

Introduction

Browsing quickly through the almost 400 pages that follow, it will become immediately clear that this book seems to have been written by mathematicians for mathematicians. And yet, the title has the catchy “High Risk Scenarios” in it. Is this once again a cheap way of introducing finance related words in a book title so as to sell more copies? The obvious answer from our, the authors’ point of view, must be no. This rather long introduction will present our case of defense: though the book is indeed written by mathematicians for a mathematically inclined readership, at the same time it grew out of a deeper concern that *quantitative risk management* (QRM) is facing problems where new mathematical theory is increasingly called for. It will be difficult to force the final product you are holding in your hands into some specific corner or school. From a mathematical point of view, techniques and results from such diverse fields as stochastics (probability and statistics), analysis, geometry and algebra appear side by side with concepts from modern mathematical finance and insurance, especially through the language of portfolio theory. At the same time, risk is such a broad concept that it is very much hoped that our work will eventually have applications well beyond the financial industry to areas such as reliability engineering, biostatistics, environmental modelling, to name just a few.

The key ingredients in most of the theory we present relate to the concepts of risk, extremes, loss modelling and scenarios. These concepts are to be situated within a complex random environment where we typically interpret complexity as high-dimensional. The theory we present is essentially a one-period theory, as so often encountered in QRM. Dynamic models, where time as a parameter is explicitly present, are not really to be found in the pages that follow. This does not mean that such a link cannot be made; we put ourselves however in the situation where a risk manager is judging the riskiness of a complex system over a given, fixed time horizon. Under various assumptions of the random factors that influence the performance of the system the risk manager has to judge today how the system will perform by the end of the given period. At this point, this no doubt sounds somewhat vague, but later in this introduction we give some more precise examples where we feel that the theory as presented may eventually find natural applications.

A first question we would like to address is

“*Why we two?*”

There are several reasons, some of which we briefly like to mention, especially as they reflect not only our collaboration but also the way QRM as a field of research and applications is developing. Both being born in towns slightly below or above sea level, Amsterdam and Antwerp, risk was always a natural aspect of our lives. For the

second author this became very explicit as his date of birth, February 3, 1953, was only two days after the disastrous flooding in Holland. In the night of January 31 to February 1, 1953, several 100 km of dykes along the Dutch coast were breached in a severe storm. The resulting flooding killed 1836 people, 72 000 people needed to be evacuated, nearly 50 000 houses and farms and over 200 000 ha of land were flooded. A local newspaper, *De Yssel- en Lekstreek*, on February 6, 1953 ran a headline “Springtij en orkaan veroorzaken nationale ramp. Nederland in grote watersnood”¹. The words of the Dutch writer Marsman from 1938 came back to mind: “En in alle gewesten, wordt de stem van het water, met zijn eeuwige rampen, gevreesd en gehoord.”² As a consequence, the Delta Project came into being with a clear aim to build up a long-lasting coastal protection through an elaborate system of dykes and sluices. Though these defense systems could never guarantee 100% safety for the population at risk, a safety margin of 1 in 10 000 years for the so-called Randstad (the larger area of land around Amsterdam and Rotterdam) was agreed upon. Given these safety requirements, dyke heights were calculated, e.g. 5.14 m above NAP (Normaal Amsterdams Peil). A combination of environmental, socioeconomic, engineering and statistical considerations led to the final decision taken for the dyke and sluice constructions. For the Dutch population, the words of Andries Vierlingh from the book *Tractaet van Dyckagie* (1578) “De meeste salicheyt hangt aen de hooghte van eenen dyck”³ summarized the feeling of the day. From a stochastic modelling point of view, the methodology entering the solution of problems encountered in the *Delta Project* is very much related to the analysis of extremes. Several research projects related to the modelling of extremal events emerged, examples of which include our PhD theses Balkema [1973] and Embrechts [1979]. Indirectly, events and discussions involving risk and extremes have brought us together over many years.

By now, the stochastic modelling of extremes, commonly referred to as *Extreme Value Theory* (EVT), has become a most important field of research, with numerous key contributors all over the world. Excellent textbooks on the subject of EVT exist or are currently being written. Moreover, a specialized journal solely devoted to the stochastic theory of extremes is available (*Extremes*). Whereas the first author (Balkema) continued working on fundamental results in the realm of heavy tailed phenomena, the second author (Embrechts) became involved more in areas related to finance, banking, insurance and risk management. Banking and finance have their own tales of extremes. So much so that Alan Greenspan in a presentation to the Joint Central Bank Research Conference in Washington D.C. in 1995 stated (see Greenspan [1996]):

“From the point of view of the risk manager, inappropriate use of the normal distribution can lead to an understatement of risk, which must be

¹“Spring tide and hurricane cause a national disaster. The Netherlands in severe water peril.”

²“And in every direction, one hears and fears the voice of the water with its eternal perils.”

³“Most of the happiness depends on the height of a dyke.”

balanced against the significant advantage of simplification. From the central bank's corner, the consequences are even more serious because we often need to concentrate on the left tail of the distribution in formulating lender-of-last resort policies. Improving the characterization of the distribution of extreme values is of paramount concern."

Also telling is the following statement taken from *Business Week* in September 1998, in the wake of the LTCM hedge fund crisis:

"Extreme, synchronized rises and falls in financial markets occur infrequently but they do occur. The problem with the models is that they did not assign a high enough chance of occurrence to the scenario in which many things go wrong at the same time – the 'perfect storm' scenario."

Around the late nineties, we started discussions on issues in QRM for which further methodological work was needed. One aim was to develop tools which could be used to model markets under extreme stress scenarios. The more mathematical consequence of these discussions you are holding in your hands.

It soon became clear to us that the combination of extremes and high dimensions, in the context of scenario testing, would become increasingly important. So let us turn to the question

"Why in the first part of the title High Risk Scenarios and Extremes?"

The above mentioned Delta Project and QRM have some obvious methodological similarities. Indeed protecting the coastal region of a country from sea surges through a system of dykes and sluices can be compared with protecting the financial system (or bank customers, insurance policy holders) from adverse market movements through the setting of a sufficiently high level of regulatory risk capital or reserve. In the case of banking, this is done through the guidelines of the Basel Committee on Banking Supervision. For the insurance industry, a combination of international guidelines currently under discussion around Solvency 2 and numerous so-called local statutory guidelines have been set up. The concept of dyke height in the Delta Project translates into the notion of risk measure, in particular into the widely used notion of *Value-at-Risk* (VaR). For instance, for a given portfolio, a 99% 10-day VaR of one million euro means that the probability of incurring a portfolio loss of one million euro or more by the end of a two-week (10 trading days) period is 1%. The 10 000 year return period in the dyke case is to be compared with the 99% confidence level in the VaR case. Both sea surges and market movements are complicated functions of numerous interdependent random variables and stochastic processes. Equally important are the differences. The prime one is the fact that the construction of a dyke concerns the modelling of natural (physical, environmental) processes, whereas finance (banking) is very much about the modelling of social phenomena. Natural events may enter as triggering events for extreme market movements but are seldom

a key modelling ingredient. An example where a natural event caused more than just a stir for the bank involved was the Kobe earthquake and its implications for the downfall of Barings Bank; see Boyle & Boyle [2001]. For the life insurance industry, stress events with major consequences are pandemics for instance. Also relevant are considerations concerning longevity and of course market movements, especially related to interest rates. Moving to the non-life insurance and reinsurance industry, we encounter increasingly the relevance of the modelling of extreme natural phenomena like storms, floods and earthquakes. In between we have for instance acts of terrorism like the September 11 attack. The “perfect storm scenario” where many things go wrong at the same time is exemplified through the stock market decline after the New Economy hype, followed by a longer period of low interest rates which caused considerable problems for the European life insurance industry. This period of economic stress was further confounded by increasing energy prices and accounting scandals.

In order to highlight more precisely the reasons behind writing these lectures, we will restrict our attention below to the case of banking. Through the Basel guidelines, very specific QRM needs face that branch of the financial industry. For a broad discussion of concepts, techniques and tools from QRM, see McNeil, Frey & Embrechts [2005] and the references therein. Besides the regulatory side of banking supervision, we will also refer to the example of portfolio theory. Here relevant references are for instance Korn [1997] and Fernholz [2002]. Under the *Basel guidelines* (www.bis.org/bcbs) for market risk banks calculate VaR; this involves a holding period of 10 days at the 99% confidence level for regulatory (risk) capital purposes and 1 day 95% VaR for setting the bank’s internal trading limits. Banks and regulators are well aware of the limitations of the models and data used so that, besides the inclusion of a so-called multiplier in the capital charge formula, banks complement their VaR reporting with so-called stress scenarios. These may include larger jumps in key market factors like interest rates, volatility, exchange rates, etc. The resulting question is of the “what if”-type. *What* happens to the bank’s market position *if* such an extreme move occurs. Other stress scenarios may include running the bank’s trading book through some important historical events like the 1987 crash, the 1998 LTCM case or September 11. Reduced to their simplest, but still relevant form, the above stress scenarios can be formalized as follows. Suppose that the market (to be interpreted within the CAPM-framework, say; see Cochrane [2001]) moves strongly against the holder of a particular portfolio. Given that information, what can be said about risk measurement features of that portfolio. Another relevant question in the same vein is as follows. Suppose that a given (smaller) portfolio moves against the holder’s interest and breaches a given risk management (VaR) limit. How can one correct some (say as few as possible) positions in that portfolio so that the limit is not breached anymore. For us, motivating publications dealing with this type of problem are for instance Lüthi & Studer [1997], Studer [1997] and Studer & Lüthi [1997].

The *high-dimensionality* within our theory is related to the number of assets in the portfolio under consideration. Of course, in many applications in finance, dimension reduction techniques can be used in order to reduce the number of assets to an effective dimensionality which often is much lower and indeed more tractable. The decision to be made by the risk manager is to what extent important information may have been lost in that process. But even after a successful dimension reduction, an effective dimensionality between five and ten, say, still poses considerable problems for the application of standard EVT techniques. By the nature of the problem extreme observations are rare. The curse of dimensionality very quickly further complicates the issue.

In recent years, several researchers have come up with high-dimensional (market) models which aim at a stochastic description of macro-economic phenomena. When we restrict ourselves to the continuous case, the multivariate normal distribution sticks out as the benchmark model par excellence. Besides the computational advantages for the calculation of various relevant QRM quantities such as risk measures and capital allocation weights, it also serves as an input to the construction of more elaborate models. For instance the widely used *Student t* model can be obtained as a random mixture of multivariate normals. Various other examples of this type can be worked out leading to the class of elliptical distributions as variance mixture normals, or beyond in the case of mean-variance mixture models. Chapter 3 in McNeil, Frey & Embrechts [2005] contains a detailed discussion of elliptical distributions; a nice summary with emphasis on applications to finance is Bingham & Kiesel [2002]. A useful set of results going back to the early development of QRM leads to the conclusion that within the class of *elliptical models*, standard questions asked concerning risk measurement and capital allocation are well understood and behave much as in the exact multivariate normal case. For a concrete statement of these results, see Embrechts, McNeil & Straumann [2002]. A meta-theorem, however, says that as soon as one deviates from this class of elliptical models, QRM becomes much more complicated. It also quickly becomes context- and application-dependent. For instance, in the elliptical world, VaR as a risk measure is subadditive meaning that the VaR of a sum of risks is bounded above by the sum of the individual VaRs. This property is often compared to the notion of diversification, and has a lot to do with some of the issues we discuss in our book. As an example we briefly touch upon the current important debate on the modelling of *operational risk* under the Advanced Measurement Approach (AMA) which is based on the Loss Distribution Approach (LDA); once more, for detailed references and further particulars on the background, we refer to McNeil, Frey & Embrechts [2005]. For our purposes it suffices to realize that, beyond the well-known risk categories for market and credit risk, under the new Basel Committee guidelines (so-called Basel II), banks also have to reserve (i.e. allocate regulatory risk capital) for operational risk. According to Basel II, *Operational Risk* is defined as the risk of loss resulting from inadequate or failed internal pro-

cesses, people or systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk; for details on the regulatory framework, see www.bis.org/bcbs/. Under the LDA, banks are typically structured into eight business lines and seven risk categories based on the type of operational loss. An example is corporate finance (business line) and internal fraud (risk type). Depending on the approach followed, one has either a 7-, 8-, or 56-dimensional problem to model. Moreover, an operational risk capital charge is calculated on a yearly basis using VaR at the 99.9% level. Hence one has to model a 1 in 1000 year event. This by every account is *extreme*. The high dimensionality of 56, or for some banks even higher, is obvious. The subadditivity question stated above is highly relevant; indeed a bank can add up VaRs business line-wise, risk type-wise or across any relevant subdivision of the 8×7 loss matrix. A final crucial point concerns the reduction of these sums of VaRs taking “diversification effects” into account. This may (and typically does) result in a rather intricate analysis where concepts like risk measure coherence (see Artzner et al. [1999]), EVT and copulas (non-linear dependence) enter in a fundamental way. Does the multivariate extreme value theory as it is presented on the pages that follow yield solutions to the AMA-LDA discussion above? The reader will not find ready-made models for this discussion. However, the operational risk issue briefly outlined above makes it clear that higher dimensional models are called for, within which questions on extremal events are of paramount importance. We definitely provide a novel approach for handling such questions in the future. Admittedly, as the theory is written down so far, it still needs a considerable amount of work before concrete practical consequences emerge. This situation is of course familiar from many (if not all) methodological developments. Besides the references above, the reader who is in particular interested in the operational risk example, may consult Chavez-Demoulin, Embrechts & Nešlehová [2006] and Nešlehová, Embrechts & Chavez-Demoulin [2006]. From information on operational risk losses available so far, one faces models that are skew and (very) heavy-tailed. Indeed, it is the non-repetitive (low-frequency) but high-severity losses that are of main concern. This immediately rules out the class of elliptical distributions. Some of the models discussed in our book will come closer to relevant alternatives. We are not claiming that the theory presented will, in a not too distant future, come up with a useful 56-dimensional model for operational risk. What we are saying, however, is that the theory will yield a better understanding of quantitative questions asked concerning extremal events for high-dimensional loss portfolios.

Mathematicians are well advised to show humbleness when it comes to model formulation involving uncertainty, especially in the field of economics. In a speech entitled “Monetary Policy Under Uncertainty” delivered in August 2003 in Jackson Hole, Wyoming, Alan Greenspan started with the following important sentence: “Uncertainty is not just an important feature of the monetary policy landscape; it is the defining characteristic of that landscape.” He then continued with some sentences

which are occasionally referred to, for instance by John Mauldin, as *The Greenspan Uncertainty Principle*:

“Despite the extensive efforts to capture and quantify these key macro-economic relationships, our knowledge about many of the important linkages is far from complete and in all likelihood will always remain so. Every model, no matter how detailed and how well designed conceptually and empirically, is a vastly simplified representation of the world that we experience with all its intricacies on a day-to-day basis. Consequently, even with large advances in computational capabilities and greater comprehension of economic linkages, our knowledge base is barely able to keep pace with the ever-increasing complexity of our global economy.”

And further,

“Our problem is not the complexity of our models but the far greater complexity of a world economy whose underlying linkages appear to be in a continual state of flux... In summary then, monetary policy based on risk management appears to be the most useful regime by which to conduct policy. The increasingly intricate economic and financial linkages in our global economy, in my judgment, compel such a conclusion.”

For many questions in practice, and in particular for questions related to the economy at large, there is no such thing as *the model*. Complementary to the quotes above, one can say that so often the road towards finding a model is far more important than the resulting model itself. We hope that the reader studying the theory presented in this book will enjoy the trip more than the goals reached so far. We have already discussed some of the places we will visit on the way.

One of the advantages of modern technology is the ease with which all sorts of information on a particular word or concept can be found. We could not resist googling “High Risk-Scenarios”. Needless to say that we did *not* check all 11 300 000 entries which we obtained in 0.34 seconds. It is somewhat disturbing, or one should perhaps say sobering, that our book will add just one extra entry to the above list. The more correct search, keeping the three words linked as in the title of our book, yielded a massive reduction to an almost manageable 717. Besides the obvious connections with the economic and QRM literature, other fields entering included terrorism, complex real-time systems, environmental and meteorological disasters, biosecurity, medicine, public health, chemistry, ecology, fire and aviation, Petri nets or software development. Looking at some of these applications it becomes clear that there is no common understanding of the terminology. From a linguistic point of view, one could perhaps query the difference between “High-Risk Scenario” and “High Risk-Scenario”. Rather than doing so, we have opted for the non-hyphenated version. In its full length “High Risk Scenarios and Extremes” presents a novel mathematical

theory for the analysis of extremes in multi-dimensional space. Especially the econometric literature is full of attempts to describe such models. Relevant for our purposes are papers like Pesaran, Schuermann & Weiner [2004], Pesaran & Zaffaroni [2004], Dees et al. [2007], and in particular the synthesis paper Pesaran & Smith [2006] on the so-called global modelling approach. From a more mathematical finance point of view, Platen [2001] and Platen [2006], and Fergusson & Platen [2006] describe models to which we hope our theory will eventually be applicable. Further relevant publications in this context are Banner, Fernholz & Karatzas [2006] and Fernholz [2002].

In the preceding paragraphs, we explained some of our motivations behind the first part of the title: “High Risk Scenarios and Extremes”. The next, more mathematical question is

“Why the second part of the title, A Geometric Approach?”

A full answer to this question will become clear as the reader progresses through the pages that follow. There are various approaches possible towards a multivariate theory of extremes, most of these being coordinatewise theories. This means that, starting from a univariate EVT, a multivariate version is developed which looks at coordinate maxima and their weak limit laws under appropriate scaling. Then the key question to address concerns the dependence between the components of the nondegenerate limit. In the pages that follow, we will explain that, from a mathematical point of view, a more geometrical, coordinate-free approach towards the stochastic modelling is not only mathematically attractive, but also very natural from an applied point of view. For this, first recall the portfolio link stated above. A portfolio is merely a linear combination of underlying risk factors X_1, \dots, X_d with weights w_1, \dots, w_d . Here X_i stands for the future, one-period value of some underlying financial instrument. The hopefully rare event that the value of the portfolio

$$V(\mathbf{w}) = \sum_{i=1}^d w_i X_i$$

is low can be expressed as $\{V(\mathbf{w}) \leq q\}$ where q is some value determined by risk management considerations. A value below q should only happen with a very small probability. Now, of course, the event $\{\sum_{i=1}^d w_i X_i \leq q\}$ has an immediate geometric interpretation as the vector (X_1, \dots, X_d) hitting a halfspace determined by the portfolio weights (w_1, \dots, w_d) at the critical level q . Furthermore, depending on the type of position one holds, the signs of the individual w_i 's will be different: in portfolio language, one moves from a long to a short position. Further, the world of financial derivatives allows for the construction of portfolios, the possible values of which lie in specific subspaces of \mathbb{R}^d . The first implication is that one would like to have a broad theory that yields the description of rare events over a wide range

of portfolio positions. The geometry enters naturally through the description of this rare event set as a halfspace. A natural question then to ask is *what do we know about the stochastic behaviour of (X_1, \dots, X_d) given that such a rare event has occurred?* Thus a theory is needed which yields results on the conditional distribution of a random vector (the risk factors) given that a linear combination of these factors (a portfolio position or market index) surpasses a high (rare) value.

The interpretation of high or rare value depends on the kind of position taken, hence in the first instance, the theory should allow for resulting halfspaces to drift to infinity in a general (non-preferred) direction. This kind of isotropic limit nevertheless yields a rich theory covering many of the standard examples in finance and insurance.

At the same time, however, one also needs to consider theories for multivariate extremes where the rare event or high risk scenario corresponds to a “drifting off” to infinity in one specific direction. This of course is the case when one is interested in one particular portfolio with fixed weights over the holding period (investment horizon) of the portfolio. Another example concerns the operational risk problem discussed above. Here the one-year losses correspond to random variables L_1, \dots, L_d where, depending on the approach used, d can stand for eight business lines, seven loss types or fifty-six combinations of these. Under Basel II, banks have to come up with a risk measure for the total loss $L_1 + \dots + L_d$ and hence a natural question to ask is the limiting behaviour of the conditional distribution of the vector (L_1, \dots, L_d) given that $L_1 + \dots + L_d$ is large. This is an example where one is interested in the conditional behaviour of the risk factors (L_1, \dots, L_d) in the direction given by the vector $(1, \dots, 1)$. The mathematics entering the theory of multivariate extremes in a particular direction in \mathbb{R}^d is different from the “isotropic” theory mentioned above and translates into different invariance properties of classes of limit laws under appropriate transformations. Examples of research papers where the interplay between geometrical thinking and the discussion of multivariate rare events are to be found include for instance Hult & Lindskog [2002] and Lindskog [2004]. The latter PhD thesis also contains a nice summary of the various approaches available to the multivariate theory of regular variation and its applications to multivariate extreme value theory. Besides the various references on this topic presented later in the text, we also like to mention Fougères [2004] and Coles & Tawn [1991]. The necessary statistical theory is nicely summarized in Coles [2001].

Perhaps an extra remark on the use of geometric arguments, mainly linked to invariance properties and symmetry arguments is in order. It is no doubt that one of the great achievements of 19th century and early 20th century mathematics is the introduction of abstract tools which contribute in an essential way to the solution of applied problems. Key examples include the development of Galois Theory for the solution of polynomial equations or Lie groups for the study of differential equations. By now both theories have become fundamental for our understanding of natural phenomena like symmetry in crystals, structures of complex molecules or quantum

behaviour in physics. For a very readable, non-technical account, see for instance, Ronan [2006]. We strongly believe that geometric concepts will have an important role to play in future applications to Quantitative Risk Management.

By now, we made it clearer why we have written this text; the motivation comes definitely from the corner of QRM in the realm of mainly banking and to some extent insurance. An alternative title could have been “Stress testing methodology for multivariate portfolios”, though such a title would have needed a more concrete set of tools for immediate use in the hands of the (financial) portfolio manager. We are not yet at that level. On the other hand, the book presents a theory which can contribute to the discussion of stress testing methodology as requested, for instance, gp in statements of the type

“Banks that use the internal models approach for meeting market risk capital requirements must have in place a rigorous and comprehensive stress testing program. Stress testing to identify events or influences that could greatly impact banks is a key component of a bank’s assessment of its capital position.”

taken from Basel Committee on Banking Supervision [2005]. Over the years, numerous applications of EVT methodology to this question of stress testing within QRM have been worked out. Several examples are presented in McNeil, Frey & Embrechts [2005] and the references therein; further references beyond these include Bensalah [2002], Kupiec [1998] and Longin [2000]. There is an enormous literature on this topic, and we very much hope that academics and practitioners contributing to and interested in this ever-growing field will value our contribution and indeed help in bringing the theory presented in this book to full fruition through real applications. One of the first tasks needed would be to come up with a set of QRM questions which can be cast in our geometric approach to high risk stress scenarios. Our experience so far has shown that such real applications can only be achieved through a close collaboration between academics and practitioners. The former have to be willing (and more importantly, capable) to reformulate new mathematical theory into a language which makes such a discussion possible. The latter have to be convinced that several of the current quantitative questions asked in QRM do require new methodological tools. In that spirit, the question

“For whom have we written this book?”

should in the first instance be answered by: “For researchers interested in understanding the mathematics of multivariate extremes.” The ultimate answer should be “For researchers and practitioners in QRM who have a keen interest in understanding the extreme behaviour of multivariate stochastic systems under stress”. A key example of such a system would be a financial market. At the same time, the theory presented here is not only coordinate-free, but also application-free. As a consequence,

we expect that the book may appeal to a wider audience of “extreme value adepts”. One of the consequences of modern society with its increasing technological skills and information technology possibilities is that throughout all parts of science, large amounts of data are increasingly becoming available. This implies that also more information on rare events is being gathered. At the same time, society cares (or at least worries) about the potential impact of such events and the necessary steps to be taken in order to prevent the negative consequences. Also at this wider level, our book offers a contribution to the furthering of our understanding of the underlying methodological problems and issues.

It definitely was our initial intention to write a text where (new) theory and (existing) practice would go more hand in hand. A quick browse through the pages that follow clearly shows that theory has won and applications are yet to come. This of course is not new to scientific development and its percolation through the porous sponge of real applications. The more mathematically oriented reader will hopefully find the results interesting; it is also hoped that she will take up some of the scientific challenges and carry them to the next stage of solution. The more applied reader, we very much hope, will be able to sail the rough seas of mathematical results like a surfer who wants to stay near the crest of the wave and not be pulled down into the depths of the turbulent water below. That reader will ideally guide the former into areas relevant for real applications. We are looking forward to discuss with both.