

Preface to the Second Edition

When I took office, only high energy physicists had ever heard of what is called the Worldwide Web . . . Now even my cat has its own page — Bill Clinton, 1996

This book gives a comprehensive survey of non-life insurance mathematics. Originally written for use with the actuarial science programs at the Universities of Amsterdam and Leuven, it is now in use at many other universities, as well as for the non-academic actuarial education program organized by the Dutch Actuarial Society. It provides a link to the further theoretical study of actuarial science. The methods presented can not only be used in non-life insurance, but are also effective in other branches of actuarial science, as well as, of course, in actuarial practice.

Apart from the standard theory, this text contains methods directly relevant for actuarial practice, for example the rating of automobile insurance policies, premium principles and risk measures, and IBNR models. Also, the important actuarial statistical tool of the Generalized Linear Models is studied. These models provide extra possibilities beyond ordinary linear models and regression, the statistical tools of choice for econometricians. Furthermore, a short introduction is given to credibility theory. Another topic that always has enjoyed the attention of risk theoreticians is the study of ordering of risks. The book reflects the state of the art in actuarial risk theory; many results presented were published in the actuarial literature only recently.

In this second edition of the book, we have aimed to make the theory even more directly applicable by using the software R. It provides an implementation of the language S, not unlike S-Plus. It is not just a set of statistical routines but a full-fledged object oriented programming language. Other software may provide similar capabilities, but the great advantage of R is that it is open source, hence available to everyone free of charge. This is why we feel justified in imposing it on the users of this book as a de facto standard. On the internet, a lot of documentation about R can be found. In an Appendix, we give some examples of use of R. After a general introduction, explaining how it works, we study a problem from risk management, trying to forecast the future behavior of stock prices with a simple model, based on stock prices of three recent years. Next, we show how to use R to generate pseudo-random datasets that resemble what might be encountered in actuarial practice.

Models and paradigms studied

The time aspect is essential in many models of life insurance. Between paying premiums and collecting the resulting pension, decades may elapse. This time element is less prominent in non-life insurance. Here, however, the statistical models are generally more involved. The topics in the first five chapters of this textbook are basic for non-life actuarial science. The remaining chapters contain short introductions to other topics traditionally regarded as non-life actuarial science.

1. *The expected utility model*

The very existence of insurers can be explained by the expected utility model. In this model, an insured is a risk averse and rational decision maker, who by virtue of Jensen's inequality is ready to pay more than the expected value of his claims just to be in a secure financial position. The mechanism through which decisions are taken under uncertainty is not by direct comparison of the expected payoffs of decisions, but rather of the expected utilities associated with these payoffs.

2. *The individual risk model*

In the individual risk model, as well as in the collective risk model below, the total claims on a portfolio of insurance contracts is the random variable of interest. We want to compute, for example, the probability that a certain capital will be sufficient to pay these claims, or the value-at-risk at level 99.5% associated with the portfolio, being the 99.5% quantile of its cumulative distribution function (cdf). The total claims is modeled as the sum of all claims on the policies, which are assumed independent. Such claims cannot always be modeled as purely discrete random variables, nor as purely continuous ones, and we use a notation, involving Stieltjes integrals and differentials, encompassing both these as special cases.

The individual model, though the most realistic possible, is not always very convenient, because the available dataset is not in any way condensed. The obvious technique to use in this model is convolution, but it is generally quite awkward. Using transforms like the moment generating function sometimes helps. The Fast Fourier Transform (FFT) technique gives a fast way to compute a distribution from its characteristic function. It can easily be implemented in R.

We also present approximations based on fitting moments of the distribution. The Central Limit Theorem, fitting two moments, is not sufficiently accurate in the important right-hand tail of the distribution. So we also look at some methods using three moments: the translated gamma and the normal power approximation.

3. *Collective risk models*

A model that is often used to approximate the individual model is the collective risk model. In this model, an insurance portfolio is regarded as a process that produces claims over time. The sizes of these claims are taken to be independent, identically distributed random variables, independent also of the number of claims generated. This makes the total claims the sum of a random number of iid individual claim amounts. Usually one assumes additionally that the number of claims is a Poisson variate with the right mean, or allows for some overdispersion by taking a negative

binomial claim number. For the cdf of the individual claims, one takes an average of the cdfs of the individual policies. This leads to a close fitting and computationally tractable model. Several techniques, including Panjer's recursion formula, to compute the cdf of the total claims modeled this way are presented.

For some purposes it is convenient to replace the observed claim severity distribution by a parametric loss distribution. Families that may be considered are for example the gamma and the lognormal distributions. We present a number of such distributions, and also demonstrate how to estimate the parameters from data. Further, we show how to generate pseudo-random samples from these distributions, beyond the standard facilities offered by R.

4. *The ruin model*

The ruin model describes the stability of an insurer. Starting from capital u at time $t = 0$, his capital is assumed to increase linearly in time by fixed annual premiums, but it decreases with a jump whenever a claim occurs. Ruin occurs when the capital is negative at some point in time. The probability that this ever happens, under the assumption that the annual premium as well as the claim generating process remain unchanged, is a good indication of whether the insurer's assets match his liabilities sufficiently. If not, one may take out more reinsurance, raise the premiums or increase the initial capital.

Analytical methods to compute ruin probabilities exist only for claims distributions that are mixtures and combinations of exponential distributions. Algorithms exist for discrete distributions with not too many mass points. Also, tight upper and lower bounds can be derived. Instead of looking at the ruin probability $\psi(u)$ with initial capital u , often one just considers an upper bound e^{-Ru} for it (Lundberg), where the number R is the so-called adjustment coefficient and depends on the claim size distribution and the safety loading contained in the premium.

Computing a ruin probability assumes the portfolio to be unchanged eternally. Moreover, it considers just the insurance risk, not the financial risk. Therefore not much weight should be attached to its precise value beyond, say, the first relevant decimal. Though some claim that survival probabilities are 'the goal of risk theory', many actuarial practitioners are of the opinion that ruin theory, however topical still in academic circles, is of no significance to them. Nonetheless, we recommend to study at least the first three sections of Chapter 4, which contain the description of the Poisson process as well as some key results. A simple proof is provided for Lundberg's exponential upper bound, as well as a derivation of the ruin probability in case of exponential claim sizes.

5. *Premium principles and risk measures*

Assuming that the cdf of a risk is known, or at least some characteristics of it like mean and variance, a premium principle assigns to the risk a real number used as a financial compensation for the one who takes over this risk. Note that we study only risk premiums, disregarding surcharges for costs incurred by the insurance company. By the law of large numbers, to avoid eventual ruin the total premium should be at least equal to the expected total claims, but additionally, there has to be a loading in

the premium to compensate the insurer for making available his risk carrying capacity. From this loading, the insurer has to build a reservoir to draw upon in adverse times, so as to avoid getting in ruin. We present a number of premium principles, together with the most important properties that characterize premium principles. The choice of a premium principle depends heavily on the importance attached to such properties. There is no premium principle that is uniformly best.

Risk measures also attach a real number to some risky situation. Examples are premiums, infinite ruin probabilities, one-year probabilities of insolvency, the required capital to be able to pay all claims with a prescribed probability, the expected value of the shortfall of claims over available capital, and more.

6. *Bonus-malus systems*

With some types of insurance, notably car insurance, charging a premium based exclusively on factors known a priori is insufficient. To incorporate the effect of risk factors of which the use as rating factors is inappropriate, such as race or quite often sex of the policy holder, and also of non-observable factors, such as state of health, reflexes and accident proneness, many countries apply an experience rating system. Such systems on the one hand use premiums based on a priori factors such as type of coverage and list-price or weight of a car, on the other hand they adjust these premiums by using a bonus-malus system, where one gets more discount after a claim-free year, but pays more after filing one or more claims. In this way, premiums are charged that reflect the exact driving capabilities of the driver better. The situation can be modeled as a Markov chain.

The quality of a bonus-malus system is determined by the degree in which the premium paid is in proportion to the risk. The Loimaranta efficiency equals the elasticity of the mean premium against the expected number of claims. Finding it involves computing eigenvectors of the Markov matrix of transition probabilities. R provides tools to do this.

7. *Ordering of risks*

It is the very essence of the actuary's profession to be able to express preferences between random future gains or losses. Therefore, stochastic ordering is a vital part of his education and of his toolbox. Sometimes it happens that for two losses X and Y , it is known that every sensible decision maker prefers losing X , because Y is in a sense 'larger' than X . It may also happen that only the smaller group of all risk averse decision makers agree about which risk to prefer. In this case, risk Y may be larger than X , or merely more 'spread', which also makes a risk less attractive. When we interpret 'more spread' as having thicker tails of the cumulative distribution function, we get a method of ordering risks that has many appealing properties. For example, the preferred loss also outperforms the other one as regards zero utility premiums, ruin probabilities, and stop-loss premiums for compound distributions with these risks as individual terms. It can be shown that the collective model of Chapter 3 is more spread than the individual model it approximates, hence using the collective model, in most cases, leads to more conservative decisions regarding premiums to be asked, reserves to be held, and values-at-risk. Also, we can prove

that the stop-loss insurance, demonstrated to be optimal as regards the variance of the retained risk in Chapter 1, is also preferable, other things being equal, in the eyes of all risk averse decision makers.

Sometimes, stop-loss premiums have to be set under incomplete information. We give a method to compute the maximal possible stop-loss premium assuming that the mean, the variance and an upper bound for a risk are known.

In the individual and the collective model, as well as in ruin models, we assume that the claim sizes are stochastically independent non-negative random variables. Sometimes this assumption is not fulfilled, for example there is an obvious dependence between the mortality risks of a married couple, between the earthquake risks of neighboring houses, and between consecutive payments resulting from a life insurance policy, not only if the payments stop or start in case of death, but also in case of a random force of interest. We give a short introduction to the risk ordering that applies for this case. It turns out that stop-loss premiums for a sum of random variables with an unknown joint distribution but fixed marginals are maximal if these variables are as dependent as the marginal distributions allow, making it impossible that the outcome of one is ‘hedged’ by another.

In finance, frequently one has to determine the distribution of the sum of dependent lognormal random variables. We apply the theory of ordering of risks and comonotonicity to give bounds for that distribution.

We also give a short introduction in the theory of ordering of multivariate risks. One might say that two random variables are more related than another pair with the same marginals if their correlation is higher. But a more robust criterion is to restrict this to the case that their joint cdf is uniformly larger. In that case it can be proved that the sum of these random variables is larger in stop-loss order. There are bounds for joint cdfs dating back to Fréchet in the 1950s and Höfdding in the 1940s. For a random pair (X, Y) , the copula is the joint cdf of the ranks $F_X(X)$ and $F_Y(Y)$. Using the smallest and the largest copula, it is possible to construct random pairs with arbitrary prescribed marginals and (rank) correlations.

8. *Credibility theory*

The claims experience on a policy may vary by two different causes. The first is the quality of the risk, expressed through a risk parameter. This represents the average annual claims in the hypothetical situation that the policy is monitored without change over a very long period of time. The other is the purely random good and bad luck of the policyholder that results in yearly deviations from the risk parameter. Credibility theory assumes that the risk quality is a drawing from a certain structure distribution, and that conditionally given the risk quality, the actual claims experience is a sample from a distribution having the risk quality as its mean value. The predictor for next year’s experience that is linear in the claims experience and optimal in the sense of least squares turns out to be a weighted average of the claims experience of the individual contract and the experience for the whole portfolio. The weight factor is the credibility attached to the individual experience, hence it is called the credibility factor, and the resulting premiums are called credibility

premiums. As a special case, we study a bonus-malus system for car insurance based on a Poisson-gamma mixture model.

Credibility theory is actually a Bayesian inference method. Both credibility and generalized linear models (see below) are in fact special cases of so-called Generalized Linear Mixed Models (GLMM), and the R function `glmm` is able to deal with both the random and the fixed parameters in these models.

9. Generalized linear models

Many problems in actuarial statistics are Generalized Linear Models (GLM). Instead of assuming a normally distributed error term, other types of randomness are allowed as well, such as Poisson, gamma and binomial. Also, the expected values of the dependent variables need not be linear in the regressors. They may also be some function of a linear form of the covariates, for example the logarithm leading to the multiplicative models that are appropriate in many insurance situations.

This way, one can for example tackle the problem of estimating the reserve to be kept for IBNR claims, see below. But one can also easily estimate the premiums to be charged for drivers from region i in bonus class j with car weight w .

In credibility models, there are random group effects, but in GLMs the effects are fixed, though unknown. The `glmm` function in R can handle a multitude of models, including those with both random and fixed effects.

10. IBNR techniques

An important statistical problem for the practicing actuary is the forecasting of the total of the claims that are Incurred, But Not Reported, hence the acronym IBNR, or not fully settled. Most techniques to determine estimates for this total are based on so-called run-off triangles, in which claim totals are grouped by year of origin and development year. Many traditional actuarial reserving methods turn out to be maximum likelihood estimations in special cases of GLMs.

We describe the workings of the ubiquitous chain ladder method to predict future losses, as well as, briefly, the Bornhuetter-Ferguson method, which aims to incorporate actuarial knowledge about the portfolio. We also show how these methods can be implemented in R, using the `glm` function. In this same framework, many extensions and variants of the chain ladder method can easily be introduced. England and Verrall have proposed methods to describe the prediction error with the chain ladder method, both an analytical estimate of the variance and a bootstrapping method to obtain an estimate for the predictive distribution. We describe an R implementation of these methods.

11. More on GLMs

For the second edition, we extended the material in virtually all chapters, mostly involving the use of R, but we also add some more material on GLMs. We briefly recapitulate the Gauss-Markov theory of ordinary linear models found in many other texts on statistics and econometrics, and explain how the algorithm by Nelder and Wedderburn works, showing how it can be implemented in R. We also study the stochastic component of a GLM, stating that the observations are independent

random variables with a distribution in a subclass of the exponential family. The well-known normal, Poisson and gamma families have a variance proportional to μ^p for $p = 0, 1, 2$, respectively, where μ is the mean (heteroskedasticity). The so-called Tweedie class contains random variables, in fact compound Poisson–gamma risks, having variance proportional to μ^p for some $p \in (1, 2)$. These mean-variance relations are interesting for actuarial purposes. Extensions to R contributed by Dunn and Smyth provide routines computing cdf, inverse cdf, pdf and random drawings of such random variables, as well as to estimate GLMs with Tweedie distributed risks.

Educational aspects

As this text has been in use for a long time now at the University of Amsterdam and elsewhere, we could draw upon a long series of exams, resulting in long lists of exercises. Also, many examples are given, making this book well-suited as a textbook. Some less elementary exercises have been marked by [♠], and these might be skipped.

The required mathematical background is on a level such as acquired in the first stage of a bachelors program in quantitative economics (econometrics or actuarial science), or mathematical statistics. To indicate the level of what is needed, the book by Bain and Engelhardt (1992) is a good example. So the book can be used either in the final year of such a bachelors program, or in a subsequent masters program in actuarial science proper or in quantitative financial economics with a strong insurance component. To make the book accessible to non-actuaries, notation and jargon from life insurance mathematics is avoided. Therefore also students in applied mathematics or statistics with an interest in the stochastic aspects of insurance will be able to study from this book. To give an idea of the mathematical rigor and statistical sophistication at which we aimed, let us remark that moment generating functions are used routinely, while characteristic functions and measure theory are avoided in general. Prior experience with regression models, though helpful, is not required.

As a service to the student help is offered, in Appendix B, with many of the exercises. It takes the form of either a final answer to check one's work, or a useful hint. There is an extensive index, and the tables that might be needed in an exam are printed in the back. The list of references is not a thorough justification with bibliographical data on every result used, but more a collection of useful books and papers containing more details on the topics studied, and suggesting further reading.

Ample attention is given to exact computing techniques, and the possibilities that R provides, but also to old fashioned approximation methods like the Central Limit Theorem (CLT). The CLT itself is generally too crude for insurance applications, but slight refinements of it are not only fast, but also often prove to be surprisingly accurate. Moreover they provide solutions of a parametric nature such that one does not have to recalculate everything after a minor change in the data. Also, we want to stress that 'exact' methods are as exact as their input. The order of magnitude of errors resulting from inaccurate input is often much greater than the one caused by using an approximation method.

The notation used in this book conforms to what is usual in mathematical statistics as well as non-life insurance mathematics. See for example the book by Bowers et al. (1986, 1997), the non-life part of which is similar in design to the first part of this book. In particular, random variables are capitalized, though not all capitals actually denote random variables.

Acknowledgments

First and most of all, the authors would like to thank David Vyncke for all the work he did on the first edition of this book. Many others have provided useful input. We thank (in pseudo-random order) Hans Gerber, Elias Shiu, Angela van Heerwaarden, Dennis Dannenburg, Richard Verrall, Klaus Schmidt, Bjørn Sundt, Vsevolod Malinovskii, who translated this textbook into Russian, Qihe Tang, who translated it into Chinese in a joint project with Taizhong Hu and Shixue Cheng, Roger Laeven, Ruud Koning, Pascal Schoenmaekers, Julien Tomas, Katrien Antonio as well as Steven Vanduffel for their comments, and Vincent Goulet for helping us with some R-problems. Jan Dhaene and Marc Goovaerts acknowledge the support of the Fortis Chair on Financial and Actuarial Risk Management.

World wide web support

The authors would like to keep in touch with the users of this text. On the internet page <http://www1.fee.uva.nl/ke/act/people/kaas/ModernART.htm> we maintain a list of all typos that have been found so far, and indicate how teachers may obtain solutions to the exercises as well as the slides used at the University of Amsterdam for courses based on this book.

To save users a lot of typing, and typos, this site also provides the R commands used for the examples in the book.

Amsterdam,
Leuven,
Louvain-la-Neuve,
March 22, 2009

Rob Kaas
Marc Goovaerts
Jan Dhaene
Michel Denuit

Modern Actuarial Risk Theory

Using R

Kaas, R.; Goovaerts, M.; Dhaene, J.; Denuit, M.

2008, XVIII, 382 p., Softcover

ISBN: 978-3-642-03407-7