# CHAPTER 1

# Creation of the Method



This first part of the book presents our experiences in operational risk modelling from a subjective point of view over the past 10 years.

## 1.1 FROM ARTIFICIAL INTELLIGENCE TO RISK MODELLING

We are engineers specialized in artificial intelligence. We have been working together for about 25 years, and early in our careers we spent a lot of time on applications of neural networks, Bayesian networks, and what was called data mining at this time. It was almost a generation before the current popularity of these techniques.

Back in the 1990s, few industries had the means to invest in artificial intelligence and data mining: mostly banking, finance, and the defense sector, as they had identified applications with important stakes. The defense usually conducts its own research, and so, quite naturally, we spent a lot of time working in research and development with banks and insurance companies, for applications such as credit rating, forecasting financial markets, and portfolio allocation. We were fortunate enough that the French Central Bank was our first client for this service, for several years. Thanks to a visionary managing director, the French Central Bank created in the early 1990s an AI team of more than 20 people working on applications ranging from natural language processing to credit scoring.

Our conclusion was mixed. Machine-learning techniques were generally not better than conventional linear techniques. This mediocre performance was not related to the techniques themselves, but to the data. When you try to predict the default of a company from its financial ratios, you will always have several companies with exactly the same profile, but that will not share the same destiny. This is because the observed data do not include all of the variables that could help predict the future. The talent and the pugnacity of the leader, the competitive environment, and so on, are not directly represented in the accounting or financial data. However, these nonfinancial indicators are the ones that will make the difference, all things being equal otherwise. Finally, in rating or classification applications, and whatever the technique used, the rates of false positives or false negatives were usually very close.

This is even more applicable when you are trying to predict the markets. We were most of the time trying to forecast the return of one particular market at various horizons, using either macroeconomic variables, or technical variables. We would have been largely satisfied with a performance just slightly better than flipping a coin. Again, the performance of nonlinear models was comparable to other techniques. In a slightly more subtle way here, the limitation was expressed through the dilemma between the complexity of the model and the stability of its performances: to get a model with stable performance, this model must be simple. The best compromise is often the linear model.

About 10 years ago, a Head of Operational Risk for a large bank asked us to think about the use of the Bayesian networks to model the operational risk, and to seek to evaluate possible extreme events. Not surprisingly, she was advised to do so by the former managing director of the French Central Bank, which we mentioned previously.

We were immediately intrigued and interested in the subject. We liked the challenge of leaving aside for a while the "big data" analysis to work on models based on "scarce data"! We thought, and continue to think, that the work of a modeler is not to look for mathematical laws to represent data, but to understand the underlying mechanisms and to gain knowledge about them. It is not surprising that one of us wrote his PhD thesis on the translation of a trained neural network into an intelligible set of rules.

Going back to operational risks, or more precisely to one of the requirements of AMA (advanced measurement approach), the problem was formulated mathematically quite simply, but seemed to require an enormous work.

The mathematical problem was to estimate an amount $M$ such that it could only be exceeded with a probability of 0.1%, regardless of the combination of operational risk events that could be observed in the forthcoming year.

In practical terms, this meant answering several questions, all of them more difficult than the other:

1. *Identification*. What are the major events that my institution could be exposed to next year? How to identify them? How to structure them? How to keep only those that are extreme but realistic (that is, how not to quantify a *Jurassic Park* scenario!).
2. *Evaluation*. How to evaluate the probability that one of them will occur? If it occurs, how to evaluate the variability of its consequences?
3. *Interdependencies*. All adverse events will not happen at the same time. However, certain events can weaken a business and make other extreme events more likely. For example, a significant natural event can weaken control capabilities and increase the risk of fraud. How to evaluate the correlations between these events?

Once we became acquainted to the problem, we did two things:

1. As consultants, we studied closely the risk management system of the bank.
2. As researchers, we studied the state of the art on the question of quantification.

We must admit that if we were impressed by the work done by our client, this was not the case on the state of the art.

This customer, which is one of the largest French banks, serves today nearly 30 million customers with more than 70,000 employees, and covers most of the banking business lines, even if it does not compare in that with the large investment banks in Europe or in the United States.

The Head of Operational Risks had put in place a set of risk mappings.

This work was based on a breakdown of the bank's activities. The breakdown did not use processes – as we found later that most banks do – but objects. The objects were of different nature: products, people, systems, buildings, and so on. This approach was consistent with the overall risk analysis approach proposed by the ARM method.[1] According to this approach, a risk is defined by a combination of Event, Object, and Consequence. A risk is therefore defined by the encounter of an event and a resource likely to be affected by this event. We will come back to this, but this approach, common to ARM and ISO 31000 and shared by most industries and research organizations working on major risks, is extremely structuring and fertile for modeling.

This mapping was not only a catalog. For each type of exposed object, the Operational Risk department of this bank had established a working group consisting of a risk manager and several experts to identify and assess risks in a simple way. Contrary to what we have seen later in sometimes more prestigious organizations, this work was not only an expert evaluation obtained during a meeting, but was a structured and well-argued document, which could be reviewed and discussed by the internal audit bodies and by the regulator. As a conclusion of each study, each of the risks identified and considered significant by the working group for the type of object considered, was the subject of a quantified evaluation. This assessment was in the form of a simple formula that evaluated the cost of risk.

The analysis was describing a mechanism by which a loss could be observed, and the indicators used made it possible to quantify it. For example, the default or disruption of a supplier could impact the business during the time needed to switch to a backup supplier. The switching time would of course depend on the quality of prior mitigation actions. This helps defining the outline of a "Supplier Failure" model: list all the critical suppliers, evaluate for each of them a probability of default, evaluate the impact of the supplier unavailability on the bank's revenue, and assess the time to return to normal operations. The combination of these different factors, all assessed with a certain degree of uncertainty, made it possible to consider building a model. We have subsequently validated this approach for all types of risks, irrespective of the type of exposed objects: people, buildings, products, stock market orders, applications, databases, suppliers, models , and so on.

## 1.2 MODEL LOSSES OR RISKS?

The other part of our preparatory research concerned the state of the art on modeling. We were surprised to find that the dominant model was called the LDA, for Loss Distribution Approach, and was actually a statistical model of past losses, not a risk model.

The point that surprised us the most is the effort statisticians made to search for laws that would fit the data, without seeking any theoretical justification for choosing the law. We had some theoretical knowledge of the modelling of financial markets, in which the use of a normal law results from the theoretical framework of efficient markets. This framework, proposed by Bachelier, demonstrates that if the markets are efficient, then the distribution of returns follows a normal distribution. This theoretical hypothesis is clear and debatable. We can accept or reject the hypothesis of efficient markets. It can be considered that there exists insider information that distorts the markets. This discussion regards the validity of the models, of the same nature

[1] The method taught by "The Institutes" to obtain the qualification of Associate in Risk Management. See https://www.theinstitutes.org (accessed 5/10/2018).

as the discussions that one may have in physics, on the fact that the hypothesis of perfect gases, or incompressible fluids is or is not acceptable, and therefore that the associated equations are applicable.

On the modelling of operational risks, nothing like that. The choice of a law did not come from a theoretical discussion, but only from its ability to fit to the data, which seemed to us contradictory to any modelling logic. Moreover, the data considered in the adjustment are not of the same nature.

The principle of the LDA is (1) to assume that the average number of losses observed in one year will also be observed in the following years although with some variance (this is represented by the use of a frequency law, for example a Poisson law), and (2) to adjust a theoretical distribution on the amounts of observed losses.

Taken literally, this approach means that the only variability of the losses lies in their number and in their arrangement (an unfavourable year can suffer several significant losses). In other words, randomness would lie only in the realizations, and not in the nature of the risk scenarios. According to this principle, a tsunami would be then only an unexpectedly big wave. Even if the adjustment of a theoretical distribution on the height of the waves makes it possible mathematically to calculate the probability of a wave of 20 or 30 meters of height, it does not account for the difference of nature of the two phenomena: tsunamis are not caused by the same process as waves.

This approach seemed to be wrong for several reasons. Regardless of the possibility of statistically adjusting a law without knowing the theoretical form that this law must take, what would be the logic to use past losses to anticipate future losses, even as technologies evolve, risks evolve, and banking activities evolve?

Why use credit card fraud loss history before EMV chips implementation, in a context where EMV chip cards are now widespread and being used? How not to see that the regulatory pressure on the risks related to the conduct of banks depends on the political climate? Would the political will to punish the banks for the economic and human disaster of the subprime crisis be applied with the same rigor if Barack Obama had not been president at that time? How about losses related to sold or obsolete activities? For example, our French client had in its accounts a significant loss related to a model error on market activities, which led the management of the bank to sell these activities: was it then justified to consider this loss in the history used to extrapolate future losses?

Of course, we know the argument of “quants” in banks, which can be summarized in a few words. Even if things change, past losses are representative of an institution, its size and its culture, and therefore they can be validly used, even to predict losses of another nature. In other words, a loss observed on market activities contains information to anticipate a possible loss on the use of cryptocurrencies, because the risk profile of a bank has a certain stability, which gives it a certain propensity to take risks, a certain appetite for risk, independent of activities and technologies. This sounds like an attempt to give a soul to a banking institution that would remain stable through all the changes. We will not engage in this metaphysical terrain of the soul of organizations, but to consider that this soul would manifest itself through the taking of operational risk, seems to us to be fanciful at best.