Ruin theory revisited: stochastic models for operational risk.

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

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

December 10, 2002

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 would come about, we review some of the tools which may be useful towards the statistical analysis of operational loss data.

Keywords: operational risk, heavy tails, ruin probability, extreme value theory, 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-, credit- 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. [21], Pezier [47, 48] for detailed, critical discussions on this and further issues underlying the Accord. In our paper we start from the preamble that a capital charge for operational risk will come about (eventually starting in 2006) and discuss 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 the financial (including insurance) industry.

In Table 1, taken from Crouhy et al. [17], 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

1. **People risk:**	• Incompetency • 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. [17]).

terminology of Pezier [48], 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 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 data base is available. We also discuss some of the mathematical research stemming from questions related to operational risk.

2 Data and preliminary stylised facts

Typically, operational risk losses are grouped in a number of categories (like in Table 1). In Pezier [48], these categories are further aggregated to the three levels nominal, ordinary and exceptional operational risks. Within each category, losses are (or better said, 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 are gathering data, besides reporting current losses, an effort is made to build up data bases 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 are becoming available. One possibility for the latter could be cross–banking pooling of loss data in order to find the main characteristics of the underlying loss distributions against which a participating 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 from a regulatory, risk management point of view. Whereas such a pooling 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 more than questionable.

For obvious reasons, operational risk data are hard to come by. One reason is no doubt the confidentiality, another the relatively short historical 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, for the purpose of this paper referred to as Types 1, 2 and 3. As stated above, these data correspond to modified real data. In Figure 4 we have pooled 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, also showing (especially for Type 3, see also Figure 5) a tendency to increase over time. This non–stationarity can partly be due to the already mentioned selection bias.

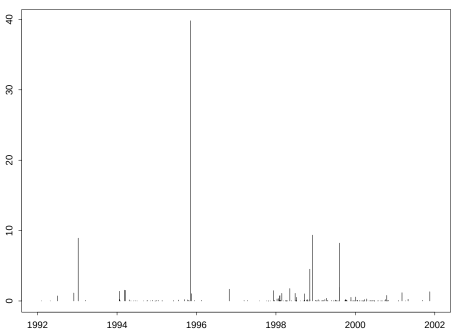

Figure 1: Operational risk losses, Type 1, $n = 162$.

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

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 primarily aimed at calculating a risk–capital charge, but more at finding a sensible quantitative summary going beyond the pure descriptive. Similar approaches are well–known from the realm of reliability (see for instance Bedford and Cooke [10]), (non–life) insurance, reinsurance and total quality control (like in Does et al. [22]).

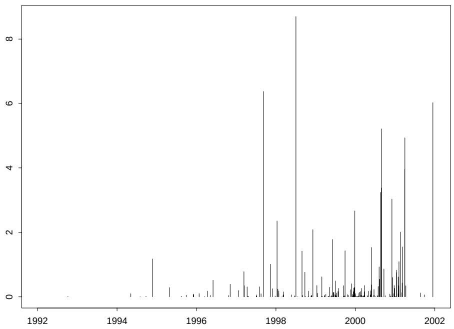

Figure 3: Operational risk losses, Type 3, $n = 175$.

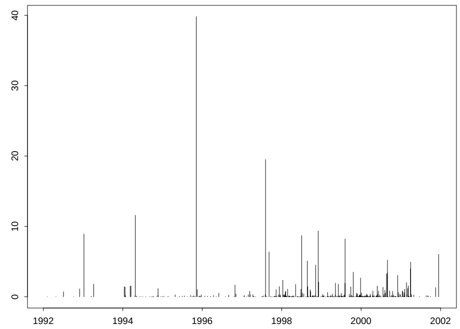

Figure 4: Pooled operational risk losses, $n = 417$.

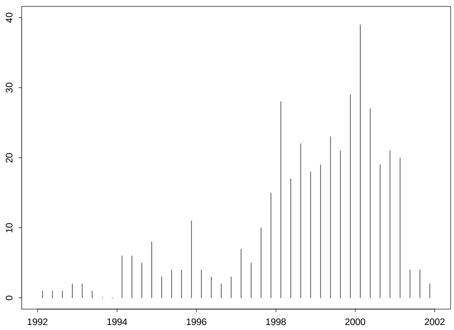

Figure 5: Quarterly loss numbers for the pooled operational risk losses.

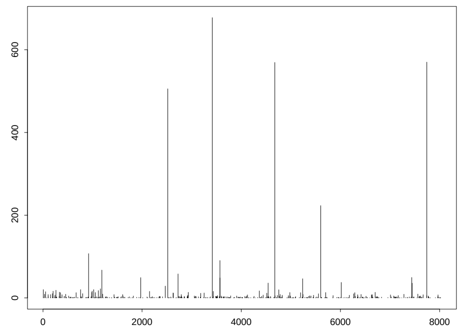

Figure 6: Fire insurance loss data, $n = 417$.

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. [29], Figure 6.2.12.

Clearly, the large losses are of 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 [18] and Medova [42]. At this point, we would like to clarify a misconception which seems to persist in the literature; see for instance Pezier [48]. In no way will EVT be able to "predict" exceptional operational risk losses like those present in the Barings case for instance. Already in the introduction to Embrechts et al. [29], it was stated very clearly that EVT is not a magical tool producing estimates out of thin air but it tries to make the best use of whatever data one may have about extreme phenomena. Moreover, and indeed equally important, EVT formulates very clearly under what conditions estimates on extreme events can be worked out. Especially concerning exceptional losses (Pezier [48]), there is very little statistical theory, including EVT, can contribute. Therefore, throughout the paper, we will only apply EVT to operational risk data which have some sort of underlying repetitiveness.

3 The problem

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

A typical operational risk data base will consist of realisations of random variables

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

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

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

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 $F_{L_{T+1}}^{\leftarrow}$ 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. [29], p. 130. For a graphical definition, see Figure 7.

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

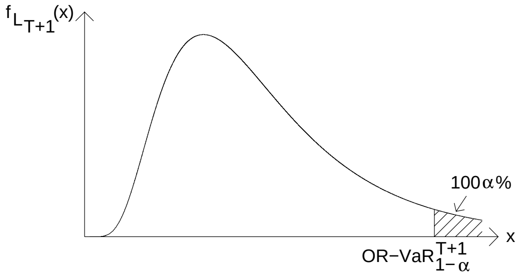

Figure 7: Calculation of operational risk VaR.

impossible. Moreover, at this aggregated loss level, across a wide range of operational risk types, no theory (including EVT) will be able to come up with any scientifically sensible estimate. As such, hoping that EVT will be helpful here is illusory. However, within quantitatively well defined sub–categories, like 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 base on it 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)

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

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

$$F_{T+1,u}(u+x) = P\left(L_{T+1} - u \leq x \mid L_{T+1} > u\right)\,, \quad x \geq 0\,,$$

where u is typically a predetermined high loss level. For instance one could take $u = \text{OR–VaR}^{T+1}_{1-\alpha}$. See Section 4.1 for 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 will be demanded. The main benefit we see lies in a bank internal modelling, rather than 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 [24] and Ebnöther et al. [25], for instance.

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 [27]. A further comparison with actuarial methodology can, for instance, be found in Duffy [23].

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 goes 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 indicates Pareto–type power behaviour of the loss distribution $P(X_1 > x) = x^{-\alpha}h(x)$; see Figure 9.

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

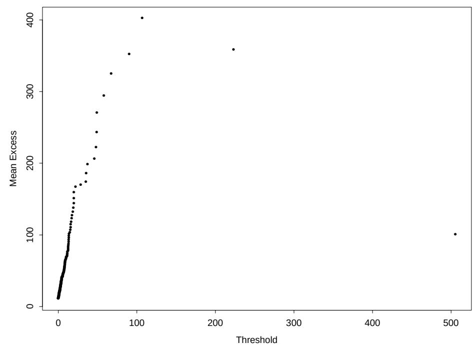

Figure 8: Empirical mean excess function for the fire loss data.

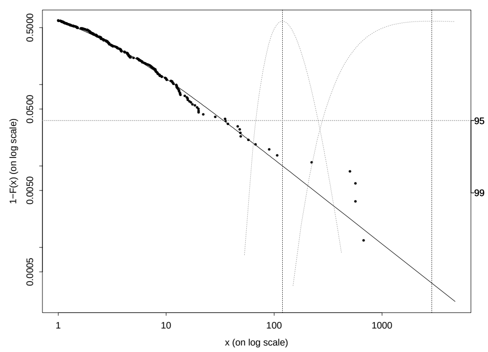

Figure 9: Empirical and fitted distribution tails on log-log scale, including estimates for VaR and C–VaR for the fire loss data.

The estimates obtained are $\widehat{\alpha} = 1.04$ with corresponding 99% VaR value of 120 and estimated 99% C–VaR of 2890. 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; due to a value of $\widehat{\alpha}$ (=1.04) close to 1, a very large 95% confidence interval is obtained. An α value less than one would correspond to an infinite mean model. A value between one and two 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 via http://www.math.ethz.ch/∼mcneil.

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, the exceedances of u form a homogeneous Poisson process and both are independent. See Leadbetter [40] 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 [26], Chavez–Demoulin and Embrechts [14] and Coles [15] for a discussion of these techniques. The asymptotic independence between exceedance times and excesses makes likelihood fitting straightforward.

We once more instruct the reader to look very carefully at the conditions needed before a POT analysis can be performed and be well aware of the "garbage in garbage out" problem; see Embrechts et al. [29], pp. 194, 270, 343. EVIS allows for several diagnostic checks on these conditions.

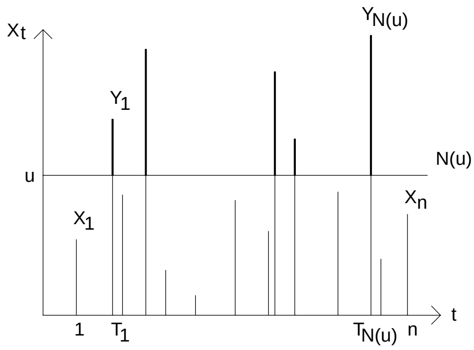

Figure 10: Stylised presentation of the POT method.

When we turn 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. As an example, we use the POT method to fit a generalised Pareto distribution to the pooled losses above $u = 0.4$. We estimate the 99% VaR and the 99% C–VaR, including their 95% confidence intervals; see 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.

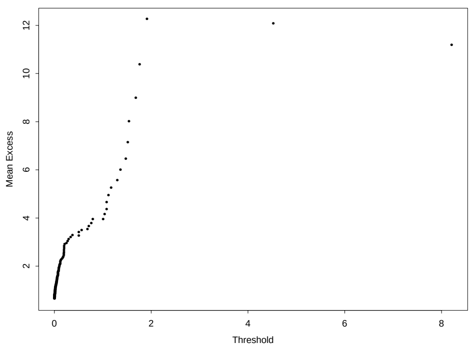

Figure 11: Mean excess plot for operational risk losses, Type 1.

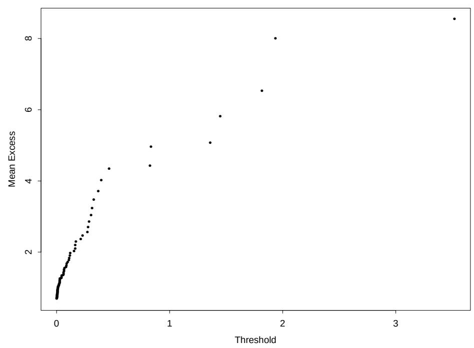

Figure 12: Mean excess plot for operational risk losses, Type 2.

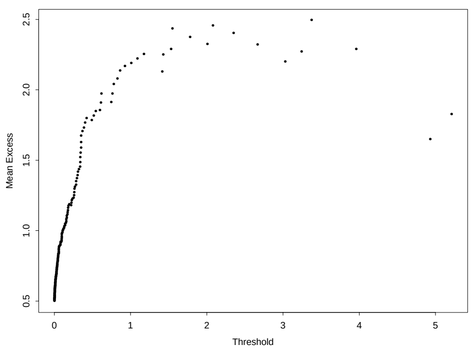

Figure 13: Mean excess plot for operational risk losses, Type 3.

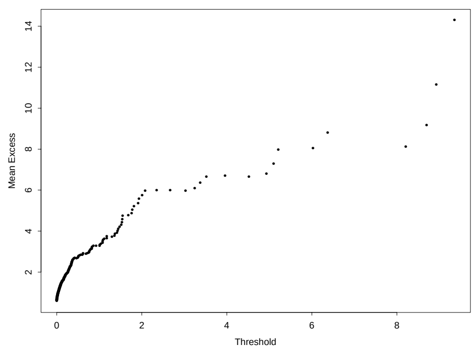

Figure 14: Mean excess plot for pooled operational risk losses.

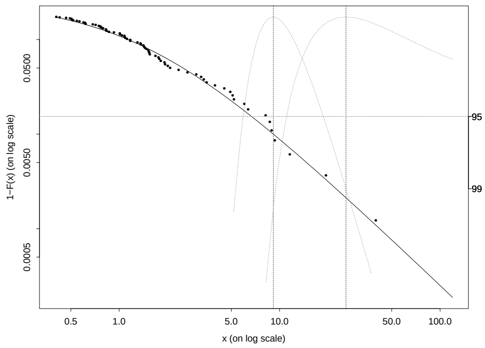

Figure 15: Empirical and fitted distribution tails for pooled operational losses on log-log scale, including estimates for VaR and C–VaR.

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 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 data bases are built up for years some distance in the past; one only "remembers" the larger losses. Our EVT analysis can (and should) be adjusted for such a switch in size and/or intensity; once more, Chavez–Demoulin and Embrechts [14] contains the relevant methodology. We will come back to this point in the next section, where we allow ourselves to make a more mathematical (actuarial) excursion in the realm of insurance risk theory.

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 studies which establish detailed features of individual processes nor their interdependencies. Note that in Ebnöther [24] and Ebnöther et al. [25] 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 various 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. Also here, the "regime switch" is clearly seen, splitting the time axis in roughly pre– and past–1998.

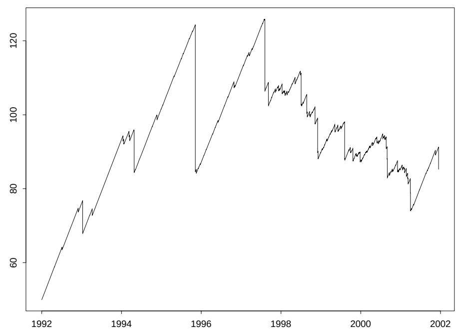

Figure 16: Risk process C_t with $u = 50$, $c = 28$ and the loss process from Figure 4.

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 over a given time horizon $[\underline{T}, \overline{T}]$ equal to ϵ:

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

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 in 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 this for a wide class 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, there exist a constant $R \in (0,\infty)$ (the so–called adjustment or Lundberg constant) and a constant $C \in (0,1)$ 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 [36], Asmussen [3] and Rolski et al. [50] for details. 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 [32], Embrechts et al. [29]):

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

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 few) large claim(s). For a more detailed discussion on this "path leading to ruin" see Embrechts et al. [29], 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, the estimate (6) holds. In the actuarial literature, the former result was first proved by von Bahr [7], the latter by Thorin and Wikstad [52]. The final version for so–called subexponential claims is due to Embrechts and Veraverbeke [32]. 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. 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 [32]. Here the claim interarrival times are still independent, but have a general distribution function, not necessarily exponential.

– Generalisations to risk processes with dependent interclaim 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 [33], Schmidli [51] and Veraverbeke [53].

– A very general result of the type (7) for the distribution of the ultimate supremum of a random walk with a negative drift is derived in Mikosch and Samorodnitsky [43]. Mathematically, these results are equivalent with ruin estimation for a related risk model.

For all of these models an estimate of the 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 already 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 [31] through a combination of the "one large claim heuristics" and the notion of operational time (time change). Below we restrict attention to the infinite horizon case $\Psi(u)$. First of all, the 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 the type (7), see Asmussen and Binswanger [4], and Asmussen et al. [5].

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 the 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 typically (but not exclusively) modelling 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 to model the stylised facts of operational risk data as discussed in Section 2. We can then look at this general time–changed process and define its corresponding infinite horizon ruin function:

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

and 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 [31], 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 in each state a random amount of time, 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 [31] may be useful to characterise interesting classes of loss processes where ruin behaves like 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 is mainly aimed at showing that tools for such problems are at hand and await the availability of more detailed loss data bases.

Some remarks are in order here. 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 are to be found in the analysis of teletraffic data; see for instance Resnick and Samorodnitsky [49]. Whereas the basic operational risk model in (1) may be of a more general nature than the ones discussed above, support seems to exist that under fairly general conditions, the tail behaviour of $P(L_{T+1} > x)$ will be power like. Further, the notion of time change may seem somewhat artificial. This technique has however been around in insurance mathematics for a long time in order to transform a complicated loss process to a more standard one; see for instance Cramér [16] or Bühlmann [13]. Within finance, these techniques were introduced through the fundamental work of Olsen and Associates on Θ–time; see Dacorogna et al. [19]. Further references are Ané and Geman [1], Geman et al. [34, 35] 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. The situation is somewhat akin to the relationship between a Brownian motion based model (like the Black–Scholes–Merton model) and the more recent models based on general semi–martingales. It is a well–known result, see Monroe [44], 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 typically a yearly operational risk variable will 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 the 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, one immediately has that

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

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. [29] 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 like 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 [45]), and Fast Fourier Transform methods (see Bertram [11]). Grübel and Hermesmeier [37, 38] are excellent review papers containing further references. The actuarial literature contains numerous publications on the subject; good places to start are Panjer and Willmot [46] and Hogg and Klugman [39].

Finally, looking at (1), several aggregation operations are going on, including the superposition of the different loss frequency processes $(N^{t,i})_{i=1,\ldots,s}$ and the aggregation of the different loss size variables $\left(Y_k^{t,i}\right)_{k=1,\ldots,N^{t,i},i=1,\ldots,s}$. For the former, techniques from the theory of point processes are available; see for instance Daley and Vere–Jones [20]. 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. [30], can no doubt be used here.

5 Final comment

As already stated in the introduction, conditional on the further development and implementation of quantitative operational risk measurement within the financial industry, tools from the realm of insurance as discussed in this paper may well become relevant. Our paper serves the goal of 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. We strongly doubt however that a full operational risk capital charge can be based solely on statistical modelling.

References

[1] Ané, T. and Geman, H. (2000). Order flow, transaction clock and normality of asset returns. *Journal of Finance*, **55**, 2259–2284.

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

[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), 297–318.

[5] Asmussen, S., Binswanger, K., and Højgaard, B. (2000). Rare events simulation for heavy-tailed distributions. *Bernoulli*, **6**(2), 303–322.

[6] Asmussen, S., Henriksen, L.F., and Klüppelberg, C. (1994). Large claims approximations for risk processes in a Markovian environment. *Stoch. Proc. Appl.*, **54**, 29–43.

[7] von Bahr, B. (1975). Asymptotic ruin probabilities when exponential moments do not exist. *Scand. Actuar. J.*, 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] Bedford, T. and Cooke, R.M. (2001). *Probabilistic Risk Analysis: Foundations and Methods.* Cambridge University Press.

[11] Bertram, J. (1981). Numerische Berechnung von Gesamtschadenverteilungen. *Blätter der DGVM*, **15**, 175–194.

[12] 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, 339–356.

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

[14] Chavez–Demoulin, V. and Embrechts, P. (2003). Smooth extremal models in finance and insurance. *Journal of Risk and Insurance*, to appear.

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

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

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

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

[19] Dacorogna, M.M., Gençay, R., Müller, U.A., Olsen, R.B., and Pictet, O.V. (2001). *An introduction to high-frequency finance.* Academic Press.

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

[21] 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.

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

[23] Duffy, P. (2002). Operational risk. *The Actuary*, October 2002, 24-25.

[24] Ebnöther, S. (2001). Quantitative aspects in operational risk. Diploma thesis ETHZ.

[25] Ebnöther, S., Vanini, P., McNeil, A.J., and Antolinez–Fehr, P. (2001). Modelling operational risk. Submitted.

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

[27] Embrechts, P. (2003). Insurance Analytics. Guest Editorial, *British Actuarial Journal*, to appear.

[28] Embrechts, P., Frey, R., and McNeil, A.M. (2003). *Quantitative Methods for Risk Management.* Book manuscript, ETH Zürich.

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

[30] Embrechts, P., McNeil, A.M., 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, 176–223.

[31] Embrechts, P. and Samorodnitsky, G. (2002). Ruin problem and how fast stochastic processes mix. *Annals of Applied Probability*, to appear.

[32] 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**, 55–72.

[33] Furrer, H. (1998). Risk processes perturbed by α-stable Lévy motion. *Scand. Actuar. J.*, 59–74.

[34] Geman, H., Madan, D.B., and Yor, M. (2001). Time changes for Lévy processes. *Mathematical Finance*, **11**(1), 79–96.

[35] Geman, H., Madan, D.B., and Yor, M. (2002). Stochastic volatility, jumps and hidden time changes. *Finance and Stochastics*, **6**(1). 63–90.

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

[37] Grübel, R. and Hermesmeier, R. (1999). Computation of compound distributions I. Aliasing errors and exponential tilting. *Astin Bulletin*, **29**, 197–214.

[38] Grübel, R. and Hermesmeier, R. (2000). Computation of compound distributions II. Discretization errors and Richardson extrapolation. *Astin Bulletin*, **30**, 309–331.

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

[40] Leadbetter, M.R. (1991). On a basis for 'Peaks over Threshold' modeling. *Statistics and Probability Letters*, **12**, 357-362.

[41] McNeil, A. (1997). Estimating the tails of loss severity distributions using extreme value theory. *Astin Bulletin*, **27**, 117–137.

[42] Medova, E. (2000). Measuring risk by extreme values. *Risk*, November 2000, S20–S26.

[43] Mikosch, T. and Samorodnitsky, G. (2000). The supremum of a negative drift random walk with dependent heavy–tailed steps. *Ann. Appl. Probab.* **10**, 1025–1064.

[44] Monroe, I. (1978). Processes that can be embedded in Brownian motion. *Annals of Probability*, **6**(1), 42–56.

[45] Panjer, H. (1981). Recursive evaluation of a family of compound distributions. *Astin Bulletin*, **12**, 22–26.

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

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

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

[49] Resnick, S. and Samorodnitsky, G. (2000). A heavy traffic limit theorem for workload processes with heavy tailed service requirements. *Management Science*, **46**, 1236-1248.

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

[51] Schmidli, H. (1999). Perturbed risk processes: a review. *Theory of Stochastic Processes*, **5**, 145–165.

[52] Thorin, O. and Wikstad, N. (1977). Calculation of ruin probabilities when the claim distribution is lognormal. *Astin Bulletin* **9**, 231–246.

[53] Veraverbeke, N. (1993). Asymptotic estimates for the probability of ruin in a Poisson model with diffusion. *Insurance: Mathematics and Economics*, **13**, 57–62.