

Interdisciplinary Institute for Innovation



**The risk of a major nuclear
accident: calculation and
perception of probabilities**
François Lévêque

Working Paper 13-ME-02

July 2013



CERNA, MINES ParisTech
60 boulevard Saint Michel
75006 Paris, France
Email: ferreira@mines-paristech.fr



The risk of a major nuclear accident: calculation and perception of probabilities

François Lévêque

The accident at Fukushima Daiichi, Japan, occurred on 11 March 2011. This nuclear disaster, the third on such a scale, left a lasting mark in the minds of hundreds of millions of people. Much as Three Mile Island or Chernobyl, yet another place will be permanently associated with a nuclear power plant which went out of control. Fukushima Daiichi revived the issue of the hazards of civil nuclear power, stirring up all the associated passion and emotion.

The whole of this paper is devoted to the risk of a major nuclear accident. By this we mean a failure initiating core meltdown, a situation in which the fuel rods melt and mix with the metal in their cladding¹. Such accidents are classified as at least level 5 on the International Nuclear Event Scale. The Three Mile Island accident, which occurred in 1979 in the United States, reached this level of severity. The explosion of reactor 4 at the Chernobyl plant in Ukraine in 1986 and the recent accident in Japan were classified as class 7, the highest grade on this logarithmic scale². The main difference between the top two levels and level 5 relates to a significant or major release of radioactive material to the environment. In the event of a level-5 accident, damage is restricted to the inside of the plant, whereas, in the case of level-7 accidents, huge areas of land, above or below the surface, and/or sea may be contaminated³.

Before the meltdown of reactors 1, 2 and 3 at Fukushima Daiichi, eight major accidents affecting nuclear power plants had occurred worldwide⁴. This is a high figure compared with the one calculated by the experts. Observations in the field do not appear to fit the results of the probabilistic models of nuclear accidents produced since the 1970s. Oddly enough the number of major accidents is closer to the risk as perceived by the general public. In general we tend to overestimate any risk relating to rare, fearsome accidents. What are we to make of this divergence? How are we to reconcile observations of the real

¹ We have adopted a broad definition of what constitutes a major accident. Strictly speaking this term is limited to a core meltdown followed by massive discharge.

² The severity of an event increases tenfold, at each level: accordingly level 2 is 10 times worse than level 1, level 3 is 10 times worse than level 2, and so on.

³ The term 'major accident' used in this part is more general than its usage in the terminology established by the INES scale. In this ranking only level-7 accidents rate as 'major'; a level-5 core melt counts as an 'accident with wider consequences', whereas its level-6 counterpart counts as a 'serious accident'.

⁴ SER (1959), Enrico Fermi 1 (1961), Chapelcross 2 (1967), Saint-Laurent A1 (1969) and A2 (1980), Three Mile Island 2 (1979), Chernobyl 4 (1986), Greifswald 5 (1989). List taken from Thomas B. Cochran, 2011, Fukushima Nuclear Disaster and Its Implication for US Nuclear Power Reactors. This figure only includes grid-connected reactors. It does not include accidents on research reactors such as the SL-1 meltdown in Idaho, USA, in 1961.

world, the objective probability of an accident and the subjective assessment of risks? Did the experts err on the side of optimism? How should risk and its perception be measured?

Calculating risk

In a companion paper⁵ we made the tentative suggestion that the cost of accidents was low, indeed negligible, when compared with the value of the electricity generated. This introductory conclusion was based on a scratch calculation drawing on upper-case assumptions, multiplying the (ill) chance of a disaster of 1 in 100,000 years of a reactor's operation by damage costing €1,000 billion. The cost of such an accident amounts to €1 per MWh generated. Setting aside the back of an envelope, which lends itself to quick, simple calculations, we shall now look at the matter in greater detail, from both a theoretical and an empirical standpoint.

At a cost of €0.10 per MWh, €1 per MWh or €10 per MWh?

Is the cost per MWh of an accident negligible, at less than 10 euro cents? Or is it just visible, at about €1, or rather a significant fraction of the several tens of euros that a MWh costs to generate? How does our scratch calculation, in the companion paper⁶, compare with existing detailed assessments and how were the latter obtained?

Risk is the combination of a random event and a consequence. To calculate the risk of an accident, the probability of its occurrence is multiplied by the damage it causes⁷. Much as many other forms of disaster, a major nuclear accident is characterized by an infinitesimal probability and huge damage. A frequently used short cut likens the former to zero, the latter to infinity. As we all know, multiplying zero by infinity results in an indeterminate quantity. The short cut is easy but idiotic. The probability of an accident is not zero; unfortunately some have already occurred. Nor yet is the damage infinite. Even the worst case accident on a nuclear reactor cannot lead to the destruction of the planet and humankind. (Would the latter outcome, following a collision with an asteroid 10 kilometres in diameter, for example, or quarks going out of control in a particle accelerator, correspond to infinite damage?⁸) Mathematically, multiplying a very small number by a very large one always produces a determinate value. So nuclear risk assessments seek to approximate the two numbers and then multiply them.

In its worst-case scenario ExternE, the major European study of the external effects of various energy sources published in 1995, estimated the cost of an accident at €83 billion.

⁵ François Lévêque, Estimating the costs of nuclear power: benchmarks and uncertainties, I3 Working Paper 13-ME-01, May 2013.

⁶ François Lévêque, Estimating the costs of nuclear power: benchmarks and uncertainties, I3 Working Paper 13-ME-01, May 2013.

⁷ Other risk dimensions such as the duration and reversibility of damage are sometimes taken into account too.

⁸ In his book on catastrophes Richard A. Posner estimates it at \$600 trillion (6 followed by 14 zeros).

This estimate was based on the hypothetical case of a core melt on a 1,250 MW reactor, followed two hours later by emissions lasting only an hour containing 10% of the most volatile elements (cesium, iodine) of the core. The population was exposed to a collective dose of 291,000 person-sieverts. This contamination ultimately caused about 50,000 cancers, one-third of which were fatal, and 3,000 severe hereditary effects. In a few days it caused 138 early diseases and nine fatalities. The impact on public health accounted for about two-thirds of the accident's cost. The study also assessed the cost of restrictions on farming (lost production, agricultural capital, etc.), and the cost of evacuating and rehousing local residents. This string of figures gives only a tiny idea of all the data required to estimate the cost of a nuclear accident. It merely lists some of the main parameters, in other words those that may double, or indeed multiply by ten, the total cost of economic damage. Let us now take a closer look.

In theory the extent of emissions may reach the release of the entire contents of the reactor core. The explosion at Chernobyl released the equivalent of 30% of the radioactive material in the reactor, a huge, unprecedented proportion. Emissions from Fukushima Daiichi's three damaged reactors are estimated to have amounted to 10 times less than the amount released in Ukraine. The collective dose reflects emissions measured in person-sieverts and the sum of the radiation absorbed by groups of people subject to varying levels of exposure. The collective dose depends on emissions, but also the weather conditions and population density. Depending on whether radioactivity is deposited by rain on an area of woodland or a city, the number of people exposed will obviously vary. The person-Sv unit is used because it is generally assumed that the biological effects of radiation follow a linear trend: the health effect of exposure of 20,000 people to 1 millisievert, or 20 people to 1 sievert is consequently taken as being the same. This approach is based on the assumption that even the lowest level of exposure is sufficient to increase the number of fatalities and diseases, in particular cancers. It is controversial because it implies that the natural radioactivity that exists in some areas, such as Brittany in France, exposes local residents to a specific hazard. For our present purposes we shall treat it as an upper-case hypothesis, in relation to one setting a threshold below which ionizing radiation has no effect. Translating the collective dose into figures for fatalities and diseases then depends on the effect it is decided to use. For example positing a 5% risk factor means that out of 100 people exposed, five will be affected by the disease under consideration. The final step in assessing health effects involves choosing a monetary value for human life. Without that it is impossible to add the damage to public health to the other consequences, such as population displacement or soil decontamination. There are several methods for calculating the value of human life, based for example on the amount allocated to reducing road accidents or the average contribution of a single individual to their country's economy, in terms of gross domestic product⁹.

The assumptions used in studies of the quantity and dispersion of emissions, the collective dose received, the risk factor, the value of human life, the number of people displaced, or indeed the amount of farmland left sterile all contribute to creating substantial disparities

⁹ These methods are open to dispute and actively debated. Readers wishing to pursue the matter may refer to Chapter 4 of W. Kip Viscusi's *Rational Risk Policy*, Oxford University Press, 1998.

in estimates¹⁰. Just looking at two indicators – the number of additional cancers and the total cost of an accident – is sufficient to grasp the scale of variations. In the ExternE study cited above, the scenario corresponding to the largest volume of emissions led to 49,739 cancers and cost €₁₉₉₅83.2 billion, whereas the scenario with the lowest emissions led to 2,380 cancers and cost €₁₉₉₅3.3 billion. In a recent German study¹¹, the low-case values calculated were 255,528 cancers and €₂₀₁₁199 billion – which corresponds to frequently quoted orders of magnitude. But the high-case figures reported by the study are far larger, with 5.3 million cancers and €₂₀₁₁5,566 billion. It is unusual for experts to produce such a high estimate, with a single accident leading to millions of cancers and total damage amounting to thousands of billions of euros. However it is close to the orders of magnitude reported in the first studies carried out in the 1980s¹² after the Chernobyl disaster. Allowance must nevertheless be made for such extreme figures, which correspond to worst-case scenarios. For example in the German study just mentioned, the weather conditions were strong wind, changing in direction, and light rain (1 millimetre per hour). The rain and wind severely contaminated an area of 22,900 square kilometres (a circle about 85 kilometres in diameter), occupied by millions of people who had to be evacuated. The most catastrophic scenarios obviously correspond to accidents at power plants in densely populated areas. Some 8.3 million people live inside a 30-kilometre radius round the nuclear power plant at Karachi, Pakistan. Worldwide there are 21 nuclear plants with more than a million people living within a 30-kilometre radius around them¹³.

Rather than picking a random number from the various damage assessments, the right approach would be to take these uncertainties into account, particularly the ones which affect the collective dose and risk factor. An accident having occurred with a given quantity of emissions, we now need to plot a curve indicating, for example, that there is a 1% probability of the event causing economic losses in excess of €1,000 billion, a 10% probability of losses ranging from €500 billion and €999 billion, or indeed a 5% probability of losses below €1 billion. In conceptual terms an exercise of this sort is easy to carry out. But in practice the problem is obtaining sufficient data on variations in the determinants of damage. This is the case for weather parameters: the wind and rainfall conditions have been statistically established for each plant. But for many other factors, the scale of their variation is unknown. In which case it must be modelled on the basis of purely theoretical considerations.

There have been too few accidents in the past with significant emissions to allow observation of the statistical variations affecting their impacts, such as the frequency of cancers. We do not even know exactly how many fatal cancers followed the Chernobyl disaster. This is not so much due to a lack of epidemiologic studies or monitoring of the local population: plenty of studies have been carried out since the accident; nor yet their manipulation inspired by on some conspiracy theory. On the contrary the uncertainty is due to the fact that cancer is a very common cause of death and that cancers caused by

¹⁰ Some recent estimates (Institut de Radioprotection et de Sûreté Nucléaire, 2013) have also sought to estimate direct and indirect macroeconomic consequences, as well as the impact on society as a whole and on private individuals, through the emotional and psychological disturbance caused by an accident.

¹¹ Versicherungsforen Leipzig, Calculating a risk-appropriate insurance premium to cover third-party liability risks that result from operation of nuclear power plants, commissioned by the German Renewable Energy Federation, April 2011.

¹² *op. cit.*

¹³ Reactors, residents and risk, Nature online, 21 April 2011, doi: 10.1038/472400a.

ionizing radiation are difficult to isolate. According to the Chernobyl Forum – a group of international organizations – of the 600,000 people who received the highest radiation doses (emergency recovery workers, or liquidators, and residents of the most severely contaminated areas), 4,000 fatal cancers caused by radiation are likely to be added to the 100,000 or so cancers normally expected for a population group of this size. Among the 5 million people exposed to less severe contamination and living in the three most affected countries (Ukraine, Belarus and the Russian Federation), the Forum forecasts that the number of additional fatal cancers will amount to several thousand, 10,000 at the very most. This estimate is one of the lowest in the literature. For the whole population of the contaminated areas in the three countries, estimates vary between 4,000 and 22,000 additional deaths. It should be noted that these figures do not take into account emissions outside the officially contaminated areas, nor in other parts of these countries, nor yet in Europe or the rest of the world. The second set of estimates is less reliable, more controversial, as radiation exposure per person is very low. Only the hypothesis of a linear relationship between dose and effect leads to additional fatal cancers outside the three areas, estimated for example at 24,700 by Lisbeth Gronlund, a senior scientist at the US Union of Concerned Scientists¹⁴.

If we accept the rough estimate suggested in my previous paper¹⁵, with a €1,000 billion loss for a major nuclear accident, we would be somewhere in the upper range of estimates. But also in relation to other accidents in the past: Three Mile Island cost an estimated \$1 billion and Chernobyl several hundreds of billions of dollars¹⁶. The provisional estimate for Fukushima Daiichi is about \$100 billion¹⁷.

Calculating the frequency of major accidents

We shall now turn to the task of putting figures on the probability of a major nuclear accident. The ExternE study reports a probability of a core melt of 5×10^{-5} per reactor-year¹⁸, in other words a 0.00005 chance of an accident on a reactor operating for one year; or alternatively, due to the selected unit, a frequency of five accidents for 100,000 years of reactor operation, or indeed a frequency of one accident a year if the planet boasted a fleet of 100,000 reactors in operation. Following a core melt, two possibilities are considered: either an 8-in-10 chance that radiation will remain confined inside the reactor containment; or a 2-in-10 chance that part of the radiation will be released into the

¹⁴ How many cancers did Chernobyl really cause? Monthly Review Magazine
<http://mrzine.monthlyreview.org/2011/gronlund070411.html>

¹⁵ François Lévêque, Estimating the costs of nuclear power: benchmarks and uncertainties, i3 Working Paper 13-ME-01, May 2013.

¹⁶ Belarus, for instance, has estimated the losses over 30 years at \$235 billion. Footnote 6, on page 33 of Chernobyl's Legacy: Health, Environmental and Socio-Economic Impacts, published by The Chernobyl Forum.

¹⁷ Estimate, ranging from \$70 billion to \$250 billion, published by the Japan Center for Economic Research. There is no estimate of the effects on public health.
<http://newsonjapan.com/html/newsdesk/article/89987.php>

¹⁸ This figure is taken from the very first studies carried out by EDF in the late 1980s on its 900 MW nuclear reactors. In its subsequent, more detailed studies, the figure was 10 times lower, at about 5×10^{-6} . But in both cases the relevant studies made no allowance for external initiating factors such as earthquakes or floods.

environment. In the first case damage is estimated at €431 million, in the second case at €83.252 billion. As we do not know which of the two scenarios will actually happen, the forecast damage is calculated using its expected value, in other words: $0.8 \times 431 + 0.2 \times 83252$, which equals roughly €17 billion. This simple example illustrates two connected concepts, which are essential to understanding probabilistic analysis of accidents: conditional probability and event trees.

Conditional probability, event trees and probabilistic safety assessment

The probability that in the event of core melt the radiation will remain confined inside the reactor containment is 0.8 (8-in-10 chance). This is a conditional probability. It is commonly denoted using a vertical bar: $p(\text{release}|\text{melt})$. In a general way, A and B being two events, it is written as $p(A|B)$, which reads as 'the probability of A given B'. Conditional probability is a key concept. In Section 3 we shall see that it gave rise to a fundamental mathematical formula, known as Bayes' rule, or theorem. It enables us to update our appraisals on the basis of new information. In the present case conditional probability is used as a tool for estimating the probability of various sequences of events, and for estimating the cost of the event among their number which leads to a major accident. For example, $p(\text{release}|\text{melt}|\text{cooling system failure}|\text{loss of emergency power source}|\text{protective wall round plant breached by wave}|\text{7.0 magnitude quake})$. All the possible sequences form an 'event' tree, with a series of forks, each branch being assigned a probability, for example p for one branch and therefore $(1-p)$ for the other¹⁹. Try to picture an apple tree growing on a trellis, and the route taken by sap to convey water to each of its extremities, one of which is diseased. The event tree maps the route taken by a water molecule which reaches the diseased part without taking any of the other possible routes.

Probabilistic assessment of nuclear accidents is based on this type of tree structure. It seeks to identify all the possible technical paths leading to an accident, then assigns a probability to the faulty branch at each fork. The starting point is given by the probability of a factor triggering an accident, for example a 1-in-1,000 chance per year of a quake resulting in peak ground acceleration of 0.5g at the site of the plant. The outcome is the occurrence of core melt, or the massive release of radiation into the environment following meltdown. There may be many intermediate forks, concerning both technical (1-in-10,000 chance that a pump will break down for each year of operation) and human failures (1-in-5,000 chance that an operator disregards a flashing light on the control panel).

¹⁹ If there are only two options – such as drawing a red ball or a black ball from an urn only containing balls of these two colours – once you know the probability of one option, the other can be deduced, the sum of the two being equal to 1. So if the probability of drawing a red ball is $1/3$, the probability of not drawing a red ball is $2/3$; as all the non-red balls are black, the probability of drawing a black ball is $2/3$.

The first large-scale probabilistic assessment²⁰ was carried out in the US in the 1970s. It was led by Norman Rasmussen, then head of the nuclear engineering department at the Massachusetts Institute of Technology. The study was commissioned by the Atomic Energy Commission, which was keen to reassure public opinion by demonstrating that the risks, albeit real were actually infinitesimal. To this end it circulated a misleading summary of the study, which over-simplified its findings. To impress readers the document made dubious comparisons – not contained in the main report – with other risks. For example it asserted that the likelihood of a person dying as a result of a nuclear accident was about the same as being hit by a meteorite. Such distortion of the report and some of the errors it contained prompted a major controversy²¹. The Nuclear Regulatory Commission, which was set up to separate nuclear safety from the other missions allocated to the AEC, rejected the contents of the report summarized in 1979. But ultimately the Rasmussen study is remembered for the work done establishing a detailed method, rather than the values it calculated for probabilities and damages. Since then probabilistic assessment has become more rigorous and an increasing number of such safety studies have been carried out.

The first probabilistic safety assessments changed several deeply rooted beliefs. They highlighted the possible input of operators, capable of either interrupting a sequence of material faults, or in some cases making it worse. Accidents and preventive measures do not only have a technical dimension. Rasmussen and his fellow scientists showed that the loss of liquid through small breaks in the primary cooling system could also be a frequent cause of accidents, largely disregarded until then. Several years later the Three Mile Island disaster prompted new interest in probabilistic safety assessment. Since then its scope has broadened and it has grown more complex. It has taken into account additional factors which may initiate an accident, both natural and human (for example the risk of falling aircraft). Probabilistic safety assessments have now been carried out on all the nuclear plants in the US and many others worldwide. Similarly reactor engineering firms carry out such studies for each reactor model while it is still in the design stage.

The main purpose of probabilistic safety assessments is not to estimate the probability of an accident on a specific plant or reactor, but rather to detect exactly what may go wrong, to identify the weakest links in the process and to understand the faults which most contribute to the risk of an accident. In short such studies are a powerful instrument for preventing and ranking priorities, focusing attention on the points where efforts are required to improve safety. But our obsession with single numbers has pushed this goal into the background and all we remember of these studies is the final probability they calculate, namely *core-melt frequency*. This bias is all the more unfortunate because the overall result is rarely weighted by any measure of uncertainty. If no confidence interval is indicated, we do not know whether the figure reported – for example one accident per 100,000 reactor-years – is very close to the mean value – for example an 8-out-of-10 likelihood that the accident frequency ranges from 0.9 to 1.1 accidents per 100,000 reactor-years – or more widely spread, with an 8-out-of-10 likelihood that the frequency ranges from 0.1 and 10 accidents per 100,000 reactor-years. Intuitively it is not the same

²⁰ US National Regulatory Commission, Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants [NUREG-75/014 (WASH-1400)], October 1975.

risk, but in both cases the mean frequency is the same. In the second case there is some likelihood that more than 10 accidents will occur per 100,000 reactor-years, whereas in the first case this risk is almost non-existent. Of the studies which have estimated the uncertainty, it is worth noting an NRC appraisal carried out in 1990 on five plants. It found, for example, that for the two PWRs at Surry (Virginia) the confidence interval for the mean frequency of 4.5×10^{-5} ranged from 1.3×10^{-4} at its upper limit to 6.8×10^{-6} at the lower limit. The first figure indicates that there is a 5% chance that the value of the frequency may be even greater; the second that there is a 5% chance it may be lower. In other words, there is a 90% chance that the core-melt frequency is somewhere between the two limits. However this interval was not calculated. It was based on the judgement of various experts who were questioned.

The widespread lack of any mention of the distribution on either side of the mean may be explained by the method employed. The probabilities assigned to the various branches of the tree used to calculate the overall core-melt frequency are selected as best estimates. The final number is single, because it is the sum of a succession of single numbers. Safety specialists naturally know how to use statistics and calculate uncertainties. They do not make do with averages. But they are concerned with the details, because it is here that there is scope for improving safety, for example using a probability density function to model the failure of a particular type of pump. At this scale, the error and distribution parameters – with barbaric names such as standard deviation, mode, variance or kurtosis – are generally entered. So why are they not systematically used to obtain an overall core-melt probability expressed as more than just a single number? The first reason is that uncertainty propagation in an event tree is far from trivial, compared to addition or multiplication. It is not just a matter of adding up or combining the standard deviations at each fork to obtain the one governing core-melt frequency. The second reason is that the prime aim of probabilistic safety assessments is not to obtain a final result. Rather they focus on the details of each branch. Only in recent years have the specialists started paying sustained, systematic attention to presenting the uncertainty affecting the aggregate probability of an accident.

Non-specialists are consequently inclined to think that probabilistic safety assessments reveal the true value for accident frequency for a given reactor, whereas in fact this value is subject to uncertainty. At best it is possible to offer a confidence interval within which the probability of an accident will fall.

It is worth noting that with advances in reactor technology the probability of a core melt has dropped. For example, on its 1,300 MW *palier*, or step, EDF estimated the core-melt frequency as 7.2×10^{-6} . On the following generation, represented by EPR, the results of safety studies carried out by Areva and vetted by the British regulator show a core-melt frequency of 2.7×10^{-7} per reactor-year²², almost 20 times lower. In the US an Electric Power Research Institute study²³ found that the mean core-melt frequency of the [US] fleet had dropped by a factor of five since the early 1990s.

²¹ See in particular the criticism voiced by the Union of Concerned Scientists, published in *The Risks of Nuclear Power Reactors: a Review of the NRC Reactor Safety Study*, Henry W. Kendall (dir); Richard B. Hubbard, Gregory C. Minor, 1977, 210 pages.

²² HSE Health and Safety Executive, Nuclear Directorate, Generic Design Assessment – New Civil Reactors Build, Step 3, PSA of the EDF and Areva UK EPR division 6, Assessment report n° AR 09/027-P, 2011.

²³ Safety and Operational Benefits of Risk-Informed Initiatives, EPRI White Paper, February 2008.

Divergence between real-world observation of accidents and their frequency as predicted by the models

The Fukushima Daiichi disaster revealed an order-of-magnitude difference between the accident frequencies forecast by probabilistic safety assessments and observed frequencies²⁴. Since the early 1960s and grid-connection of the first nuclear reactor, 14,400 reactor-years have passed worldwide. This figure is obtained by adding up all the years of operation of all the reactors ever built that generated kWhs, whether or not they are still in operation, were shut down earlier than planned or not. In other words the depth of observation currently at our disposal is equivalent to a single reactor operating for 14,400 years. Given that the global fleet currently numbers about 500 reactors, it may make more sense to say that this depth is equivalent to 28.8 years for 500 reactors. At the same time, since grid-connection of the first civil reactor, 11 partial or total core-melt accidents have occurred, of which three at Fukushima Daiichi. So the recorded core-melt frequency is 11 over 14,400, or 7.6×10^{-4} , or indeed an accident for every 1,300 reactor-years. Yet the order of magnitude reported by probabilistic safety studies ranges from 10^{-4} to 10^{-5} , or an accident every 10,000 to 100,000 reactor years. Compared to 1,300 that means a ten to hundred-fold divergence between calculated and observed probabilities.

How can such a large divergence be justified? The reasons are good or bad, trivial or complicated.

The first possible reason is simply bad luck. Just because you score six, five times running with the same die, it does not mean it is loaded. There is a 1-in-7,776 chance of this sequence with a perfectly balanced die. So experts firmly convinced of the accuracy of their models or passionate advocates of nuclear power may set aside the suggestion that the calculated frequencies are erroneous, despite them being much lower. Much as with the die, 14,400 reactor-years is not enough to obtain a true picture. This reason is legitimate in principle but it does mean ignoring observations if there are only a limited number. All in all it is not very different from the opposite standpoint which consists in discarding probabilistic safety assessments and only accepting observations. The right approach is to base our reasoning on data obtained from both observation and modelling. Faced with uncertainty, all data should be considered, whether obtained from the field or from laboratories. We shall examine this approach in greater depth in the section devoted to Bayesian revision of probabilities.

A variation on the bad luck theory is to point out that the observed frequency actually falls within the range predicted by probabilistic assessments. As we saw above, in their forecast for the Surry plant, the experts estimated a 5% likelihood that the core-melt frequency would exceed 1.3×10^{-4} , in other words an accident every 769 reactor-years. The observed value of an accident every 1,300 reactor-years is actually lower than this limit value. So convinced experts have no reason to review their position: there is no divergence between observations and the model. This stance might carry some weight if core-melt frequencies were reported with a confidence interval, but this is not the case. Moreover the previous comment still holds true: the upper and lower limits on uncertainty must move to accommodate fresh observations.

²⁴ See Cochran and Repussard.

The second possible reason is that the probabilistic assessments are not exhaustive. The event trees they examine and assess do not cover all the possible scenarios. The first safety studies only focused on internal initiating events, such as a device failure. The sequences of faults initiated by an earthquake or flood which might lead to core-melt have only gradually been taken into account. The validity of calculated frequencies is restricted to the perimeter under study. If no allowance is made for the risk of a falling aircraft the frequency is lower. The studies which do take it into account estimate that it is lower than 10^{-7} per reactor-year²⁵. On its own this figure is too small to significantly change core-melt frequency which is much higher. All this example shows is that by adding scenarios the frequency gradually increases. Little streams make big rivers. The Fukushima Daiichi accident is a concrete illustration of missing scenarios. It made people realise that spent-fuel pools could cause a massive release of radioactive material into the atmosphere. Probabilistic safety assessments do not usually register a break in the water supply to these pools as a possible initiating event. At Fukushima Daiichi, much as other Japanese nuclear plants, the possibility that two risk factors – an earthquake and a tsunami – might coincide was apparently not studied either. Readers may be surprised by this oversight. Tidal waves and quakes are frequent in Japan and the two events are connected: one triggers the other. In fact the scenario which was not considered (nor its probability assessed) was the failure of the regional electricity network, knocked out by the quake, combined with flooding of the plant, due to the tsunami. With the surrounding area devastated, the backup diesel pumps underwater and the grid down, the power plant was left without an electricity supply for 11 days. Safety assessments generally assume power will be restored within 24 hours.

But is it possible for probabilistic studies to take into account all possible scenarios? Obviously not. It is impossible to imagine the unimaginable or to conceive the inconceivable: joking apart, there is an intrinsic limit to probabilistic analysis: it can be applied to risk and uncertainty, not to situations of incompleteness.

²⁵ EDF studied the risk of falling aircraft for the Flamanville 3 reactor. The probability of an attack on one of its safety functions is 6.6×10^{-8} per reactor-year. See EDF PSA <http://www.edf.com/html/epr/rps/chap18/chap18.pdf>.

Risk, uncertainty and incompleteness

These basic concepts may be explained using the example of an urn containing different coloured balls. We shall start from what is a certainty. It may be described using the case of an urn only containing balls of the same colour, red for instance. We know that if we pick one ball out of the urn, it will inevitably be red. The outcome is a foregone conclusion. Risk corresponds to an urn of which the content is known, 30 red balls and 60 white, for instance. We can no longer be sure of picking out a red ball, but we do know that we have a 1-in-3 chance of picking a red, a 2-in-3 chance of picking a white. In the theoretical jargon, we would say that all the states of the world (or indeed the universe of events) are known with certainty, and for each state or event there is a corresponding probability, also known with certainty. Uncertainty may be represented by an urn known to contain 30 red balls, but we do not know whether the 60 others are black or white. So only the probability of picking a red ball (1-in-3) is known. On the other hand all the states of the world (picking a red, black or white ball) are known. There is no possibility of a surprise, such as picking a blue ball. Lastly, incompleteness corresponds to an urn full of balls of unspecified colours. We may pick a white or a purple ball, perhaps even a multicoloured one. Unlike risk and uncertainty, in a situation of incompleteness all the states of the world are not known. So probability theory cannot apply. A probability cannot be assigned to an unknown event²⁶.

This presentation makes a distinction between uncertainty and risk. However this vocabulary is not universally accepted and must be handled with care. The term 'uncertainty' is often used with a broader sense, which encompasses the notion of risk. The part of uncertainty which is not covered by the term 'risk' is then referred to as 'ambiguity' or 'non-specific uncertainty'²⁷. But such quibbles are of secondary importance, the priority being to draw a line between situations in which we have probabilities for all the events under consideration, and situations in which we do not. In the second case it is necessary to make assumptions in order to assign probabilities to events for which they are unknown. Returning to the urn containing 30 red and 60 black or white balls. A simple way of assigning a probability to picking a black ball and a white ball would be to posit that the two events are of equal likelihood, namely a 1-in-2 chance of picking a black or white ball from among the 60 which are not red²⁸. The probability of picking a black, white or red ball, from among the 90 balls in the urn, is 1-in-3 (30/90). In other words, ignorance is treated by assuming equiprobability: if n outcomes are possible, and if we have no idea of their chances of occurring, we consider them to be equiprobable (equally probable) and equal to $1/n$. This approach provides a way of treating non-specific uncertainty as a risk, thus making it possible to apply the calculation of probabilities to situations of uncertainty in general (in other words including non-specific uncertainty).

²⁶ Nor is it possible to assign probabilities to the share of known events in an incomplete whole. The sum of all probabilities must be 1, but as the probability associated with the sub-set of disregarded events is not known, it cannot be subtracted from 1 to obtain the probability assigned to known events.

²⁷ In an even more comprehensive definition of uncertainty, this term also encompasses incompleteness. The term 'radical uncertainty' may be used to distinguish this form of uncertainty from non-specific uncertainty.

²⁸ The intuition underpinning this hypothesis is that all the distributions of black and white balls (1 white, 59 black; 2 white, 58 black; ...; 30 white, 30 black; 59 white, 1 black) are possible, and that there is nothing to suggest that one is more probable than another. It is no more likely that the distribution contains more black than white balls, than the opposite.

The third reason is that every event is unique, so probability theory cannot apply. The divergence between the observed frequency of accidents and their calculated probability is not a matter of bad luck, but results from the impossibility of applying probability theory to exceptional or one-off events. This reason is intuitive but must nevertheless be discarded.

In our minds the concept of probability is associated with that of frequency, and the law of large numbers too. We all learned at school that probability is the ratio of favourable or unfavourable outcomes to the number of possible cases. We all remember having to apply this definition to the observation of rolling dice or drawing cards. We also recall that calculating a probability requires the operation to be repeated a large number of times. We need to toss a coin several dozen times to grasp that it will land on one or the other side roughly the same number of times. This frequency-based approach to calculating probabilities is the best known and it does not work without data. But there are other ways of, or theories for, analyzing probabilities, which do away with the need for repeated experiments, and consequently a large number of observations, to calculate frequencies. It is quite possible to carry out probabilistic analysis without any observation at all. The reader may recall the wager made by Pascal. The thinker was puzzled about the right approach to adopt regarding uncertainty of the existence of God. Here there could be no repeated events. The existence of a Supreme Being was a singular proposition, which Pascal subjected to probabilistic reasoning.

So probabilistic logic can be applied to one-off events. The concept of probability refers to a degree of certainty regarding the veracity of a proposition. It applies for instance to the reasoning of a court judge who takes a different view of the guilt of the accused if it is known that the latter lacks an alibi or that traces of his or her DNA have been found on the body of the victim. John Maynard Keynes defined probability theory thus: 'Part of our knowledge we obtain direct; and part by argument. The theory of probability is concerned with that part which we obtain by argument, and it treats of the different degrees in which the results so obtained are conclusive or inconclusive.' This reference to the author of *The General Theory of Employment, Interest and Money* [1936] may surprise readers unfamiliar with the work of the Cambridge economist. *A Treatise on Probability* [1921] was one of the first works published by Keynes, but is still well worth reading even now.

The theory of subjective probability opens up a second approach to probability, not based on frequency. It has been advocated and developed by three leading figures in economic science: Britain's Frank Plumpton Ramsey, already cited with reference to discount rates; Bruno de Finetti, from Italy; and an American, Leonard Jimmie Savage. These authors understand the concept of probability as the belief which an individual invests in an event, regardless of whether or not the event recurs. According to Finetti, probability is the degree of confidence of a given subject in the realization of an event, at a given time and with a given set of data²⁹. So probability is not an objective measurement, because it depends on the observer and his or her knowledge at that time. Probability is thus assimilated to the odds a punter will accept when betting on a given outcome, odds at which he or she will neither lose nor win money. Imagine, for example, that two nuclear experts are asked to bet on the likelihood of a nuclear accident occurring in Europe in the

²⁹ Quoted by Simona Morini, Bruno de Finetti: l'origine de son subjectivisme, *Journal Electronique d'Histoire des Probabilités et de la Statistique*, Vol. 3, n°2, December 2007, p. 1-16.

course of the next 30 years. One agrees to odds of 100 to 1, the other 120 to 1, or 200 to 1. So according to subjective probability theory anyone can bet on anything. However we should not be misled by the terms 'belief' and 'subjective'. The odds accepted do not depend on a person's mood or state of mind, but are supposed to be based on their knowledge. Furthermore the person making the wager is deemed to be rational, being subject to the rules governing the calculation of probabilities. For example he or she cannot bet 3 to 1 for and 1 to 2 against a given outcome. The theory of decision under uncertainty, a monument elaborated by Savage in the middle of the last century, assumes that economic agents comply with all the axioms for calculating probabilities: they must behave as perfect statisticians.

The fourth reason for the divergence between observations and forecasts is that the models may be faulty or use the wrong parameters. It is vital to avoid mistakes when assigning a probability to the known states of the world, and consequently when measuring the probability and associated uncertainty. Returning to what happened at Fukushima Daiichi, the six reactors at the plant were commissioned in the 1970s. We do not have access to the probabilistic studies carried out or commissioned by the plant operator, Tepco, at the time of the plant's construction or afterwards. But we do know some of the values taken into account by the generating company or the regulator for the risk of an earthquake or tsunami. The figures were largely underestimated. The plant was designed to withstand a magnitude 7.9 earthquake and a 3.1 metre tidal wave. On 11 March 2011 at 14:46 JST it was subjected to a magnitude 9 tremor, then swamped by a wave more than 10 metres high. So, much as Tepco, the nuclear industry is purportedly inclined to play down risks by picking values or models favourable to their growth. Unless one is an adept of conspiracy theories, this explanation for the divergence between observations and forecasts is barely convincing. For many years safety authorities and independent experts have scrutinized such probabilistic studies.

The wrong values taken into account for Fukushima Daiichi

When the nuclear power plant at Fukushima Daiichi was built the risk of an earthquake exceeding magnitude 8 on the Richter scale was estimated³⁰ as less than 2×10^{-5} per year. This value was taken from the work of modelling and numerical simulation carried out for each plant in Japan by the National Research Institute for Earth Science and Disaster Prevention (NIED). However historical research has identified six major quakes which have occurred on the Sanriku coast since 869. That was the year of the Jogan undersea earthquake, probably the most devastating ever known on this stretch of coastline until March 2011. The various pieces of evidence which have been gathered suggest that these quakes all exceeded magnitude 8. This means the observed annual frequency should be about 5×10^{-3} , in other words 100 times higher than the results calculated by the NIED³¹.

³⁰ Woody Epstein, A Probabilistic Risk Assessment Practitioner looks at the Great East Japan Earthquake and Tsunami, A Ninokata Laboratory White Paper, 29 April 2011.

³¹ The estimated frequencies for the other NPPs on the same coastline do not display such a large difference. At Fukushima Daini the estimated frequency is 10 times lower than the historic frequency; at Onagawa the two frequencies converge. The differences are due to the choice of too fine a mesh to differentiate seismic risks in the Sanriku area. See the article by Geller published in 2011 in Nature.

The protective seawall at Fukushima Daiichi was built in 1966. It was six metres high, a dimension decided in line with a classical deterministic principle: take the most severe historic event and add a twofold safety margin. The three-metre wave which struck the coast of Chile in 1960 was taken as a baseline value. This was a surprising choice, the Jogan quake having triggered a four-metre wave locally, a fact that was already known when the plant was built³². Be that as it may, 40 years later, our understanding of past tsunamis had made significant progress but the initial height of the wall was considered to comply with the guidelines for tsunami assessment set by the regulatory authority³³. Such compliance was based on an annual probability of less than 10^{-4} . Yet historical studies have established that waves exceeding eight metres have been recorded on the Sanriku coast. Witness the stones set into the ground marking the points furthest inland reached by the flood. Some of these stones are more than 400 years old. The inscriptions on them urge residents not to build homes lower down the slope. Core samples have revealed sedimentary deposits left by previous tsunamis. At the Onagawa plant, on the basis of remains found in the hills one kilometre inland, it was estimated that the 1611 quake had caused a six-to-eight-metre wave³⁴. On the evidence of old records Epstein estimated that the average frequency of a tsunami of eight metres or more on the Sendai plain, behind the Fukushima Daiichi plant, was about one every 1,000 years. He estimated as 8.1×10^{-4} the probability of an earthquake of a magnitude equivalent to or greater than a seismic intensity of 6³⁵ or more followed by a tidal wave of over eight metres. Due to the layout of the plant, this dual shock would almost certainly lead to flooding of the turbine building, destruction of the diesel generators, loss of battery power and station blackout lasting at least eight hours. This sequence of events would correspond, according to Epstein, to a probability of core-melt of about 10^{-4} , five times higher than the authorized limit. Tepco based its various probabilistic studies³⁶ on very low values, very probably taken from a badly designed historical database or unsuitable simulation models.

Two points are of particular note in this long list of possible reasons for the divergence between the observed frequency and calculated probability of an accident. Firstly, with regard to method, we should make use of all the available instruments, combining empirical and theoretical data. Assessments of the risk of a major nuclear accident based exclusively on either data from past observations or theoretical probabilistic simulations lead to a dead end. Secondly, in practical terms, the limitations of probabilistic assessment lead to the deployment of deterministic safety measures. As we do not know all the states of the world, nor yet the probability of all those we do know, it is vitally important to install successive, redundant lines of defence as a means of protection against

³² W. Epstein p. 52.

³³ 2002 guidelines. Tsunami Assessment Method for Nuclear Power Plants in Japan, published by the Tsunami Evaluation Subcommittee, Nuclear Civil Engineering Committee, Japan Society of Civil Engineers (JSCE). 'We have assessed and confirmed the safety of the nuclear plants [at Daiichi] based on the JSCE method published in 2002.' [Tsunami Study for Fukushima 1 and 2, p. 14]. Written statement by Tepco submitted to the regulator, cited by Woody Epstein p. 24.

³⁴ W. Epstein.

³⁵ On Japan's Shindo scale of seismic intensity. Seismic scales not being strictly speaking comparable with one another, seismic intensity 6 on the Shindo scale is roughly equivalent to between magnitude 6 and 7 on the Richter scale.

unforeseeable or ill appraised events. In short we must protect ourselves against the unknown. Among others³⁷ we may raise protective walls or quake-resistant structures able to withstand events twice as severe as the worst recorded case, install large numbers of well protected backup diesel generators, and build nesting containments – the first one in the fuel cladding itself, forming an initial barrier between radioactive elements and the environment – the largest one, made of reinforced concrete encasing the reactor and steam-supply system.

Perceived probabilities and aversion to disaster

The concept of subjective probability can be a misleading introduction to individuals' perception of probabilities. Perceived probability, much as its subjective counterpart, varies from one person to the next but it does not take the same route in our brain. The former expresses itself rapidly, effortlessly, in some sense as a reflex response; the latter demands time, effort and supervision. Perceived probability is based on experience and routine, whereas subjective probability is rooted in reason and optimization. In response to the question, 'Which is the most dangerous reactor: the one for which the probability of an accident over the next year is 0.0001; or the one which has a 1 in 10,000 chance of having an accident over the next year?', most people will spontaneously pick the second answer. Yet reason tells us that the two reactors are equally dangerous (i.e. $1:10,000 = 0.0001$). So how do we perceive risks? Are individuals poor statisticians, or perhaps not statistically minded at all? In which case, what use is the theory of decision under uncertainty, given that it requires us to calculate probabilities accurately? How do individuals make a choice if they do not optimize their decisions? These questions are central to 40 years' work by experimental cognitive psychology on how individuals assess the probability of events. Understanding this work is essential, for almost all its findings contribute to amplifying the perception of nuclear risk.

Biases in our perception of probabilities

What is the connection between cognitive psychology – which is an experimental science – and economics? In fact, there is a significant link. In 2002, in Stockholm, Daniel Kahneman was awarded the most coveted distinction to which an economist can aspire: the Sveriges Riksbank prize in Economic Sciences in memory of Alfred Nobel. His work was distinguished for 'for having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms'. That year, the winner of the Nobel prize for economics was not an economist, but a psychologist!

Economic analysis focuses on many subjects and lends itself to many definitions: decision theory is among these subjects; it may, among others, be defined as the science of human behaviour. Economics investigates how humans seek the best means of achieving their aims (Lionel Robins, 1932). When an individual needs to take a decision under uncertainty, he or she assigns more or less weight to the probabilities associated with each

³⁶ It is of little concern here whether this was done knowingly or not. This point is addressed in Part III.

³⁷ For a detailed discussion of the concept of defence in depth see Part III of the book.

choice. A rational being faced with an alternative between an action yielding a satisfaction of 100 associated with a probability of 0.3, and an action yielding a satisfaction of 105 associated with a probability of 0.29, will choose the second option because its mathematical expectation is greater ($100 \times 0.3 < 105 \times 0.29$). The expected utility of the payout is thus optimized.

The Swiss mathematician Daniel Bernoulli laid the foundations of expected utility theory. In his 1738 essay he raised the question of how to give formal expression to the intuition according to which the rich were prepared to sell insurance to the poor, and that the latter were prepared to buy it. He also sought to resolve an enigma of great interest at the time, subsequently known as the Saint Petersburg paradox: why does a gamble which holds an infinite expectation of gain³⁸ not attract players prepared to bet all they own? The answer to both these questions is to be found in our aversion to risk. This psychological trait means that we prefer a certain gain of 100 to an expected gain of 110. This is explained by the fact that a rich person attaches less value to a sum of €100 than a poor person does, which is translated into the contemporary jargon of economics as the declining marginal utility of income. In mathematical terms it is represented by the concave form of the utility function. (The curve, which expresses our satisfaction depending on the money we own, gradually flattens out. Bernoulli thus used a logarithmic function to represent utility).

The connection between risk aversion and the form of the utility function may not be immediately apparent to the reader. In which case, illustrating the point with figures should help. Let us assume that €100 yields a satisfaction of 1; €200 yields a proportionately lower satisfaction, 1.5 for example; €220 less still, say 1.58. I offer you €100 which you either accept or we toss a coin with the following rule: heads, I keep my €100; tails, you take the €100 and I add €120 more. Which option would you choose? The certainty of pocketing €100, or a one in two chance of winning €220?. In the first case your gain is €100, whereas in the second case there is the expectation of €110 ($€220 \times 1/2$). But what matters to you is not the money, but the satisfaction – or utility – it yields. The first case yields a utility of 1, the second an expected utility of 0.79 ($1.58 \times 1/2$). So you choose the first option, which does not involve any risk. With the resolution of the Saint Petersburg paradox, by altering the form of the utility function, Bernoulli opened the way for progress towards decision theory, which carried on to Kahneman. This was achieved through a back-and-forth exchange between economic modelling and psychological experimentation. The latter would, for instance, pick up an anomaly – in a particular instance people's behaviour did not conform to what the theory predicted – and the former would repair it, altering the mathematical properties of the utility function or the weighting of probabilities. The paradoxes identified by Allais and Ellsberg were two key moments in this achievement.

During a university lecture in 1952 Maurice Allais, a Professor at the Ecole des Mines in Paris, handed out a questionnaire to his students, asking them to choose between various simple gambles, arranged in pairs. He then collected their answers and demonstrated that

³⁸ Imagine the following game in a casino: a coin is tossed a number of times, with an initial kitty of \$1, which doubles each time the coin comes up tails. The game stops and the player pockets the kitty when the coin lands heads up. At the first toss, the winner pockets \$1 if the coin lands heads up; if it is tails, the game continues, with a second toss. If the coin then lands heads up, the player pockets \$2, otherwise it is tossed a third time and so on. The expectation of gain equals $1 \times 1/2 + 2 \times 1/4 + 4 \times 1/8 + 8 \times 1/16 + \dots$ or $1/2 + 1/2 + 1/2 + 1/2 + \dots$, making an infinite amount provided the casino has unlimited resources.

they had contradicted expected utility theory in what was then its most advanced form, as developed by Leonard Jimmie Savage. Their responses violated one of the theory's axioms, which, along with the others, was supposed to dictate the rational decision under uncertainty. In simple terms the Allais paradox may be expressed as follows³⁹: the first gamble offers the choice between (A) the certainty of winning €100 million, and (B) winning €500m with a probability of 0.98, or otherwise nothing. The second gamble offers the choice between (C) receiving €100 million with a probability of 0.01, or otherwise nothing, and (D) €500 million with a probability of 0.0098, or otherwise nothing. The paired gambles are therefore the same, but with a probability 100 times lower. Most of the students chose A, not B, but D rather than C. It seemed to them that the probability of winning €500 million with D (0.0098) was roughly the same as that of winning €100 million with C (0.01), whereas in the previous case the same difference of 2% between probabilities seemed greater. Yet, according to Savage, rational behaviour would have dictated that if they chose A rather than B, they should also pick C, not D. Ironically, Savage also attended the lecture and handed in answers to the questionnaire which contradicted his own theory.

A common solution to the Allais paradox is to weight the probabilities depending on their value, with high coefficients for low probabilities, and vice versa. Putting it another way, the preferences assigned to the probabilities are not linear. This is more than just a technical response. It makes allowance for a psychological trait, which has been confirmed since Allais' time by a large number of experimental studies: people overestimate low probabilities and underestimate high probabilities. In other words they tend to see rare events as more frequent than they really are, and very common events as less frequent than is actually the case.

To our preference for certainty, rather than risk, and our varying perception of probabilities depending on their value, we must add another phenomenon well known to economists: our aversion to ambiguity. This characteristic was suggested by Keynes, and later demonstrated by Daniel Ellsberg in the form of a paradox. In his treatise on probabilities the Cambridge economist posited that greater weight is given to a probability that is certain than to one that is imprecise. He illustrated his point by comparing the preferences for a draw from two urns containing 100 balls. One contains black and white balls, in a known, half-and-half proportion; the other urn also contains black and white balls but in an unknown proportion. In the second urn, all distributions are possible (0 black, 100 white; 1 black, 99 white; ...; 100 black, 0 white) with equal probability ($p = 1/101$). So the expected probability of drawing a white ball is also $1/2$ ⁴⁰, the same value as for the probability of drawing a white ball from the first urn. But we would rather win, by drawing white (or black) balls from the first urn.

In 1961 Ellsberg revisited Keynes' example, developing it and making experiments. We shall now look at his experiment with an urn containing balls of three different colours. In all it contains 90 balls, of which 30 are red and the remaining 60 either black or yellow. So all we accurately know is the probability of drawing a red ($1/3$) and that of drawing a black or yellow – in other words one that is not red – ($2/3$). You are presented with two pairs of gambles. In the first pair, (A) you win €100 if you draw a red ball from the urn, (B)

³⁹ Gigerenzer, *Heuristics and the Law*, p. 32.

⁴⁰ It equals $(1/101)(100/100 + 99/100 + \dots + 1/100 + 0/100)$.

you win €100 if you draw a black ball from the urn. Which option do you choose? If, like most people, you are averse to ambiguity, you choose A. The other pair is more complicated: (C) you win €100 if you draw a red or a yellow ball from the urn, (D) you win €100 if you draw a black or a yellow ball from the urn. The question, however, is the same: which option do you choose?. The answer is also the same: aversion to ambiguity prompts most people to pick D rather than C. You know very well that you have a 2/3 chance of drawing a black or yellow ball (one that is not red). The paradox resides in the fact that the preference for A and D is inconsistent on the part of a rational person, as modelled in the classical theory of expected utility. Preferring A to B implies that the player reckons subjectively that the probability of drawing a black ball is lower than 1/3, whereas for a red ball it is higher than 1/3. Knowing that there is a 2/3 probability of drawing a black or yellow ball, the player deduces that the probability of drawing a yellow ball is higher than 1/3. As the probability of drawing a yellow or black ball is higher than 1/3 in both cases, their sum must exceed 2/3. According to Savage, the player picks C. But in the course of experiments, most players who choose A also choose D, which suggests that there is an anomaly somewhere, which the aversion to ambiguity corrects.

Just as there is a premium for taking risks, some compensation must be awarded to players for them to become indifferent to gain (or loss) with a 1/2 probability or an unknown probability with an expected value of 1/2. Technically speaking there are several solutions for remedying this problem, in particular by using the utility function, yet again⁴¹. It is worth understanding Ellsberg's paradox because ambiguity aversion with regard to a potential gain, has its counterpart with regard to a loss: with the choice between a hazard associated with a clearly defined probability – because the experts are in agreement – and a hazard of the same expected value – because the experts disagree – people are more inclined to agree to exposure to the first rather than the second hazard. Putting it another way, in the second instance people side with the expert predicting the worst-case scenario. Simple intuition enables us to better understand this result. If players prefer (A) the prospect of drawing a red ball associated with the certainty of a 1/3 probability, rather than that (B) of drawing a black ball, it is because they are afraid that in the latter case the person operating the experiment may be cheating. The latter may have put more yellow balls in the urn than black ones. The experimental proposition is suspect⁴². Out of pessimism, the players adjust their behaviour to suit the least favourable case, namely the absence of black balls in the urn, comparable to the worst-case scenario among those proposed by the experts.

Kahneman's work followed on from previous research, but it also diverged in two respects.

He and his fellow author, Amos Tversky⁴³, introduced two changes to the theory of expected utility⁴⁴. Firstly, individuals no longer base their reasoning on their absolute


⁴¹ The Choquet integral.

⁴² The aim is not to explain Ellsberg's paradox: an individual who finds this behaviour suspect for the first pair of gambles, should subsequently prefer C rather than D, unless he or she also suspects that the experimenter has changed the content of the urns in the meantime!

⁴³ Amos Tversky and Daniel Kahneman co-authored a great many academic papers, in particular those awarded the prize in economics by Sweden's central bank. But Tversky died at the age of 59 and was consequently not awarded the Nobel prize for economics alongside Kahneman.

⁴⁴ See Daniel Kahneman and Amos Tversky, 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, XLVII (1979), p. 263-291; and Daniel Kahneman and Amos Tversky, 'Advances in

wealth, but in relative terms with regard to a point of reference. For example, if your boss gives you a smaller rise than your fellow workers, you will perceive this situation as a loss of utility, rather than a gain. The value function does not start from zero, corresponding to no wealth, subsequently rising at a diminishing rate as Bernoulli indicated with his model. A whole new part of the curve, located to the left of the zero, represents the value of losses in relation to the status quo. This part is convex (see Figure 1): just as we derive less satisfaction from our wealth rising from 1,000 to 1,100, than when it rises from 100 to 200, so our perception of the loss between -1,000 and 1,100 is less acute than between -100 and -200.



Secondly individuals are more affected by loss than gain. For example if a teacher randomly hands out mugs marked with their university shield to half the students in a lecture theatre, giving the others nothing, the recipients, when asked to sell their mugs, will set a price twice as high as the one bid by those who did not received a gift⁴⁵. Reflecting such loss aversion the slope of the value function is steeper on the loss side than on the gain side. Regarding the perception of probabilities Kahneman and Tversky revisit the idea of distortion, in particular for extreme values (overweighting of low probabilities, underweighting of high probabilities). Armed with these value functions and probability weighting, the decision-maker posited by the two researchers assesses the best option out of all those available and carries on optimizing the outcome.

Kahneman has also experimented widely and published on heuristics, the short-cuts and routines underpinning our decisions. In so doing he distances himself from previous work on optimal decision-making under uncertainty. The decision-maker no longer optimizes nor maximizes the outcome. Anomalies in behaviour with regard to expected utility theory are no longer detailed to enhance this theory and enlarge its scope, but rather to detect the ways we react and think. Here the goal of research is exclusively to describe and explain: in the words of the philosophers of science it is positive. The aim is no longer to propose a

prospect theory: Cumulative representation of uncertainty'. *Journal of Risk and Uncertainty* 5 (4), 1992, p. 297-323.

⁴⁵ This type of experiment, simplified here, is presented in detail in Daniel Kahneman, Jack L. Knetsch and Richard H. Thaler, Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias, *Journal of Economic Perspectives*, Vol. 5, n°1, 1991, p. 193-206.

normative theoretical framework indicating how humans or society should behave. Observing the distortion of probabilities becomes a way of understanding how our brain works. This line of research is comparable to subjecting participants to optical illusions to gain a better understanding of sight. For example a 0.0001 probability of loss will be perceived as lower than a 1/10,000 probability. Our brain seems to be misled by the presentation of figures, much as our eyes are confused by an optical effect which distorts an object's size or perspective. This bias seems to suggest that the brain takes a short-cut and disregards the denominator, focusing only on the numerator.

Psychology has also done a great deal of experimental work revealing a multitude of micro-reasoning processes, which sometimes overlap, their denomination changing from one author to the next. Economists may find this confusing, much as psychologists are often thrown by the mathematical formalism of decision theory. Readers may feel they are faced with a choice between two unsatisfactory options. On the one hand, an economic model has become increasingly complex as it has been built up with the addition of post hoc hypotheses, the theory of expected utility having gradually taken onboard aversion to risk, ambiguity, loss, while nevertheless remaining a schematic, relatively unrealistic representation. On the other hand, a host of behavioural regularities, identified through observation, throw light on the countless facets of the decision-making process but, lacking a theoretical basis, simply stand side by side. In one case we have a discipline which is still basically normative, in the other a science obsessed by detailed description. Building bridges between economics and psychology does not eliminate their differences. However in a recent book targeting the general public, *Thinking, Fast and Slow*, Kahneman sets out to reconcile the two disciplines. In a synthesis which rises above doctrinal differences he draws a distinction between two modes of reasoning which direct our thoughts and decisions, one automatic, the other deliberate. The first mode is swift, and requires no effort on our part nor supervision, being mainly based on association and short-cuts. The second one is quite the opposite, slow, demanding effort and supervision, largely underpinned by deduction and rules we have learnt. Supposing we submit the following problem to a group of mathematically adroit students⁴⁶: a baseball bat and ball cost €110, the bat costs €100 more than the ball, so how much does the ball cost? The spontaneous answer is €10. But in that case the bat would cost €110, and the whole kit €120! Leave the students to think for a while or to write out the equations on paper and they will find the right answer.

By revisiting and extending the classic psychological distinction between the two cognitive systems Kahneman stops seeing heuristics and calculation as mutually exclusive. It is no longer a matter of deciding whether humans are rational or irrational: they are both.

Perception biases working against nuclear power

The overall biases in our perception of probabilities, discussed above, amplify the risk of a nuclear accident in our minds. To this we must add other forms of bias, with which economists are less familiar, but which contribute to the same trend.

⁴⁶ Kahneman submitted this problem to his students at Princeton (**Vérifier**).

A major nuclear accident is a rare event. Its probability is consequently overestimated. Much as smallpox or botulism, the general public thinks its frequency is higher than it really is. Nuclear technology is thus perceived as entailing a greater risk than other technologies. Paul Slovic, an American psychologist who started his career working on perception of the dangers of nuclear power, asked students and experts to rank 30 activities and technologies according to the risk they represented. The students put nuclear power at the top of the list (and swimming last), whereas the experts ranked it in 20th position (predictably putting road accidents first). Faced with a low probability individuals are inclined to over-reassure themselves and demand higher protection or compensation. Protecting the personnel of nuclear power plants against workplace accidents costs more than in any other field⁴⁷. In a general way we over-invest in protection against events with a low probability, the incremental improvement being seen as more beneficial than it really is, in contradiction with the fact that the risks at Fukushima Daiichi were under-estimated.

The risk of a nuclear accident is ambiguous. Expert appraisals diverge depending on whether they are based on probabilistic assessments or the observation of accidents. Furthermore probabilistic assessments produce different figures depending on the reactors or initiating events they consider; the same applies to observations of accidents, for which the definitions and consequently lists differ. Lastly there is a divergence of views between industry experts working for operators and engineering firms, and scientists hostile to the atom, such as the members of Global Change in France or the American Concerned Scientists Association. The effect suspected by Keynes and demonstrated by Ellsberg is clearly at work here. In the face of scientific uncertainty we are inclined to opt for the worst-case scenario. The highest accident probability prevails. The same phenomenon affects the controversial estimation of damages in the event of a major accident. The highest estimates tend to gain the most widespread credence.

The asymmetry between loss and gain highlighted by Kahneman and Tversky is behind a widespread effect which is less decisive and specific for nuclear power⁴⁸. The main consequence of this commonplace, third bias is that it favours the status quo, the drawbacks of change being seen as greater than the benefits (there is a bend in the curve of the value function shown in Figure 1 near the point of reference separating the feeling of gain from that of loss). As a result local residents will tend to oppose the construction of a nuclear power plant in their vicinity, but on the other hand they will block plans to close a plant which has been operating for several years. Obviously what holds true for nuclear power is equally valid for other new facilities, be they gas-fired plants or wind farms.

The risks of a nuclear accident are also distorted by the scale of potential damages and their impact on public opinion. With low frequency and a high impact they are among the various dread risks, along with plane crashes and terrorist attacks targeting markets, hotels or buses, or indeed cyclones. The perception of the consequences of such events is such that their probability is distorted. It is as if the denominator had been forgotten. Rather than acknowledging the true scale of the accident, attention seems to focus exclusively on the accident itself. Disregard for the denominator, mentioned very briefly above, is connected to several common routines of varying similarity which have been

⁴⁷ W. Kip Viscusi, *Rational Risk Policy*, Oxford University Press, 1998.

⁴⁸ Moreover it is not strictly speaking a biased perception of probability. Rather the utility function is distorted, depending on whether there is a loss or gain.

identified by risk psychologists⁴⁹. Let us look at the main routines at work. The availability heuristic is a short-cut which prompts us to answer questions on probabilities on the basis of examples which spontaneously spring to mind. The frequency of murders is generally perceived, quite wrongly, to be higher than that of suicides. Violent or disastrous events leave a clearly defined, lasting mark, in particular due to the media attention they attract. It is consequently easy to latch onto them. Chernobyl reminds us of the accident at Three Mile Island, and in turn Fukushima Daiichi brings Chernobyl to mind. It matters little that each case is different, with regard to the established causes, the course of events or the consequences in terms of exposure of the population to radiation. The spontaneous mode of thought works by analogy and does not discriminate. The representativeness heuristic is based on similarity to stereotypes. It originated with a well known experiment by Kahneman and Tversky. Test participants were told that Linda was 31, single, outspoken and very bright. She had majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also took part in anti-nuclear demonstrations. Then they were asked to pick the job Linda did, in descending order of likelihood: elementary schoolteacher; bookstore salesperson, attending yoga classes; active feminist; psychiatric social worker; member of the League of Women Voters; bank teller; insurance salesperson; bank teller and an active feminist. Most of the participants placed the last answer ahead of the antepenultimate one. Yet logic tells us that there is a higher probability of being a bank teller, than being a bank teller and a militant feminist, the latter category being a sub-set of the female bank-teller population. The representativeness heuristic prompts us to confuse frequency and plausibility. We are consequently able to find regularities and trends where they do not exist. Two successive accidents separated by a short lapse of time will be interpreted as a substantial deterioration in nuclear safety, whereas their close occurrence is merely a matter of chance. This heuristic resembles another bias which often misleads us, namely our tendency to generalize on the basis of low numbers. Rather than waiting to see how a low-probability event repeats itself, tending towards infinity or even with 100 recurrences, we immediately deduce the probability of its recurrence.

To round off this list of biases in our perception of frequencies unfavourable to nuclear power, a word on very low probabilities. It is difficult to grasp values such as 0.00001, or even smaller. Much as it is hard to appreciate huge figures, in the billion billions, often due to the lack of tangible points of reference to which to connect them, dividing 1 by several hundred thousands makes little sense to the average person. Cass Sunstein, a Harvard law professor, carried out an interesting experiment. He asked his students how much they were prepared to pay, at the most, to eliminate a cancer risk of 1 in 100,000 or 1 in 1,000,000. He also presented the problem in slightly different terms, adding that it 'produces gruesome and painful death as the cancer eats away at the internal organs of the body'. In the first case his students were prepared to spend more on saving one individual in 100,000 than on one in a million. On the other hand, playing on their emotions narrowed the difference. They forgot the denominator, despite being students with a basic understanding of fractions and probabilities! To an economist, the biased perception of frequencies is very odd. The way individuals react is light years away from the marginalist reasoning which underpins economic models. As an economist it is disconcerting to

⁴⁹ See Kahneman; and above all, Cass R. Sunstein and Richard J. Zeckhauser, *Overreaction to Fearsome Risks*, 2008.

discover that cutting a risk from 5×10^{-4} to 10^{-4} can pass unnoticed, even if it took time and means to achieve this shift.

In practice if the authorities want to reassure citizens worried by nuclear risk, they will have difficulty basing their arguments on progress in reactor safety. It would apparently be more effective to emphasize the benefits. Slovic and a fellow author have noted a negative correlation between perception of risks and benefits. If the benefits of a given activity are thought to be high, the corresponding risks are perceived as being low and vice versa; on the other hand if the risks are thought to be high, the associated benefits will be underestimated. In more prosaic terms, if I like something, I will play down the risks involved; if I dislike something, I will minimize the benefits. It seems that shale gas currently benefits from this negative correlation in the US. Deeply attached to energy independence and wholly convinced of the economic benefits of cheap energy, Americans underestimate the risks entailed by exploiting this resource.

Unlike very low probabilities, everyone understands zero. By playing on people's emotions it seems possible to persuade them to spend large amounts to eradicate risk completely. In the field of nuclear power, the decision by Germany to shut down its reactors following the Fukushima Daiichi disaster is a case in point. Economists have estimated the loss resulting from the decision to retire the plants earlier than planned at between €15 billion and €40 billion⁵⁰.

In short, however psychologists address perception biases, the latter go against nuclear power. Their effect is cumulative, one consolidating another: in this way they amplify the perceived probability of an accident. At a practical level this has two major consequences. Firstly, there is a risk of over-investing in safety. For this to happen all that is required is for the authorities to follow the trend, either because policy-makers have an interest in giving way to the electorate's demands, or because their perception of probabilities is no different from that of private individuals. Secondly, the options for alternative investment are distorted. On account of its dread nature, the perceived risk of a nuclear accident has been much more exaggerated than the risk inherent in other generating technologies. In 2010 the Organization for Economic Cooperation and Development published a study⁵¹ on the comparative assessment of serious accidents (causing more than five deaths). The survey covered the whole world, from 1969 to 2000. In terms of fatalities and accidents coal came top of the list, a very long way ahead of nuclear power⁵². Does this correspond to your perception, given that the list of accidents during this period of time includes the Chernobyl disaster? Very probably not, much as most other people. If the authorities base their action on perceived probabilities, their decisions may ultimately detract from the goal of reducing risks and loss of human life.

⁵⁰ See S. Douquet and F. Lévêque, The economic loss of the early retirement of nuclear power plants, energypolicyblog.com, 18 February 2012

⁵¹ Comparing Nuclear Accident Risks with Those from Other Energy Sources, OECD, 2010.

⁵² In all 80,250 people lost their lives in 1,870 accidents. Which accident caused the largest number of fatalities? The Chernobyl disaster? No! The failure of the Banquiao and Shimantan dams in China in 1975. This accident, disregarded by all but the local population, claimed 30,000 lives. Bear in mind that, as we saw at the start of this Part, the number of fatalities at Chernobyl did not exceed 60 and that 20,000 was an upper-range estimate of subsequent early deaths. Even if the OECD study had considered these deaths, serious accidents in the coal industry during this period caused more deaths. Moreover, in all fairness, the early deaths caused by coal should also be taken into account in any comparison with nuclear power.

The decisions taken following the terrorist attacks on 11 September 2001 provide a simple illustration of the dual effect of our aversion to dread risk. During the three months following the disaster Americans travelled less by plane, more by car. The well known German psychologist Gerd Gigerenzer has demonstrated that this behavioural shift resulted in an increase in the number of road deaths, exceeding the 265 passengers who perished in the hijacked planes. As the fearsome pictures of the collapse of the twin towers of the World Trade Center receded into the past, the sudden distortion of probabilities must have been partly dissipated. On the other hand we are all affected by the long-term impact of 9-11 whenever we travel by plane. Air passengers all over the world undergo a series of trials –queuing to be checked before entering the boarding area, taking off belts and shoes, emptying pockets, taking electronic devices out of hand luggage, trashing water bottles, putting other liquids into plastic bags, forgetting keys and setting off the alarm, and being handled by complete strangers. Not to mention the rise in airport taxes to cover the extra security costs. In view of the infinitely low risk actually involved, they do seem to have been taken too far⁵³.

In conclusion there is good reason to suppose that the discrepancy between the perceived probability of a nuclear accident and the figures advanced by the experts will persist. The work by experimental psychologists on the amplification of risks for low-probability events seems convincing. The distorting effect of the Fukushima Daiichi disaster is likely to be a durable feature, particularly as we receive regular reminders, if only when progress is achieved or shortcomings are observed in the resolution of the many outstanding problems (soil decontamination, public health monitoring, dismantling of reactors, waste processing among others). It will also require great political determination to avoid pitfalls similar to tightening up checks on airline passengers.

The magic of Bayesian analysis

An English Presbyterian minister and a French mathematician forged a magic key which enables us to update our probabilistic judgements, use probability as a basis for reasoning without being a statistician, reconcile objective and subjective probability, combine observed frequency and calculated probability, and predict the probability of the next event.

The Bayes-Laplace rule

Thomas Bayes was the first person to use the concept of conditional probability, for which Pierre-Simon Laplace found a more widespread application. We presented this concept briefly in the opening section on the escalating severity of nuclear accidents, more

⁵³ In addition many observers maintain that they have not been very effective.

specifically in relation to the probability of massive release of radioactive material in the event of a core melt, or in other words given that core meltdown has occurred. This probability is denoted using a vertical bar: $p(\text{release}|\text{melt})$. In a general way, A and B being two events, the conditional probability is written as $p(A|B)$, which reads as ‘the probability of A given B’. Or, to be more precise, we should refer to two conditional probabilities; as we are focusing on two events, we may also formulate the conditional probability $p(B|A)$, the probability of B given A. In the case of a core melt and the release of radioactivity, $p(\text{release}|\text{melt})$ is almost equal to 1: with just a few exceptions (water loss from a spent-fuel pool) a massive release is impossible unless the core has melted.

Bayes’ rule and conditional probabilities

The Bayes-Laplace rule, more commonly known as Bayes’ rule, enables us to write an equation linking two conditional probabilities. It is written as:

$$p(A|B) = p(A)[p(B|A)/p(B)]$$

It reads: ‘the probability of A given B is equal to the probability of A multiplied by the probability of B given A divided by the probability of B’.

It may also be written as:

$$(2) p(A|B) = [p(A)p(B|A)]/[p(A)p(B|A)+p(\text{not } A)p(B|\text{not } A)]$$

This formula is typically illustrated with examples from medicine. We shall use the following data for the problem:

The probability that a patient undergoing an X-ray examination has a cancer is 0.01, so $p(\text{cancer}) = 0.01$

If a patient has a cancer, the probability of a positive examination is 0.8, so $p(\text{positive}|\text{cancer}) = 0.8$

If a patient does not have a cancer, the probability of a positive examination is 0.1, so $p(\text{positive}|\text{no cancer}) = 0.1$

What is the probability that a patient has cancer if the examination is positive, so $p(\text{cancer}|\text{positive})$?

Using equation (2), we find 0.075⁵⁴.

Bayes’ rule, set forth in the box above, is the key to inductive reasoning. Let us look at how it works.

As the reader is probably more familiar with the concept of deduction rather than induction we shall start by confronting the two approaches. Deduction generally proceeds from the general to the specific: ‘All dogs have four legs; Fido, Rover, Lump and Snowy are dogs; so they have four legs.’ Induction generalizes on the basis of facts or empirical data: ‘Fido, Rover, Lump and Snowy are dogs; they have four legs; so all dogs have four legs.’

⁵⁴ $p(\text{cancer}|\text{positive}) = [p(\text{cancer})p(\text{positive}|\text{cancer})]/[p(\text{positive}|\text{cancer})p(\text{no cancer}) + p(\text{positive}|\text{no cancer})p(\text{no cancer})] = [0.01 \times 0.8]/[0.01 \times 0.8 + 0.99 \times 0.1] = 0.075$.

A simple equation enables us to grasp the inductive power of Bayes' rule. Let us suppose that event A is a hypothesis, H, and event B is an observed data, d. So $p(A|B)$ is the probability that H is true, given that d is observed. Bayes' rule is written thus:

$$p(H|d) = p(H)[p(d|H)/p(d)]$$

The conditional probability $p(d|H)$ encapsulates the essence of deductive reasoning. It indicates the probability of a data d being observed when hypothesis H is true. The reasoning starts from a hypothesis, rooted in theory, to arrive at the probability of an observation. For example, the modern theory of physics implies the existence of a massless particle, the neutrino. Since the work of the physicist Wolfgang Pauli, much research has been devoted to proving its existence. In simple terms, if the theory is right, the particle should produce a particular effect. The conditional probability $p(H|d)$ expressing inductive reasoning takes the same route but in the opposite direction. It indicates the veracity of a hypothesis, if a data d has been observed. Here, we start from the observation to infer the degree of certainty of a hypothesis. If our name is Isaac Newton, an apple falling on our head may lead to the hypothesis of universal gravity. We work back from the effect to the cause.

The equation above also shows how Bayes' rule provides a means of reviewing an appraisal in the light of new information.

At the outset only the general, or prior, probability $p(H)$ is known. It may be based on objective data from the past, on scientific theory or indeed on subjective belief. New data, which changes the picture, then comes to notice. Thanks to Bayes' rule we can calculate a new, posterior probability, namely that H is true now that data d has been brought to light. It depends on the prior probability $p(H)$ and the multiplier $[p(d|H)/p(d)]$. The inverse probability, $p(d|H)$, is the probability of observing data d given that H is true. $p(d)$ is the probability of observing data d; and is often referred to as the likelihood. The multiplier determines how much the prior probability is updated. It is worth noting, for example, that if the observation yields no additional information, then $p(d|H)$ equals $p(d)$ and the multiplier equals 1, so there is no reason to update the prior probability. Intuition would suggest that if there is no relation between d and H, observing d should not change the degree of certainty assigned to H. On the other hand, when the observation is more probable if the hypothesis is true, $p(d|H)$ being greater than $p(d)$ and the multiplier greater than 1, the probability will be revised upwards. In other words, if d and H are linked I must upgrade my appreciation of the hypothesis. But in the opposite case, when the observation is less probable if the hypothesis is true, $p(d|H)$ being smaller than $p(d)$ and the multiplier less than 1, the probability will be revised downwards. Or in other words, if d and H conflict, I must downgrade my appreciation of the hypothesis. Lastly, it should be borne in mind that a posterior probability calculated on the basis of a new item of information may in turn be used as a prior probability for calculating another posterior probability taking into account yet another new item of information. And so on. As the new inputs accumulate, the initial prior probability exerts less and less influence over the posterior probability.

Bayes' rule thus makes it possible to calculate probabilities with only a limited number of observations. It even provides for a complete lack of data, in which case the prior probability is unchanged. For example you did not toss the coin your adversary offered, or perhaps you have never tossed a coin. If you do not know whether the coin is loaded and

your adversary displays no criminal tendencies, you will choose 1/2 as the subjective initial odds and you will stick to these odds if the game is cut short before it even starts.

Are we naturally good at statistics?

In the above we presented Bayesian analysis as a relatively intuitive mechanism for inductive reasoning and updating earlier appraisals. However applying the rule on which it is based is tricky. In the previous section we explained that Ivy League students sometimes came up with the wrong answers to much simpler probability problems. Disregard for the denominator, apparently so common, is such that one may wonder whether the average person can really grasp the concept of probability. Probability is commonly defined by a frequency, the ratio of observed cases to possible cases. Omitting the denominator renders a probability completely worthless.

Thanks to experimental psychology the focus of the classical debate on whether decision-making by humans has a rational or irrational basis has shifted to the statistical mind, raising the question of whether we are capable of reasoning in terms of probabilities⁵⁵. The response by most experimental psychologists is categorical. Kahneman and Tversky sum up their position thus: 'People do not follow the principles of probability theory in judging the likelihood of uncertain events [...] Apparently, people replace the laws of chance by heuristics which sometimes yield reasonable estimates and quite often do not.' According to Slovic our lack of probabilistic skills is a proven fact: humans have not developed a mind capable of grasping uncertainty conceptually. This view was endorsed by the well known evolutionary biologist Stephen Jay Gould. After examining the work of the winner of the 2002 Nobel prize in economics and his main co-author, Gould concluded that: 'Tversky and Kahneman argue, correctly I think, that our minds are not built (for whatever reason) to work by the rules of probability, though these rules clearly govern our universe.'

Under these circumstances Bayesian logic seems even further beyond our reach. The case of the medical examination described above was presented to American physicians and medical students. An overwhelming majority answered that the chances of being ill if the examination is positive ranged from 70% to 80%, 10 times higher than the right answer, 7.5%. A survey in Germany revealed that HIV-Aids counsellors were just as confused by Bayesian reasoning. After being given the data to calculate the probability that a person who had tested positive was actually a carrier of the virus, test participants produced answers very close to 1, whereas the right answer is closer to 50%. If you think you are ill and decide to have an examination, make sure you choose a doctor who is trained in statistical analysis.

In short it seems to be an open-and-shut case. As Kahneman and Tversky assert: '[In] his evaluation of evidence, man is apparently not a conservative Bayesian: he is not Bayesian

⁵⁵ NB: we are not talking about innumeracy, a shortcoming similar to illiteracy but for numbers. The experiments related above involved educated participants, students or graduates.

at all⁵⁶.’ Participants in tests go so far as to resist attempts by experimenters to correct and instruct them. This reaction is comparable to the impact of optical effects: although we know they are optical illusions we persist in taking a bump for a hole, or in seeing an area as smaller or larger than it really is.

In the 1990s Gerd Gigerenzer firmly opposed this standpoint. According to the German psychologist humans are not deaf to statistical reasoning. Unlike Gould he upholds the idea that the human mind evolved, while integrating Bayesian algorithms. However to be of any use such algorithms must be expressed in practical terms, close to their internal format, as is apparently the case with frequencies in their natural form, but not probabilities.

Presenting Bayes’ rule in other terms

People make mistakes in Bayesian calculations because the problem is not clearly stated. They make far fewer mistakes if data are set forth in a ‘natural frequency format’, instead of being given as frequencies expressed as percentages or values ranging from 0 to 1.

This approach presents the medical examination in the following way:

- (a) 10 out of 1,000 patients who undergo an X-ray examination have cancer;
- (b) 8 out of 10 patients with cancer test positive;
- (c) 99 out of the 990 patients who do not have cancer test positive.

If we take another group of patients, who have also tested positive, how many would you expect to have cancer?

Presented in these terms one may easily reason as follows: 8 patients who tested positive have cancer and 99 patients who tested positive do not, so in all 107 patients tested positive; the proportion of patients with cancer among those who tested positive is therefore $8/107$, in other words a probability of 0.075.

This problem and numerous variants were experimented by Gigerenzer and his fellows on groups comprising doctors, students and ordinary people. They also developed a method for learning Bayes’ rule in under two hours, which is now widely used in Germany to train medical staff.

Cognitive science has made further progress, since the experimental work of Kahneman, Tversky and Gigerenzer, in particular with regard to language learning. Recent work seems to favour the idea of a Bayesian brain⁵⁷, endorsing the conviction expressed by Laplace two centuries earlier in the introduction to his philosophical treatise on probabilities: ‘[...] the theory of probabilities is basically just common sense reduced to calculus; it makes one appreciate precisely something that informed minds feel with a sort of instinct, often without being able to account for it.’

⁵⁶ Daniel Kahneman and Amos Tversky, ‘Subjective probability: a judgment of representativeness’, in *Judgements under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic and A. Tversky (ed.), Cambridge University Press, 1982, p. 46.

⁵⁷ See Stanislas Dehaene, *Le cerveau statisticien: la révolution Bayésienne en sciences cognitives*, a course of lectures at the Collège de France, in 2011-12.

Choosing the right prior probability

We explained above that the degree of prior certainty we choose to assign to a hypothesis may be based on objective data from the past, scientific theory or indeed subjective belief. It is now time to tighten up that statement and show how this prior choice influences the degree of posterior certainty.

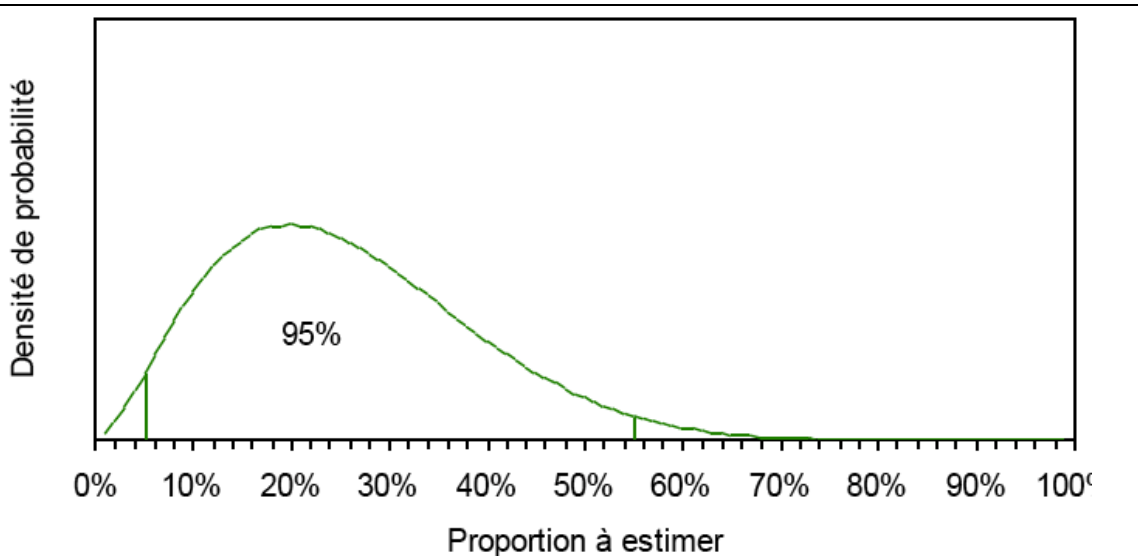
Bayes' rule applies equally well to statistical reasoning as to probabilistic logic. In the first case the prior value we take is based on observation. For example a die has been rolled 10 times, the 6 has come up twice, and we are about to roll the same die, starting a new series. The prior probability of the 6 coming up again (or, more simply, the prior) is $2/10$. Before rolling any dice at all, we may also choose a theory such as the one which posits the equiprobability of any of the faces coming up. This can no longer be treated as an observation, rather a belief. It assumes that the die is perfectly balanced, which is never the case. The strict equiprobability of $1/6$ is rarely observed, even if a die is rolled several thousand times. Some people have taken the trouble to check this. So by choosing the prior, we are engaging in probabilistic logic. Equiprobability is the most plausible initial hypothesis, lacking any knowledge of the properties of the die on the board. Adopting a more subjective stance we may also opt to assign a prior probability of $1/10$ to the 6, because we know we are unlucky or having a bad day.

The effect of the prior is described visually in the following box. It shows that only in two cases does it have no effect: either there are an infinite number of observations or measurements, or the beliefs are vague and all the outcomes equiprobable. Between the two the prior carries a varying amount of weight.

Illustrating Bayesian revision

Figure 2 describes a probability density function⁵⁸. It expresses your prior judgement on the proportion of your compatriots who hold the same opinion as you do on nuclear power. According to you, the proportion is probably (95% chance of being right) between 5% and 55%, and its most likely value is 20%. (95% because the area below the curve has been normalized to 1 and the segments at the two extremities add up to 5%; 20% because the highest point of the curve coincides with this value on the x-axis.) Of course, someone else may be more sure than you are, in which case the curve would be more pointed; on the other hand a third party might only have a very vague prior conviction, making the curve very flat.

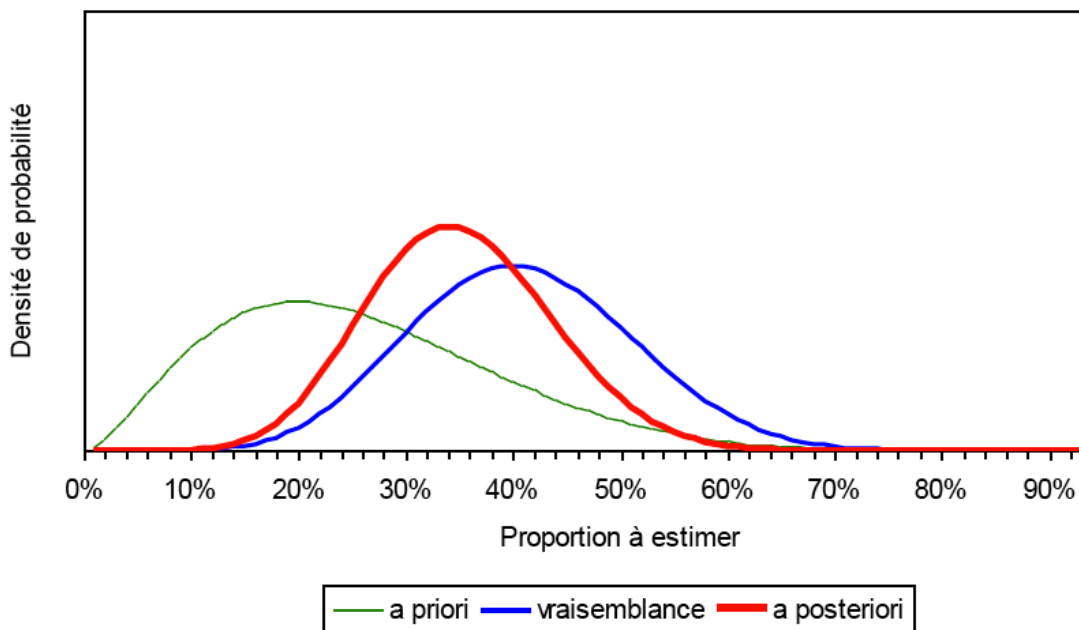
⁵⁸ Figure taken from Bayes Fastoche, by Emmanuel Grenier.



Probability density; Estimated proportion

Let us now assume that one of your friends has carried out a mini-survey of acquaintances. Of the 20 people she has questioned, eight (40%) shared your point of view. With some knowledge of statistics you will be able to define what is known as the likelihood function, summarizing the survey's findings. The blue curve in Figure 3 plots its density. Had a higher number of people been surveyed, the curve would have had a steeper peak. As we all know the larger the number of people polled, the greater the certainty associated with poll results.

You are now ready to update your initial judgement by applying Bayes' rule. Just multiply the two previous functions. The red curve in Figure 3 plots the new density function which expresses the proportion of your compatriots who hold the same opinion as you do on nuclear power, given the results of your friend's mini-survey.



Probability density; Estimated proportion; prior; likelihood; posterior

It is immediately apparent that the posterior is more certain than the prior – the red curve is not as flat as the green one. This is hardly surprising, because the new data confirms your judgement. It is also worth noting that observation carries more weight than the prior: the posterior peaks at 34%, a value closer to 40% than to 20%. Finally you should bear in mind that if the survey had polled a larger number of people, the red posterior curve would have been even closer to the blue curve, which would – others things being equal – also rise to a steeper peak. Likelihood would have weighed even more heavily on the prior. Ultimately, if observations are carried out on very large samples, the event is repeated a large number of times, or a quasi-exhaustive survey is conducted, the choice of the prior no longer has any bearing on the posterior. It makes no difference.

On the other hand, what happens if the prior is changed? If your prior is subject to greater uncertainty (a flatter green curve), it will weigh less heavily on the posterior, which will resemble more closely the likelihood (the blue curve). The prior exerts less influence. If, at the outset, you have absolutely no idea of how many people share your views (a horizontal line from one end to the other), the prior will exert no influence on the posterior; the latter will follow the likelihood exactly – a horizontal straight line does not distort, through multiplication, the likelihood density function. The weightless prior is obtained in the same way as equiprobability when rolling dice. If you know nothing about its lack of balance, the posterior will only be influenced by each roll of the dice, not your initial belief.

In this example Bayes' rule worked one of its many magic tricks, reconciling perceived probability with its objective counterpart. It combines an initial perception, regarding the plausibility of a hypothesis or the probability of an event, with a calculated probability based on new knowledge.

Predicting the probability of the next event

Will the sun rise tomorrow? According to Laplace the answer is uncertain, having estimated the probability of it not rising to be $1/1.826215!$ In this example – which earned him much mockery and criticism – he applied a general formula which he had worked out himself. His rule of succession is based on Bayes' rule.

The problem of succession may be stated in the following abstract terms: what is the probability of picking a red ball out of an urn at the $n+1$ th attempt, given that for the n previous attempts, k red balls were picked? Or alternatively, posing a more concrete question addressed in the following section, given that the global reactor fleet has experienced 11 core-melt events in 14,400 reactor-years, what is the probability of another core-melt occurring tomorrow?

Your intuitive answer to the question on the next draw is very probably k/n . Previously k red balls have been picked from the urn in the course of n draws, so you conclude that the same proportion will also be valid next time. It is not a bad solution, but there is a better

one. From a mathematical point of view your answer does not work if $n = 0$ ⁵⁹. So you cannot use this formula to determine the probability of the first event, for example a core-melt before an accident of this type ever happened. Using the formula means you can only base your judgement on observation. Much as Doubting Thomas, you only believe what you can see. You may however have carried out previous observations (though this is fairly unlikely in the case of picking balls out an urn) or perhaps you fear that the person who set the urn on the table wants to cheat you. In short you are not basing your decision on your prior; you are not using Bayesian reasoning.

To solve the succession problem Laplace resorted to Bayes' rule, though it was not yet known as such. Indeed Laplace reinvented it, very probably never having heard of the work done by the British minister. He started from an *a priori* judgement, assuming that only two outcomes were possible: picking a red ball or a ball that was not red. Each had an equal chance of occurring, so the probability was $1/2$. He thus obtained a formula $(k+1)/(n+2)$. For example, if five draws are made and a red ball is only picked once, the probability predicted by the formula for the sixth draw is $(1+1)/(5+2)$, or $2/7$, or indeed 0.286. Alternatively, if we return to the example of the sun⁶⁰, Laplace went back to the earliest time in history, 5,000 years or 1,826,214 days earlier, making 1,826,213 successful sunrises out of 1,826,213. The probability that the sun would also rise on the following day was consequently $(1,826,213+1)/(1,826,213+2)$ and the probability of the sun not rising was $1-(1,826,213+1)/(1,826,213+2)$, or one chance in 1.8 million.

Laplace proceeded as if two virtual draws had been made, in addition to the real ones, one yielding a red ball, the other a non-red ball. Hence the simple, intuitive explanation of the origin of his formula: he adds 2 to n in the denominator, because there are two virtual draws; and adds 1 to k in the numerator, to account for the virtual draw of a red ball. From this point of view the choice of the sun is unfortunate. The merit of the prior depends on it being selected advisedly. Laplace was a gifted astronomer, yet he chose a prior as if he knew nothing of celestial mechanics, as if he lacked any understanding of the movement of the sun apart from the number of its appearances.

The choice of the prior to predict picking a red ball at the next draw is more successful. It anticipates Keynes' principle of indifference: when there is a priori reason to suppose that one outcome is more probable than another, equiprobability is the only option. If there is nothing to indicate a bias, the prior probability that a coin will land heads up is $1/2$, the prior probability of scoring a 3 with a roll of a die is $1/6$, and so on. A $1/n$ probability makes complete sense if no prior knowledge nor expertise suggests that a specific outcome among the n possible outcomes is more probable than another.

Thanks to progress in probability theory we now have a better instrument than the Laplace formula for solving the succession problem. The solution found by the French mathematician, $(k+1)/(n+2)$, has become the particular case in a more general formula. This formula is still the result of Bayesian reasoning, but it has the key advantage of introducing a parameter which expresses the strength of the prior in relation to new

⁵⁹ If n equals zero, k equals zero; zero divided by zero is indeterminate.

⁶⁰ 'Thus we find that an event having occurred successively any number of times, the probability that it will happen again the next time is equal to this number increased by unity divided by the same number, increased by two units. Placing the most ancient epoch of history at 5,000 years ago, or at 1,826,213 days, and the sun having risen constantly in the interval at each revolution of 24 hours, it is a bet of 1,826,214 to 1 that it will rise again tomorrow.' p. 23 *Essai Philosophique sur les Probabilités*, 1825.

observations and knowledge. For the two extreme values of this parameter, either your confidence in the prior is so strong that nothing will change your mind – so the posterior cannot diverge from the prior – or your confidence in the prior is so weak that the tiniest scrap of new data will demolish it – leaving the posterior exclusively dependent on the new data.

In this modern version the expected probability of picking a red ball in the $n+1$ th draw is written as $(k+st)/(n+s)$, where t is the prior expected probability – for example 0.5 for equiprobability – and s is the parameter measuring the strength of the prior. The two virtual draws, of which one is successful, added by Laplace, are replaced here by st virtual draws (where $t < 1$) of which s are successful. In other words we return to Laplace's formula if $s=1$ and $t=1/2$. Bear in mind that if $s=0$, the posterior, in other words the probability of drawing a red ball in the $n+1$ th draw becomes k/n . Only the observations count. On the other hand, if s tends towards infinity, the posterior tends towards the prior⁶¹. Only the prior counts.

The general formula, $(k+st)/(n+s)$, corresponds to the summit of a curve similar to the blue curve in Figure 3. Similarly the value t marks the top of the green curve in the same figure. So Bayes' rule does apply to functions. The prior probability and the observed probability are both random variables for which only the distribution and classical parameters (expectation, variance) are known⁶². The exact value is not known, but simply estimated by statements such as 'there is a 95% chance that the target value is between 0.3 and 0.6'. The parameter s which expresses the strength of the prior may therefore be interpreted as the uncertainty of the prior. The greater the value of s , the steeper the curve; the lower the value of s , the flatter the curve (in mathematical terms $1/s$ is proportional to the variance which measures the distribution on either side of the mean value of the prior function). Expert opinions can be finely quantified, thanks to this property. They can be queried too, for the most probable prior measurement of a phenomenon (t), but also the strength of their opinion (s). In this way two experts can state that they estimate the probability of radioactive emissions in the event of a core melt as 0.1, but one may be confident, assigning a value of 10 to s , whereas the other is much less certain, only crediting s with a value of 0.5.

An interesting feature of this way of calculating the probability of the next event is that it enables us to grasp the connection between perceived probability and acquired objective probability. At the outset all you have is your subjective prior (the expected probability, t). Then you receive the objective information that k red balls have been picked in n draws (an expected probability of k/n). So your perception of the next draw is given by the posterior (the expected probability of $(k+st)/(n+s)$) which combines the two previous elements. In other words the perceived probability is seen as the updated prior probability taking into account the acquired objective probability. This interpretation is fruitful because it provides a theoretical explanation⁶³ for our biased perception of low and high probabilities. There is a linear relation between the expectation of the perceived

⁶¹ If s is very high, the denominator n is negligible compared to s while the numerator k becomes negligible compared to st , the ratio is close to t , the initial expectation of probability.

⁶² The general formula is obtained by choosing a beta distribution with parameters $[(st, s(1-t))]$ for the prior function, and a binomial distribution for the likelihood function. In this case a beta distribution with parameters $[st+k, s(1-t)+n-k]$ is used for the posterior function.

⁶³ See W. Kip Viscusi, op cit.

probability $(k+st)/(n+s)$ and the expectation of the acquired objective probability k/n . If $k/n < t$, the perceived probability is less than the objective probability.

What is the global probability of a core melt tomorrow?

Following the Fukushima Daiichi disaster two opponents of nuclear power published their estimate of the risk of a major accident in a French daily⁶⁴. Both have a scientific background: Bernard Laponche is a nuclear physicist; Benjamin Dessus is an engineer and economist. Both started their career at France's Atomic Energy Commission (CEA). They asserted that, 'in the next 30 years the probability of a major accident is [...] over 100% in Europe'. The exact figure cited was 129%. Quite something, a probability of more than 1. It was as if the certainty of an accident was not enough to impress public opinion, so the two authors invented a super-certain hazard. Perhaps they were inspired by the manufacturers of washing powders which wash 'whiter than white'.

Joking apart, what caused this mistake? Certainly not the data they cited, which were correct. Dessus and Laponche's calculations were based on a global fleet of 450 reactors (of which 143 in Europe), which has suffered four core melt events followed by a massive release of radioactive material (Chernobyl 4, Fukushima Daiichi 1, 2 and 3) with total operations spanning 14,000 reactor-years. So the observed accident frequency is $4/14,000$, or about 0.0003. The mistake is due to confusion between the number of accidents expected over a 30-year period and the probability of an accident. The expected number of accidents for this period, with an observed frequency of 0.0003 is indeed 1.29 ($30 \times 143 \times 0.0003$). But this number is not a probability. Confusing it with a probability is tantamount to saying that in a family with four children the probability of having a daughter is 2 (0.5×4), whereas in fact it is $15/16$ ($1 - 1/2 \times 2 \times 2 \times 2$)!

The mistake, if it is deliberate, is all the more stupid as an error-free calculation produces a figure which is sufficiently impressive in itself: the probability of a major accident in Europe over the next 30 years is 0.72. Were Dessus and Laponche afraid the general public would underestimate such a high probability? Were they trying to rectify this perception bias? The figure of 0.72 is calculated by modelling the probability of a major accident using a binomial distribution function [30×143 ; 0,0003]. Statistically speaking, this involves presenting the problem to students in these terms: 14,000 reactor-years have been polled ('Have you had a major accident?'); four answered yes; what is the probability of an accident over the next 30 years for a fleet of 143 units?⁶⁵

However the mistake discussed above is harmless compared to a crippling fault in the approach which the rule of succession reveals. Dessus and Laponche behave as if the probability of another accident tomorrow can only be elucidated and parametered by the past observed frequency, k/n . They choose to focus exclusively on observation, completely disregarding all the other learning that has accumulated on the subject. Yet tens of

⁶⁴ B. Dessus and B. Laponche, *Libération*, 5 June 2011, Accident nucléaire: une certitude statistique.

⁶⁵ To simplify matters, we shall assume here that the number of accidents follows a binomial distribution. The probability of there not being a major accident in Europe over the next 30 years is $1 - 0.0003$ per reactor-year, or $(1 - 0.0003)^{30 \times 143}$, or about 0.28. The probability of there being a major accident in Europe over the next 30 years is therefore 0.72. The choice of a 30-year period is largely arbitrary; the calculation can be reduced to the probability of an accident in Europe next year, for which we find 0.042, or a nearly 4-in-1,000 chance.

thousands of engineers, researchers, technicians and regulators have been working on precisely this topic for the past 60 years. It is as if the only hope of salvation is in major accident data, despite their number being extremely limited, unlike other forms of knowledge. By considering that the parameter indicating the strength of the prior (previously designated as s) has a value of zero, these two opponents of nuclear power are behaving in exactly the same way as their counterparts, the obtuse advocates of the atom, who stick rigidly to their prior assumptions, refusing to attach any credit to observations of past accidents to predict the probability of an accident in the future (in other words, leaving s to tend towards infinity).

A simple example may be used to illustrate how observation and other knowledge can be combined to predict future nuclear risks. There are no particular difficulties regarding observations, apart from their very limited number. This constraint may, however, be slightly reduced if we focus on the risk of a core-melt accident, of which there are 11 cases⁶⁶, rather than the risk of core melt leading to release, of which there are only four instances. As ever, choosing the prior is tricky, but we may take advantage of the availability of probabilistic safety assessments, the results of which summarize all the existing knowledge on accidents, apart from observations. On the basis of assessments carried out in the 1990s on US reactors, we may choose to add to the data on 11 accidents in 14,000 reactor-years a virtual observation of 1.6 accidents in 25,000 reactor-years. This virtual observation is derived from the expected value of the probability of core melt estimated by the experts as 6.5×10^{-5} per reactor-year associated with an uncertainty measured by a prior strength of 25,000⁶⁷. We thus obtain values for t and s and can apply the formula $(k+st)/(n+s)$, or $[11+(25,000 \times 6.5 \times 10^{-5})]/14,000+25,000$, which equals 3.2×10^{-4} . In other words the experts propose a mean probability of core-melt per reactor-year of 6.5×10^{-5} ; and observed evidence leads to 7.8×10^{-4} . By advisedly combining the two sources of knowledge we obtain a probability of 3.2×10^{-4} . To make this figure easier to grasp, we may translate it by calculating that the probability of a core-melt accident next year in Europe is 4.4 in 1,000. This probability is less than half the value obtained from the observed frequency of core-melt accidents on its own.

It is nevertheless very high, because the approach used in this example disregards progress towards greater safety. It assumes that the accident probability we want to measure has remained constant over time. But the design of reactors is steadily changing. Reactors built today are on the whole safer than their first-generation counterparts. Furthermore safety standards are increasingly strict and more efficiently applied, which in turn improves safety performance. Indeed these gains are reflected in the risk of core melt. A large number of the accidents considered occurred during the early years of the development of civil nuclear power⁶⁸.

Nor is the above approach ideal, because it assumes that events are unconnected. Yet the European, or global, nuclear fleet in 2012 is fairly similar to the one that existed in 2011, in 2010, and so on. In other words, we are not dealing with a situation akin to rolling dice, in which the result of one roll has absolutely no bearing on the following one. On the

⁶⁶ The eight core-melt accidents previous to Fukushima Daiichi (see footnote 4 on page #), plus the meltdown of reactors 1,2 and 3 at this NPP.

⁶⁷ Specify source.

⁶⁸ See also the changes in US probabilistic safety studies highlighted by the Electric Power Research Institute.

contrary the occurrence of one major accident may point to others, because we are still dealing with the same reactors⁶⁹. For example, an accident may indicate that a similar problem affects other reactors of the same design or subject to comparable natural risks. A recent study⁷⁰ sought to remedy these two shortcomings. The model it used weights accident observations according to the date when they occurred. It calculates a sort of discount rate which, with time, tends to reduce the influence of old accidents over the probability of an accident now. With this model the probability of a core-melt accident next year in Europe is 0.7 in 1,000, six times lower than with the previous approach.

How may we conclude this second part devoted to risk? By emphasizing the pointless opposition between experts and the general public, with both parties accusing the other of rank stupidity. The experts, with their sophisticated calculations, have allegedly made massive mistakes estimating the frequency of major accidents and the associated damage. Locked up in their laboratories, their certainty has purportedly blinded them to events that are glaringly obvious, namely the recurrence of disasters since Three Mile Island. On the other hand the general public is supposedly guilty of yielding exclusively to its fear of disaster, stubbornly clinging to its rejection of calculated probabilities. Brainwashed by a stream of disaster movies and incapable of basing the slightest decision on statistical evidence, people's reactions are ruled by emotion and nothing else; they are ostensibly quite unable to see that coal or hydroelectric dams constitute risks every bit as serious as nuclear power plants. So you may choose sides, denigrating either the experts or the general public! However there is no contradiction between analysing on paper sequences of events that may potentially lead to a core melt, and observing accidents and their causes. Furthermore theory provides ways of combining and associating such knowledge. It enables us to find rational explanations for the irrational, which becomes much less or not at all irrational. We have a better grasp of how probabilities are distorted, either through simple heuristics or gradual acquisition of Bayesian methods, depending on one's school of thought. What was once perceived as irrational loses that connotation once it is seen to conform to identified mechanisms. A fuzzy border now separates rational decision-making from other forms of choice. Is it irrational to take decisions under uncertainty without allowing for risk aversion? Surely it would in fact be irrational to go on denying the existence of such aversion? When Maurice Allais presented Jimmie Savage with a puzzle the latter's answer contradicted the results of his own theory. When he appears as a guest-speaker in Israel, Daniel Kahneman avoids travelling on Tel Aviv buses for fear of an attack. In so doing he is fully aware his actions are dictated by perceived, rather than calculated probability.

The loss of bearings which fuels discord between experts and the general public, opposing the rational and irrational, is most uncomfortable. It raises a terrifying question: should government base its decisions on perceived probabilities or on those calculated by experts? In the case of nuclear power, the former are largely overestimated, a situation which is likely to last. It would be foolish to treat the attitude of the general public as the

⁶⁹ At least with respect to their location and type. Improvements to safety through changes to equipment and operating procedures mean that a reactor is not quite the same over time.

⁷⁰ Lina Escobar Rangel and François Lévêque, How did Fukushima Daiichi core meltdown change the probability of nuclear accidents?, i3 Working Paper 12-ME-06, October 2012.

expression of fleeting fears which can quickly be allayed, through calls to reason or the reassuring communication of the 'true' facts and figures. If government and the nuclear industry did share an illusion of this sort, it would only lead to serious pitfalls. The reality test, in the form of hostile demonstrations or electoral reversals, may substantially add to the cost for society of going back on past decisions based exclusively on expert calculations. On the other hand the airline security syndrome will inevitably spread to nuclear power unless the authorities pay attention to public perception of probabilities. If the propensity to invest in nuclear safety is not brought under control, many expensive new protective measures will accumulate, without necessarily having any effect. Which brings us to the subjects of the next two papers: safety regulation and nuclear policy.