European Finance Review **2:** 113–124, 1999.

Learning about Risk: Some Lessons from Insurance

HÉLYETTE GEMAN
University Paris IX Dauphine and ESSEC, Finance Department, 95021, Cergy Pointoise, France; E-mail: P_geman@edu.essec.fr

Abstract. This paper argues that in the fundamental subject of financial risk analysis, some valuable lessons may be drawn from insurance. The probability of ruin, defined as a first passage time, carries a dynamic element whose absence in Value at Risk is one liability, among others. Extreme value theory, which has been successfully applied to insurance shortly after it was introduced in probability, may offer a coherent framework for analyzing the extreme moves such as the ones observed in recent foreign exchange and financial crises. Lastly, we show that the genuine hazards generated by global capital markets and illustrated by the events of summer 1998, generate a market incompleteness that existing models of defaultable bonds do not fully address. In contrast, the long experience of risk premium analysis in the insurance and reinsurance industry, as well as the existence of historical data on natural disasters, render the valuation of catastrophe bonds less perilous than that of defaultable bonds.

Key words: Value at Risk; Probability of Ruin; Extreme Value Theory; Defaultable Bonds; Catastrophe Bonds.

For centuries, the function of insurers and reinsurers has been to sell risk coverage. In the last few decades, they have been joined in this activity by banks and financial institutions, previously dedicated primarily to the intermediation of funds and increasingly to the intermediation of risks. Today both groups face the same challenge: collect and manage risks by looking for markets where these risks may be hedged or unbundled. When hedges do not exist, for instance due to the structural incompleteness of the corresponding market, a careful measure of the institution exposure needs to be established and, if necessary, prudent reserves built up. When the risk can be unbundled, the task is to devise the right structured contract or security. The goal of this note is to show that in these two situations, lessons may be drawn from insurance, especially for matters related to catastrophic risk.

Advances in financial theory, increasing market volatility, changing regulation, and major shifts in the structure of the financial services industry have propelled market and credit risk to the forefront of the debates in public policy, in academic research and in corporate risk management. Aggregation of market risk from different portfolios in an investment bank or different lines of business within a large firm is being driven by the necessity of controlling risk at the corporate level and manage them most efficiently. Another area of advance is the joint study of market risk and credit risk which, at the analytical level, means acknowledging the relationships between the probabilities of loss due to a market and credit event and the subsequent impact on the firm value.

In the wake of several high-profile failures of risk management such as Metallgesellschaft, Long Term Capital Management or Nomura bank, a need has emerged for better quantification of the financial risks facing corporations and financial services institutions. Prominently featured in this quest for a standardized risk measure is the Value at Risk or VaR. In its various forms, VaR has gained strong support from industry and regulatory bodies such as the Group of Thirty or the Bank for International Settlements. The European Union's Capital Adequacy Directive makes the VaR of a bank's trading book a major input in the calculation of its capital reserve requirement.

However, one should not overstate the novelty of the VaR concept nor underestimate its limits. In the "pre-VaR days", portfolio managers were routinely computing the potential Profit and Loss of their distributions at different horizons. And the Value at Risk at a given level of confidence, 0.95 for example, is nothing but the left 5% quantile of the portfolio Profit and Loss distribution for a date H. As in the case of VaR, the analysis of this quantity requires one to exhibit the *state variables* X_i, i.e., the sources of randomness in the position value, as well as the linear versus nonlinear dependence with respect to these state variables in addition to the crucial discussion of the distribution of the changes of the X_i and, in particular, of their extreme moves. The same issues have been discussed in the insurance industry when studying the so-called probability of ruin which is a more stringent – and probably better – measure of risk exposure.

1. From the Probability of Ruin to VaR

Insurers have been computing for at least a century the probability of ruin of the insurance company. In fact, not only does this quantity present interesting resemblances to VaR, after adjustments of the sources of probability and risk from the insurance industry to the banking industry but it also has the merit, conveyed in its name, of addressing the key issue, namely the probability of insolvency. The foundation of actuarial risk theory is classically identified with the Cramer–Lundberg model, which has, in its simplified version, the following structure: the aggregate claim process $(L(t))_{t\geq 0}$ of the insurance company is defined as

$$L(t) = \sum_{i=1}^{N(t)} S_i$$

where S_i is the amount of claim i, the S_i are random variables taking positive values and supposed to be independent and identically distributed and $N(t)$ is a homogeneous Poisson process with intensity λ, independent of the S_i.

Hence the aggregate claim distribution at time t is

$$G_t(x) = P(L(t) \leq x) = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)}{n!} F^{n^*}(x), \tag{1}$$

where

$$F^{n^*}(x) = P(L(t) \leq x) = P\left(\sum_{i=1}^{n} S_i \leq x\right),$$

is the n-fold convolution of the distribution F of the claim S_i, usually supposed to be a gamma variable with density

$$f(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}, \qquad \alpha \text{ and } \beta > 0$$

or a truncated normal distribution of the type

$$f(x) = \sqrt{\frac{2}{\pi}} e^{x^2/2}.$$

The surplus of the insurance company at time t is defined as

$$U(t) = u + ct - L(t), \tag{2}$$

where $u > 0$ denotes the initial capital and $c > 0$ the amount of insurance premia received per unit of time. The probability of ruin within the finite horizon T for an initial surplus u is defined as

$$\begin{aligned} r(u, T) &= P(u(t) < 0) \qquad \text{for some } t \leq T \\ &= P(L(t) - ct > u). \end{aligned} \tag{3}$$

The probability of ruin in infinite time is $r(u) = r(u, \infty)$, $u > 0$.

Hence, a key problem in insurance, to which a vast amount of literature has been devoted over time, is the determination of the initial capital u and the suitable insurance premium rate u. Obviously, this choice depends on the measure of solvency one wants to emphasize at a given horizon T. The first step is to require that $r\ (u, T) < 1$ for all u and T positive, but this is insufficient and we are naturally led to a condition of the type:

(Prob of ruin at the horizon T) $< \alpha$ (where $\alpha = 0.05$ for instance)

which implies that

$$L(T) - cT - u < \mathrm{VaR}_\alpha(T), \tag{4}$$

but also that the same inequality holds for all dates $t < T$. As defined in Equation (3), ruin is a *first-passage time* and hence the corresponding required capital will always be greater that the value at risk at a fixed horizon (whether that horizon be a day or three months).

A first necessary condition for $r(u)$ to be bounded is an economic constraint on the premium rate c relative to the intensity λ of the claim process and the mean μ of the claim distribution, namely that

$$\frac{c}{\lambda\mu} = 1 + \rho \qquad \text{where } \rho > 0. \tag{5}$$

This 'net profit condition' which reduces the claim distribution to its mean is obviously not sufficient for $r(u)$ to be bounded. Assuming it is satisfied, a mathematical condition – called 'small claim condition' and established in the first part of the 20th century by Cramer and Lundberg – requires *exponentially small probabilities* of occurrence of large claims in order to obtain an upper bound for the ruin probability in terms of the initial capital u. Not surprisingly, this condition bears only on the tail of the claim distribution. In practice, claims are mostly modelled by log-normal, log-gamma or even more heavy-tailed distributions and refinements of the Cramer–Lundberg results need to be established in order to verify that the probability of ruin is still reasonably low. Moreover, changes in the world climatic conditions as well as in the value of property, structures and equipment cannot allow to safely state that the claim distribution is identical over time. These are some of the main problems which make VaR analysis a perilous exercise, namely the misspecification of the distributions and their non-stationarity or at least the non-stationary of the correlations. The good news from the insurers perspective however is that correlations between claims located in different regions of the world can legitimately be viewed as stationary and close to zero. Recent world events have shown that neither of these two key properties holds for distributions of asset returns in the financial markets and that for instance credit risk diversification in a portfolio of sovereign bonds from countries in different parts of the world may well turn out to be in fact risk concentration when all correlations become suddenly close to one. Besides the problem of correlations, the misspecification of market moves and changes in the portfolio value impinges upon the reliability of VaR estimates.

2. Is Risk due to Model Errors or to Erroneous Models?

One of the standard assumptions made by conventional VaR models such as the original form of J.P. Morgan's Riskmetrics is that the returns on a portfolio have a normal distribution, assumption which allows an elementary computation of all quantiles, in particular of the ninety-ninth percentile. In fact, the Riskmetrics Technical manual (fourth edition, 1997) states 'An important advantage of assuming that changes in asset prices are normally distributed is that we can make predictions about what to expect to occur'. The distribution is not defined by the data but chosen in order to make some statistical tools available. In practice, returns are rarely normally distributed and outliers (or crashes) more common than the Gaussian distribution would suggest. Though there is variation from market to market, distributions of daily returns are mostly both sharp-peaked and fat-tailed, and even more so when measured on shorter horizons. Normality may be recovered through a change of probability measure or a change of time scale (see Ané-Geman, 1997) but this change is asset-specific and does not extend easily to several instruments. Even in the case of a simple portfolio comprised of n stocks, the return r_p is a weighted-average of the returns on the n individual stocks; hence normality for r_p requires multivariate normality of the n components, a very strong assumption.

Returns on options and option portfolios are inherently non Gaussian. This is more generally the case with instruments that contain embedded options and have a nonlinear payout profile with respect to the underlying market variables. For instance, the prepayment option in mortgage-backed securities, collateralized mortgage obligations and even more exotic mortage-related products have played a key role in the recent period of persistently falling interest rates. The importance of this nonlinearity increases as the option becomes closer to the money or when the time to maturity increases. To avoid a possibly significant exposure to gamma risk in large option portfolios, a January 1996 amendment to the Basle Capital Accord stipulates that banks that write options are required to measure delta, gamma and vega risk when calculating the capital to be held against them.

Britten Jones and Shaefer (1998) address the issue of these nonlinear terms in the computation of the Value at Risk and show the immediate increase in complexity it generates: even in the simple situation of normally distributed changes in the state variables X_i, the squares $(\Delta X_i)^2$ become non-centered χ^2-distributions and a linear combination of these quantities with possibly negative coefficients is no longer a χ^2 variable. Consequently, the quantile probabilities cannot be read from a table of χ^2 as soon as there is more than one market variable at stake. Britten Jones and Shaefer propose to use an approximation (the so-called Salomon–Stephen approximation) for the two sums of terms of the same sign and then recombine their difference through Monte-Carlo. An alternative technique to handle this signed combination of χ^2 variables consists in inverting the characteristic function of the algebraic sum, using the semi-closed form expression provided by Provost and Rudiuk (1995). In the two approaches, the consideration of nonlinear terms in VaR is not straightforward.

At an econometric level, one problem related to VaR is to decide whether the processes we observe are stationary. If they are at least weakly stationary, then ergodic theory states that we can estimate parameters with a confidence level in some proportion to the sample size. But if we introduce higher dimensional processes to account for the numerous sources of risk, or if we assume Markov switching distributions, then the amount of necessary data increases exponentially. Statements made after the 1987 crash and the summer 1998 financial turmoil that these were 'ten sigma events' would require millions of observations to give the claims a scientific meaning.

A well-known problem of inductive learning in any scientific instance is the so-called bias/variance dilemma. One may decide to choose a simple model with few parameters, but there will be structural errors in the estimators or in their classification. In our risk analysis, it will translate into a high variance of VaR due to the parsimony of the model. Or a more complex representation is chosen, in which case there may not be enough data to calibrate the parameters, leading to the risk of overfitting. An elementary illustration of this risk is provided by the misleading identification of a unique cubic spline using four points or a polynomial of degree n using $(n + 1)$ points. The increasing mathematical sophistication and complexity of the techniques may mask in a subtle manner a similar misspecification of the problem.

In order to address the fact that traditional parametric methods based on the estimation of the entire density are not well suited to the determination of extreme quantiles, a number of authors have proposed first in insurance (Embrechts and Schmidli, 1994), then in finance (Longin, 1997) the use of 'extreme value theory'. This field of probability and statistics provides a theoretical framework to estimate the tail of the distribution on the basis of extreme event data rather than all the data and arguably appears appropriate for risk management in finance. Most of the literature on the subject starts with the crucial assumption that returns X_i are in the maximum domain of attraction of a Frechet distribution, or in other words that,

$$P(X_i > x) = k(x)x^{-a} \tag{6}$$

where $k(x)$ is often assumed to be constant and α a key parameter to be estimated. However, as observed by Diebold et al. (1997), the bias-variance tradeoff regarding the size of the sample is analogous to the one discussed earlier: the statistical results in extreme value theory are much less sophisticated than the probabilistic ones and, in particular do not provide at this point any rule on the choice of the sample. Moreover, a classical hypothesis in this literature is the independence of returns; and the empirical evidence of financial asset returns being conditionally heteroskedastic, in particular over short horizons, has been widely documented. Lastly, it is not clear that the assumption of Frechet limiting distribution for returns as expressed in Equation (9) is necessarily the correct one. A paper by Inoue (1998) raises this issue and some recent empirical studies suggest that, in fact, the Weilbull distribution may be more appropriate for extreme returns observed on the financial markets (the third possibility provided by extreme value theory, the Gumbel distribution – obtained as the limiting distribution for normal and lognormal variables – being rejected in all empirical findings).

Let's now turn to a subject where extreme value theory has been applied in practice, namely the domain of catastrophic risk securitization.

3. Catastrophe Bonds: Less Risky than Defaultable Bonds

Catastrophe bonds, sometimes called Act of God bonds, are bonds issued by an insurance or reinsurance company (or sometimes by a state, as did the state of California when insurance companies decided after the 12.5 billion insured losses generated by the Northridge earthquake in 1994 not to write any more coverage against earthquake). The coupon (and part of the principal, depending on the tranche) is not paid if some well-specified event occurs or if a sinistrality index in a designated geographical area (e.g., the PCS Property Claim Services index to which are tied the pay-offs of catastrophe derivatives launched in December 1993 by the Chicago Board of Trade) reaches a certain threshold during a specified time period.

The credit risk is practically eliminated thanks to the cash transfer of the proceeds of the bond issue into a special purpose vehicle, in turn mostly invested in Treasury securities. Consequently, catastrophe bonds look essentially like defaultable bonds, except that they are better instruments in terms of portfolio diversification and easier to price:

(a) the correlation between catastrophe bond return and market return is very close to zero, hence the diversification effect is outstanding.
(b) the fraction of the principal at risk is not a random quantity as in the case of defaultable bonds, but fully specified at issuance
(c) the hazard rate, i.e., the probability of an event over the next accounting period, can be derived from observable variables provided by weather and meteorology departments.

Let us recall that in the simple Poisson case – which has been, since Lundberg's work in 1903, the central model for insurance losses – the hazard rate is identical to the intensity λ of the Poisson process; i.e., the probability of an event occuring in the interval $[t, t+dt]$ is equal to $\lambda\, \mathrm{d}t$ and the probability of no event up to time t is $\mathrm{e}^{-\lambda t}$.

In the so-called reduced form models for the valuation of defaultable bonds, default is also represented by an unpredictable stopping time and the terminology from insurance, 'hazard rate', has been adopted by most authors (Madan-Unal 1993, Duffie-Singleton 1994 and many others). This hazard rate λ in the case of catastrophe bonds is adjusted deterministically over time according to meteorology information. In the risky debt market, λ_t evolves stochastically over time (as recent world events have shown) and the recovery rate $(1-\ell)$ is also a random variable; intensity and magnitude of losses may be exogeneous or derived from some equilibrium model. One central result, derived by arbitrage arguments, is that the credit spread over Treasurys may be written as $s = \lambda.\ell$ and the value at time t of a cash-flow ϕ to be paid at time T is

$$V(t) = E_Q\left[\phi \exp\left(-\int_t^T (r(u) + \lambda.\ell)\mathrm{d}u\right) / F_t\right]. \tag{7}$$

where Q is the risk-adjusted probability measure and F_t the information available at time t. In the case of deterministic parameters λ and ℓ, an intuitive proof of the expression of the spread in terms of λ and ℓ can be obtained by writing in two different manners the value at date t of a risky dollar to be paid at time $t + \mathrm{d}t$. The first one involves a risk-adjusted discount factor, namely a spread s added to the (locally risk-free) short-term rate $r(t)$; the other a probability-weighted average of the payouts discounted at the risk-free rate, i.e.,

$$\frac{1}{1 + (r + s)\,\mathrm{d}t} = 1.(1 - \lambda\,\mathrm{d}t) + (1 - \ell)\lambda\,\mathrm{d}t. \tag{8}$$

Formula (7) is very elegant and in total continuity with the actuarial methodology described above; but it raises some questions. The equivalent martingale measure Q is not unique since the completeness of the defaultable bond market is questionable, in particular when the intensity λ is supposed to be stochastic. The identification of the two key parameters, λ and ℓ, is a very difficult problem, under any probability measure and in particular under the measure Q since it should incorporate the risk premia.

Financial institutions do not necessarily have a database big enough to extract reliable estimators. Also it is not clear that rating agencies which provide not only credit ranking but also migration probability matrices for different grade bonds do not work as well on quicksand since the current situation of globally integrated financial markets has not been experienced before. Lastly, when the loss ratio ℓ and (or) the intensity λ are stochastic, the use of the T-forward neutral probability measure Q_T (see Geman 1989) which is a powerful tool to price random cash-flows under stochastic interest rates does not provide an immediate answer because of the non independence of the level of interest rates and *the credit* spread $\lambda\ell$.

In a somewhat circular theoretical argument, formula (6) is frequently used to infer the value of the spread as a function of the maturity of the bond and rating of the firm. This methodology obviously assumes that markets price credit in a coherent manner, i.e., understand the risk and know the recovery value. The recent shrinking of emerging market departments in a number of institutions hints that this assumption may not necessarily be founded. The question 'how to price' then remains open. The fundamental lesson derived from option theory is that the only safe answer resides in the price of the (self-financing) hedging portfolio. If no hedge can be found, the existence of an arbitrage price is not obvious. The valuation of the security then requires either the use of the utility function of a representative agent or of a risk-minimization criterion appropriate for the type of incompleteness at stake and leads in all cases to the relevance of prudent reserves.

Formula (6) applies to catastrophe bonds in a much safer manner because of the greater reliability of the estimation of the key parameters. Data have been accumulated over time and their treatment has received renewed attention over the last few years because of the crucial role they play in the valuation of catastrophe derivatives (traded on the Chicago Board of Trade or in OTC transactions). Historical data on insured losses are used to estimate the probabilities of exceeding certain trigger levels in a defined area; then, adjustments for changes in demographics and property values need to be made. It is important to notice that, at variance with the random evolution over time of credit worthiness in the financial markets, these adjustments are deterministic. Several engineering firms have developed a real expertise in providing loss estimates. These firms collect and maintain data at the zip code, county, state and regional level on industry exposures, building inventories and construction codes. They then employ sophisticated software to simulate structural damage to existing homes and business under a variety of natural disasters in different regions.

If we take as an example the recent CAT bonds issued by Swiss Reinsurance on behalf of the insurance company Tokio Marine, the coupon (or the principal, for some riskier tranches) is nullified or reduced if the magnitude of a seismic event on the Richter scale in the Grand Kanto area of Japan reaches a given threshold. The available information on the dollar (or yen) amount of losses generated by an earthquake in terms of the distance of its epicenter to major cities is bewildering by its quality. As for the probabilities of occurrence of such seismic events, they have been studied in departments of civil engineering by academics who were routinely using Monte Carlo simulations thirty years ago.

Consequently, the price at time t of a catastrophe bond can be written as

$$V(t) = E_Q\left[\sum_{k=1}^{n} \phi_k \exp\left(-\int_t^{t+k} r(s)\,\mathrm{d}s\right) 1_{I_k<A}/F_t\right], \tag{9}$$

where

- the cash flows ϕ_k, $k = 1, \ldots, n-1$, represent the fixed coupon payments at dates $t + k$ and ϕ_n includes the principal repayment (if a coupon of the type Libor plus basis points is paid, a swap may be added to the structure to recover the fixed-coupon framework)
- I_k is the loss index over the year $(t + k - 1, t + k)$ and A the trigger level
- Q is the risk-adjusted probability measure whose unicity depends on the 'primitive assets' available in the market to which the index is related as well as the number of sources of randomness in the index representation. For instance, in the case of the CBOT insurance derivatives, Geman (1994) proposed to extract from layers of reinsurance the construction of the probability measure Q.

We can observe that, using the reasonable assumption of independence between interest rates and natural risks, formula (8) can be reduced to

$$V(t) = \sum_{k=1}^{n} B(t, t+k)\phi_k E_Q[1_{I_k<A}/F_t], \tag{10}$$

and no choice of interest rate model is necessary. If the coupon ϕ_k is not fixed but depends on the Treasury yield curve observed at date $(t+k)$, the same formula prevails with ϕ_k replaced by $E_{Q_{t+k}}(\phi_k)$ where Q_{t+k} is the forward-adjusted probability measure relative to date $t+k$ and defined by its Radon-Nikodym derivative

$$\frac{\mathrm{d}Q_{t+k}}{\mathrm{d}Q} = \frac{\exp\left(-\int_t^{t+k} r(s)\,\mathrm{d}s\right)}{B(t,t+k)}.$$

Then, the only quantity which needs to be computed is

$$E_Q[1_{I_k<A}/F_t] = Q[(I_k < A)/F_t],$$

and this obviously requires the specification of the index I and its dynamics over the period $(t+k-1, t+k)$.

The index I can be an aggregate loss index such as the one published by institutions like Insurance Statistical Office or Property Claims Service; it can also be a number on the Richter scale or another measure of a catastrophic event, not subject to possible manipulations in order to avoid moral hazard problems. In all cases, the index I should be represented by a stochastic process including a jump component (this modelling may also be appropriate for temperature indexes, such as those measured by the National Weather Service in the United States and to which heating and cooling degree-days weather derivatives are tied).

The dynamics of I under the probability measure Q may be described by the stochastic differential equation

$$\mathrm{d}I_t = I(t^-)[\alpha\,\mathrm{d}t + \sigma\,\mathrm{d}W_t] + k\,\mathrm{d}N_t \tag{11}$$

where

- α and σ are constants accounting for the diffusion part of I, i.e., for the small moves
- (W_t) is a Brownian motion under Q
- (N_t) is a Poisson process with intensity λ, where λ denotes the frequency of the jumps
- (W_t) and (N_t) are assumed to be independent.

The solution of equation (10) is the Doléans–Dade exponential

$$I(t) = \mathrm{e}^{X(t)}\left[S(0) + k\int_0^t \mathrm{e}^{-X(u)}\right], \tag{12}$$

where

$$X(t) = \left(\alpha - \frac{\sigma^2}{2}\right)t + \sigma N(t).$$

The values of parameters α, σ, λ, k can be derived (see Cummins and Geman, 1995) from the market prices of insurance derivatives. The computation of the probabilities $Q(I_k < A)$ requires Monte-Carlo simulations of the trajectories of the index I as described in Equation (12), since an exact closed-form solution only exists today for the cases of a pure diffusion ($k = 0$) or a pure jump ($\alpha = \sigma = 0$). These quantities incorporated in Equation (10) give the price of the catastrophe bond.

Among many other properties which could be discussed, an important one needs to be emphasized at this point: the market value of a catastrophe bond, as long as the proceeds of the bond issue are mostly secured in Treasuries through a Special Purpose Vehicle (and this is the case in practice), does not depend on the *size* of the issue. This is another key difference between catastrophe bonds and sovereign or corporate bonds, with major consequences in terms of risk management.

4. Conclusion

Extreme value theory, which researchers in insurance were the first to introduce for the analysis of catastrophic risk, may become a useful tool for financial risk management. The recent financial turmoil in the emerging markets confirms that a probability quantile cannot solve all problems. When dealing with rare events, historical time series will always be lacking part of the relevant information and it becomes necessary to remodel these data by accounting for events not yet observed. It may also be useful to put a tighter norm (in a mathematical and regulatory sense) on the admissible portfolio strategies in order to avoid insolvency not only at the horizon of analysis but also at any intermediate date.

Regarding catastrophic risk securitization, this procedure allows to transfer in thin layers natural disaster risk into the financial markets with a great diversification effect, not only on the investors' portfolios but also in terms of nature risk sharing worldwide. Given the cyclical lack of capacity in the insurance and reinsurance industry, this transfer should arguably improve the general welfare.

References

Ané, T and Geman, H. (1997) Order flow, transaction clock and normality of asset returns, Manuscript, CEREG, University Paris IX Dauphine.

Britten Jones, M. and Shaefer, S. (1998) Nonlinear value at risk, forthcoming, *European Finance Review*.

Culp, G. L., Miller M. H. and Neves, A. M. (1998) Value at risk: Uses and abuses, *Journal of Applied Corporate Finance*, Winter.

Cummins, J. D. and Geman, H. (1995) Pricing catastrophe insurance futures and call spreads: An arbitrage approach, *Journal of Fixed Income*, March.

Diebold, F. X., Schuermann T. and Stroughair, J. D. (1998) Pitfalls and opportunities in the use of extreme values theory in risk management, Manuscript, the Wharton School.

Duffie, D. and Singleton, K. J. (1994) Modeling term structures of defaultable bonds, Manuscript, Graduate School of Business, Stanford University.

Embrechts, P. and Schmidli, H. (1994) Modelling of extremal events in insurance and finance, *Math. Meth. Op. Res.* **39**.

Geman, H. (1989) The importance of the forward neutral probability measure in a stochastic approach to interest rates, ESSEC Working Paper.

Geman, H. (1994) Castastrophe calls, *RISK*, September.

Hill, B. M. (1975) A simple general approach to inference about the tail of a distribution, *Annals of Statistics* **3**.

Inoue, A. (1998) A conditional goodness-of-fit test for time series, Manuscript, Department of Economics, University of Pennsylvannia.

Lamarre, M., Townsend, B. and Shah, H. C. (1992) Application of the bootstrap method to quantify uncertainty in seismic hazard estimates, *Bull. Seis. Soc. Am.* **82**, 1.

Leadbetter, M. R., Lindgren, G. and Rootzen, H. (1983) *Extremes and related properties of random sequences and processes*. Springer-Verlag, New York.

Madan, D. and Unal, H. (1993) Pricing the risks of default, Manuscript, College of Business, University of Maryland.

Provost, P. and Rudiuk, E. (1996) The exact distribution of indefinite quadratic forms in no-central normal vectors, *Annals of the Institute of Statistical Mathematics* **48**(2).

Pye, G. (1974) Gauging the default premium, *Financial Analysts Journal*, Jan–Feb.

Resnick, S. I. (1987) *Extreme Value, Regular Variation and Point Processes*, Springer-Verlag, New York.

Taleb, N. (1997) Against value at risk, Working Paper, Taleb Research.