Part II: Quantitative Risk Analysis

In Part I of this series on risk management[1], we presented a historical perspective of our understanding of risks and uncertainties and introduced the main concepts of a comprehensive risk management process. In particular, we emphasized the importance of an appropriate quantification of risks, and we briefly addressed the prominent role that probability theory plays in this context.

In the present article, we discuss in more detail how uncertainties and risks are represented in mathematical terms and what kind of information can be extracted from such a quantitative description. It turns out that there is no unique way to measure risks, and we analyze and compare some of the risk measures that are used in present-day risk management approaches.

Finally, we address some issues and problems in connection with a quantitative assessment of uncertainties and risks in real-world situations, and we discuss the importance of scenario simulations and sensitivity analyses.

Characterization of uncertainties and risks

In mathematical terms, uncertainties are represented as random variables. These are variables for which the value they will take cannot be predicted unambiguously. Random variables are only characterized in terms of the probabilities with which they take their different values. A well-known example of a random variable is the number of points in a throw of a fair die. This number can be 1, 2, 3, 4, 5, or 6, and all six values are equally likely, ie, they all have the same probability, 1/6. If such a die is thrown a large number of times, this implies, for example, that a '6' will occur in about one sixth of all throws.

Random variables are thus completely characterized by a so-called 'probability distribution'. This specifies the occurrence probability for each value the variable can assume. In the case of continuous random variables, the corresponding probability distribution becomes a continuous function and is then often called 'probability density' (see below).

Risks arise because of uncertainties, and the quantity with which a particular risk is associated is usually a function of many uncertain (random) variables. As an example, we may consider the risks associated with the annual revenue from generating and selling electricity. If we are, in particular, interested in the way this revenue is affected by a reduced availability of the generator, the relevant uncertainties are the number of

1) See ABB Review 2/2004 pp 66–70

outages, the duration of these outages, and the market price of electricity at the time of the outages. Using methods from probability theory, we can then determine the probability distribution of the annual revenue, ie the information needed to estimate the corresponding risks, provided that the probability distributions of the underlying uncertainties are known. (In some cases, we also have to take into account possible correlations between the individual uncertainties)

Probability distributions

To discuss the information contained in a probability distribution, we consider three typical examples (see 1). The Poisson distribution 1a refers to the number of events per time period that are observed in a so-called Poisson process. Poisson processes are used as an accurate model for many practical problems, eg, to describe the random occurrence of defects or outages in an industrial process or the number of cars that arrive at an intersection in a given time interval. The example shown in 1a is based on an average number of 3 events (eg, 3 outages per year). In this case, for example, the probability of zero events is 5 percent, about the same as that of observing six events, and the probability that we observe two, three, or four events is about 60 percent.

The Poisson distribution is an example of a discrete probability distribution, ie the underlying random variable (number of events) can only assume discrete values (0, 1, 2, 3, etc.). Often, however, we are concerned with continuous random variables that can assume any value in a given interval. The properties of such random variables are then described by a so-called 'probability density function' as shown in the examples of 1b and 1c. The exponential distribution 1b is often used to model the random characteristics of the duration of outages or of the time to failure of system components (measured in hours, for example). 1b also illustrates the interpretation of such probability densities: The shaded area is equal to the probability that the duration D of an outage is between 5 and 7.5 hours. The total area under the curve is thus equal to one (the probability that the outage duration has any value between zero and infinity is 100 percent).

The perhaps best-known probability density is the normal (or Gaussian) distribution shown in 1c. It accurately represents the uncertainty observed in many biological and technical systems. The probability distribution of the value of a portfolio of assets or of an industrial project, for example, is often also very close to a normal distribution - at least if it depends on a not too small number of independent random variables. In the example of 1c, the mean (eg, project) value is assumed to be 10 (eg, US$ 10 million), and the indicated area under the curve (from minus infinity to zero) is equal to the probability that the project will lead to a loss (has a negative value).

1 Examples of probability distributions (probability densities)

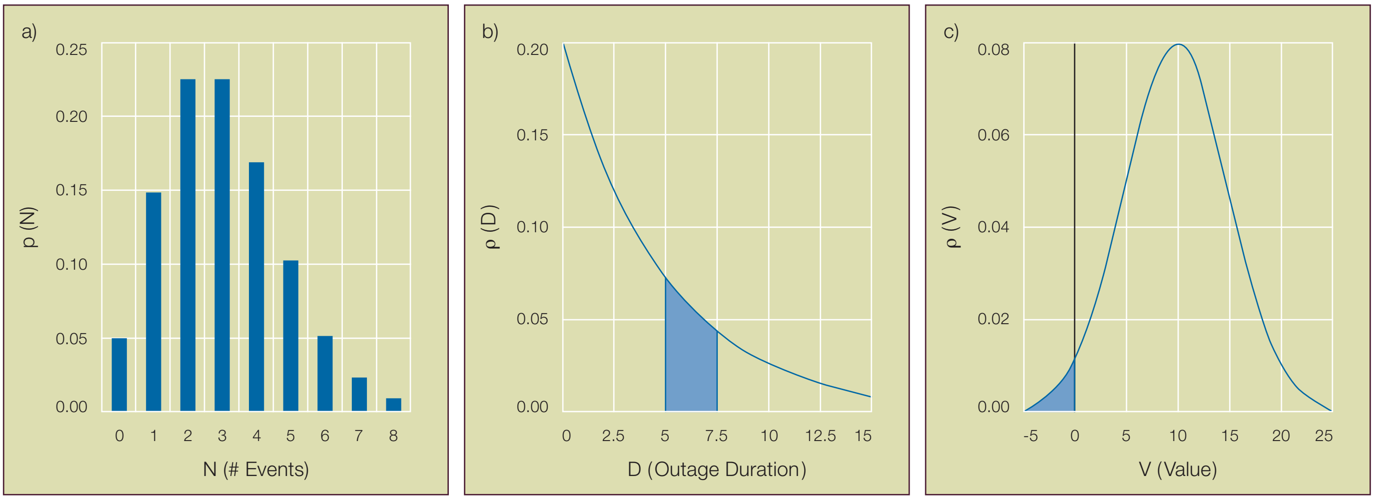

Poisson distribution | Exponential distribution | Normal (Gaussian) distribution

In mathematical terms, such areas (between minus infinity and a given value V) represent the corresponding integral over the probability density function. The resulting function is called 'cumulative probability distribution' 2. For any value of V, it specifies the probability that the (project) value is smaller than V.

2 Cumulative normal probability distribution (corresponding to the probability density of Figure 1(c)).

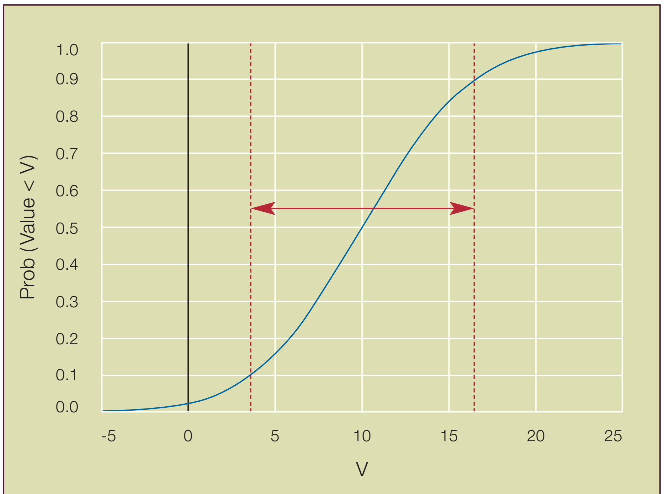

The dashed vertical lines indicate the 10 percent and 90 percent confidence limits for V.

From the cumulative probability function, we can read off most of the information required to quantitatively characterize the risks associated with the random nature of the corresponding variable. In 2, we have, for example, indicated the 10 percent and 90 percent confidence limits for the random variable V. These tell us that with a probability of 10 percent, V will be below 3.6, and with a probability of 90 percent below 16.4 (ie, above 16.4 with a probability of 10 percent). The corresponding 80 percent confidence interval is thus given by [3.6 < V < 16.4].

The most well-known characteristics of a random variable, its mean value μ and its variance σ^2, however, can in general not be evaluated from the cumulative probability distribution. They have to be determined with the help of the probability density. For the Poisson distribution of 1a, for example, μ is calculated by summing $N \cdot p(N)$ over all values of N, and σ^2 is then obtained by summing $(N - \mu)^2 \cdot p(N)$. For the examples shown in 1b and 1c, these sums are replaced by the corresponding integrals.

Risk measures

The main purpose of a quantification of risks and uncertainties is to obtain a sound basis for our respective decisions. As we have seen, all available information on an uncertain (random) variable is, in principle, contained in the corresponding probability distribution (or probability density). A probability distribution, however, is not a practical form of information for a decision maker, and there is thus a need for concise (preferably 'single number') risk measures.

Suppose we are concerned with the risk that the value of an investment portfolio (or of a planned project) changes in an unfavorable direction. The simplest measure of this risk is given by the mean (expected) value of the potential losses. Outside the finance industry, this is often the only quantitative risk measure that is estimated. The expected loss, however, is only one of the risk measures that have to be considered in a comprehensive risk management process. It represents an inherent cost of any business activity and thus affects the expected net revenue, but it gives no information about the probability and magnitude of larger than expected losses. Corresponding information, however, is important for determining the capital required to cover such losses.

More informative risk measures in this respect, eg, confidence limits, can be evaluated from the cumulative probability distribution of the relevant quantity 2. One of these risk measures, called Value at Risk (VaR) [1][2], has recently received a lot of attention and has now reached the status of a generally accepted risk measurement standard.

Value at Risk is defined as the expected *maximum* loss (over a given time period) that will not be exceeded with a given probability. To define a VaR-measure completely, we thus have to specify the time period over which the changes in value, ΔV, are considered, as well as the probability α (confidence level) with which a potential loss should not exceed VaR. In mathematical terms, VaR is determined by the following implicit equation:

$$\text{Prob}\ [\ \Delta V < -\ \text{VaR}\ ] = 1 - \alpha,$$

where Prob refers to the cumulative probability distribution of ΔV, the change in portfolio or project value over a given time period. The definition of VaR, for a confidence level of 95 percent, is illustrated in 3.

Value at Risk has become an important risk management instrument in the finance industry. According to the 'Capital Accord' of the Basel Committee on Banking Supervision [3], banks are now required to calculate VaR on a daily basis, for a time interval (holding period) of 10 days and a confidence level of 99 percent. The VaR-concept, however, also plays an increasingly important role in the non-finance industry, and analogous measures have been defined for a variety of business risks, eg, Profit at Risk and Credit Value at Risk.

The VaR-concept, on the other hand, also has some limitations. In particular, it generally provides no information on the expected magnitude of losses that are larger than the VaR-limit. Although such losses only occur with a small probability (equal to 1 – α), it may be very dangerous not to know 'how bad it can become if something goes wrong'.

Aggregation of risks

Most business risks depend on a number of different uncertain factors. If we now assume that we can determine or estimate the probability distributions of the individual risk factors, we are still faced with the problem of aggregating these into the probability distribution of the total risk.

A generally applicable and commonly used approach to solve such problems is Monte Carlo Simulation. In the Monte Carlo simulation method, the probability distribution of the total risk is determined by using a large sample of randomly generated values for the individual risk factors. These values are drawn from the known or estimated probability distributions of the different risk factors, or directly from corresponding historical data. Monte Carlo methods have the advantage that no specific form for the individual probability distributions has to be assumed and that correlations can easily be taken into account. The drawback, however, is that a reasonably accurate estimation of VaR usually requires a very large number of Monte Carlo steps.

If we are satisfied, however, with an approximate evaluation of the aggregated risk probability distribution, we can use analytical methods and do not have to perform time-consuming Monte Carlo simulations. A widely used approach is based on the central limit theorem. This tells us that the probability distribution of a sum of independent random variables can be approximated by a normal distribution if the number of individual random variables is large enough. A normal distribution, however, is completely characterized by its mean value and its variance, and for a sum of random variables, these are simply calculated by summing the respective contributions of the individual terms.

The approximation of an aggregated risk probability distribution by a normal (Gaussian) distribution is illustrated in 4. It refers to a sum of ten binary risk factors, each of which assumes a value of 10 with a probability of 0.2 and a value of zero with a probability of 0.8. In general, the 'Gauss approximation' is already sufficiently accurate for a sum of only about five risk factors, provided that their probability distributions are not too asymmetric, and if very precise information about the tails of the aggregated distribution is not required.

Assessment of probability distributions

For some types of risk, we can base our assessment of the corresponding probability distribution on historical data. This is certainly the case for market risks such as, eg, currency and interest rate risks. For banks, the Basel Committee on Banking Supervision indeed sets a number of binding standards [3]) for the measurement of such risks, eg:

- 'The choice of historical observation period (sample period) for calculating value-at-risk will be constrained to a minimum length of one year.'
- 'Banks should update their data sets no less frequently than once every three months and should also re-

3 Illustration of the definition of Value at Risk (VaR) for a confidence level of 95 percent.

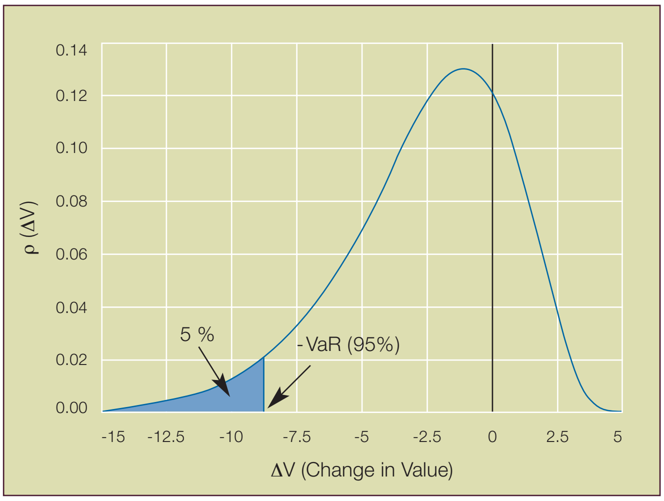

4 Cumulative distribution function for the aggregated risk of ten binary risk factors. Comparison with the approximation by a normal distribution (continuous curve).

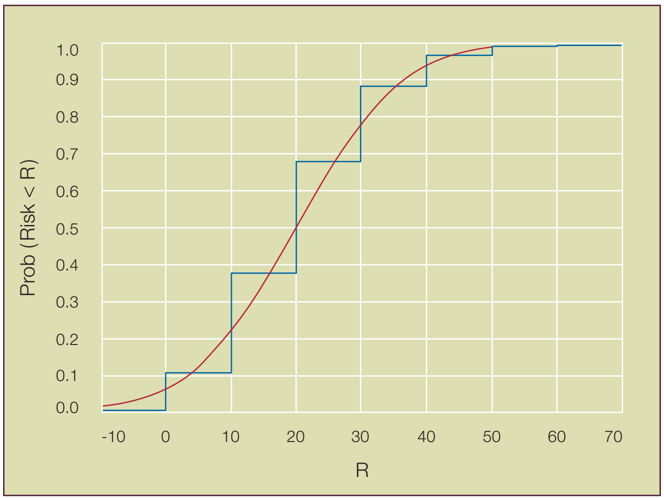

assess them whenever market prices are subject to material changes.'

As an example, 5 shows a histogram of the relative changes of the USD/CHF closing rates during 1993. To facilitate the calculation of VaR-limits, such a histogram of historical data is often approximated by a normal distribution with the same mean value and standard deviation. As the example shows, however, this approximation underestimates the very small and, more importantly, the very large changes. It is an established fact that the probability distributions of market risks have so-called 'fat tails', ie, large changes occur more frequently than estimated by a normal distribution, and there are indications that some technical risks (eg, power blackouts) exhibit the same phenomenon. To estimate the probability of large losses or damages, it is therefore important that these observations are adequately taken into account.

5 Histogram of the relative USD/CHF day-to-day changes during 1993.

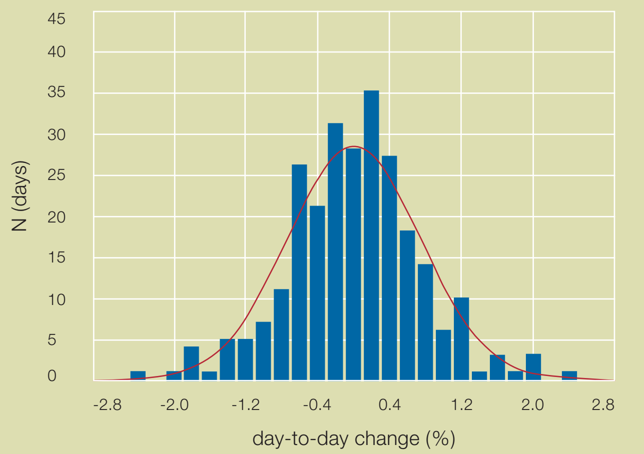

The continuous curve indicates the approximation by a normal probability distribution with the same mean value and standard deviation.

Outside the finance industry, historical data may be available in some specific cases (eg, outage and performance data for certain components or systems), but much more often the probabilistic characteristics of the different risks have to be assessed without reliable statistical data. In such cases, we have to base our analysis primarily on expert judgment. Experts, however, will usually not specify risks in terms of a probability distribution but, eg, in terms of an exposure and of a probability with which the corresponding costs will actually materialize (or in terms of the minimum, most likely, and maximum impact of a given risk). From these characteristics, we then have to infer, perhaps with some additional information, an adequate probability distribution, or at least an estimate for the mean value and variance of the different risk factors.

Scenario simulations

Given the fact that often only a rough and not very reliable estimate of the different risks is available, is it then still useful to perform a detailed quantitative analysis?

Most risk managers will agree that it is always better to have some quantitative information than none at all. However, in situations where we only have rough estimates or guesses about our risks, it obviously makes little sense to perform a very accurate mathematical analysis. It is then much more important that a corresponding evaluation is supplemented by appropriate scenario simulations. With the help of such simulations, we can, eg, analyze the effects of modeling uncertainties and determine the sensitivity of our results with respect to different assumptions.

As the inherent uncertainties associated with an extrapolation of risk estimations to future time periods can never be completely eliminated, scenario simulations play an important role in any quantitative risk analysis procedure. Moreover, scenario simulations often represent the only alternative to determine the impact of extreme events ('stress testing') or, eg, to examine the relative effects of different hedging strategies.

The use of scenario simulations will be demonstrated in part III of this series, where the possibilities and limitations of a quantitative risk analysis are illustrated on the basis of a fictitious but representative industrial project.

Jakob Bernasconi
ABB Switzerland Ltd.
Corporate Research
jakob.bernasconi@ch.abb.com

References

[1] see, eg, 'Value at Risk', by Thomas J. Linsmeier and Neil D. Pearson, Financial Analysts Journal, Vol. 56, No. 2, March/April 2000, pp. 47-67

[2] A good source for information about VaR and quantitative risk management in general is the GloriaMundi website ('All about Value at Risk™'): http://www.gloriamundi.org

[3] Basel Committee Publications No. 24, 'Amendment to the Capital Accord to Incorporate Market Risks', January 1996, modified in September 1997,http://www.bis.org/publ/bcbs24a.htm