Ruin Theory Revisited: Stochastic Models for Operational Risk

Paul Embrechts
Roger Kaufmann*
Department of Mathematics
ETH Zürich
CH–8092 Zürich
Switzerland

Gennady Samorodnitsky†
School of ORIE
Cornell University
Ithaca, NY
USA

April 2, 2004

Abstract

The new Basel Capital Accord has opened up a discussion concerning the measurement of operational risk for banks. In our paper we do not take a stand on the issue of whether or not a quantitatively measured risk capital charge for operational risk is desirable; however, given that such measurement will come about, we review some of the tools which may be useful towards the statistical analysis of operational loss data. We also discuss the relevance of these tools for foreign reserves risk management of central banks.

Keywords: central banks, extreme value theory, heavy tails, operational risk, risk management, ruin probability, time–change.

1 Introduction

In [9], the following definition of operational risk is to be found: "The risk of losses resulting from inadequate or failed internal processes, people and

*Research supported by Credit Suisse Group, Swiss Re and UBS AG through RiskLab, Switzerland.

†Research partially supported by NSF grant DMS–0071073 at Cornell University.

systems or from external events." In its consultative document on the New Basel Capital Accord (also referred to as Basel II or the Accord), the Basel Committee for Banking Supervision continues its drive to increase market stability in the realms of market risk, credit risk and, most recently, operational risk. The approach 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 finally Pillar 3 concerns Market Discipline. Applied to credit and operational risk, within Pillar 1, quantitative modelling techniques play a fundamental role, especially for those banks opting for an advanced, internal measurement approach. It may well be discussed to what extent a capital charge for operational risk (estimated at about 12% of the current economic capital) is of importance; see Daníelson et al. [23], Goodhart [46] and Pezier [60, 61] for detailed, critical discussions on this and further issues underlying the Accord.

In our paper we start from the premise that a capital charge for operational risk will come about (eventually starting in 2007), and we examine some quantitative techniques which may eventually become useful in discussing the appropriateness of such a charge, especially for more detailed internal modelling. Independent of the final regulatory decision, the methods discussed in our paper have a wider range of applications within quantitative risk management for central banks, the financial (including insurance) industry and supervising authorities. In particular, we are convinced that these methods will play a role in the construction of quantitative tools for integrated risk management, including foreign reserves risk management for central banks. First of all, in an indirect way, central banks have a keen interest in the stability of the global banking system and as such do follow up on the quality/diversity of tools used by the financial industry. Some of these tools used for the modelling of operational risk are discussed in the present paper. Secondly, whether for the demand on foreign reserves one adheres to the intervention model or an asset choice model (see Batten [10]), central banks face risk management decisions akin to commercial banks, albeit under different political and economical constraints. Finally, the role and function of central banks is no doubt under discussion (see Goodhart [45]), and therefore risk management issues which were less relevant some years ago may become important now. In particular, operational risk ought to be of great concern to any central bank. As discussed in Batten [10], a central bank typically confronts two types of economic phenomena – expected and unexpected – to which it makes policy responses. Faced with unanticipated economic but also external (e.g. catastrophic environmental) events to which it may wish to respond, it must hold additional reserves. Furthermore, as any institution, central banks face operational losses owing to system failure and fraud, for instance. How such losses impact on foreign reserve policy very much depends on the portfolio model chosen (Batten [10]).

In Table 1, taken from Crouhy et al. [19], we have listed some typical types of operational risks. It is clear from this table that some risks are difficult to quantify (like incompetence under people risk), whereas others lend themselves much easier to quantification (as for instance execution error under transaction risk). As already alluded to above, most of the techniques discussed in this paper will have a bearing on the latter types of risk. In the terminology of Pezier [61], this corresponds to the ordinary operational risks. Clearly, the modelling of the latter type of risks is insufficient to base a full capital charge concept on.

The paper is organised as follows. In Section 2 we first look at some stylised facts of operational risk losses before formulating, in a mathematical

1. **People risk:**	• Incompetence • Fraud
2. **Process risk:**	
A. Model risk	• Model/methodology error • Mark–to–model error
B. Transaction risk	• Execution error • Product complexity • Booking error • Settlement error • Documentation/contract risk
C. Operational control risk	• Exceeding limits • Security risks • Volume risk
3. **Technology risk:**	• System failure • Programming error • Information risk • Telecommunications failure

Table 1: Types of operational risks (Crouhy et al. [19]).

form, the capital charge problem for operational risk (Pillar 1) in Section 3. In Section 4 we present a possible theory together with its limitations for analysing such losses, given that a sufficiently detailed loss database is available. We also discuss some of the mathematical research stemming from questions related to operational risk. Most of our discussion will use language close to Basel II and commercial banking. At several points, we shall highlight the relevance of the tools presented for risk management issues for central banks.

2 Data and Preliminary Stylised Facts

Typically, operational risk losses are grouped in a number of categories (as in Table 1). In Pezier [61], these categories are further aggregated to the three levels of nominal, ordinary and exceptional operational risks. Within each category, losses are (or better, have to be) well–defined. Below we give an example of historical loss information for three different loss types. These losses correspond to transformed real data. As banks gather data, besides reporting current losses, an effort is made to build up databases going back about 10 years. The latter no doubt involves possible selection bias, a problem one will have to live with till more substantial data warehouses on operational risk become available. One possibility for the latter could be cross–bank pooling of loss data in order to find the main characteristics of the underlying loss distributions against which a particular bank's own loss experience can be calibrated. Such data pooling is well–known from non–life insurance or credit risk management. For Basel II, one needs to look very carefully into the economic desirability of such a pooling arrangement from a regulatory, risk management point of view. Whereas this would be most useful for the very rare large losses (exceptional losses), at the same time, such losses are often very specific to the institution and hence from that point of view make pooling somewhat questionable.

For obvious reasons, operational risk data are hard to come by. This is to some extent true for commercial banks, but considerably more so for central banks. One reason is no doubt the issue of confidentiality, another the relatively short period over which historical data have been consistently gathered. From the quantifiable real data we have seen in practice, we summarise below some of the stylised facts; these seem to be accepted throughout the industry for several operational risk categories. By way of example, in Figures 1, 2 and 3 we present loss information on three types of operational losses, which are for the purpose of this paper referred to as Types 1, 2 and 3. As stated above, these data correspond to modified real data. Figure 4 pools these losses across types. For these pooled losses, Figure 5 contains quarterly loss numbers. The stylised facts observed are:

– loss amounts very clearly show extremes, whereas

– loss occurrence times are definitely irregularly spaced in time, and also show (especially for Type 3, see also Figure 5) a tendency to increase over time. This non–stationarity can partly be explained by the already mentioned selection bias.

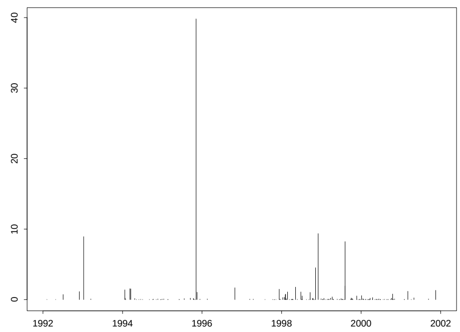

Figure 1: Operational risk losses, Type 1, $n = 162$.

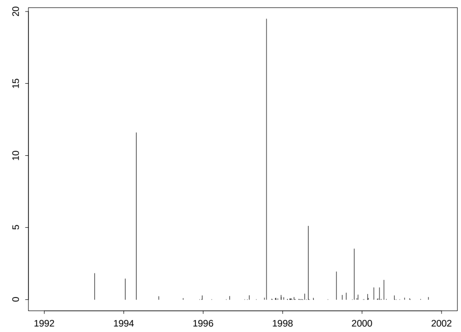

Figure 2: Operational risk losses, Type 2, $n = 80$.

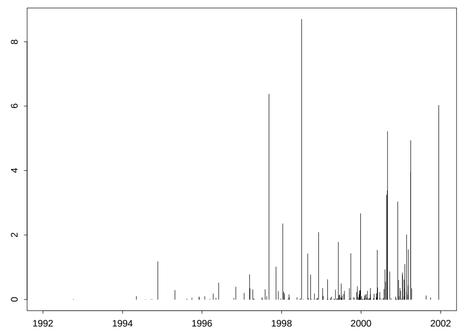

Figure 3: Operational risk losses, Type 3, $n = 175$.

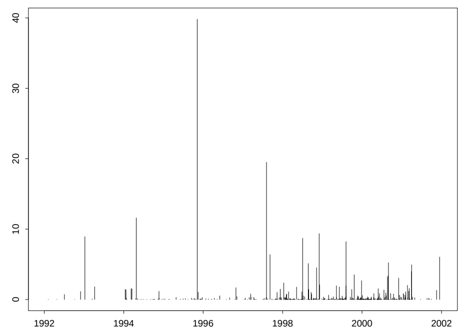

Figure 4: Pooled operational risk losses, $n = 417$.

Any serious attempt of analytic modelling will at least have to take the above stylised facts into account. The analytic modelling referred to is not primar-

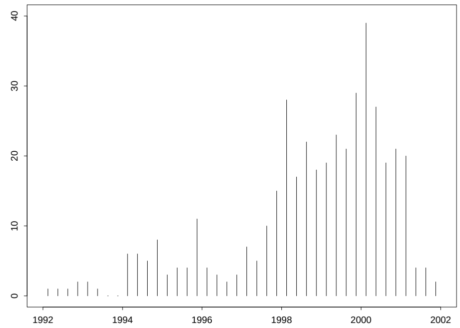

Figure 5: Quarterly loss numbers for the pooled operational risk losses.

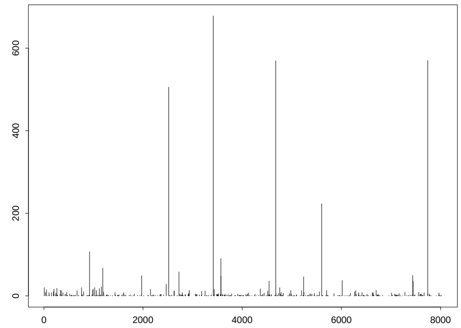

Figure 6: Fire insurance loss data, $n = 417$.

ily aimed at calculating a risk–capital charge, but more at finding a sensible quantitative summary that goes beyond the purely descriptive. Similar approaches are well–known from the realm of reliability (see for instance Bedford and Cooke [11]), (non–life and re–)insurance (Hogg and Klugman [51]) and total quality control (as in Does et al. [26]).

In order to show some similarities with property insurance loss data, in Figure 6 we present $n = 417$ losses from a fire insurance loss database. For the full set of data, see Embrechts et al. [36], Figure 6.2.12.

Clearly, the large losses are a main concern, and hence, extreme value theory (EVT) can play a major role in analysing such data. Similar remarks have been made before concerning operational risk; see for instance Cruz [20] and Medova [55]. At this point, we would like to clarify a misconception which seems to persist in the literature; see for instance Pezier [61]. In no way will EVT be able to "predict" exceptional operational risk losses such as those present in the Barings case, for instance. Indeed, the introduction to Embrechts et al. [36] states very clearly that EVT is not a magical tool that can produce estimates out of thin air, but is instead one that tries to make the best use of whatever data exist on extreme phenomena. Moreover, and indeed equally important, EVT formulates very clearly under what conditions estimates on extreme events can be worked out. Especially with regard to exceptional losses (Pezier [61]), there is very little that statistical theory, including EVT, can contribute. On the other hand, EVT is very useful when it comes to analysing extreme events such as catastrophic storms or floods, where these events occur within a specific physical or environmental model and where numerous observations on "normal" events exist; see Finkenstädt and Rootzén [40]. Clearly a case like Barings falls outside the range of EVTs applicability. On the other hand, when data on sufficient "normal" and a few extreme events within a well–defined class of risks exist, then EVT offers a very powerful statistical tool allowing one to extrapolate from the "normal" to the extreme. Numerous publications within financial risk management exemplify this; see for instance Embrechts [30]. Specific applications of EVT to risk management questions for central banks can be found in De Brandt and Hartman [25] and Hartman et al. [50]. Relevant problems where EVT technology could be applied are discussed in Borio et al. [14]. These papers mainly concern spillover of crises between financial markets, contagion, systemic risk and financial stability. An example of an EVT analysis related to the interest rate crisis of 2000 in Turkey involving interventions at the currency level (the lira) by the Turkish central bank is discussed in Gençay and Selçuk [44]. Embrechts [31] discusses the broader economic issues underlying the application of EVT to financial risk management.

In the next sections we concentrate on the calculation of an operational risk charge based on EVT and related actuarial techniques.

3 The Problem

In order to investigate the kind of methodological problems one faces when trying to calculate a capital charge or reserve for (quantifiable) operational risks, we introduce some mathematical notation.

A typical operational risk database will consist of realisations of random variables

$$\left\{ Y_k^{t,i} : t = 1, \dots, T\,, \quad i = 1, \dots, s \quad \text{and} \quad k = 1, \dots, N^{t,i} \right\}$$

where

- T stands for the number of years ($T = 10$, say);
- s corresponds to the number of loss types (for instance $s = 6$), and
- $N^{t,i}$ is the (random) number of losses in year t of type i.

Note that in reality $Y_k^{t,i}$ is actually thinned from below, i.e.

$$Y_k^{t,i} = Y_k^{t,i} I_{\left\{ Y_k^{t,i} \geq d^{t,i} \right\}}$$

where $d^{t,i}$ is some lower threshold below which losses are disregarded. Here $I_A(\omega) = 1$ whenever $\omega \in A$, and 0 otherwise. Hence, the total loss amount for year t becomes

$$L_t = \sum_{i=1}^{s} \sum_{k=1}^{N^{t,i}} Y_k^{t,i} = \sum_{i=1}^{s} L_{t,i}\,, \quad t = 1, \dots, T\,. \tag{1}$$

Denote by F_{L_t}, $F_{L_{t,i}}$ the distribution functions of L_t, $L_{t,i}$, $i = 1, \dots, s$. One of the capital charge measures discussed by the industry (Basel II) is the Value–at–Risk (VaR) at significance α (typically $0.001 \leq \alpha \leq 0.0025$ for operational risk losses) for next year's operational loss variable L_{T+1}. Hence

$$\text{OR–VaR}_{1-\alpha}^{T+1} = F_{L_{T+1}}^{\leftarrow}(1-\alpha)\,,$$

where $F^{\leftarrow}_{L_{T+1}}$ denotes the (generalised) inverse of the distribution function $F_{L_{T+1}}$, also referred to as its quantile function. For a discussion of generalised inverses, see Embrechts et al. [36], p. 130. Figure 7 provides a graphical definition of this.

Figure 7: Calculation of operational risk VaR.

It is clear that, with any realistically available number T years worth of data, an in–sample estimation of VaR at this low significance level α is very difficult indeed. Moreover, at this aggregated loss level, across a wide range of operational risk types, any statistical theory (including EVT) will have difficulties in coming up with a scientifically sensible estimate. However, within quantitatively well–defined sub–categories, such as the examples in Figures 1–4, one could use EVT and come up with a model for the far tail of the loss distribution and calculate a possible out–of–sample tail fit. Based on these tail models, one could estimate VaR and risk measures that go beyond VaR, such as Conditional VaR (C–VaR)

$$\text{OR–C–VaR}^{T+1}_{1-\alpha,i} = E\left(L_{T+1,i} \mid L_{T+1,i} > \text{OR–VaR}^{T+1}_{1-\alpha,i}\right) , \quad i = 1, \ldots, s ,$$

or more sophisticated coherent risk measures; see Artzner et al. [2]. Furthermore, based on extreme value methodology, one could estimate a conditional loss distribution function for the operational risk category (categories) under investigation,

$$F^{i}_{T+1,u_i}(u_i + x) = P\left(L_{T+1,i} - u_i \leq x \mid L_{T+1,i} > u_i\right) , \quad x \geq 0 , \quad i = 1, \ldots, s ,$$

where u_i is typically a predetermined high loss level specific to loss category i. For instance one could take $u_i = \text{OR–VaR}_{1-\alpha,i}^{T+1}$. Section 4.1 provides more details on this.

We reiterate the need for extensive data modelling and pooling before risk measures of the above type can be calculated with a reasonable degree of accuracy. In the next section we offer some methodological building blocks which will be useful when more quantitative modelling of certain operational risk categories is demanded. The main benefit we see lies in a bank internal modelling, rather than providing a solution towards a capital charge calculation. As such, the methods we introduce have already been tested and made operational within a banking environment; see Ebnöther [28] and Ebnöther et al. [29].

4 Towards a Theory

Since certain operational risk data are in many ways akin to insurance losses, it is clear that methods from the field of (non–life) insurance can play a fundamental role in their quantitative analysis. In this section we discuss some of these tools, also referred to as Insurance Analytics. For a discussion of the latter terminology, see Embrechts [32]. A further comparison with actuarial methodology can, for instance, be found in Duffy [27]. As mentioned in the Introduction, we have also made some references to EVT applications to specific risk management issues for central banks; the comments made below also apply to these.

4.1 Extreme Value Theory (EVT)

Going back to the fire insurance data (denoted $X_1, \ldots, X_n$) in Figure 6, a standard EVT analysis is as follows:

(EVT-1) Plot the empirical mean excess function

$$\widehat{e}_n(u) = \frac{\sum_{k=1}^{n} (X_k - u)^+}{\sum_{k=1}^{n} I_{\{X_k > u\}}}$$

as a function of u and look for (almost) linear behaviour beyond some threshold value. For the fire insurance data, $\widehat{e}_n(u)$ is plotted in Figure 8. A possible threshold choice is $u = 1$, i.e. for this case, a value low in the data.

(EVT-2) Use the so–called Peaks over threshold (POT) method to fit an EVT model to the data above $u = 1$; plot the data (dots) and the fitted model (solid line) on log–log scale. Linearity supports Pareto–type power behaviour of the loss distribution $P(X_1 > x) = x^{-\alpha} h(x)$; see Figure 9. Here h is a so–called slowly varying function, i.e. for all $x > 0$, $\lim_{t\to\infty} \frac{h(tx)}{h(t)} = 1$. For $h \equiv c$, a constant, a log–log plot would be linear.

(EVT-3) Estimate risk measures such as a 99% VaR – and 99% C–VaR – and calculate 95% confidence intervals around these risk measures; see Figure 9.

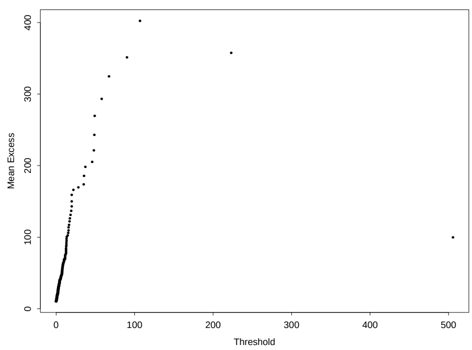

Figure 8: Empirical mean excess function for the fire loss data.

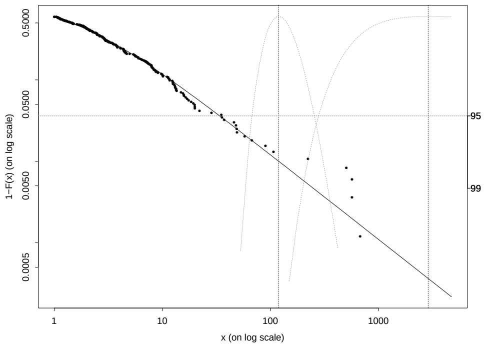

Figure 9: Empirical and fitted distribution tails on log–log scale, including estimates for VaR and C–VaR for the fire loss data.

The estimates obtained are $\widehat{\alpha} = 1.04$ with a corresponding 99% VaR value of 120 and an estimated 99% C–VaR of 2890 (note the huge difference). Figure 9 contains the so–called profile likelihood curves with maximal values in the estimated VaR and C–VaR. A 95% confidence interval around the 99% VaR 120 is given by $(69, 255)$. The right vertical axis gives the confidence interval levels. The interval itself is obtained by cutting the profile likelihood curves at the 95% point. A similar construction (confidence interval) can be obtained for the C–VaR; owing to a value of $\widehat{\alpha}$ (=1.04) close to 1, a very large 95% confidence interval is obtained which hence puts great uncertainty on the point estimates obtained. An α value less than 1 would correspond to an infinite mean model. A value between 1 and 2 yields an infinite variance, finite mean model. By providing these (very wide) confidence intervals in this case, EVT already warns the user that we are walking very close (or even too close) to the edge of the available data.

The software used, EVIS (Extreme Values in S–Plus) was developed by Alexander McNeil and can be downloaded from http://www.math.ethz.ch/~mcneil. It is no doubt a main strength of EVT that, under precise underlying model assumptions, confidence intervals for the risk measures under consideration can be worked out. The techniques used belong to the realm of maximum likelihood theory. We would however like to stress "under precise model assumptions". In Embrechts et al. [35] a simulation study by McNeil and Saladin [54] is reported which estimates, in the case of independent and identically distributed (iid) data, the number of observations needed in order to achieve a pre–assigned accuracy. For instance, in the iid case and a Pareto loss distribution with tail index 2 (a realistic assumption), in order to achieve a reasonable accuracy for the VaR at $\alpha = 0.001$, a minimum number of 100 observations above a 90% threshold u is needed (corresponding to an original 1,000 observations).[1]

The basic result underlying the POT method is that the marked point process of excesses over a high threshold u, under fairly general (though very precise) conditions, can be well approximated by a compound Poisson process (see Figure 10):

$$\sum_{k=1}^{N(u)} Y_k \delta_{T_k}$$

where (Y_k) iid have a generalised Pareto distribution and $N(u)$ denotes the number of exceedances of u by (X_k). The exceedances of u form (in the limit) a homogeneous Poisson process and both are independent. See Leadbetter [52] for details. A consequence of the Poisson property is that inter–exceedance times of u are iid exponential. Hence such a model forms a good first guess. More advanced techniques can be introduced taking, for instance, non–stationarity and covariate modelling into account; see Embrechts [30], Chavez–Demoulin and Embrechts [16] and Coles [17] for a discussion of these techniques. The asymptotic independence between exceedance times and excesses makes likelihood fitting straightforward.

[1] See also Embrechts et al. [36], pp. 194, 270 and 343 for the need to check conditions for the underlying data before an EVT analysis can be used. EVIS allows for several diagnostic

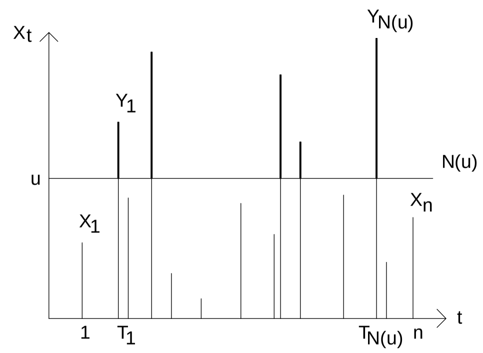

Figure 10: Stylised presentation of the POT method.

Turning to the mean excess plots for the operational risk data from Figures 1–3 (for the type–specific data) and Figure 4 (for the pooled data), we clearly see the typical increasing (nearly linear) trends indicating heavy–tailed, Pareto–type losses; see Figures 11–14 and compare them with Figure 8. As a first step, we can carry out the above extreme value analysis for the pooled data, though a refined analysis, taking non–stationarity into account, is no doubt necessary. Disregarding the possible non–stationarity of the data, one could be tempted to use the POT method and fit a generalised Pareto distribution to the pooled losses above $u = 0.4$, say. Estimates for the 99% VaR and the 99% C–VaR, including their 95% confidence intervals, are given in Figure 15. For the VaR we get a point estimate of 9.1, and a 95% confidence interval of $(6.0, 18.5)$. The 99% C–VaR beyond 9.1 is estimated as 25.9, and the lower limit for its 95% confidence interval is 11.7. Since, as in the fire insurance case, the tails are very heavy ($\widehat{\alpha} = 1.63$), we get a very large estimate for the upper confidence limit for C–VaR.

As already discussed before, the data in Figure 4 may contain a transition from more sparse data over the first half of the period under investigation to more frequent losses over the second half. It also seems that the early

checks on these conditions.

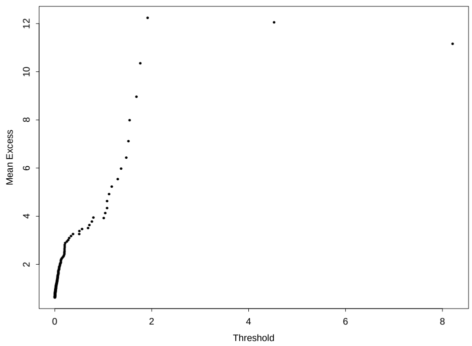

Figure 11: Mean excess plot for operational risk losses, Type 1.

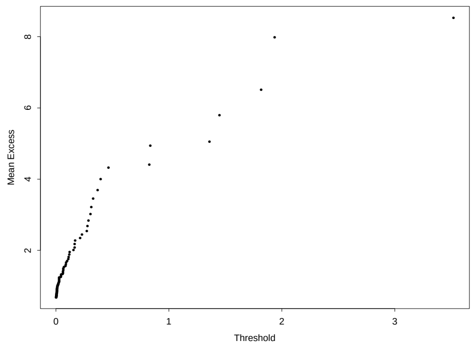

Figure 12: Mean excess plot for operational risk losses, Type 2.

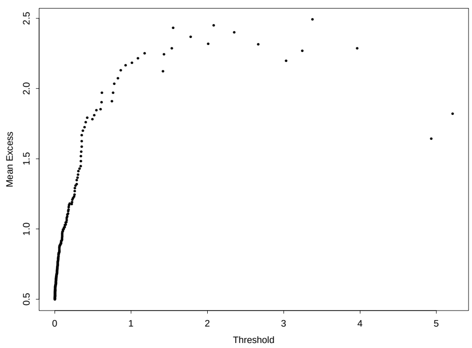

Figure 13: Mean excess plot for operational risk losses, Type 3.

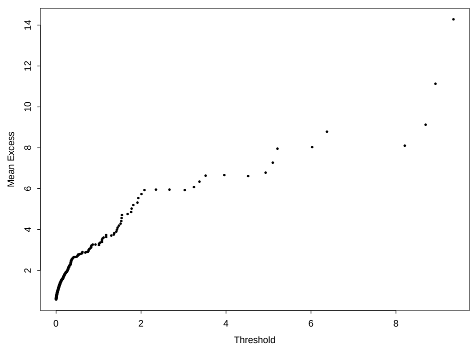

Figure 14: Mean excess plot for pooled operational risk losses.

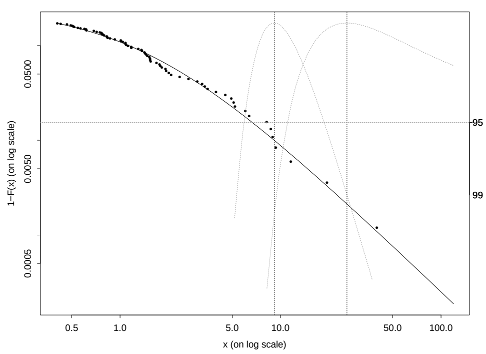

Figure 15: Empirical and fitted distribution tails for pooled operational losses on log–log scale, including estimates for VaR and C–VaR.

losses (in Figure 4 for instance) are not only more sparse, but also heavier. Again, this may be due to the way in which operational loss databases are built up. When gathering data for years some distance in the past, one only "remembers" the larger losses. Our EVT analysis should be adjusted for such a switch in size and/or intensity. Though Chavez–Demoulin and Embrechts [16] contains the relevant methodology, one should however realise that for such more advanced modelling, many more observations are needed.

In the next section, we make a more mathematical (actuarial) excursion in the realm of insurance risk theory. Risk theory provides models for random losses in a dynamic setting and yields techniques for the calculation of reserves given solvency constraints; see for instance Daykin et al. [24]. This theory has been applied in a variety of contexts and may also yield a relevant toolkit, especially in combination with EVT, if central banks were to manage their foreign reserves more actively and hence be more exposed to market, credit, and operational risk. In particular, ruin theory provides a longer–term view on solvency for a dynamic loss process. We discuss some of the key aspects of ruin theory with a specific application to operational risk below.

4.2 Ruin Theory Revisited

Given that (1) yields the total operational risk loss of s different sub–categories during a given year, it can be seen as resulting from a superposition of several (namely s) compound processes. So far, we are not aware of any studies which establish detailed features of individual processes or their interdependencies. Note that in Ebnöther [28] and Ebnöther et al. [29], conditions on the aggregated process are imposed: independence, or dependence through a common Poisson shock model. For the moment, we summarise (1) in a stylised way as follows:

$$L_t = \sum_{k=1}^{N(t)} Y_k \,,$$

where $N(t)$ is the total number of losses over a time period $[0, t]$ across all s categories and the Y_k's are the individual losses. We drop the additional indices.

From an actuarial point of view, it would now be natural to consider an initial (risk) capital u and a premium rate $c > 0$ and define the cumulative risk process

$$C_t = u + ct - L_t \,, \quad t \geq 0 \,. \tag{2}$$

In Figure 16 we have plotted such a risk process for the pooled operational risk losses shown in Figure 4. Again, the "regime switch" is clearly seen, splitting the time axis into roughly pre– and past–1998.

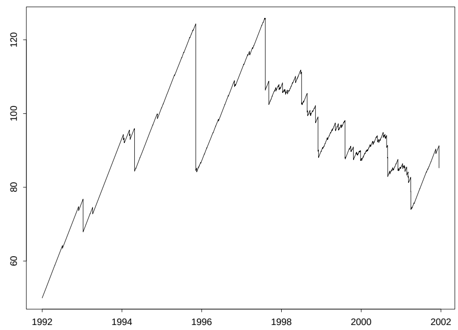

Figure 16: Risk process C_t with $u = 50$, $c = 28$ and the loss process from Figure 4.

Given a small $\epsilon > 0$, for the process in (2), a risk capital u_ϵ can then be calculated putting the so–called ruin probability, i.e. the probability for the surplus process C_t to become negative, over a given time horizon $[\underline{T}, \overline{T}]$ equal to ϵ, small:

$$\Psi\left(u_\epsilon; \underline{T}, \overline{T}\right) = P\left(\inf_{\underline{T} \leq t \leq \overline{T}} (u_\epsilon + ct - L_t) < 0\right) = \epsilon \,. \tag{3}$$

For $c = 0$, $\underline{T} = T+$, $\overline{T} = T + 1$

$$u_\alpha = \text{OR–VaR}_{1-\alpha}^{T+1}.$$

The level of insolvency 0 is just chosen for mathematical convenience. One could, for instance, see c as a premium rate paid to an external insurer taking (part of) the operational risk losses or as a rate paid to (or accounted for by) a bank internal office. The rate c paid and the capital u_ϵ calculated would then be incorporated into the unit's overall risk capital.

Classical actuarial ruin theory concerns estimation of $\Psi(u; \underline{T}, \overline{T})$ in general and $\Psi(u,T) = \Psi(u;0,T)$, $0 < T \leq \infty$ in particular, and for a wide range of processes. The standard assumption in the famous Cramér–Lundberg model is that $(N(t))$ is a homogeneous Poisson(λ) process, independent of the losses (Y_k) iid with distribution function G and mean $\mu < \infty$. Under the so–called net–profit condition (NPC), $c/\lambda > \mu$, one can show that, for "small claims" Y_k, a constant $R \in (0,\infty)$ (the so–called adjustment or Lundberg constant) and a constant $C \in (0,1)$ exist, so that:

$$\Psi(u) = \Psi(u,\infty) < e^{-Ru}\,, \quad u \geq 0\,, \tag{4}$$

and

$$\lim_{u\to\infty} e^{Ru}\Psi(u) = C\,. \tag{5}$$

The small claims condition leading to the existence of $R > 0$ can be expressed in terms of $E(e^{RY_k})$ and typically holds for distribution functions with exponentially bounded tails. The constant C can be calculated explicitly; see for instance Grandell [47], Asmussen [3] and Rolski et al. [63] for details. For the case $\underline{T} = 0$, $\overline{T} = \infty$ and a process for which (4) holds, we can solve for u_ϵ in (3), obtaining

$$u_\epsilon = \frac{1}{R}\log\frac{1}{\epsilon}\,,$$

a quantity which can statistically be estimated given sufficient loss data.

For operational risk losses, the small claims condition underlying the so–called Cramér–Lundberg estimates (4) and (5) are typically not satisfied. Operational risk losses are heavy–tailed (power tail behaviour) as can be seen from Figures 11–14. Within the Cramér–Lundberg model, the infinite–horizon ($T = \infty$) ruin estimate for $\Psi(u) = \Psi(u, \infty)$ becomes (see Embrechts and Veraverbeke [39], Embrechts et al. [36]):

$$\Psi(u) \sim (\frac{c}{\lambda} - \mu)^{-1} \int_u^\infty (1 - G(x))\, dx\,, \quad u \to \infty\,. \tag{6}$$

Hence the ruin probability $\Psi(u)$ is determined by the tail of the loss distribution $1 - G(x)$ for x large, meaning that ruin (or a given limit excess) is caused by typically one (or a few) large claim(s). For a more detailed discussion on this "path leading to ruin", see Embrechts et al. [36], Section 8.3 and the references given there. The asymptotic estimate (6) holds under very general conditions of heavy tailedness, the simplest one being $1 - G(x) = x^{-\alpha} h(x)$ for h slowly varying and $\alpha > 1$. In this case (6) becomes

$$\Psi(u) \sim C\, u^{1-\alpha} h(u)\,, \quad u \to \infty\,, \tag{7}$$

where $C = [(\alpha - 1)(\frac{c}{\lambda} - \mu)]^{-1}$. Hence ruin decays polynomially (slow) as a function of the initial (risk) capital u. Also for lognormal claims, estimate (6) holds. In the actuarial literature, the former result was first proved by von Bahr [7], the latter by Thorin and Wikstad [65]. The final version for so–called subexponential claims is due to Embrechts and Veraverbeke [39]. In contrast to the small claims regime estimates (4) and (5), the heavy–tailed claims case (6) seems to be robust with respect to the underlying assumptions of the claims process. The numerical properties of (6) are however far less satisfactory.

Besides the classical Cramér–Lundberg model, an estimate similar to (6) also holds for the following processes:

– Replace the homogeneous Poisson process $(N(t))$ by a general renewal process; see Embrechts and Veraverbeke [39]. Here the claim inter–arrival times are still independent, but have a general distribution function, not necessarily an exponential one.

– Generalisations to risk processes with dependent inter–claim times, allowing for possible dependence between the arrival process and the claim sizes are discussed in Asmussen [3], Section IX.4. The generalisations contain the so–called Markov–modulated models as a special case; see also Asmussen et al. [6]. In these models, the underlying intensity model follows a finite state Markov chain, enabling for instance the modelling of underlying changes in the economy in general or the market in particular.

– Ruin estimates for risk processes perturbed by a diffusion, or by more general stochastic processes, are for instance to be found in Furrer [41], Schmidli [64] and Veraverbeke [66].

– A very general result of type (7) for the distribution of the ultimate supremum of a random walk with a negative drift is derived in Mikosch and Samorodnitsky [56]. Mathematically, these results are equivalent with ruin estimation for a related risk model.

For all of these models an estimate of type (7) holds. Invariably, the derivation is based on the so–called "one large claim heuristics"; see Asmussen [3], p. 264. These heuristics may eventually play an important role in the analysis of operational risk data.

As discussed above, as yet there is no clear stochastic model available for the general operational risk process (1). Consequently, it would be useful to find a way to obtain a broad class of risk processes for which (7) holds. A solution to this problem is presented in Embrechts and Samorodnitsky [38] through a combination of the "one large claim heuristics" and the notion of operational time (time–change). Below we restrict our attention to the infinite horizon case $\Psi(u)$. First of all, estimate (7) is not fine enough for accurate numerical approximations. It rather gives a benchmark estimate of ruin (insolvency), deliminating the heavy–tailed ("one claim causes ruin") situation from the light–tailed estimates in (4) and (5) where most (small) claims contribute equally and ruin is remote, i.e. has an exponentially small probability. For a discussion on numerical ruin estimates of type (7), see Asmussen and Binswanger [4], and Asmussen et al. [5]; the keywords here are rare event simulation.

Suppose that we are able to estimate ruin over an infinite horizon for a general stochastic (loss) process (L_t), a special case of which is the classical Cramér–Lundberg total claim process in (2) or the risk processes listed above. Suppose now that, for this general loss process (L_t), we have a ruin estimate of form (7). From (L_t), more general risk processes can be constructed using the concept of time–change $(\Delta(t))$. The latter is a positive, increasing stochastic process that typically (but not exclusively) models economic or market activity. The more general process $(L_{\Delta(t)})$ is the one we are really interested in, since its added flexibility could allow us to model the stylised facts of operational risk data as discussed in Section 2. The step from the classical Cramér–Lundberg model (L_t) to the more general process $(L_{\Delta(t)})$ is akin to the step from the classical Black–Scholes–Merton model to more general stochastic volatility models. We can now look at this general time–

changed process and define its corresponding infinite horizon ruin function:

$$\Psi_\Delta(u) = P\left(\sup_{t\geq 0}\left(L_{\Delta(t)} - ct\right) > u\right).$$

We then ask for conditions on the process parameters involved, as well as for conditions on $(\Delta(t))$, under which

$$\lim_{t\to\infty} \frac{\Psi(t)}{\Psi_\Delta(t)} = 1\,, \tag{8}$$

meaning that, asymptotically, ruin is of the same order of magnitude in the time–changed (more realistic) process as it is for the original (more stylised) process. These results can be interpreted as a kind of robustness characterisation for general risk processes so that the polynomial ruin probability estimate (7) holds. In Embrechts and Samorodnitsky [38], besides general results for (8) to hold, specific examples are discussed. Motivated by the example of transaction risk (see Table 1), Section 3 in the latter paper discusses the case of mixing through Markov chain switching models, also referred to as Markov–modulated or Markov renewal processes. In the context of operational risk, it is natural to consider a class of time–change processes $(\Delta(t))$ in which time runs at a different rate in different time intervals, depending on the state of a certain underlying Markov chain. The Markov chain stays a random amount of time in each state, with a distribution that depends on that state. Going back to the transaction risk case, one can think of the Markov chain states as resulting from an underlying market volume (intensity) index. These changes in volumes traded may for instance have an effect on back office errors. The results obtained in Embrechts and Samorodnitsky [38] may be useful to characterise interesting classes of loss processes where ruin behaves as in (7). Recall from Figure 5 the fact that certain operational risk losses show periods of high (and low) intensity. Future dynamic models for sub–categories of operational risk losses will have to take these characteristics into account. The discussion above mainly aims to show that tools for such problems are at hand and await the availability of more detailed loss databases.

Some remarks should be made at this point. Within classical insurance risk theory, a full solution linking heavy–tailedness of the claim distribution to the long–tailedness of the corresponding ruin probability is discussed in Asmussen [3]. Alternative models leading to similar distributional conclusions can be found in the analysis of teletraffic data; see for instance Resnick and Samorodnitsky [62] and Finkenstädt and Rootzén [40]. Whereas the basic operational risk model in (1) may be of a more general nature than the ones discussed above, support seems to exist for the supposition that under fairly general conditions, the tail behaviour of $P(L_{T+1} > x)$ will be power–like. Furthermore, the notion of time–change may seem somewhat artificial. This technique has however been around in insurance mathematics for a long time and is used to transform a complicated loss process into a more standard one; see for instance Cramér [18] or Bühlmann [15]. Within finance, these techniques were introduced through the fundamental work of Olsen and Associates on Θ–time; see Dacorogna et al. [21]. Further work has been done by Ané and Geman [1], Geman et al. [42, 43] and more recently Barndorff–Nielsen and Shephard [8]; they use time–change techniques to transform a financial time series with randomness in the volatility to a standard Black–Scholes–Merton model. As stated above, the situation is somewhat akin to the relationship between a Brownian motion–based model (such as the Black–Scholes–Merton model) and more recent models based on general semi–martingales. It is a well–known result, see Monroe [57], that any semi–martingale can be written as a time–changed Brownian motion.

4.3 Further Tools

In the previous section, we briefly discussed some (heavy–tailed) ruin type estimates which, in view of the data already available on operational risk, may become useful. From the realm of insurance, several further techniques may be used. Below we mention some of them without entering into details. Recall from (1) that a yearly operational risk variable will typically be of the form:

$$L = \sum_{k=1}^{N} Y_k \tag{9}$$

where N is some discrete random variable counting the total number of claims within a given period across all s loss classes, say, and Y_k denotes the kth claim. Insurance mathematics has numerous models of type (9) starting with the case where N is a random variable independent of the iid claims (Y_k) with distribution function G, say. In this case, the probability of a loss exceeding a certain level equals

$$P(L > x) = \sum_{k=1}^{\infty} P(N = k) \left(1 - G^{*k}(x)\right), \tag{10}$$

where G^{*k} denotes the kth convolution of G. Again, in the case that $1 - G(x) = x^{-\alpha}h(x)$, and the moment generating function of N is analytic in 1, it is shown in Embrechts et al. [36] that

$$P(L > x) \sim E(N)x^{-\alpha}h(x), \quad x \to \infty.$$

Several procedures exist for numerically calculating (10) under a wide range of conditions. These include recursive methods such as the Panjer–Euler method for claim number distributions satisfying $P(N = k) = (a + \frac{b}{k})P(N = k-1)$ for $k = 1, 2, \ldots$ (see Panjer [58]), and Fast Fourier Transform methods (see Bertram [12]). Grübel and Hermesmeier [48, 49] are excellent review papers containing further references. The actuarial literature contains numerous publications on the subject; good places to start are Panjer and Willmot [59] and Hogg and Klugman [51].

Finally, looking at (1), several aggregation operations are taking place, including the superposition of the different loss frequency processes $(N^{t,i})_{i=1,\dots,s}$ and the aggregation of the different loss size variables $\left(Y_k^{t,i}\right)_{k=1,\dots,N^{t,i},i=1,\dots,s}$. For the former, techniques from the theory of point processes are available; see for instance Daley and Vere–Jones [22]. The issue of dependence modelling within and across operational risk loss types will no doubt play a crucial role; copula techniques, as introduced in risk management in Embrechts et al. [37], can no doubt be used here.

5 Final Comments

As stated in the Introduction, tools from the realm of insurance as discussed in this paper may well become relevant, conditional on the further development and implementation of quantitative operational risk measurement within the financial industry. Our paper aims at encouraging a better exchange of ideas between actuaries and risk managers. Even if one assumes full replicability of operational risk losses within the several operational risk sub–categories, their interdependence will make detailed modelling difficult. The theory presented in this paper is based on specific conditions and can be applied in cases where testing has shown that these underlying assumptions are indeed fulfilled. The ongoing discussions around Basel II will show at which level the tools presented will become useful. However, we strongly doubt that a full operational risk capital charge can be based solely on statistical modelling.

Some of the tools introduced and mainly exemplified through their application to the quantitative modelling of operational risk are no doubt useful well beyond in more general risk management. Further research will have to look more carefully at the risk management issues underlying the central bank landscape. In particular, the issue of market liquidity under extreme events, together with the more active management of foreign reserves and the longer time view will necessitate tools that complement the existing ones in risk management for commercial banks. Insurance analytics as presented in this paper, and discussed more in detail in Embrechts [33], will no doubt form part of this.

References

[1] Ané, T. and Geman, H. (2000). Order Flow, Transaction Clock and Normality of Asset Returns. *Journal of Finance*, **55**, pp. 2259–84.

[2] Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent Measures of Risk, *Mathematical Finance*, **9**, pp. 203–28.

[3] Asmussen, S. (2000). *Ruin Probabilities*. World Scientific.

[4] Asmussen, S. and Binswanger, K. (1997). Simulation of Ruin Probabilities for Subexponential Claims. *Astin Bulletin*, **27**(2), pp. 297–318.

[5] Asmussen, S., Binswanger, K., and Højgaard, B. (2000). Rare Events Simulation for Heavy–tailed Distributions. *Bernoulli*, **6**(2), pp. 303–22.

[6] Asmussen, S., Henriksen, L. F., and Klüppelberg, C. (1994). Large Claims Approximations for Risk Processes in a Markovian Environment. *Stochastic Processes and their Applications*, **54**, pp. 29–43.

[7] von Bahr, B. (1975). Asymptotic Ruin Probabilities when Exponential Moments Do Not Exist. *Scandinavian Actuarial Journal*, pp. 6–10.

[8] Barndorff–Nielsen, O. E. and Shephard, N. (2002). Financial Volatility and Lévy–based Models. Book project.

[9] Basel Committee on Banking Supervision (2001). Consultative Document. Overview of the New Basel Capital Accord.

[10] Batten, D. S. (1982). *Central Banks' Demand for Foreign Reserves under Fixed and Floating Exchange Rates*. Federal Reserve Bank of St. Louis.

[11] Bedford, T. and Cooke, R. M. (2001). *Probabilistic Risk Analysis: Foundations and Methods.* Cambridge University Press.

[12] Bertram, J. (1981). Numerische Berechnung von Gesamtschadenverteilungen. *Blätter der DGVM*, **15**, pp. 175–94.

[13] Blum, P., Dias, A., and Embrechts, P. (2002). The ART of Dependence Modelling: The Latest Advances in Correlation Analysis. In: *Alternative Risk Strategies* (ed. M. Lane). Risk Waters Group, London, pp. 339–56.

[14] Borio, C., Furfine, C., and Lowe, P. (2001) Procyclicality of the Financial System and Financial Stability; Issues and Policy Options. BIS Papers No. 1.

[15] Bühlmann, H. (1970). *Mathematical Methods in Risk Theory.* Springer.

[16] Chavez–Demoulin, V. and Embrechts, P. (2004). Smooth Extremal Models in Finance and Insurance. *Journal of Risk and Insurance*, to appear.

[17] Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values.* Springer.

[18] Cramér, H. (1930). *On the Mathematical Theory of Risk.* Skandia Jubilee Volume, Stockholm.

[19] Crouhy, C., Galai, D., and Mark, R. (2000). *Risk Management.* McGraw–Hill, New York.

[20] Cruz, M. G. (2002). *Modelling, Measuring and Hedging Operational Risk.* Wiley.

[21] Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B., and Pictet, O. V. (2001). *An Introduction to High–frequency Finance.* Academic Press.

[22] Daley, D. J. and Vere–Jones, D. (1988). *An Introduction to the Theory of Point Processes.* Springer.

[23] Daníelsson, J., Embrechts, P., Goodhart, C., Keating, C., Muennich, F., Renault, O., and Shin, H. S. (2001). An Academic Response to Basel II. Financial Markets Group, London School of Economics, Special Paper No. 130.

[24] Daykin, C. D., Pentikäinen, T., and Pesonen, M. (1994). *Practical Risk Theory for Actuaries.* Chapman and Hall/CRC, London.

[25] De Brandt, O. and Hartman, P. (2000). Systemic Risk: A Survey. Working Paper No. 35, European Central Bank, Frankfurt.

[26] Does, R. J. M. M., Roes, K. C. B., and Trip, A. (1999). *Statistical Process Control in Industry.* Kluwer Academic.

[27] Duffy, P. (2002). Operational Risk. *The Actuary*, October 2002, pp. 24–25.

[28] Ebnöther, S. (2001). Quantitative Aspects in Operational Risk. Diploma thesis ETH Zürich.

[29] Ebnöther, S., Vanini, P., McNeil, A. J., and Antolinez–Fehr, P. (2003). Operational Risk: A Practitioner's View. *Journal of Risk*, **5**(3), pp. 1–15.

[30] Embrechts, P. (ed.) (2000). *Extremes and Integrated Risk Management.* Risk Waters Group, London.

[31] Embrechts, P. (2003). Extremes in Economics and the Economics of Extremes. In: *Extreme Values in Finance, Telecommunications, and the Environment* (eds. B. Finkenstädt, H. Rootzén). Chapman and Hall/CRC, London, pp. 169–83.

[32] Embrechts, P. (2003). Insurance Analytics. Guest Editorial, *British Actuarial Journal*, **8**(IV), pp. 639–41.

[33] Embrechts, P. (2003). Insurance Analytics: Actuarial Tools for Financial Risk Management. Opening Lecture at the XXXIV Astin Colloquium in Berlin.

[34] Embrechts, P., Frey, R., and McNeil, A. J. (2004). *Quantitative Risk Management: Concepts, Techniques and Tools.* Book manuscript, ETH Zürich.

[35] Embrechts, P., Furrer, H., and Kaufmann, R. (2003). Quantifying Regulatory Capital for Operational Risk. *Derivatives Use, Trading and Regulation*, **9**(3), pp. 217–33.

[36] Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance.* Springer.

[37] Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and Dependence in Risk Management: Properties and Pitfalls. In: *Risk Management: Value at Risk and Beyond* (ed. M. Dempster). Cambridge University Press, pp. 176–223.

[38] Embrechts, P. and Samorodnitsky, G. (2003). Ruin Problem and How Fast Stochastic Processes Mix. *The Annals of Applied Probability*, **13**, pp. 1–36.

[39] Embrechts, P. and Veraverbeke, N. (1982). Estimates for the Probability of Ruin with Special Emphasis on the Possibility of Large Claims. *Insurance: Mathematics and Economics*, **1**, pp. 55–72.

[40] Finkenstädt, B. and Rootzén, H. (eds.) (2003). *Extreme Values in Finance, Telecommunications, and the Environment*. Chapman and Hall/CRC, London.

[41] Furrer, H. (1998). Risk Processes Perturbed by α–stable Lévy Motion. *Scandinavian Actuarial Journal*, pp. 59–74.

[42] Geman, H., Madan, D. B., and Yor, M. (2001). Time Changes for Lévy Processes. *Mathematical Finance*, **11**(1), pp. 79–96.

[43] Geman, H., Madan, D. B., and Yor, M. (2002). Stochastic Volatility, Jumps and Hidden Time Changes. *Finance and Stochastics*, **6**(1), pp. 63–90.

[44] Gençay, R. and Selçuk, F. (2001). Overnight Borrowing, Interest Rates and Extreme Value Theory. Working Paper, Bilkent University.

[45] Goodhart, C. (1988). *The Evolution of Central Banks*. MIT Press.

[46] Goodhart, C. (2001). Operational Risk. Special Paper 131, Financial Markets Group, London School of Economics.

[47] Grandell, J. (1991). *Aspects of Risk Theory.* Springer.

[48] Grübel, R. and Hermesmeier, R. (1999). Computation of Compound Distributions I. Aliasing Errors and Exponential Tilting. *Astin Bulletin*, **29**, pp. 197–214.

[49] Grübel, R. and Hermesmeier, R. (2000). Computation of Compound Distributions II. Discretization Errors and Richardson Extrapolation. *Astin Bulletin*, **30**, pp. 309–31.

[50] Hartman, P., Straetmans, S., and de Vries, C.G. (2001). Asset Market Linkages in Crisis Periods. Working Paper No. 71, European Central Bank, Frankfurt.

[51] Hogg, R. V. and Klugman, S. A. (1984). *Loss Distributions.* Wiley.

[52] Leadbetter, M. R. (1991). On a Basis for 'Peaks over Threshold' Modeling. *Statistics and Probability Letters*, **12**, pp. 357–62.

[53] McNeil, A. J. (1997). Estimating the Tails of Loss Severity Distributions Using Extreme Value Theory. *Astin Bulletin*, **27**, pp. 117–37.

[54] 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.

[55] Medova, E. (2000). Measuring Risk by Extreme Values. *Risk*, November 2000, pp. 20–26.

[56] Mikosch, T. and Samorodnitsky, G. (2000). The Supremum of a Negative Drift Random Walk with Dependent Heavy–tailed Steps. *The Annals of Applied Probability*, **10**, pp. 1025–64.

[57] Monroe, I. (1978). Processes that Can Be Embedded in Brownian Motion. *Annals of Probability*, **6**(1), pp. 42–56.

[58] Panjer, H. (1981). Recursive Evaluation of a Family of Compound Distributions. *Astin Bulletin*, **12**, pp. 22–26.

[59] Panjer, H. and Willmot, G. (1992). *Insurance Risk Models*. Society of Actuaries, Schaumburg, Illinois.

[60] Pezier, J. (2002). A Constructive Review of Basel's Proposals on Operational Risk. Working Paper, ISMA Centre, University of Reading.

[61] Pezier, J. (2002). Operational Risk Management. Working Paper, ISMA Centre, University of Reading.

[62] Resnick, S. and Samorodnitsky, G. (2000). A Heavy Traffic Limit Theorem for Workload Processes with Heavy–tailed Service Requirements. *Management Science*, **46**, pp. 1236–48.

[63] Rolski, R, Schmidli, H., Teugels, J., and Schmidt, V. (1999). *Stochastic Processes for Insurance and Finance*. Wiley.

[64] Schmidli, H. (1999). Perturbed Risk Processes: A Review. *Theory of Stochastic Processes*, **5**, pp. 145–65.

[65] Thorin, O. and Wikstad, N. (1977). Calculation of Ruin Probabilities when the Claim Distribution Is Lognormal. *Astin Bulletin*, **9**, pp. 231–46.

[66] Veraverbeke, N. (1993). Asymptotic Estimates for the Probability of Ruin in a Poisson Model with Diffusion. *Insurance: Mathematics and Economics*, **13**, pp. 57–62.