1

Introduction

1.1 A view on the evaluation of risk

1.1.1 The role of mathematics

How is evaluation of risk influenced by modern computing? Consider the way we use mathematics, first as a vendor of models of complicated risk processes. These models are usually stochastic. They are in general insurance probability distributions of claim numbers and losses and in life insurance and finance, stochastic processes describing lifecycles and investment returns. Mathematics is from this point of view a *language*, a way risk is expressed, and it is a language we must master. Otherwise statements of risk cannot be related to reality, it would be impossible to say what conclusions mean in any precise manner and nor could analyses be presented effectively to clients. Actuarial science is in this sense almost untouched by modern computational facilities. The basic concepts and models remain what they were, notwithstanding, of course, the strong growth of risk products throughout the last decades. This development may have had something to do with computers, but not much with computing per se.

However, mathematics is also *deductions* with precise conclusions derived from precise assumptions through the rules of logic. That is the way mathematics is taught at school and university. It is here that computing enters applied mathematical disciplines like actuarial science. More and more of these deductions are implemented in computers and carried out there. This has been going on for decades. It leans on an endless growth in computing power, a true technological revolution opening up simpler and more general computational methods which require less of users.

1.1.2 Risk methodology

An example of such an all-purpose computational technique is **stochastic simulation**. Simplified versions of processes taking place in real life are then reproduced

The world of risk	*In the computer*	*Sources for* $\widehat{M}$:
		Historical experience
$M \rightarrow X$	$\widehat{M} \rightarrow X^*$	*The implied market view*
	$\uparrow$	*Deductions from no-arbitrage*
	Assumed mechanism	*Judgement, physical modelling*

Figure 1.1 The working process: Main steps in risk evaluation.

in the computer. Risk in finance and insurance is future uncertain gains and losses, designated in this book by letters such as X and Y. Typical examples are compensations for claims in general insurance, pension schemes interrupted upon death in life insurance and future values of shares and bonds in finance. There are also secondary (or derived) products where values and payoffs are channelled through contract clauses set up in advance. Such agreements are known as **derivatives** in finance and **reinsurance** in insurance.

The mathematical approach, unanimously accepted today, is through probabilities with risks X and Y seen as **random variables**. We shall know their values eventually (after the event), but for planning and control and to price risk-taking activities we need them in advance and must fall back on their probabilities. This leads to a working process such as the one depicted in Figure 1.1. The real world on the left is an enormously complicated mechanism (denoted M) that yields a future X.

We shall never know M, though our paradigm is that it does exist as a well-defined stochastic mechanism. Since it is beyond reach, a simplified version $\widehat{M}$ is constructed in its place and used to study X. Its expected value is used for valuation and the percentiles for control, and unlike in engineering we are rarely concerned with predicting a specific value. Note that everything falls apart if $\widehat{M}$ deviates too strongly from the true mechanism M. This issue of **error** is a serious one indeed. Chapter 7 is an introduction.

What there is to go on when $\widehat{M}$ is constructed is listed on the right in Figure 1.1. Learning from the past is an obvious source (but not all of it is relevant). In finance, current asset prices bring market opinion about the future. This so-called **implied view** is introduced briefly in Section 1.4, and there will be more in Part III. Then there is the theory of **arbitrage**, where riskless financial income is assumed impossible. The innocent-looking no-arbitrage condition has wide implications, which are discussed in Chapter 14. In practice, some personal judgement behind $\widehat{M}$ is often present, but this is not for general argument, and nor shall we go into the physical modelling used in large-claims insurance where hurricanes, earthquakes or floods are imitated in the computer. This book is about how $\widehat{M}$ is constructed from the first three sources (historical data above all), how it is implemented in the computer and how the computer model is used to determine the probability distribution of X. Note that $\widehat{M}$ is inevitably linked to the past even though perceived trends and changes may have been built into it, and there is no way of knowing how well it captures a future which sometimes extends over decades. While this doesn't make the mathematical approach powerless, it does suggest simple and transparent models and a humble attitude towards it all.

1.1.3 The computer model

The real risk variable X will materialize only once. The economic result of a financial investment in a particular year is a unique event, as is the aggregated claim against an insurance portfolio during a certain period of time. With the computer model, that is different. Once it has been set up it can be played as many times as we please. Let $X_1^*, \dots, X_m^*$ be realizations of X revealing which values are likely and which are not, and how bad things might be if we are unlucky. The * will be used to distinguish computer simulations from real variables and m will always denote the number of simulations.

The method portrayed on the left of Figure 1.1 is known as the **Monte Carlo** method or **stochastic simulation**. It belongs to the realm of numerical integration; see Evans and Schwarz (2000) for a summary of this important branch of numerical mathematics. Monte Carlo integration dates back a long way. It is computationally slow, but other numerical methods (that might do the job faster) often require more expertise and get bogged down for high-dimensional integrals, which are precisely what we often need in practice. The Monte Carlo method is unique in handling many variables well.

What is the significance of numerical speed anyway? Does it really matter that some specialized technique (demanding more time and know-how to implement) is (say) one hundred times faster when the one we use only takes a second? If the procedure for some reason is to be repeated in a loop thousands of times, it would matter. Often, however, slow Monte Carlo is quite enough, and, indeed, the practical limit to its use is moving steadily as computers become more and more powerful. How far have we got? The author's portable computer from 2006 (with T60p processor) took three seconds to produce ten million Pareto draws through Algorithm 2.13 implemented in Fortran. This is an insurance portfolio of 1000 claims simulated 10 000 times! Generating the normal is even faster, and if speed is a priority, try the table methods in Section 4.2.

One of the aims of this book is to demonstrate how these opportunities are utilized. Principal issues are how simulations programs are designed, how they are modified to deal with related (but different) problems and how different programs are merged to handle situations of increasing complexity with several risk factors contributing jointly. The versatility and usefulness of Monte Carlo is indicated in

Section 1.5 (and in Chapter 3 too). By mastering it you are well equipped to deal with much that comes your way and avoid getting stuck when pre-programmed software doesn't have what you need. What platform should you go for? Algorithms in this book are written in a pseudo-code that goes with everything. Excel and Visual Basic are standard in the industry and may be used even for simulation. Much higher speed is obtained with C, Pascal or Fortran, and in the opinion of this author people are well advised to learn software like those. There are other possibilities as well, and the open-source R-package is used with the exercises. Much can be achieved with a platform you know!

1.2 Insurance risk: Basic concepts

1.2.1 Introduction

Property or **general insurance** is economic responsibility for incidents such as fires or accidents passed on (entirely or in part) to an insurer against a fee. The contract, known as a **policy**, releases indemnities (**claims**) when such events occur. A central quantity is the total claim X amassed during a certain period of time (typically a year). Often $X = 0$ (no events), but on rare occasions X is huge. An insurance company copes whatever happens, if properly run. It has a **portfolio** of many such risks and only a few of them materialize. But this raises the issue of controlling the total uncertainty, which is a major theme in general insurance.

Life insurance is also built up from random payments X. **Term insurance**, where beneficiaries receive compensation upon the death of the policy holder, is similar to property insurance in that unexpected events lead to payoffs. **Pension schemes** are the opposite. Now the payments go on as long as the insured is alive, and they are likely, not rare. Yet the basic approach remains the same, with random variables X expressing the uncertainty involved.

1.2.2 Pricing insurance risk

Transfers of risk through X do not take place for free. The fee (or **premium**), charged in advance, depends on the market conditions, but the expectation is a guideline. Introduce

$$\pi^{\text{pu}} = E(X), \tag{1.1}$$

which is known as the **pure premium** and defines a break-even situation. A company receiving π^{pu} for its services will, in the absence of all overhead cost and all financial income, neither earn nor lose in the long run. This is a consequence of the law of large numbers in probability theory; see Appendix A.2.

Such a pricing strategy is (of course) out of the question, and companies add **loadings** γ on top of π^{pu}. The premium charged is then

$$\pi = (1 + \gamma)\pi^{\text{pu}}, \tag{1.2}$$

and we may regard $\gamma\pi^{\text{pu}}$ as the cost of risk. It is influenced thoroughly by the market situation, and in many branches of insurance is known to exhibit strong fluctuations; see Section 11.5 for a simple model. There have been attempts to determine γ from theoretical arguments, see Young (2004) for a review, but these efforts are not used much in practice and will not be considered.

The loading concept separates the market side from the insurance process itself, but another issue is whether the pure premium is known. Stochastic models for X always depend on unknown quantities such as parameters or probability distributions. They are determined from experience or even assessed informally if historical data are lacking, and there is a crucial distinction between the true π^{pu} with perfect knowledge of the underlying situation and the $\hat{\pi}^{\text{pu}}$ used for analysis and decisions. The discrepancy between what we seek and what we get is a fundamental issue of **error** that is present everywhere (see Figure 1.1), and there is special notation for it. A parameter or quantity with a ˆ such as $\hat{\psi}$ means an estimate or an assessment of an underlying, unknown ψ. Chapter 7 offers a general discussion of errors and how they are confronted.

1.2.3 Portfolios and solvency

A second major theme in insurance is **control**. Companies are obliged to set aside funds to cover future obligations, and this is even a major theme in the *legal* definition of insurance. A company carries responsibility for many policies. It will lose on some and gain on others. In property insurance policies without accidents are profitable, those with large claims are not. Long lives in pension schemes lead to losses, short ones to gains. *At the portfolio level, gains and losses average out*. This is the beauty of a large agent handling many risks simultaneously.

Suppose a portfolio consists of J policies with claims $X_1, \dots, X_J$. The total claim is then

$$\mathcal{X} = X_1 + \cdots + X_J, \tag{1.3}$$

where calligraphic letters like $\mathcal{X}$ will be used for quantities at the portfolio level. We are certainly interested in $E(\mathcal{X})$, but equally important is its distribution. Regulators demand sufficient funds to cover $\mathcal{X}$ with high probability. The mathematical formulation is in terms of a percentile q_ϵ, which is the solution of the equation

$$\Pr(\mathcal{X} > q_\epsilon) = \epsilon \tag{1.4}$$

where ϵ is a small number (for example 1%). The amount q_ϵ is known as the **solvency capital** or **reserve**. Percentiles are used in finance too and are then often called value at risk (or **VaR** for short). As elsewhere, the true q_ϵ we seek is not the same as the estimated $\hat{q}_\epsilon$ we get.

1.2.4 Risk ceding and reinsurance

Risk is ceded from ordinary policy holders to companies, but companies do the same thing between themselves. This is known as **reinsurance**, and the ceding company is known as the **cedent**. The rationale *could* be the same; i.e., that a financially weaker agent is passing risk to a stronger one. In reality even the largest companies do this to diversify risk, and financially the cedent may be as strong as the reinsurer. There is now a chain of responsibilities that can be depicted as follows:

$$\begin{array}{ccccc} \text{original clients} & \longrightarrow & \text{cedent} & \longrightarrow & \text{reinsurer} \\ \mathcal{X}\ \textit{(primary)} & & \mathcal{X}^{\text{ce}} = \mathcal{X} - \mathcal{X}^{\text{re}} & & \mathcal{X}^{\text{re}}\ \textit{(derived)} \end{array}$$

The original risk $\mathcal{X}$ is split between cedent and reinsurer through two separate relationships, where the cedent part $\mathcal{X}^{\text{ce}}$ is net and the difference between two cash flows. Of course $\mathcal{X}^{\text{re}} \leq \mathcal{X}$; i.e., the responsibility of the reinsurer is always *less* than the original claim. Note the calligraphic style that applies to portfolios. There may in practice be several rounds of such cedings in complicated networks extending around the globe. One reinsurer may go to a second reinsurer, and so on. Modern methods provide the means to analyse risk taken by an agent who is far away from the primary source. Ceding and reinsurance are tools used by managers to tune portfolios to a desired risk profile.

1.3 Financial risk: Basic concepts

1.3.1 Introduction

Gone are the days when insurance liabilities were insulated from assets and insurance companies carried all the financial risk themselves. One trend is ceding to customers. In countries like the USA and Britain, insurance products with financial risk integrated have been sold for decades under names such as unit link or universal life. The rationale is that clients receive higher expected financial income in exchange for carrying more risk. Pension plans today are increasingly **contributed benefits** (CB), where financial risk rests with individuals. There is also much interest in investment strategies tailored to given liabilities and how they distribute over time. This is known as **asset liability management** (ALM) and is discussed in

Chapter 15. The present section and the next one review the main concepts of finance.

1.3.2 Rates of interest

An ordinary bank deposit v_0 grows to $(1 + r)v_0$ at the end of one period and to $(1 + r)^K v_0$ after K of them. Here r, the **rate of interest**, depends on the length of the period. Suppose interest is compounded over K segments, each of length $1/K$, so that the total time is one. If interest per segment is r/K, the value of the account becomes

$$\left(1 + \frac{r}{K}\right)^K v_0 \rightarrow e^r v_0, \quad \text{as} \quad K \rightarrow \infty,$$

after one of the most famous limits of mathematics. Interest earnings may therefore be cited as

$$rv_0 \quad \text{or} \quad (e^r - 1)v_0,$$

depending on whether we include ‘interest on interest’. The second form implies continuous compounding of interest and higher earnings ($e^r - 1 > r$ if $r > 0$), and now $(e^r)^k = e^{rk}$ takes over from $(1 + r)^k$. It doesn’t really matter which form we choose, since they can be made equivalent by adjusting r.

1.3.3 Financial returns

Let V_0 be the value of a financial asset at the start of a period and V_1 the value at the end of it. The relative gain

$$R = \frac{V_1 - V_0}{V_0} \tag{1.5}$$

is known as the **return** on the asset. Solving for V_1 yields

$$V_1 = (1 + R)V_0, \tag{1.6}$$

with RV_0, the financial income. Clearly R acts like interest, but there is more to it than that. Interest is a fixed benefit offered by a bank (or an issuer of a very secure bond) in return for making a deposit and is risk free. Shares of company stock, in contrast, are fraught with risk. They may go up (R positive) or down (R negative). When dealing with such assets, V_1 (and hence R) is determined by the market, whereas with ordinary interest r is given and V_1 follows.

The return R is the more general concept and is a random variable with a probability distribution. Take the randomness away, and we are back to a fixed rate of interest r. As r depends on the time between V_0 and V_1, so does the distribution of

R; as will appear many times in this book. Whether the rate of interest r really *is* risk free is not as obvious as it seems. True, you do get a fixed share of your deposit as a reward, but that does not tell its worth in **real** terms when price increases are taken into account. Indeed, over longer time horizons risk due to inflation may be huge and even reduce the real value of cash deposits and bonds. Saving money with a bank at a fixed rate of interest may also bring **opportunity cost** if the market rate after a while exceeds what you get. These issues are discussed and integrated with other sources of risk in Part III.

1.3.4 Log-returns

Economics and finance have often constructed stochastic models for R directly. An alternative is the **log-return**

$$X = \log(1 + R), \tag{1.7}$$

which by (1.5) can be written $X = \log(V_1) - \log(V_0)$; i.e., as a difference of logarithms. The modern theory of financial derivatives (Section 3.5 and Chapter 14) is based on X. Actually, X and R do not necessarily deviate that strongly since the Taylor series of $\log(1 + R)$ is

$$X = R - \frac{R^2}{2} + \frac{R^3}{3} + \cdots,$$

where R (a fairly small number) dominates so that $X \doteq R$, at least over short periods. It follows that the distributions of R and X must be rather similar (see Section 2.4), but this is not to say that the discrepancy is unimportant. It depends on the amount of random variation present, and the longer the time horizon the more X deviates from R; see Section 5.4 for an illustration.

1.3.5 Financial portfolios

Investments are often spread over many assets as **baskets** or financial **portfolios**. By intuition this must reduce risk; see Section 5.3, where the issue is discussed. A central quantity is the portfolio return, denoted $\mathcal{R}$ (in calligraphic style). Its relationship to the individual returns R_j of the assets is as follows. Let $V_{10}, \ldots, V_{J0}$ be investments in J assets. The portfolio value is then

$$\mathcal{V}_0 = \sum_{j=1}^{J} V_{j0} \quad \text{growing at the end of the period to} \quad \mathcal{V}_1 = \sum_{j=1}^{J}(1 + R_j)V_{j0}.$$

Subtract $\mathcal{V}_0$ from $\mathcal{V}_1$ and divide by $\mathcal{V}_0$, and you get the portfolio return

$$\mathcal{R} = \sum_{j=1}^{J} w_j R_j \quad \text{where } w_j = \frac{V_{0j}}{\mathcal{V}_0}. \tag{1.8}$$

Here w_j is the **weight** on asset j and

$$w_1 + \cdots + w_J = 1. \tag{1.9}$$

Financial weights define the distribution on individual assets and will, in this book, usually be normalized so that they sum to 1.

The mathematics allow negative w_j. With bank deposits this corresponds to borrowing. It is also possible with shares, known as **short selling**. A loss due to a negative development is then carried by somebody else. The mechanism is as follows. A short contract with a buyer is to sell shares at the end of the period at an agreed price. At that point we shall have to buy at market price, gaining if it is lower than our agreement, losing if not. Short contracts may be an instrument to lower risk (see Section 5.3) and require **liquidity**; i.e., assets that are traded regularly.

1.4 Risk over time

1.4.1 Introduction

A huge number of problems in finance and insurance have time as one of the central ingredients, and this requires additional quantities and concepts. Many of these are introduced below. The emphasis is on finance, where the number of such quantities is both more plentiful and more complex than in insurance. Time itself is worth a comment. In this book it will be run on equidistant sequences, either

$$\underbrace{T_k = kT}_{\text{time scale for evaluation}} \qquad \text{or} \qquad \underbrace{t_k = kh}_{\text{time scale for modelling}} \tag{1.10}$$

for $k = 0, 1, \ldots$ On the left, T is an accounting period (e.g., year, quarter, month) or the time to expiry of a bond or an option. Financial returns R_k, portfolio values $\mathcal{V}_k$ and insurance liabilities $\mathcal{X}_k$ are followed over T_k. The present is always at $T_0 = 0$, whereas $k > 0$ is the future which requires stochastic models to portray what is likely and what is not. *Negative* time will sometimes be used for past values.

The time scale h is used for modelling. It may coincide with T, but it may well be smaller so that $T = Kh$ for $K > 1$. How models on different time scales are related is important; see Section 5.7, where this issue is discussed. There is also much scope for very short time increments where $h \to 0$ (so that $K = T/h \to \infty$). This is known as **continuous-time modelling** and is above all a trick to find simple mathematical solutions. Parameters or variables are then often cited as **intensities**, which are quantities per time unit. An example is interest rates, which will on several occasions be designated rh with r an intensity and not a rate as in Section 1.3. Claim frequencies in property insurance (Chapter 8) and mortalities in life insurance (Chapter 12) are other examples of using intensities for modelling. This section is concerned with the macro time scale T only.

1.4.2 Accumulation of values

If v_0 is the original value of a financial asset, by $T_K = KT$ it is worth

$$V_K = (1 + R_1)(1 + R_2) \cdots (1 + R_K)V_0 = (1 + R_{0:K})v_0,$$

where $R_1, \ldots, R_K$ are the returns. This defines $R_{0:K}$ on the right as the K**-step** return

$$\underbrace{R_{0:K} = (1 + R_1) \cdots (1 + R_K) - 1}_{\text{ordinary returns}} \quad \text{and also} \quad \underbrace{X_{0:K} = X_1 + \cdots + X_K,}_{\text{log-returns}} \tag{1.11}$$

where $X_k = \log(1 + R_k)$ and $X_{0:K} = \log(1 + R_{0:K})$. The log-returns on the right are accumulated by adding them. Interest is a special case (an important one!) and grows from $T_0 = 0$ to T_K according to

$$r_{0:K} = (1 + r_1)(1 + r_2) \cdots (1 + r_K) - 1, \tag{1.12}$$

where $r_1, \ldots, r_K$ are the future rates. This reduces to $r_{0:K} = (1+r)^K - 1$ if all $r_k = r$, but in practice r_k will float in a way that is unknown at $T_0 = 0$.

Often V_K aggregates economic and financial activity beyond the initial investment v_0. Let B_k be income or expenses that surface at time T_k, and suppose the financial income or loss coming from the sequence $B_1, \ldots, B_K$ is the same as for the original asset. The total value at T_K is then the sum

$$V_K = (1 + R_{0:K})v_0 + \sum_{k=1}^{K}(1 + R_{k:K})B_k, \tag{1.13}$$

where $R_{k:K} = (1 + R_{k+1}) \cdots (1 + R_K) - 1$ with $R_{K:K} = 0$. Later in this section $B_1, \ldots, B_K$ will be a fixed cash flow, but further on (for example in Section 3.6) there will be huge uncertainty as to what their values are going to be, with additional random variation on top of the financial uncertainty.

1.4.3 Forward rates of interest

Future interest rates like r_k or $r_{0:K}$ are hopeless to predict from mathematical models (you will see why in Section 6.4), but there is also a market view that conveys