Holger Rootzén and Claudia Klüppelberg

A Single Number Can't Hedge Against Economic Catastrophes

Mathematics and statistics have transformed day-to-day trading in the world's financial markets. This has led to new ways to reduce (or "hedge") risks, which provide an important service to society, but also a temptation to very big gambles, with a potential for extreme losses. This paper discusses some of the ways statistics and mathematics can be used to understand and protect against very large, "catastrophic" financial risks. We argue that means don't mean anything for catastrophic risk, that separate large financial risks often are better handled by separate companies, and that the mathematical aspects of risk can't be summarized into one number. We also believe that there is a large potential for improved risk management in financial institutions, where extreme value theory, a speciality of the present authors, may be a useful tool. Improvements, however, will not come for free, but require long and hard work, where mathematics is only one part of the total effort.
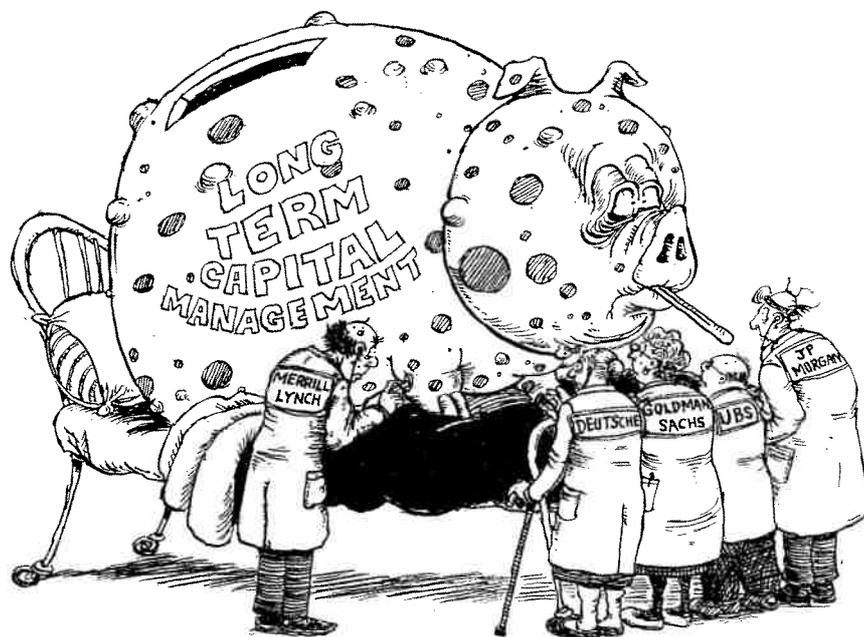

Figure 1. Long-term sickness? Reprinted from the *Economist* (3), with permission of the artist, D. Simmonds.

INTRODUCTION

On 23 September 1998, a consortium of banks and investment firms decided to spend USD 3.65 billion to rescue the Connecticut based hedgefund LTCM, Long Term Capital Management. The consortium included the Union Bank of Switzerland, which chipped in USD 300 million, and may have lost more than USD 650 mill. through direct dealings with LTCM, and investment giants Merrill Lynch, J.P. Morgan and Goldman Sachs, and the decisions were presumably taken under heavy pressure from the Federal Reserve Bank of New York.

LTCM had started in 1994 and used extensive computer modelling and sophisticated "hedging" strategies to trade in financial derivatives. It had produced spectacular profits, 43% in 1995 and 41% in 1996, and investors were lining up to be allowed to enter. The fund had almost completely free hands to invest as it chose without even insight from investors (or the public) and money couldn't be withdrawn until after 3 years. Instead, unfettered use of the most advanced financial technology in existence was promised. Parts of the fund's attractiveness undoubtedly was that two winners of the Nobel Prize in economics, Merton and Scholes, were partners, and were rumored to have invested some of their own (Nobel Prize?) money. However, in a few months in the beginning of the fall of 1998 the fund came from glory to the verge of bankruptcy, and the Federal Reserve stepped in to hinder a collapse which it believed could threaten the stability of the world financial system. For more of the story, see Coy and Wooley (1), Coy (2), Anon. (3), and Wardle (4).

LTCM traded in a wide range of derivatives and other financial instruments. Derivatives perform a very important service by providing companies and individuals with protection against future changes in prices and exchange rates. Often the risks in trading with individual derivatives are small, and the possible gains are correspondingly small. Hedgefunds, nevertheless, sometimes are able to make spectacular gains by investing not only their own capital, but also borrowed capital—LTCM had borrowed more than 50 times its own capital. Such borrowing, to get increased "leverage", is common, but the size of LTCM's leverage was not. However, borrowing doesn't only increase the size of possible gains, risks also become much larger. LTCM, as other hedgefunds, was betting on temporary disparities in the prices of related assets. This time a main factor simply was playing the 4% spread between the ruble-dominated Russian Treasury bills and the lower cost of borrowing rubles from banks. For LTCM disaster struck when the Russians halted trading in their domestic government debt market.

Table 1. Examples of financial crises 1618–1900. For a complete list until 1990 and more details see Kindleberger (5).

| Time	Related to:	Preceding speculation in:	Countries
1618–1623	Thirty-years War	subsidiary coin	Holy Roman Empire
1636–1637	boom in war against Spain	tulip bulbs	Dutch Republic
1720	Treaty of Utrecht	South Sea Company stock	England
	death of Louis XIV	Mississippi Company	France
1793	reign of terror in France	canal mania	England
1825	success of Baring loan,	Latin American bonds,	England
	decline of interest rates	mines, cotton	
1848	potato blight, wheat failure	railways, wheat	European Continent
1890	Sherman Silver Act	silver, gold	US
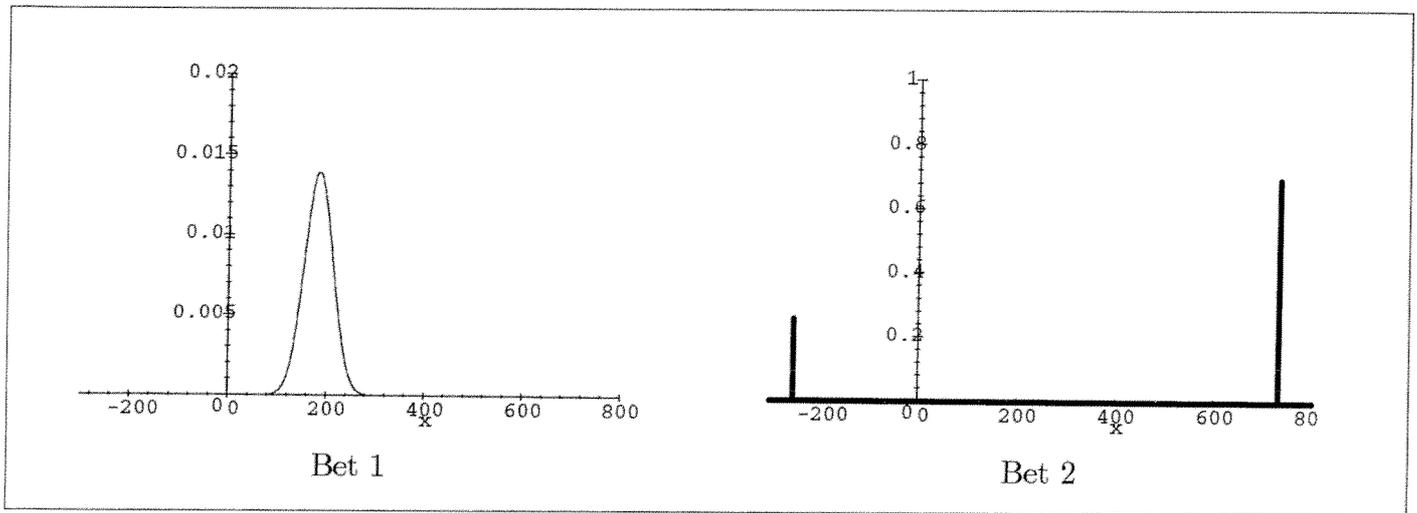
Figure 2. The distribution of total gain in two bets with the same mean.

In the aftermath, headlines like "Can you devise surefire ways to beat the markets? The rocket scientists thought they could. Boy were they wrong" (1, *Business Week*) appeared. However, later on the article also says "In a sense, maybe the problem wasn't too much rocket science but too little". In the quotes, "rocket science" means mathematical modelling, but also alludes to the substantial number of university physicists employed by the financial industry in recent years.

Other papers, among the 6 presented in this issue of Ambio, deal with risks which pose direct threats to health and life. Financial crises could be shrugged off with "it's only money". However, in reality they make our entire society poorer, and may in the final analysis pose more serious threats to health and life than technical or natural disasters.

LTCM is not the first financial disaster, but one in a very long row (cf. Table 1), and not the last one either. In this article we discuss relations between mathematics and statistics and the practical problems of risk management in financial institutions, with a hope to indicate some ways that mathematics can be useful, and when it can't, and to give examples of the interesting interplay between mathematics and the real world. Our main assertions are that risk is complex and cannot be summarized into one number, and that society isn't always better served by having very large enterprises handling big risks. The first assertion is accepted practice in many areas of risk analysis, and also often understood by people in the financial industry. Still, in our belief, it is poorly reflected in current practise and regulation. The second one, we think is much less understood, both practically and mathematically.

Better handling of financial risk will require long, determined and costly work on many fronts, mathematics being only one of those. In particular, better training and better systems for handling and disseminating information, both within firms and between firms and the public, are in our opinion of vital importance. There are no quick fixes.

CATASTROPHES DO HAPPEN

Even very unlikely events will occur, given enough time. Of course, we all know this. Time after time the catastrophes happen which were thought could never be. This in fact also is a theorem of mathematics (see 6).

Paul Lévy's zero-one law: *If the risk of a "catastrophe" tomorrow is greater than p percent (where p may be as small as 0.1% or even smaller) every day, then the catastrophe will happen (sooner or later).*

A much discussed measure of financial risk is VaR, Value-at-Risk, defined as *the p-th quantile of the distribution of the possible loss*. The so-called Basle accord (7, 8) proposed that capital requirement for individual trading desks should be set as follows. VaR for the profit/loss for for the next 10 days and $p = 1\%$, based on a historical observation period of at least 1 year (220 days) of data, should be computed and then multiplied by 3. Here the "safety factor" 3 is because it was realized that the commonly used normal distribution is unrealistic, and that 220 days history is altogether too little. This is well documented by risk managers and regulators through backtesting procedures, which are regularly performed; see e.g. Overbeck and Stahl (9).

The philosophy is that it is OK to take small enough risks even though they in the end (but hopefully only after a long while) by Paul Levy's zero-one law will end in a catastrophe. Indeed this is both right and inevitable. It is still useful to try to understand the consequences. Rather recent mathematics give a quite precise guide (see e.g. Leadbetter et al. (10) and Embrechts et al. (11)).

Compound Poisson Limit Theorem: *Let the probability that a large loss happens tomorrow be p% each day, and assume large losses occur in small clusters of mean length l. Then the probability that at least one large loss occurs in a period of T days is* $1-\exp(-Tp/l)$.

Suppose one each day computes a VaR correctly and sets aside enough money to cover VaR. E. g. using the Basle "VaR" might in a (hypothetical) situation correspond to a real one-day VaR at $p = 0.1\%$, and a typical cluster length, say a period corresponding to large fluctuations in exchange rates, might be 3 days. Then the formula gives that the risk that the capital set aside is exceeded by the loss at least once in one year is 7.6% (= 1–exp(–220 x 0.001/3)).

MEANS DON'T MEAN ANYTHING

The omnipresence of means is based on the law of large numbers—observed averages will be close to means if an experiment is repeated very many times. However, this is precisely not the case for catastrophic losses. Once incurred then the experiment which led to the loss will not be repeated.

Consider a choice between two investments. *Bet 1* is that you win 2 with probability 0.5 (of course, here 2 might mean USD 2 million), and lose 1 with probability 0.5, for each of 365 days. In *Bet 2* you instead gain 2 with probability 0.9985 and lose 1000 with probability 0.0015 each day. The gains have (up to round-

ing errors) the same mean, 365 x 0.5 in both situations, and for *Bet 1* the distribution is nicely concentrated around the mean 182.5 (Fig. 2). However, the second strategy will lead to a very different conclusion, either one loses 272 or (much) more, which happens with probability 42% or wins 730 with probability 58%. Clearly, for *Bet 2* the mean provides very little information. It is this situation which is the typical one for large-scale financial risk.

As illustrated by a variant of the celebrated Saint Petersburg paradox, risk may not even have a finite mean. Assume you are in a casino and play roulette, where red and black both have probability 1/2 (i.e. there is no 0 or 00). Now you always play "red" according to the following rule until the first time "red" happens: The first time you set $2 and win $2, if "red" comes up, if not you lose your stake. If this happens, next time you double your bet, i.e. $4, next time $8, etc. This means that at the n-th bet you may win 2^n with a probability 2^{-n}. What is the result of the game? Your mean gain is $\sum_{n=1}^{\infty} 2^{-n}2^n = \infty$.
Thus, from the view of the casino, the mean loss is infinite. However, this is not a worry for real casinos who do not have infinitely rich customers.

Even if means exist, they may be quite unstable. A very widely used model (the Pareto model) intended to capture catastrophic behavior is that the risk of a loss which exceeds the (extreme) level x has probability cx^{-a}, for some parameters c, a. Then the mean is finite if the parameter a is larger than 1, but infinite for $a \leq 1$. However, for $a > 1$ but close to 1 quite similar models may have very different means. This is illustrated in Figure 3, which shows two Pareto distributions which are very similar indeed, but where the means differ by about a factor 10.

Hence, for catastrophic risk, means don't mean anything. Of course, this is not a new observation, but nevertheless often forgotten. At a time when financial risk was measured solely by the mean square error as suggested by Gauss in the 18th century (Gauss used the method in astronomy, where measurements can be repeated many times, a situation where means make sense), Nobel Prize winner Harry Markowitz (12) proposed the so-called semivariance, where only one-sided deviations above a certain target value were considered. Fishburn (13) generalized this to so-called partial moments, which in present day terminology are conditional moments. They are based on the understanding that big losses should be measured by quantities related to high quantiles, such as, in the technical jargon of the area, expected target shortfall, target semivariance and perhaps target semi-skewness. All these measures concentrate on the probability and conditional means of high exceedances.

AN AXIOMATIC APPROACH

Artzner et al. (14) consider the problem of risk measurement from an axiomatic point of view. The authors want to determine properties which a "good" risk measure should have, and also aim at giving rules on which risks are acceptable to investment managers and regulators. To this end, they introduce certain axioms which a "good" risk measure R should satisfy.

In a way, this approach is "typical" for a mathematician. We believe it is an important contribution which serves the good purpose of thinking carefully about the problem and formalizing it in a clear way. However, as will be seen below, we also don't agree with the conclusions drawn from from these axioms, and thus not with

the axioms either. Our arguments are illustrated by examples, some of them taken from Artzner et al. (14).

Among the axioms considered as the most important ones in Artzner et al. (14) are, for given risks X and Y

Axiom 1. Monotonicity: *If $X \leq Y$, then $R(X) \leq R(Y)$.*
This is definitely, what you expect from a risk measure, that it assigns a higher number to a bigger risk.

Axiom 2. Translation invariance: $R(X + a) = R(X)-a$
If you add some real money, then the risk should decrease by the same amount.

Axiom 3. Positive homogeneity: $R(tX) = tR(X)$ *for any positive real number t.*
If you buy the same risk twice, this should simply double your risk.

Axiom 4. Subadditivity: $R(X + Y) \leq R(X) + R(Y)$
A big portfolio of different risks diversifies risk: although one asset may have a negative return, there may be assets to compensate for this.

If a risk measure R satisfies these four axioms, then it is called a *coherent risk measure*. The conclusion of Artzner et al. (14) is the following result, which in its present version requires some additional conditions.

The Artzner, Delbaen, Eber, Heath theorem: *All coherent risk measures may be obtained as means of "generalized scenarios".*
A generalized scenario is a set of posssible outcomes, with attached probabilities. The mathematical formulation is given by a set of probability measures P. Then, in formulas, for a risk X,

$$R(X) = \sup\{Ep(X)|P \in \mathsf{P}\}$$

is a coherent risk measure.
For example, for P, which only contains one probability measure P the mean satisfies Axioms 1–4 and is hence a coherent risk measure. On the other hand, the more scenarios are considered, the more conservative (i.e. larger) is the risk measure obtained.

Figure 3. The probability of a loss larger than x, for two different Pareto distributions, with c = 0.105, a = 1.1 and c = 0.99, a = 1.01. The two distributions have means 1.2 and 10, respectively.
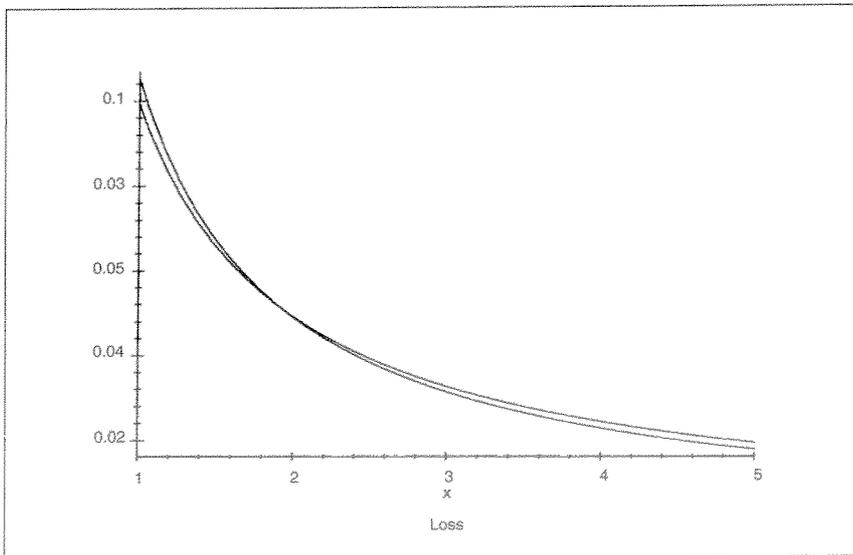
Will the supremum over all possible probability measures $P \in \mathbf{P}$, i.e. over all possible scenarios provide the correct measure of risk? Even if it does, it may be very difficult to think up all relevant scenarios for our specific risk. To give an example, sources say that LTCM worst case scenario was only about 60% as bad as the one that actually occurred (1).

One conclusion in Artzner et al. (14) (which we don't agree with) is that VaR is not a good risk measure, since it sometimes can violate the subadditivity required by Axiom 4. An example of a risk measure which comes from a generalized scenario and at the same time is one of Fishburn's high quantile risk measures is the expected shortfall with target VaR:

$$R(X) = E(X|X \leq \text{Var}(X)).$$

It fits into the above framework of the theorem by choosing the set, and hence is a coherent risk measure.

$$\mathbf{P} = \{P_A(\cdot) = P(\cdot|A) ; A \subset \Omega, P(A) \geq \infty\}$$

BIG ISN'T ALWAYS BEAUTIFUL

A standard useful way to reduce (ordinary) risks is to diversify. In the axiomatic theory of Artzner et al. (14) this is translated into Axiom 4, which says "big is beautiful". In this section we argue that for really catastrophic risks the opposite may often be the truth.

Again we first consider an example where a firm has a choice between two propositions, or "bets", both with zero mean gain. In *Bet 1* the firm gains 1 with probability 0.99 and loses 99 with probability 0.01 (and then is ruined). *Bet 2* simply consists of two independent bets of the same kind, which means that the gain is 2 with probability $0.99^2 \approx 0.98$ and that 99 or more is lost if at least one of the bets fails, which happens with probability approximately equal to 0.02. The probability of ruin is doubled in the second situation.

Next, let us again consider the Pareto model: the probability that the risk X exceeds a high level x is modelled by $P(X > x) = cx^{-a}$ so that the mean is finite for $a > 1$, but infinite for $a \leq 1$. Assume we have two independent Pareto risks, e.g. two hedge fonds, oil platforms or tankers. The subadditivity axiom suggests that a portfolio of two risks in one company always is preferable to having these risks in two different companies.

A mathematical property of two such Pareto risks X and Y is that, if independent and added, then, for large values x,

$$P(X + Y > x) \approx P(X > x) + P(Y > x),$$

i.e. the probability that the two risks together exceed some value x is approximately equal to the sum of the probabilities that the first risk is bigger than x and the probability that the second risk is bigger than x. If we take the $p\%$-quantile as a risk measure, then $R(X) = \text{VaR}(X,p)$ is given by that value x_p such that $1-p = P(X > x_p) = cx_p^{-a}$. This is just the simplest definition of VaR. Easy calculation gives

$$\text{VaR}(X,p) = (c/(1 - p))^{1/a}.$$

Checking Axioms 1–4 we see that Axioms 1–3 are satisfied— in fact,they are satisfied for VaR in general, independent of the model for X. Axiom 4 is more problematic:

$$\text{VaR}(X + Y,p) \approx (2c/(1 - p))^{1/a} = 2^{1/a}(c/(1 - p))^{1/a}$$

and

$$\text{VaR}(X,p) + \text{VaR}(Y,p) = (2c/(1 - p))^{1/a}$$

If $a < 1$ then $2^{1/a} > 2$ and hence VaR then doesn't satisfy

Axiom 4. For instance for $c = 1$ and $a = 1/2$ we obtain $\text{VaR}(X + Y,p) \approx 2(\text{VaR}(X,p) + \text{VaR}(Y,p))$ and for $c = 1$ and $a = 1/4$ we have that $\text{VaR}(X + Y,p) \approx 8(\text{VaR}(X,p) + \text{VaR}(Y,p))$. Further numerical illustration is given in Table 2. The same picture would emerge if other values of p were considered, or indeed for the entire loss distribution function.

In practical terms, this means that for Pareto risks with $a < 1$ overall risk is increased more than proportionally by taking on two independent risks. This might happen if an insurance company admits two oil platforms from completely different parts of the world into its portfolio.

These examples argue against the general usefulness of the subadditivity Axiom (15). In Artzner et al. (14) the authors claim that it is a natural requirement that "a merger does not create extra risk". They support their statement by some examples:
  i) If a firm were forced to meet a requirement of extra capital, which did not satisfy this property, the firm might be motivated to break up into two separately incorporated affiliates.
  ii) Bankruptcy risk inclines society to require less capital from a group without "firewalls" between various business units than it does require when one "unit" is protected from liability attached to failure of another "unit".
  iii) Suppose that two desks in a firm compute in a decentralized way the measures $R(X)$ and $R(Y)$ of the risks they have taken. If the function R is subadditive, the supervisor of the two desks can count on the fact that $R(X) + R(Y)$ is a feasible guarantee relative to the global risk $X + Y$.

We believe that these arguments are wrong for catastrophic business risk. The arguments for *i)* and *iii)* seem to be that it would be convenient for firms and supervisors if risk actually behaves according to Axiom 4. However, what is convenient and what is true isn't necessarily the same.

In the past, we have experienced dangers of bankruptcy with Baring's Bank, Metallgesellschaft, LTCM, and many others. If each such big risk is taken on by a different company, then if the catastrophe actually occurs only one of the companies is in the risk zone. Thinking of the economic and social consequences of the bankruptcy of a very large firm this may often be preferable. Moreover, whereas the bankruptcy of a very large firm is politically intolerable, so that the taxpayer has to come to the rescue, a smaller firm may well disappear from the market, in particular, if it becomes obvious that it has failed to set up a proper risk management.

The claim *ii)* goes against long experience in the insurance industry of the desirability of firewalls between different parts of the businesses of handling risk.

On a more technical level we have learned above that the VaR violates the subadditivity axiom of Section 4 above. This has also been indicated in Artzner et al. (14). Does this mean that one should abandon VaR completely? We instead believe that several quantiles for different levels and time horizons give a realistic and very useful picture of the risk.

ONE NUMBER DOESN'T SUFFICE

In the previous sections, we have argued that for catastrophic business risks, means don't mean anything, and that big isn't nec-

Table 2. Comparison of VaR for the sum of two Pareto risks with c = 1 and p = 95%.

a	VaR(X,95%) = VaR(Y,95%)	VaR(X + Y,95%)
1/2	400	1600
1/3	8000	64 000
1/4	160 000	2 560 000

Figure 4. DAX closing prices during 29.8.95–26.8.96 (250 data points in total). The corresponding differences, which are the daily price changes, are plotted in the right-hand graph.
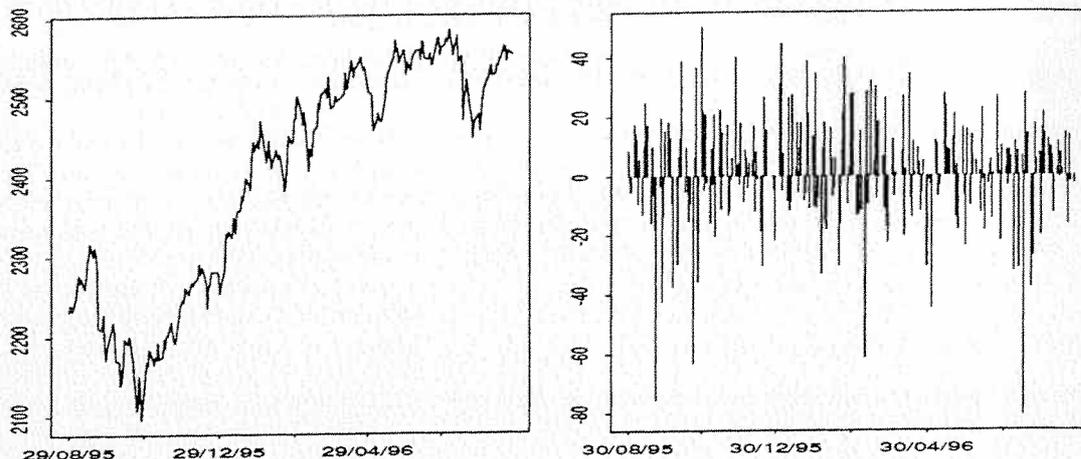
essarily beautiful. However, our main disagreement with Artzner et al. (14) is the basic underlying assumption that all (mathematical) aspects of risk can be measured by a single number. We simply don't believe that this can be done in a useful way.

As an illustration of calculations we consider a Pareto model for a real example. (A similar analysis in an insurance context may be found in Rootzén and Tajvidi (16)). The use of Pareto models is not an *ad hoc* idea, but based on a mathematical theory for the extreme values and on considerable practical testing; see Leadbetter et al. (10) or Embrechts et al. (11) for mathematical and statistical background. Mathematical thinking compensates the unavoidable drawback of little empirical experience from extremes (only the largest losses of the sample are responsible for the extreme risk) by a plausible model. In particular it provides ways to extrapolate beyond the data at hand, which is required of any realistic way of treating rare and extreme events. It is a completely constant experience—and in accordance with the Paul Lévy zero one law that sooner or later risks exceeding any past risk do occur.

The example is built on the DAX (German share index) data of Figure 4. The investor has to worry about losses, which are here the negative daily price changes. The estimated VaR is based on the so-called *peaks over threshold method*; see Embrechts et al. (11, Chapter 6). The data are not catastrophic in our sense, but might have been, at higher aggregation levels and longer time horizons. Table 3 presents the calculation of VaR for different percentages *p* and 2 different time horizons. The 10-day VaR corresponds to the regulator requirement, which is based on the idea that reserves should be built to cover losses for a time horizon needed to restructure a portfolio. This example has been treated in detail in Emmer et al. (17).

Obvious problems are that the time period covered by the data might be too short to include any real extreme event, and that the possibility of future structural changes may limit the usefulness of historical data. However, as for other risk measures, these difficulties may, at least to some extent, be dealt with by using "scenarios". VaRs like those of Table 3 in our opinion give a clear and useful measure of risk, and, if needed, a good basis for regulation.

Finally, we want to compare these figures to today's normal practice in industry. We do this for the 5% VaR corresponding to Figure 5. Taking the normal quantile as an estimate gives VaR(*p*%, *1 day*) = 29.90 which is much smaller than our estimate VaR(*p*%, *1 day*) = 42.86. On the other hand, the 10-day VaR would be estimated by the scaling property of the (zero mean) normal distribution, by multiplying the one-day VaR by

√10. This gives VaR(*p*%, *10 days*) = 94.55 which is much larger than our estimate VaR(*p*%, *10 days*) = 65.79. A bank trying to fit the proper model of a Pareto distribution to the daily data and then use the scaling constant to get the 10-day VaR would even get VaR(p%, *10 days*) = 135.53 giving a conservative upper bound for VaR(p%, 10 *days*). This has been pointed out by Stahl (18) and Danielsson et al. (19).

MARKETS CAN BE MADE SAFER

The Black-Scholes theory of option pricing and its descendants have transformed trading in financial risk. The importance is hard to overestimate. However, its main concern is with individual risks from a micro-perspective and it has little to say about risk on a macro scale, where bankruptcy of major firms may be at stake. The present paper is aimed at the latter kind of risks. Our main assertions are that

– means don't mean anything for catastrophic risks;
– catastrophic financial risks don't become smaller if collected in one big company and, the most important one;
– a single number can't capture all the different aspects of catastrophic financial risk.

To elaborate on the last part, the final comprehensive picture of financial risk would be to have the joint distributions of the possible losses at many different levels of aggregation and time horizons. This is never completely possible. However, condensed versions, like Table 3 above, and much more sophisticated future versions, based on data analysis, theoretical reasoning and analysis of alternative scenarios, are quite feasible, and in our belief will be very useful.

We also believe that quantitative methodology for financial risk management is capable of substantial improvement in the near future. Part of the improvement may be provided by actuarial science and by a speciality of the authors of this article—statistical extreme value theory. A comprehensive account of the present state of the art may be found in Embrechts et al. (11).

Table 3. Calculation of the Value-at-Risk for different percentages and time horizons.		
p	VaR(p%, 1 day)	VaR(p%, 10 days)
10	31.42	48.23
5	42.86	65.79
1	74.83	114.86

Figure 5. Comparison of different 5%-quantile estimates, with fitted normal density and generalized Pareto density in the left tail region.

Labels in figure: -30.654 (empirical); -29.823 (Normal); -42.856 (GPD); axis "DAX daily returns"
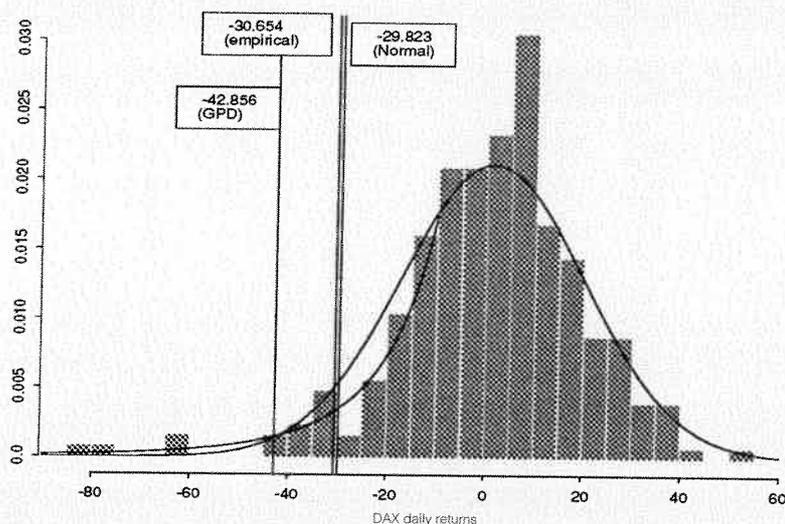
This theory is specifically aimed at modelling and analyzing rare and extreme events, such as those involved in catastrophic financial risk. However, much research remains to be done, e.g. in development of better methods for multivariate and dependent situations, in understanding the effect of different levels of aggregation and in adapting the methods to the specifics of financial risk.

In the previous sections, we tried to illustrate how mathematics and statistics are put to work in a practical situation. One face is general theory founded on theorems and axioms as a means to bring understanding to a complex and bewildering reality. The other face, and often the most important one, is statistical analysis of data and computer modelling to understand the detailed aspects of the problem.

Both general mathematics and detailed statistical modelling involve dangers: Improperly understood or inappropriate theory may become a straightjacket rather than liberator for clear thought, and statistics instead of giving understanding and detailed knowledge can drown in a flood of formulas and numbers and may put you at the mercy of complex and dubious computer programs. There are no automatic benefits from using mathematics and statistics in financial risk management. However, careful, thoughtful and persistent work on improvement of procedures and methods will make risks less risky.

It is important to understand that catastrophic risk always involves an element of gambling. What mathematics can do is to help find the odds of the gamble. Whether to gamble or not is outside the realm of mathematics.

Quantitative measures, of course, are only part of the management of financial risk. Good procedures and systematic ways of containing and reducing risk are of central importance. Similar risk-handling problems permeate manufacturing and energy producing industries. Reliability technology has been spectacularly successful in ensuring high safety in extremely complex systems, such as the world telephone network. The systematic approaches developed in reliability theory to manage risks have a large potential for reducing financial risk. Examples of such methods are FMEA (Failure Mode and Effects Analysis), a systematic approach to detecting risks before they have occurred and FTA (Fault Tree Analysis), which also attempts to provide numerical estimates of risk (cf. 20).

Financial risk offers a host of important challenges, both for quantitative methods and for the many other basic issues in risk management. We hope these challenges will be taken up in long and hard work by financial institutions and by university researchers.

References

1. Coy, P. and Wooley, S. 1998. Failed wizards of Wall Street. *Business Week*, Sept. 21, 1998, 54–59.
2. Coy, P. 1998. How the game was played—and why long-term lost. *Business Week*, Oct. 12, 1998, 40–41.
3. Anon. 1998. The rescue of long-term capital management raises disturbing questions about the health of America's and the world's leading banks. *The Economist*, Oct. 3, 1998, 93–97.
4. Wardle, I. 1998. LTCM, a hedge fund above suspicion. Le Monde diplomatique. English internet edition, November 1998.
5. Kindleberger, C.P. 1996. *Manias, Panics and Crashes. A History of Financial Crises*, 3rd Ed. Wiley, New York.
6. Chung, K.L. 1974. *A Course in Probability Theory*. 2nd Ed. Springer, Berlin.
7. Basle Committee on Banking Supervision. 1995. *The Supervisory Treatment of Market Risks*. Bank for International Settlements. Basle, Switzerland.
8. Basle Committee on Banking Supervision. 1996. *Amendment to the Capital Accord to Incorporate Market Risks*. Bank for International Settlements. Basle, Switzerland.
9. Overbeck, L. and Stahl, G. 1999. Backtesting. In: *Handbuch des Risikomanagements*. Rudolph, B. (ed.). Uhlenbruch Verlag. (In print).
10. Leadbetter, M.R., Lindgren, G. and Rootzén, H. 1983. *Extremes and Related Properties of Random Sequences and Processes*. Springer, Berlin.
11. Embrechts, P., Klüppelberg, C. and Mikosch, T. 1997. *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin.
12. Markowitz, H. 1959. *Portfolio Selection—Efficient Diversification of Investments*. Wiley, New York.
13. Fishburn, P.C. 1977. Mean-risk analysis with risk associated with below-target returns. *Am. Econ. Rev. 67*, 116–126.
14. Artzner, P., Delbaen, F., Eber, J.M. and Heath, D. 1999. Coherent measures of risk. Preprint. *Math. Finance 9*, 203–228.
15. Bernstein, L.P. 1996. *Against the Gods*. Wiley, New York.
16. Rootzén, H. and Tajvidi, N. 1997. Extreme value statistics and wind storm losses: a case study. *Scand. Actuarial J. 1*, 70–94.
17. Emmer, S., Klüppelberg, C. and Trüstedt, M. 1998. VaR—ein Mass für das extreme Risiko. *Solutions 2*, 53–63. English version available at http://www.ma.tum.de/stat/Papers/papers.html
18. Stahl, G. 1997. Three cheers. *Risk*, May 1997, 67–69.
19. Danielsson, J. Hartmann, P. and de Vries, C. 1998. The cost of conservatism. *Risk*, Jan. 1998, 101–103.
20. Høyland, A. and Rausand, M. 1994. *System Reliability Theory: Models and Statistical Methods*. Wiley, New York.

Holger Rootzén is professor of mathematical statistics at Chalmers University of Technology, Gothenburg. His main interests are stochastic processes and statistics, extreme values, image analysis, applications in engineering, medicine and industry. His address: Department of Mathematical Statistics, Chalmers University of Technology, SE-41296 Gothenburg, Sweden. www.math.chalmers.se/rootzen

Claudia Klüppelberg holds the Chair of mathematical statistics at the Department of Mathematics of Munich University of Technology. Her main interests are modelling and analysis of financial time series for risk management. Her address: Center for Mathematical Sciences, Munich University of Technology, D-80290 Munich, Germany. www.ma.tum.de/stat/