

COPYRIGHT NOTICE:

Riccardo Rebonato: Plight of the Fortune Tellers

is published by Princeton University Press and copyrighted, © 2007, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@pupress.princeton.edu

CHAPTER 1

WHY THIS BOOK MATTERS

But what I... thought altogether unaccountable was the strong disposition I observed in [the mathematicians of Laputa] toward news and politics, perpetually enquiring into public affairs, giving their judgement on matters of state, and passionately disputing every inch of a party opinion. I have indeed observed the same disposition among the mathematicians I have known in Europe, although I could never observe the least analogy between the two sciences, unless those people suppose that because the smallest circle hath as many degrees as the largest, therefore the regulation and management of the world require no more abilities than the handling and turning of a globe.

Jonathan Swift, writing about the mathematicians of Laputa

STATISTICS AND THE STATE

This book is about the quantitative use of statistical data to manage financial risk. It is about the strengths and limitations of this approach. Since we forget the past at our own peril, we could do worse than remind ourselves that the application of statistics to economic, political, and social matters is hardly a new idea. The very word “statistics” shares its root with the word “state,” a concept that, under one guise or another, has been with us for at least a few centuries. The closeness of this link, between compilations of numbers, tables of data, and actuarial information on the one hand and the organization and running of a state on the other

may today strike us as strange. But it was just when the power of this link became evident that statistics as we know it today was “invented.” So, in its early days probability theory may well have been the domain of mathematicians, gamblers, and philosophers. But even when mathematicians did lend a helping hand in bringing it to life, from the very beginning there was always something much more practical and hard-nosed about statistics.

To see how this happened, let us start at least close to the beginning, and go back to the days of the French Revolution. In the first pages of *Italian Journey* (1798), Goethe writes:

I found that in Germany they were engaged in a species of political enquiry to which they had given the name of *Statistics*. By statistical is meant in Germany an inquiry for the purpose of ascertaining the political strength of a country, or questions concerning matters of state.

By the end of the eighteenth century, when Goethe explained his understanding of the word “statistics,” the concept had been around in its “grubby” and practical form for at least a century. It is in fact in 1700 that we find another German, Leibniz, trying to forward the cause of Prince Frederick of Prussia who wanted to become king of the united Brandenburg and Prussia. The interesting point for our discussion is that Leibniz offers his help by deploying a novel tool for his argument: statistics. Prince Frederick of Prussia was at a disadvantage with respect to his political rivals, because the population of Prussia was thought to be far too small compared with that of Brandenburg to command a comparable seat at the high table of power. If at the time the true measure of the power of a country was the size of its population, the ruler could not be a Prussian. What Leibniz set out to prove was that, despite its smaller geographical size, Prussia was nonetheless more populous than was thought, indeed almost as populous as Brandenburg—and hence, by right and might, virtually as important. How did he set out to do so in those pre-census days? By an ingenious extrapolation *based solely on the Prussian*

register of births, which had been started seventeen years earlier and carefully updated since.*

The details of how the estimate was reached need not concern us here—probably so much the better for Leibniz, because the jump from the birth data collected over seventeen years to the total size of the population was, by modern statistical standards, flawed. What *is* of great current interest is the logical chain employed by Leibniz, i.e., the link between some limited information that we *do* have but that, per se, we may consider of little importance—What do we care about births in Prussia?—to information that we *would* desperately like to have but is currently beyond our reach: for Leibniz and Prince Frederick, ultimately, what is the might of the Prussian army? If anyone were ever to doubt that there is real, tangible power in data and in data collection, this first example of the application of the statistical line of argument to influence practical action should never be forgotten. The modern bank that painstakingly collects information about failures in the clearing of cheques, about minute fraud, about the delinquency of credit card holders and mortgagors (perhaps sorted by age, postcode, income bracket, etc.) employs exactly the same logic today: data give power *to actions and decisions*. To the children of the Internet age it may all seem very obvious. But, at the beginning of the eighteenth century, it was not at all self-evident that, in order to gain control over the running of the state, looking at hard, empirical, and “boring” data might be more useful than creating engaging fictions about “natural man,” “noble savages,” social contracts between the king and the citizens, etc. The first statisticians were not political philosophers or imaginative myth-makers: they were civil servants.

The parallels between these early beginnings and today’s debates about statistics run deeper. As soon as the power of basing decisions on actual data became apparent, two schools of thought quickly developed, one in France (and Britain, Scotland

*Hacking, *The Taming of Chance*.

in particular) and one in Prussia. Generalizing greatly,* the French school advocated an interpretation of the data on the basis of the “regularities of human nature”: deaths, births, illness, etc., were, according to the French and British schools of thought, no less regular, and therefore no less amenable to rigorous quantitative analysis, than, say, floods or other “natural” phenomena. Ironically, the Prussian school, that had founded the statistics bureau, failed to reap the full advantage of its head start because it remained suspicious of the French theoretical notions of “statistical law” when applied to human phenomena—and, predictably, derided the French statisticians: “What is the meaning of the statement that the average family has 2.33 children? What does a third of a child look like?”

Perhaps it is not surprising that the country of the fathers of probability theory (Descartes, Pascal, Bernoulli, Fermat, etc.) should have been on the quantitative side of the debate. Indeed, 100 years before statistics were born, Bernoulli was already asking questions such as, “How can a merchant divide his cargo between ten ships that are to brave the pirate-infested seas so as to minimize his risk?” In so doing, he was not only inventing and making use of the eponymous probability distribution, he was also discovering risk aversion, and laying the foundations of financial risk management. Probability and statistics therefore seemed to be a match made in heaven: probability theory would be the vessel into which the “hard” statistical data could be poured to reach good decisions on how to run the state. In short, the discipline of probability, to which these French minds contributed so much, appeared to offer the first glimpses of an intriguing promise: a quantitative approach to decision making.

The Prussian–French debate was not much more constructive than many of the present-day debates in finance and risk management (say, between classical finance theorists and behavioral financiers), with both parties mainly excelling in caricaturing their opponent’s position. Looking behind the squabbling, the

*For a more nuanced discussion of the two schools, again see Hacking, *The Taming of Chance*.

arguments about the applicability and usefulness of quantitative techniques to policy decisions have clearly evolved, but reverberations of the 200-year-old Franco-Prussian debate are still relevant today. The French way of looking at statistics (and of using empirical data) has clearly won the day, and rightly so. Perhaps, however, the pendulum has swung too far in the French direction. Perhaps we have come to believe, or assume, that the power of the French recipe (marrying empirical data with a sophisticated theory of probability) is, at least in principle, boundless.

This overconfident extrapolation from early, impressive successes of a new method is a recurrent feature of modern thought. The more elegant the theory, the greater the confidence in this extrapolation. Few inventions of the human mind have been more impressive than Newtonian mechanics. The practical success of its predictions and the beauty of the theory took a hold on Western thought that seemed at times almost impossible to shake off. Yet two cornerstones of the Newtonian edifice, the absolute nature of time and the intrinsically deterministic nature of the universe, were ultimately to be refuted by relativity and quantum mechanics, respectively. Abandoning the Newtonian view of the world was made more difficult, not easier, by its beauty and its successes.

It sounds almost irreverent to shift in one paragraph from Newtonian physics and the absolute nature of time to the management of financial risk. Yet I think that one can recognize a similar case of overconfident extrapolation in the current approach to statistics applied to finance. In particular, I believe that in the field of financial risk management we have become too emboldened by some remarkable successes and have been trying to apply similar techniques to areas of inquiry that are only superficially similar. We have come to conclude that we simply have to do “more of the same” (collect more data, scrub our time series more carefully, discover more powerful statistical theorems, etc.) in order to answer *any* statistical question of interest. We have come to take for granted that while some of the questions may be hard, they are always well-posed.

However, if this is not the case but the practice and the policy to control financial risk remain inspired by the nonsensical answers to ill-posed questions, then we are all in danger. And if the policies and practices in question are of great importance to our well-being (as is, for instance, the stability and prudent control of the financial system), we are all in great danger.

WHAT IS AT STAKE?

Through financial innovations, a marvelously intricate system has developed to match the needs of those who want to borrow money (for investment or immediate consumption) and of those who are willing to lend it. But the modern financial system is far more than a glorified brokerage of funds between borrowers and lenders. The magic of modern financial engineering truly becomes apparent in the way *risk*, not just money, is parceled, repackaged, and distributed to different players in the economy. Rather than presenting tables of numbers and statistics, a simple, homely example can best illustrate the resourcefulness, the reach, and the intricacies of modern applied finance.

Let us look at a young couple who have just taken out their first mortgage on a small house with a local bank on Long Island. Every month, they will pay the interest on the loan plus a (small) part of the money borrowed. Unbeknownst to them, their monthly mortgage payments will undergo transformations that they are unlikely even to imagine. Despite the fact that the couple will continue to make their monthly mortgage payments to their local bank, it is very likely that their mortgage (i.e., the rights to all their payments) will be purchased by one of the large federal mortgage institutions (a so-called “government-sponsored agency”) created to oil the wheels of the mortgage market and make housing more affordable to a large portion of the population. Once acquired by this institution, it will be pooled with thousands of other mortgages that have been originated by other small banks around the country to advance money to similar home buyers. All these

mortgages together create a single, diversified pool of interest-paying assets (loans). These assets then receive the blessing of the federal agency who bought them in the form of a promise to continue to pay the interest even if the couple of newlyweds (or any of their thousands of fellow co-mortgagors) find themselves unable to do so. Having given its seal of approval (and financial guarantee), the federal institution may create, out of the thousands of small mortgages, new standardized securities that pay interest (the rechanneled mortgage payments) and will ultimately repay the principal (the amount borrowed by the Long Island couple).

These new securities, which have now been made appealing to investors through their standardization and the financial guarantee, can be sold to banks, individuals, mutual funds, etc. Some of these standardized securities may also be chopped into smaller pieces, one piece paying only the interest, the other only the principal when (and if) it arrives, thereby satisfying the needs and the risk appetite of different classes of investors. At every stage of the process, new financial gadgets, new financial instruments, and new market transactions are created: some mortgages are set aside to provide investors with an extra cushion against interruptions in the mortgage payments; additional securities designed to act as “bodyguards” against prepayment risk are generated; modified instruments with more predictable cash flow streams are devised; and so on.

So large is this flow of money that every rivulet has the potential to create a specialized market in itself. Few tasks may appear more mundane than making sure that the interest payments on the mortgages are indeed made on time, keeping track of who has repaid their mortgage early, channeling all the payments where they are due just when they are due, etc. A tiny fraction of the total value of the underlying mortgages is paid in fees for this humble servicing task. Yet so enormous is the river of mortgage payments, that this small fraction of a percent of what it carries along, its flotsam and jetsam, as it were, still constitutes a very large pool of money. And so, even the fees earned for the administration of the various cash flows become tradeable instruments in

themselves, for whose ownership investment banks, hedge funds, and investors in general will engage in brisk, and sometimes furious, trading.

The trading of all these mortgage-based securities, ultimately still created from the payments made by the couple on Long Island and their peers, need not even be confined to the country where the mortgage originated. The same securities, ultimately backed by the tens of thousands of individual private mortgage borrowers, may be purchased, say, by the Central Bank of China. This body may choose to do so in order to invest some of the cash originating from the Chinese trade surplus with the United States. But this choice has an effect back in the country where the mortgages were originated. By choosing to invest in these securities, the Central Bank of China contributes to making their price higher. But the interest paid by a security is inversely linked to its price. The international demand for these repackaged mortgages therefore keeps (or, actually, pushes) down U.S. interest rates and borrowing costs. As the borrowing cost to buy a house goes down, more prospective buyers are willing to take out mortgages, new houses are built to meet demand, the house-building sector prospers and employment remains high. The next-door neighbor of the couple on Long Island just happens to be the owner of one small enterprise in the building sector (as it happens, he specializes in roof tiling). The strong order book of his roof tiling business and his optimistic overall job outlook make him feel confident about his prospects. Confident enough, indeed, to take out a mortgage for a larger property. So he walks into the local branch of his Long Island bank. *Au refrain.*

There is nothing special about mortgages: a similarly intricate and multilayered story could be told about insurance products, the funding of small or large businesses, credit cards, investment funds, etc. What all these activities have in common is the redirection, protection from, concentration, or diversification of some form of risk. All the gadgets of modern financial engineering are simply a bag of tricks to reshape the one and only true underlying

“entity” that is ultimately being exchanged in modern financial markets: risk.

But there is more. All these pieces of financial wizardry must perform their magic while ensuring that the resulting pieces of paper that are exchanged between borrowers and lenders enjoy an elusive but all-important property: the ability to flow smoothly and without interruptions among the various players in the economy. All the fancy pieces of paper are useful only if they can freely flow, i.e., if they can be readily exchanged into one another (thereby allowing individuals to change their risk profile at will), or, ultimately, into cash. Very aptly, this all-important quality is called “liquidity.” By itself, sheer ingenuity in inventing new financial gadgets is therefore not enough: the pieces of paper that are created must not only redirect risk in an ingenious way, they must also continually flow, and flow without hindrance.

As the mortgage example suggests, if an occlusion occurs anywhere in the financial flows that link the payments made to the local Long Island bank to the foreign exchange management reserve office of the Bank of China, the repercussions of this clogging of the financial arteries will be felt a long way from where it happens (be it on Long Island or in Beijing). It is because of this interconnectedness of modern financial instruments that financial risk has the potential to take on a so-called “systemic” dimension.

Fortunately, like most complex interacting systems, the financial system has evolved in such a way that compensating controls and “repair mechanisms” become active automatically (i.e., without regulatory or government intervention) when some perturbation occurs in the orderly functioning of the network. Luckily, the smooth functioning of this complex system does not have to rely on the shocks not happening in the first place:

[T]he interesting question is not whether or not risk will crystallize, as in one form or another risks crystallize every day. Rather the important question is whether ... our capital markets can absorb them.*

*P. Tucker, Executive Director for Markets and member of the Monetary Policy Committee, Bank of England, *Financial Stability Review* December 2005:73.

And indeed, most of the time, surprises are readily and “automatically” absorbed by the system. But if the shock is too big or if, for any reason, some of the adjusting mechanisms are prevented from acting promptly and efficiently, the normal repair mechanisms may not be sufficient to restore the smooth functioning of the financial flows. In some pathological situations, they may even become counterproductive. If this happens, the effects can quickly be felt well beyond the world of the whirling pieces of paper and can begin to affect the “real economy.” It is exactly because, in a modern economy, the boundaries between the financial plumbing and the real economy have become so blurred and porous that the management of financial risk is today vitally important—and vitally important to everyone, from Wall Street financiers down to the young couple on Long Island.

One of the salient characteristics of the modern economy is that an amazingly complex web of interactions can develop among agents who barely, if at all, know each other, to produce results that give the impression of a strong and conscious coordination (an “intelligent design”). This is, of course, in essence nothing but Adam Smith’s invisible hand in action, but the complexity of modern-day financial arrangements would have been unthinkable when *The Wealth of Nations* was written. Accepting that “no one is in charge” but that an intricate network of transactions will, most of the time, work very well remains to this day difficult to accept from an intuitive point of view. “Please understand that we are keen to move toward a market economy,” a senior Soviet official whose responsibility had been to provide bread to the population of Saint Petersburg in its communist days told economist Paul Seabright,* “but we need to understand the fundamental details of how such a system works. Tell me, for example, who is in charge of the supply of bread to the population of London.” In the industrialized Western world we may have forgotten how astonishing the reply is (“nobody is in charge”), but

*P. Seabright (2004), *The Company of Strangers* (Princeton University Press).

this does not make it any less so. In a similar vein, looking at a small portion of butter offered as part of an in-flight meal, Thomas Schelling (rhetorically) asks:

How do all the cows know how much milk is needed to make the butter and the cheese and the ice cream that people will buy at a price that covers the cost of maintaining and milking the cow and getting each little piece of butter wrapped in aluminum foil with the airlines's insignia printed on it?*

Of all these complex, resilient, self-organizing systems, the financial industry is probably one of the most astonishing. Nobody was “in charge” to ensure that when the couple from Long Island walked into their local bank branch they would get a competitive quote for their mortgage in hours; nobody was in charge to make sure that the depositor who provided the funds needed by the couple to purchase their house would be forthcoming just at the right moment; nobody was in charge when the same depositor changed his mind a week later and withdrew the money he had deposited; nobody was in charge to ensure that a buyer would be willing to part with his money within hours, or minutes, of the security from the pool of mortgages being created and offered on the market; and when, thousands of miles away, a Bank of China official decided to invest some of the Chinese trade surplus in the purchase of the same security, nobody had communicated his intentions and nobody had arranged for a market maker to buy the security itself from the primary investor and hold it in his inventory. Nobody is in charge and yet, *most of the time*, all of these transactions, and many more, flow without any problems.

There, however, in that innocent “most of the time,” is the rub. The financial system is robust, by virtue of literally hundreds of self-organizing corrective mechanisms, but it is not *infinitely* robust. This is not surprising because the cornerstone institution of the financial system, the bank itself, is precariously

*T. Schelling (1978), *Micromotives and Macrobehavior* (New York and London: Norton).

perched on the border of stability: just because nobody is in charge, every day a bank owes its existence to the law of large numbers. How does this happen? The most fundamental activity of a bank is the so-called “maturity transformation”: accepting money from depositors who may want it back tomorrow, and lending the “same money” to borrowers who may want to hold on to it for thirty years. This is all well and good unless all of the depositors simultaneously want their money back. In normal conditions, of course, it is virtually impossible for this to happen, and the statistical balance between random withdrawals from depositors and equally random repayments from borrowers requires a surprisingly small safety buffer of “liquid cash.” And, by the way, many other arrangements in our nobody-is-in-charge industrialized world rely on a similar (apparently fragile and precarious) statistical balance: from the provision of food to the distribution of electricity, the use of roads, telephone lines, petrol stations, etc. As long as all the users make their decisions close-to-independently, only a relatively small safety margin (of spare electricity, spare phone line capacity, petrol in filling stations, food on supermarket shelves, etc.) has to be kept.

Conditions, however, are not always “normal.” For a variety of reasons, unpredictable coordination of actions can occur that break the independence in the decisions and disrupt the apparently most robust systems: in the 1980s, for instance, the simultaneous switching on of millions of kettles during the commercial break of a TV program broadcasting a British royal wedding caused chaos across the electricity grid.* Similarly, fears about the robustness of supply can, in next to no time, empty supermarket shelves and create queues at petrol stations. The rumors at the root of these unusual coordinated behaviors need not be true for the effect to be highly disruptive and, in some circumstances, potentially devastating: in a state of imperfect information a set of rationally calculating agents looking after their self-interest will

*Seabright, *The Company of Strangers*.

all reach the same conclusions, and close-to-simultaneously head for the supermarket or the petrol station.*

The amazing efficiency and resilience of the supply chain in the industrialized world therefore contains a potentially fragile core for many goods and services. Banks, are, however, different, and not in a way that makes a bank regulator sleep any more peacefully. If, a few minutes after midnight on New Year's eve, we find it difficult to make a telephone call to wish a happy new year to our loved ones, we desist from calling for ten minutes and, by doing so, avoid clogging the line even more. We do so spontaneously, not because some regulation tells us to behave this way, and it is rational for us to do so. But if a rumor spreads that a bank is not solvent, the individually rational thing to do is to rush to the head of the queue, not to come back tomorrow when the queue, which currently circles two blocks, will be shorter: by tomorrow the queue may well be shorter, but only because there may be no money left in the vault. This is how runs on banks occur: given the asymmetry of information between insiders and depositors, the same response (run for the till!) applies in the case of the perfectly healthy bank as in the case of the one that really has got itself into financial difficulties.

How can these panics be prevented when the rational behavior is to do what everyone else is trying to do—and when this “rational” behavior exacerbates, rather than alleviates, the problem? In most industrialized societies, governments have stepped in by creating deposit insurance. The soothing promise of the government to depositors that it stands behind the bank's commitment to repay (with the implication that their monies are safe no matter what) can prevent, or stem, the flood of withdrawals. In the case of a healthy bank, this arrangement can prevent the run from occurring in the first place.

*For an interesting discussion about the consequences of spontaneously organized collective behaviors, see, for example, C. P. Chamley (2004), *Rational Herds* (Cambridge University Press). A possible explanation of how these collective behaviors may become ordered can be found in D. Watts (1999), *Small Worlds—The Dynamics of Networks between Order and Randomness* (Princeton University Press).

It is important to understand that this underwriting of risk by the government is beneficial for the public but it is also extremely beneficial for the banks. For this reason this insurance is not offered “for free.” Via the deposit guarantee the government develops a keen, first-hand financial interest in the good working of banks, but only shares in the downside, not in the profit the bank may make. It therefore imposes another stipulation to this otherwise too-good-to-be-true arrangement: the government, the covenant goes, will indeed step in to guarantee the deposits of the public even if the bank were to fail; but, in exchange for these arrangements, which are extremely useful for the public *and for the banks*, it will acquire a supervisory role over banking operations. (Of course it is well within the remit of the government to care about systemic financial risk even in the absence of deposit insurance, but this commitment certainly focuses its mind.)

The management of financial risk therefore acutely matters, not only to the financial institutions in the long chain from the small Long Island bank to the Bank of China, but to the public at large and, as the agent of the public and the underwriter of deposit insurance, to the government. So, the long chain of financial institutions that sell, buy, and repackage risk care intensely about controlling financial risk—if for no other reason, then simply as an exercise in self-preservation. But the national and international regulators who have been appointed to take a broader, system-wide, view of financial risk also strongly care and justifiably see themselves as the ultimate safeguard of the soundness of the real economy (“ultimate,” in that they come into play if all other self-correcting mechanisms fail).

Shocks and disturbances of small-to-medium magnitude can be efficiently handled by the inbuilt feedback mechanisms of the financial markets. It is the “perfect storms” that have the potential to create havoc and destroy wealth on a massive scale. Given this state of affairs, it is not surprising that the regulatory bodies that are mandated to control the stability and proper functioning of the international financial system are basing more and more of their regulatory demands on the estimation of the probability

of occurrence of extremely rare events. In statistical terms, these requirements are expressed in terms of what are called, in technical parlance, “extremely high percentiles.” For instance, the regulators are requesting more and more frequently, and in wider and wider contexts, that a 99.9th, one-year percentile of some loss distribution be calculated for regulatory capital purposes. This is just the definition of Value-at-Risk (VaR), which we will discuss in detail in the following chapters. At this stage, we can take this expression to mean that financial institutions are required to estimate the magnitude of an adverse market event so severe that an even worse one should only be expected to occur once every thousand years.

Let us pause and “feel the magnitude” of this statement. A more familiar (although not strictly identical) question is the following: “How bad must a flood/snow fall/drought be for you to be confident that we have experienced a worse one only once since the battle of Hastings in 1066?” Again, let us pause and feel the magnitude of the statement. How would you try to estimate such a quantity? What data would you want to have? How many data records would you need? And where would you find them?

As we saw, it is not just the regulators who have a strong interest in keeping the financial cogs turning smoothly. Banks and other financial institutions to a large extent share the same concerns—if for no other reason than because surviving today is a pretty strong precondition for being profitable tomorrow. There is far more to the risk management of a financial institution than avoiding perfect storms, though. Modern (i.e., post-Markowitz) portfolio theory sees the trade-off between expected return and risk (variance of returns) as the cornerstone of modern finance. In this light, possibly the most common and important questions faced almost every day at different levels of a bank are therefore of the following type: What compensation should we, the bank, require for taking on board a certain amount of extra risk? Should we lend money to this fledgling company, which promises to pay a high interest but which may go bankrupt soon? Should we begin underwriting, distributing, and making a market in bonds

or equities, thereby earning the associated fees and trading revenues but accepting the possibility of trading losses? Should we retain a pool of risky loans, or are we better off repackaging them, adding to them some form of insurance, and selling them on?

Of course, how a financial institution should choose among different risky projects has been the staple diet of corporate finance textbooks for decades. Down in the financial trenches, it is a pressing question for all the decision makers in a bank, from the most junior relationship manager to its CEO and finance director. The risk policy of the bank (what is these days called its “risk appetite”) is typically set at a very senior level in rather broad and generic terms. But it is only after the (often rather vague) risk appetite policy has been articulated that the real questions begin: How should these lofty risk principles be applied to everyday decisions and, most importantly, to strategic choices?

The standard textbook answer to these questions used to be disarmingly simple: “Just discount the expected cash flows from the new prospective projects at the ‘appropriate’ risky rate. Compare how much, after this discounting, each new project is worth today. Choose the one with the greatest present value today. Next problem please.”

Even glossing over the difficulty in assessing the “appropriate” rate of discount, this way of looking at the budgeting and capital allocation process appears to suffer from at least one important shortcoming: unless the bank already has an “optimal” mix of risk,¹ the approach appears to neglect the *interaction* between the new projects and the existing business activities. Diversification or concentration of risk (i.e., the undesirability of putting all of one’s eggs in one basket, no matter how strong the basket looks) is not readily handled by the discounted-cash flow method.*

To overcome this shortcoming, a new approach, termed the “economic-capital project,” has recently been suggested. Unfortunately, there is no consensus about the precise meaning and applicability of the term, as it is often used in different, and even

*A brief caveat in passing: it may not be optimal for the banks to diversify, but for investors to do so. More about this later.

contradictory, ways. In chapter 8 I will therefore try to explain what it promises (a lot), and what it can deliver (alas, I fear somewhat less). I will frame that discussion in the wider context of the use and misuse of traditional statistical tools for financial risk management. For the moment, what is of relevance in this introductory chapter is that the make-or-break condition for the economic-capital project to succeed, and for a lot of current quantitative risk management to make sense, is that it should be possible to estimate the probability of extremely remote events—or, in statistical jargon, that it should be possible to estimate extremely high percentiles.

“NOT EVEN WRONG”: SENSE AND NONSENSE IN FINANCIAL RISK MANAGEMENT

One of the points that I make in this book is that the very high percentiles of loss distributions (i.e., roughly speaking, the probability of extremely unlikely losses) *cannot* be estimated in a reliable and robust way. I will argue that the difficulty is not simply a technical one, which could “simply” be overcome by collecting more data or by using more sophisticated statistical techniques, but is of a fundamental nature. Of course, more data and more powerful techniques *will* help, but only up to a point. Asking what the 99.975th percentile of a one-year loss distribution is is not a tough question, it is a close-to-meaningless one. Famously, Wolfgang Pauli, one of the founders of quantum mechanics, dispatched the theory of a student of his with the ultimate putdown: “Your theory is not correct. In fact, your theory is not even wrong. Your theory makes no predictions at all.” This is not quite the case with some utterances from hardcore quantitative risk managers. These statements *do* contain predictions. Unfortunately, they are untestable.

This claim is in itself an important one, since so much effort appears to be devoted to the estimation of these exceedingly low probabilities. This, however, is not the main “message” of my book. Apart from the feasibility issue, I believe that there is a

much deeper and more fundamental flaw in much of the current quantitative approach to risk management. There is a “sleight of hand” routinely performed in the risk-management domain that allows one to move from observed frequencies of a certain event to the probabilities of the same event occurring in the future. So, for instance, we observe the historical default *frequency* of AA-rated firms, and equate this quantity to the *probability* that a AA firm may default in the future. This association (frequency = probability) implies a view and an application of the concept of probability that, at the very least, should not be taken for granted. I believe that this “sleight of hand” is indicative of the way we look at probabilities in the whole of the risk-management arena. The philosophy that underpins the identification of frequencies with probabilities is defensible. In some scientific areas it is both sensible and appropriate. It is not, however, the only possible way of understanding probability. In the risk-management domain I believe that the view of probability-as-frequency often underpins a singularly unproductive way of thinking. At the risk of spoiling a point by overstressing it, I am tempted to make the following provocative claim. According to the prevailing view in risk management:

We *estimate* the probabilities, and from these we *determine* the actions.

For most of the types of problems that financial risk management deals with, I am tempted to say that the opposite should apply:

We *observe* the actions, and from these we *impute* the probabilities.

This statement may strike the reader as a puzzling one. What kind of probabilities am I talking about? Surely, if I want to advertise my betting odds on a coin-tossing event, *first* I should determine the probability of getting heads and *then* I should advertise my odds. The estimation of probability comes first, and action follows. What could be wrong with this approach?

In the coin-flipping case, probably nothing (although in the next chapter I will raise some reasonable doubts as to whether, even in this stylized case, the approach really always works as well as we may think). The real problem is that very few problems in risk management truly resemble the coin-flipping one. It may come as a surprise to the reader, but there are actually many types of probabilities. Yet the current thinking of risk managers appears to be anchored, without due reflection, to the oldest and most traditional view. It is a view of probability (called the “frequentist view”) that applies when we enter a statistical “experiment” with random outcomes with absolutely no prior beliefs about the likelihood of its possible outcomes; when our subsequent beliefs (and actions) are fully and only determined by the outcome of the experiment; and when the experiments can be repeated over and over again under identical conditions. Unfortunately, I will show in the next chapter that none of these requirements really applies to most of the financial-risk-management situations that are encountered in real life.

Fortunately, there *are* views of probability that are better suited to the needs of risk management. They are collectively subsumed under the term “Bayesian probability.” Bayesian probability is often seen as a measure of degree of belief, susceptible of being changed by new evidence.* The frequentist probability is not dismissed, but simply becomes a limiting, and unfortunately rather rare, case. If the conditions of applicability for the frequentist limiting case apply, a Bayesian is likely to be delighted—if for no other reason, because there are thousands of books that teach us very well how to compute probabilities when we flip identical coins, draw blue and red balls from indistinguishable urns, etc. But the abundance of books on this topic should not make us believe that the problems they (beautifully) solve are quite as numerous.

Unfortunately, the “Bayesian revolution” has almost completely bypassed the world of risk management. Sadly, quantitative risk managers, just when they claim to be embracing the

*See chapter 3 for a fuller description.

more sophisticated and cutting-edge statistical techniques, actually appear to be oblivious to the fundamental developments in the understanding of the nature of probability that have taken place in the last half century. This cannot be good, and it is often dangerous. This, in great part, is what this book is about, and why what it says matters.

WHAT DO WE DO WITH ALL THESE PROBABILITIES?

There is another aspect that current risk-management thinking does not seem to handle in a very satisfactory way: how to go from probabilistic assessments to decisions. Managing risk in general, and managing financial risk in particular, is clearly a case of decision making under uncertainty. Well-developed theories to deal with this problem do exist, but, rightly or wrongly, they are virtually ignored in professional financial risk management. Most of the effort in current risk management appears to be put into arriving at the probabilities of events (sometimes very remote ones), with the implicit assumption that, once these probabilities are in front of our eyes, the decision will present itself in all its self-evidence. So, these days it is deemed obvious that a quantitative risk manager should be well-versed in, say, extreme value theory or copula theory (two sophisticated mathematical tools to estimate the probabilities of rare or joint events, respectively). However, judging by the contents pages of risk-management books, by the fliers of risk-management conferences, or by the questions asked at interviews for risk-management jobs, no knowledge, however superficial, is required (or even deemed to be desirable) about, say, decision theory, game theory, or utility theory. A potpourri of recommendations, rules of thumb, and “self-evident truths” do exist, and sometimes these are cloaked under the rather grand term of “industry best practice.” Yet, it only takes a little scrutiny to reveal that the justifications for these decisional rules of thumb are shaky at best, sometimes even contradictory. Alas,

I will argue that some of the lore around the economic-capital project is no exception.

What can we do instead? To answer this question, I will try to explain what the theoretically accepted tools currently at the disposal of the decision maker have to offer. I will also try to suggest why they have met with so little acceptance and to explain the nature of the well-posed criticisms that have been leveled against them. I would like to suggest that some of the reasons for the poor acceptance of these decision tools are intrinsic to these theories and do reflect some important weaknesses in their setup. However, some of the reasons are intimately linked to the “type of probability” relevant in risk management. Therefore, understanding one set of issues (what type of probability matters in risk management) will help with solving the other (how we should use these probabilities once we have estimated them). This link will also provide a tool, via subjective probabilities and how they are arrived at, for reaching acceptable risk-management decisions. In essence, I intend to show that subjective probabilities are ultimately *revealed* by real choices made when something that matters is at stake. In this view probabilities and choices (i.e., decisions) are not two distinct ingredients in the decisional process, but become two sides of the same coin. Decisional consistency, rather than decisional prescription, is what a theory of decisions in risk management should strive for: *given that* I appear comfortable with taking the risk (and rewards) linked to activity A, should I accept the risk–reward profile offered by activity B?

In this vein, I will offer some guidance in the last part of this book as to how decisions about risk management in low-probability cases can be reached. Clearly, I will not even pretend to have “solved” the problem of decision making under uncertainty. I will simply try to offer some practical suggestions, and to show how the line of action they recommend can be analyzed and understood in light of the theoretical tools that we do have at our disposal.