFLIS	80822
UIFUIS	116044
UNFLIS	101002
UNFUIS	114989
UNFUNS	83857

Table 3: Sum of the differences between the observed (future) data and the estimated VaR (99.9%)

References

[1] Alexander, C., 2003 Operational Risk. Regulation, Analysis and management. In: Alexander, C. (Ed.), Financial Times Prentice Hall, London, pp. 130-170.

[2] Artzner, P., Delbaen, F., Eber, J., Heath, D., 1999 Coherent measures of risk. Mathematical Finance 9 (3), 203-228.

[3] Behrens, C. N., Lopes, H.F. and D. Gamerman, 2006 Bayesian analysis of extreme events with threshold estimation. Statistical Modelling 6, 251-263.

[4] Bernardo, J.M., Smith, A.F.M., 1994 Bayesian Theory.Wiley, Chichester.

[5] Bilotta A., Giudici P. 2004 Modelling operational losses: a bayesian approach. Quality and Reliability Engineeering International, 20, pp. 407-417.

[6] Bonafede E., P. Giudici 2007 Bayesian networks for enterprise risk assessment. Physica A, 382, pp 22-28

[7] Coles SG, Tawn JA, 1996 A Bayesian analysis of extreme rainfall data. Applied Statistics 45,463-78.

[8] Chavez-Demoulin V., Embrechts P., Neslehova, E.P., 2006 Quantitative models for Operational Risk: Extremes, Dependence and Aggregation. Journal of Banking and Finance 30, 2635-2658.

[9] Cornalba C., P. Giudici, 2004 Statistical models for operational risk management. Physica A: statistical mechanics and its applications 338, 166-172.

[10] Cruz, M. G. 2002 Modeling, Measuring and Headging Operational Risk. John Wiley and Sons, New York, Chichester.

[11] Dalla Valle L., P. Giudici, 2008 A Bayesian approach to estimate the marginal loss distributions in operational risk management. Computational Statistics and Data Analysis 52, 3107-3127.

[12] Diebolt J., Mhamed-Ali El-Aroui, Myriam Garrido and Stephane Girard, 2005. Quasi-conjugate Bayes estimates for GPD parameters and application to heavy tails modelling. Extremes 8, 57-78.

[13] Fantazzini, D. , Dalla Valle, L. and Giudici P. 2008. Copulae and Operational risks. International Journal of Risk assessment and management 9, 238-257.

[14] Giudici P. and Politou D. 2009 Modelling operational losses with graphical models and copulae. Methodology and computing in applied probability 11, 65-94.

[15] Figini S., Giudici P. 2011 Statistical merging of rating models. Journal of the operational research society 62, 1067-1074.

[16] Gamerman, D., 1997 Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference. Chapman and Hall, London.

[17] Panjer H.H. 2006 Operational Risk: Modeling Analytics, Wiley.

[18] Pickands J., 1975 Statistical inference using extreme order statistics, Annals of Statistics 3, 119-131.

[19] Robert, C.P., Casella, G., 1999 Monte Carlo Statistical Methods. Springer, NewYork.