Chapter 16

A Conversation with Paul Embrechts

Can you describe your field briefly?

First and foremost, I am a mathematician with strong research and teaching interests in actuarial mathematics. At the Department of Mathematics of the ETH Zürich, I am responsible for the education of actuaries, i.e. insurance mathematicians. Based within a very strong research department, my teaching and research in this more applied area also mirrors that environment. Our position on the educational/research scale within the worldwide actuarial community is hence to be located at the more mathematical, quantitative end. So essentially, my current research field is Quantitative Risk Management (QRM) within insurance and finance, with a particular interest in the modelling of extreme events and in gaining a better understanding of concepts related to risk aggregation and risk diversification.

Present a paradigmatic example of a model in your field, describing it in terms that are accessible to non-experts.

The Gaussian Copula Model (GCM) entered Wall Street around 2000 as THE model for credit derivative markets everybody had been waiting for. Credit derivatives are financial products that can be used both as investment tools as well as hedges (safeguards) in financial markets centred on credit, i.e. bank loans. In particular, the GCM offered bankers a way to model interdependencies (also referred to as default correlation) within and between loan portfolios. By 2009, its original star-like (some even predicted Nobel level) status was downgraded to junk level when it became known as 'A recipe for disaster. The formula that killed Wall Street.' Its original aim was to come up with a pricing and hedging tool addressing the elusive notion of default correlation. The latter technical term refers to the problem corresponding to the crucial task of modelling and risk managing credit markets where several credit (i.e. loan) positions may turn sour at the same moment in time. Building on GCM, its subsequent relatively huge financial markets, with nominal volumes into trillions of dollars, ballooned. GCM had its origin in actuarial mathematics for the modelling of so-called "joint lifes," and indeed was used with success in survival analysis when addressing the joint survival times of patients in clinical trials or epidemiological studies like "the Broken Heart Syndrome," a temporary heart condition that is often brought on by stressful situations, such as the death of a loved one.

With the help of this example, could you explain why a model is needed? And could you describe what a model is? In doing so, please answer the following sub-questions:

What is a model in your field? What are the key steps in a modelling process?

In my area of research, models typically appear as output or consequences of mathematical theorems in answer to concrete questions from practice. For instance, copula models were introduced to answer questions related to the construction of multivariate models with given marginal distributions and with a specific dependence structure. The key steps consequently include: (1) listening to practitioners, trying to achieve a clear understanding of the practical question at hand, (2) the translation of discussions from step (1) into a clear scientific (in my case, most often mathematical) context, (3) model fine-tuning after initial feedback, (4) output production and verification on whether or not this output really answers the question posed; if yes, fine, if no, some further iterations restarting with (1) may have to take place.

What is the role of mathematics in modelling?

First of all, mathematics provides a language in which a given applied, quantitative question can be clearly understood and precisely formulated. A very important part of this (re)formulation is pointing out possible misunderstandings stemming from often vague original formulations. Then it provides conditions under which certain technical consequences (like the calculation of prices) can be derived. In the latter process, it also points at possible limitations of the mathematical tools provided. And though Goethe is quoted as having said: "Mathematicians are like Frenchmen: whatever you say, they translate into their own language and forthwith it is something entirely different," looking more closely at his writings on the topic reveals however how close he was to my interpretation above of the role of mathematics in modelling.

What constituents besides a mathematical formalism are part of a model?

A first, and in my mind crucial, constituent of the above process is a sufficient understanding on behalf of the mathematical modeller of the underlying field to which his/her techniques are to be applied. Further, any model needs to take into account to what extent it can be justified, or for that matter falsified, on the basis of (statistical) data. Increasingly, computer implementation has become a necessity, but at the same time it carries the risk that the end-user becomes more and more distant from the underlying mathematical modelling process. As a consequence, any truly important model should be based on a regular exchange or feedback from practical experience in the model's day-to-day use. One has to avoid as much as possible a black-box status of models crucial within the business environment. Finally, output from a model needs to be communicated to a broader, possibly none(or less)-technical audience. This puts clear standards of quality on model-output and model-documentation.

How important is notation in modelling?

Correct notation is absolutely critical, if only to safeguard clear understanding of the underlying practical issues.

What is the role of language in modelling? Are there qualitative aspects in modelling?

As already stressed above, the role of "language" is eminently important! All too often, a model is as good for business as its flexibility and ease of communication. This bears the risk that in this Darwinian contest, the easy-models become the surviving ones, as they are often perceived as "easy to communicate to management." And as a consequence, the more involved, possibly a better model, may fall out because communication of its structure was deemed "too difficult for management to grasp." One may replace in the statements above "management" with "end-user" depending on the application.

An excellent example from the realm of regulation for banking and insurance is the difficulty encountered, since the mid-1990s, in convincing industry and regulation to move from an if-measure, Value-at-Risk (VaR), to a much more informative what-if measure, Expected Shortfall (ES). The first risk measure, VaR, addresses the frequency with which future extreme events (losses) may occur, for instance, a 1-in-100-year or 1-in-10-day event; the second, ES, addresses the much more relevant severity question: what happens IF such a rare event takes place, how much does one stand to lose? The 2011 9.0 Tohoku earthquake off the coast of North-Eastern Japan was a 1-in-10,000-year event (a VaR-like frequency), much more important information at the time was (would have been) available in an ES-like severity estimate, "What is the expected tsunami height near the Sendai area coastal atomic reactors, GIVEN that such a (very) rare event takes (or would take) place?" And yet, for far too long VaR was defended (especially in the world of banking) as much easier to communicate. In my opinion, and also from practical experience, such statements yield a grossly unjustified perception of the intellect of "management" and "end-user"! Despite academic criticism on the use of VaR, going back to the mid-1990s (VaR was "born" around 1994), only now (2015) are financial regulation and industry warming up to the transition from VaR to ES. Nevertheless, in a recent study, the International Association of Insurance Supervisors (June 2015) released a document revealing that a majority of companies still favour VaR over ES on the basis of (1) easier to understand/communicate, (2) less difficult to calculate, and (3) the shareholder view: based on a limited liability structure of companies, once a company is insolvent (i.e. the rare VaR event happened), the extra information on "severity of insolvency" (encoded in ES) becomes less relevant to shareholders. Whereas worries (1) and (2) can easily be overcome (at least in time), worry (or better said, observation (3)) is more serious and very much hinges on a still prevalent mood on Wall Street of privatising gains versus socializing losses. I am not saying that a change from VaR to ES will come any way near to changing this attitude, but a more consistent reflection on "what if" rather than (just) "if," and this at all levels of a company, would be an important step in the right direction.

Models are often said to represent a target system (typically a selected part or aspect of the world). Does this characterisation describe what happens in your field? If so, could you say how a model represents its target? In other words, how do you understand the model–world interface?

As already alluded to above, the model-to-world interface has to be of feedback-type. The target system aspect is relevant. For instance, when an apple falls from a tree I am sitting under, I may be interested in the event that the apple falls on my head (for which Newtonian mechanics will do just fine). I may, however, be interested in the precise path that a given electron in the apple follows along its path downwards (in this case, quantum mechanics becomes relevant). The target has changed. Similarly, in the GCM example above, the target could be a very specific, small, isolated credit portfolio for which the GCM would work just fine. If, however, one wants to target credit products with nominal values in the trillions affecting broad financial markets, then the GCM is clearly a far too oversimplified model. So, a clear understanding of a model's target is of prime importance, as is the feedback along the evolutionary path of a model in its applications to the outside world. When do we leave the target space for which the model was originally designed? Are we sufficiently aware of and indeed do we understand the broader target regions? All too often, these important broader model issues are not, or at best insufficiently, addressed. One of the catch words along these lines is model robustness, a theme of increasing importance to regulators and industry alike.

What is the relation between a model and theory?

This question can be answered from numerous angles, my answer only gives one. I personally see, in the examples I have considered, a model as being embedded in a theory, the latter yielding a broad model environment in which models can be formulated and tested. For instance, in finance, the theory could be that of rational markets operating under some kind of Efficient Market Hypothesis, whereas a specific model would result in the Black–Scholes–Merton price for a European Call. Alternatively, one could be more interested in Behavioural Economic Theory with a model based on Prospect Theory. Hence, I see models as being grounded in a broader theory or theoretical framework. The usefulness of this "grounding" can be illustrated on the basis of the GCM: recall that at the height of the financial crisis, 2008, say, banks and some insurance companies suffered huge (multibillion) losses from their credit portfolios, like Credit Default Swaps and Collateralised Debt Obligations. By some, in a very naïve way, the GCM was blamed for the downfall: it did not work as expected, especially when markets were under stress. A Financial Times quote at the time stated, "Why did no one notice the formula's Achilles' heel?" The simple answer is: academics (mathematicians) who understood the broader theory within which the GCM was just one little model, understood perfectly why it would backfire in moments of stress. This message was published, communicated at conferences, voiced in discussions *well before* the crisis. No one in industry cared to listen: "As long as the music is playing, you've got to get up and dance" became the market participants' excuse later on. There does (and unfortunately most often) exist a fundamental dislocation in practice between a model being used and a basic understanding of the theory within which this model is scientifically embedded. Famous examples of such an embedding are, for instance, GPS tools within the theory of general relativity, or e-banking (in)security within cryptography and, hence, to a large degree, within the mathematical theory of prime numbers.

What is the aim/use of the model: e.g. learning/exploration, optimisation/exploitation?

In broad terms, models are there to help towards a better understanding of the (not just physical) world around us. In that sense, learning/

exploration/optimisation all play an important role along the path towards this ultimate goal. Models should never be advanced to the point where their formulation and analysis become THE ultimate stand-alone goals. One may perhaps do so momentarily, but ultimately we need to descend back to the level of practical understanding, i.e. to the level of 'exploitation' mentioned in this subsection's title.

In case you use computer simulations, what is the relationship between simulations and the model?

At the basic level, given a model, computer power allows us to obtain answers to questions by simulating numerous realisations of the same model and somehow counting the number of ways in which a certain event occurs. Like testing whether a coin is fair by tossing it over and over again and counting the frequency of heads, say. Physically proving that the coin is biased may be difficult, tossing it is not. Similarly, a model used in practice may be far too complicated to allow us to analytically calculate the occurrence probabilities of certain events, say, but once the model is written down and its parameters calibrated, one can simulate numerous realisations of the process, and it also becomes possible to gauge the consequences of (small) changes to the underlying parameters and/or assumptions. The latter is also referred to as stress testing and relates to the question of model robustness as mentioned above. It is no coincidence that this methodology is often referred to as a Monte Carlo simulation. Perhaps one of the key examples in everyday life is to be found in weather prediction which, perhaps contrary to common belief, has achieved significant quality progress over the recent decennia. All this due to an optimal symbiosis of model theory and model simulation. It is hard to think of any branch of applications where (model) simulation does not play a fundamental role. But note that, typically, a model is lurking in the background. At this point, I ought to drop the notion of Big Data, though a better name is Data Science. It should be clear that most models in use out there can be falsified given huge amounts of data, even the fair coin hypothesis. With a zillion observations, we start to discern even the smallest (presumably irrelevant) deviations from a hypothesised model. Modern data scientists even brag that models are things from the past, let the (big) data speak for itself. I personally am highly sceptical, just for the simple reason that 'big data' does not necessarily mean 'large information content'…the future will tell. I personally strongly believe that even in the Big Data circus, models have an important role to play. It is a bit like the clowns in a circus, they often bring out the truth, the real sentiments in a performance as an important yardstick or mirror of our own identities, our own strengths and weaknesses.

What has been the impact of the development of new technologies or tools in your field? (e.g. telescope in cosmology, etc.)

Clearly the advent of advanced computational tools combined with ever more powerful computers are no doubt key to my field.

What is a good model?

Here, I want to start with the often (mis)quoted statement of the statistician George E.P. Box: "All models are wrong, some are useful." The correct original statement from his 1976 publication in the *Journal of the American Statistical Association* (Vol. 71, No. 356, p. 792), addressing the issue of parsimony, reads as follows, "Since all models are wrong, the scientist cannot obtain a "correct" one by excessive elaboration. On the contrary, following William of Occam he should seek an economical description of natural phenomena. Just as the ability to devise simple but evocative models is the signature of the great scientist, so overelaboration and overparameterisation is often the mark of mediocrity." In the next paragraph, he continues with, "Since all models are wrong, the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad." For those interested in the present volume, I strongly advise a careful (re)reading of Box's original paper. There is not much more that I can add beyond the fact that, mainly in the statistical, econometric, and insurance/finance literature, behind the following terminology, several tools and techniques are to be found: model selection, model validation, model adequacy, model uncertainty, model misspecification, model robustness… It would lead me too far to discuss (some of) these topics here more at length.

How would you define uncertainty? And how does the model help us understand uncertainty?

Benjamin Franklin is quoted as having said, "…in this world nothing can be said to be certain except death and taxes." For the moment disregarding both (death and taxes), everything in life is about uncertainty. This was scientifically, strongly stressed with the advent of Quantum Physics, to the extent that Albert Einstein exclaimed in disagreement, "God does not play dice with the universe!" At some point he added, "The Lord is subtle, but not malicious." This then provoked the following reply from Niels Bohr: "Einstein, stop telling God what to do." Whichever way one looks back at these early discussions between the scientific giants behind General Relativity Theory and Quantum Theory, modern science clearly puts randomness (hence, uncertainty) at the heart of it all. I personally do not want to enter into the philosophical discussions surrounding uncertainty (and risk, which we later will do), nor do I have the background to enter into a discussion on Quantum Theory. The interested reader may want to search for the plenary lecture I gave at the 30th International Congress of Actuaries in Washington D.C., April 2, 2014, with the title "Uncertainty." As stated in that lecture, for me uncertainty is about incomplete knowledge, an incompleteness I try to model to a high (though not complete) degree via a scientific (mathematical) edifice created for that purpose by A. N. Kolmogorov in 1933, when he wrote his path breaking *Grundbegriffe der Wahrscheinlichkeitsrechnung* (*Foundations of the Theory of Probability*). I realise that in doing so, I do dodge various important issues, which we later will come back to. I am convinced that by using Kolmogorov's triplet of "Sample Space," "Event Space," and "Probability Measure," we can do a lot in describing uncertainty in numerous applications. Within this, by now, classical theory of probability and statistics, one can (and has) achieve(d) considerable progress concerning the modelling of random phenomena. Let me just give a pedagogical example I often use in my introductory lectures on probability and statistics in order to discuss the notion of uncertainty and what to highlight when theory and models become important. Suppose you have a group of students divided up into two roughly equal subgroups. To students in Subgroup 1, you instruct to toss a coin 200 times at home and report the results (110010110001…) as these are produced. The second group goes home and writes down such numbers the way they think a fair coin produces such a sequence (111010010001…). Ask the students to write down the numbers, ordered as produced, on a card together with their names (but not their subgroup) on the back. I claim that by just looking at the card one can separate both groups almost perfectly (here the 200 plays a role). Indeed, one can prove (using the standard mathematical model for fair coin tossing) that Subgroup 1 (the coin tossing students) produce a longest sequence of 1s (heads, say), with very high probability, between 5 and 10. Subgroup 2 (the coin toss thinkers) will typically not come up with sequences of subsequent 1s (or 0s for that matter) of length more than 5. The precise mathematical model calculations are not so easy. One can, however, easily simulate this experiment and come to the same conclusion. I leave it to the reader to make the small step from this experiment to (high frequency) financial markets with stock prices moving up (1), down (–1), or no change (0). In a perfect random stock market, sequences of subsequent ups may turn out to be surprisingly long, just like for instance subsequent runs of black in roulette. On August 18, 1913, a staggering sequence of 26 successive blacks were observed in Monte Carlo. Model calculations can however explain that, in all the roulette games witnessed all over the world over a longer period of time, this is not such a rare event. A more interesting and telling (real) story is that of the statistician who beat online casino because the computer programme underlying the online roulette wheel simulations produced more switches from black to red and backwards than proper randomness allows for. After a short winning streak, our statistician was forbidden to play, and subsequently the computer programme was altered so as to generate spins of the wheel which resemble more closely truly random spins. This I find a compelling example of the power of modelling and I urge readers to consult on this and further examples from the realm of uncertainty in *Significance*, December 2013, 10(6), published jointly by the *American Statistical Association* and the *Royal Statistical Society*. D.J. Hand's *The Improbability Principle: Why Coincidences, Miracles and Rare Events Happen Every Day*, Scientific American, 2014, yields further food for thought. Finally, coming back to the stock market translation from long streaks of ups or downs for a coin or a roulette wheel to ups and downs on financial markets, entertaining texts on the topic carry titles like *A Random Walk Down Wall Street* (Malkiel) and *A Non-Random Walk Down Wall Street* (Lo and MacKinlay).

How would you define risk? And how does the model help us understand risk?

According to common usage, risk entails both uncertainty and exposure, i.e. possible consequences. In our textbook on *Quantitative Risk Management* (McNeil, Frey, Embrechts, Princeton University Press, Revised Edition, 2015), we give the following definition of risk: "Any event or action that may adversely affect an organisation's ability to achieve its objectives and execute its strategies." First of all, this rather restrictive definition only looks at the downside, and neglects the equally (if not more) important upside. Also, it implicitly links risk to measurable objectives and clearly formulated strategies. Reality allows for many shades of grey here: from purely analytical formulations to much more soft interpretations. And hence, also a multitude of names occur in the academic literature and relevant practice (especially within banking and insurance). For instance, in Asian countries the preferred word for risk is resilience with its more proactive interpretation. So, an important distinction is to be made between (1) aleatory risk, as the risk coming from random fluctuations, as in a coin-tossing-based game, and (2) epistemic risk, as the risk due to our incomplete understanding of a problem. For the first type of risk, probability and statistics offer excellent tools, for the second, the key advice is, "be aware of this not-full understanding and learn more." In my own experience, most real problems in practice are a combination of both types of risks, (1) and (2). A typical situation presents itself when one wants to price and hedge a complicated financial or insurance deal: besides the random fluctuations at the level of the underlying data used and model assumptions made, typically customer and market behaviour, the influence of changing accounting rules, the legal environment, and changes in the political landscape all play an important role. On the latter, politics, Donald Rumsfeld's statement, "There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we know we don't know. But there are also unknown unknowns. There are things we don't know we don't know," at first may sound a bit hyperbolic, but at the time he made the statement it surely set risk and (un)certainty free for a broader public debate. On a more profound academic level, in his famous 1921 treatise on Risk and Uncertainty, Frank Knight equates risk to measurable uncertainty and reserves the term uncertainty for unmeasurable uncertainty. This distinction relates risk to objective probabilities, whereas uncertainty corresponds to subjective probabilities. Numerous discussions by many authors emerged, including John Maynard Keynes and Bruno de Finetti, the former a key figure in economics and finance, the latter taking centre stage in insurance. Leaving some of the more philosophical discussions aside, when we assume a middle-of-the-road definition of risk as any event where both uncertainty and exposure act together, then it becomes clear where modelling plays an important role. And this not only in attempts to quantify both ingredients, but also, and in my mind very importantly, in describing clear boundaries beyond which a given model output is not anymore justifiable. Here, we enter the increasingly important field of model uncertainty and robustness. For instance, if we want to calculate a risk measure of a given aggregate position of risk factors, but we have little information on the interdependence between the risk factors, it is better to give best–worst bounds on this risk measure and not try to force upon the end-user a single number. This kind of functional model uncertainty compounds the always existing parameter uncertainty, which is of a purely statistical nature. The story around the use of the Gaussian Copula Model mentioned at the beginning can perfectly be framed in this context: besides the parameter uncertainty in the underlying models for the default probabilities for the underlying credit-based positions, there was great functional (macroeconomic) uncertainty with respect to the interdependence between these credit positions, and this especially in moments of market stress. The latter occurred with catastrophic consequences for the world at large as soon as American house prices did something they never did before…they started to fall!

What is the role of stress testing and sensitivity analysis in the understanding of risk and/or uncertainty? (It might be helpful to clarify what stress testing and sensitivity analysis involve in your field).

A proper stress test and/or sensitivity analysis belongs to most applied problems I have worked on. This may be either disguised under the form of parameter and model uncertainty or explicitly demanded for, like for instance, in banking and insurance regulation. An example of the latter case consists, for instance, of the broadly publicised, and in some notable cases, not very successful stress tests by governments on the stability of their respective banking systems. A further example, this time from the insurance side, consists of the various stress tests demanded for by the regulators responsible for the (private) life insurance business. Such a test might include downturns in the stock market and even the calculation of the consequences on current balance sheets of past extreme events like the 1987 crash or the 2007–2009 financial crisis. Also relevant are the potential consequences of a pandemic.

What do you consider to be the work/result that has had the most significant impact on your field? And why?

I occasionally ask students or visitors passing by my office to single out from my rather extensive personal library on mathematics the couple of books that look most worn out. Almost invariably, the books chosen are William Feller's two volumes on *Probability Theory* and Walter Rudin's *Real and Complex Analysis*, as well as his *Functional Analysis*. The fact that these books go back to my days as a student of course helped in the wear and tear process, but even so. These two categories of mathematical texts exemplify my own attitude to the field of modelling. In Feller's books, one not only learns the more technical aspects of probabilistic thinking, but also the importance of intuition in the field of uncertainty. On the other hand, Rudin's texts are prime examples of the beauty and relevance of mathematical rigour. In my opinion, many (if not all) problems in practice need a combination of both skills: intuition and a basic understanding of the broader issues together with care for clear definitions and an in-depth understanding of the underlying conditions of the model within a more mathematical framework. The GCM once more is an example where this symbiosis of thoughts should have played more strongly. At least, we as mathematicians should foster both! It is impossible to single out a specific scientific contribution in my field of research with main impact on my work as such a choice so much depends on the type of (applied) problems I am currently working on. In my career, I have been uniquely blessed with working environments and intellectual impulses from numerous fields: starting with pure and applied mathematics as a student, early research in applied probability with actuarial applications, economics, finance, statistics, operations research, and industry experience from working together on concrete problems, to membership of boards of independent directors in the financial and insurance industry. It is this combination of experiences that has by far had the strongest impact on me as a(n academic) risk modeller. From these contacts and influences emerged the two (by now classical) books I co-authored with former students on *Modelling of Extremal Events for Insurance and Finance* (Springer, 1997) and *Quantitative Risk Management* (Princeton University Press, 2005, 2015). This educational path is difficult to copy; at all stages of life, uncertainty and risk, in the form of decisions and their consequences, have to be taken into account. I have been fortunate both at a personal as well as a professional level to have taken the right (or at least challenging, interesting) turns whenever the road ahead forked. By far the most important skills I learned are an openness of mind and an eagerness to learn both at a societal as well as academic level. And whenever I work on a truly applied problem, the words of Hamlet to Horatio are not far away: "There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy."