

Model Risk Management

This book provides the first systematic treatment of model risk, outlining the tools needed to quantify model uncertainty, to study its effects, and, in particular, to determine the best upper and lower risk bounds for various risk aggregation functionals of interest.

Drawing on both numerical and analytical examples, this is a thorough reference work for actuaries, risk managers, and regulators. Supervisory authorities can use the methods discussed to challenge the models used by banks and insurers, and banks and insurers can use them to prioritize the activities on model development, identifying which ones require more attention than others.

In sum, it is essential reading for all those working in portfolio theory and the theory of financial and engineering risk, as well as for practitioners in these areas. It can also be used as a textbook for graduate courses on risk bounds and model uncertainty.

Ludger Rüschendorf is Professor of Mathematics at the University of Freiburg. He is author of more than 200 research papers and a number of textbooks, in a variety of subjects in probability, statistics, and analysis of algorithms, as well as in risk analysis and mathematical finance. A main topic in his research is the modeling and analysis of dependence structures.

Steven Vanduffel is Professor in Risk Management at the Solvay Business School at Vrije Universiteit Brussel. He has authored papers for leading journals including *Journal of Risk and Insurance*, *Finance and Stochastics*, *Mathematical Finance*, and *Journal of Econometrics*. He has won prizes including the Robert I. Mehr Award (2022), the Robert C. Witt Award (2018), and the Redington Prize (2015).

Carole Bernard is Professor in Finance at Grenoble Ecole de Management and Vrije Universiteit Brussel. She has published articles in leading journals in finance, insurance, operations research, and risk management, including *Management Science*, *Journal of Risk and Insurance*, *Journal of Banking and Finance*, and *Mathematical Finance*.

Model Risk Management

Risk Bounds under Uncertainty

LUDGER RÜSCHENDORF

University of Freiburg

STEVEN VANDUFFEL

Vrije Universiteit Brussel

CAROLE BERNARD

Grenoble Ecole de Management



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press & Assessment
978-1-009-36716-5 — Model Risk Management
Ludger Rüschendorf, Steven Vanduffel, Carole Bernard
Frontmatter
[More Information](#)



Shaftesbury Road, Cambridge CB2 8EA, United Kingdom
One Liberty Plaza, 20th Floor, New York, NY 10006, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India
103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467
Cambridge University Press is part of the University of Cambridge.

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence

www.cambridge.org
Information on this title: www.cambridge.org/9781009367165
DOI: 10.1017/9781009367189

© Ludger Rüschendorf, Steven Vanduffel and Carole Bernard 2024

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press & Assessment

First published 2024

A catalogue record for this publication is available from the British Library

Library of Congress Cataloging-in-Publication Data

Names: Rüschendorf, Ludger, 1948- author. | Vanduffel, Steven, author. | Bernard, Carole, author.

Title: Model risk management : risk bounds under uncertainty / Ludger Rüschendorf, Steven Vanduffel, Carole Bernard.

Description: New York, NY : Cambridge University Press, [2023] | Includes bibliographical references and index.

Identifiers: LCCN 2023032170 (print) | LCCN 2023032171 (ebook) | ISBN 9781009367165 (hardback) | ISBN 9781009367189 (epub)

Subjects: LCSH: Risk management. | Management. | Computer simulation.

Classification: LCC HD61 .R867 2023 (print) | LCC HD61 (ebook) |

DDC 658.15/52–dc23/eng/20230810

LC record available at <https://lcn.loc.gov/2023032170>

LC ebook record available at <https://lcn.loc.gov/2023032171>

ISBN 978-1-009-36716-5 Hardback

Cambridge University Press & Assessment has no responsibility for the persistence or accuracy of URLs for external or third-party Internet Web sites referred to in this publication and does not guarantee that any content on such Web sites is, or will remain, accurate or appropriate.

Contents

Preface	<i>page</i> ix
Introduction	xi
Part I Risk Bounds for Portfolios Based on Marginal Information	1
1 Risk Bounds with Known Marginal Distributions	5
1.1 Some Basic Notions and Results of Risk Analysis	5
1.2 Standard Bounds, VaR Bounds, and Worst Case Distributions	9
1.3 Worst Case Risk Vectors: The Conditional Moment Method and the Mixing Method	15
1.4 Mixability and Convex Minima of Portfolios	18
2 The Rearrangement Algorithm	23
2.1 Formulation of Dependence Problems as Rearrangement Problems	24
2.2 The Rearrangement Algorithm	29
2.3 The Block Rearrangement Algorithm	37
3 Dual Bounds	47
3.1 Hoeffding–Fréchet Functionals	47
3.2 Dual Representation of Hoeffding–Fréchet Functionals	48
3.3 Sharpness of the Dual Bounds for the Tail Risk	55
4 Asymptotic Equivalence Results	59
4.1 Introduction and Motivation	59
4.2 Asymptotics for Conservative Capital Charges: The Homogeneous Finite Mean Case	61
4.3 Asymptotics for VaR in the Infinite Mean Case	64
4.4 Asymptotic Equivalence in the Inhomogeneous Case	67

vi	Contents	
	Part II Additional Dependence Constraints	71
	5 Improved Standard Bounds	75
	5.1 Standard and Improved Standard Bounds for Monotonic Functions	75
	5.2 Computation of Improved Standard Bounds	80
	5.3 VaR Bounds with Two-Sided Dependence Information	84
	6 VaR Bounds with Variance Constraints	92
	6.1 VaR Bounds with an Upper Bound on the Variance	93
	6.2 Algorithm to Approximate VaR Bounds	99
	6.3 Generalization to Constraints on Higher Moments on Default Probabilities	106
	6.4 Discussion of Consequences for Credit Risk Modeling	115
	7 Distributions Specified on a Subset	122
	7.1 Improved Hoeffding–Fréchet Bounds	122
	7.2 Distribution Functions or Copulas with Given Bounds on a Subset	125
	7.3 Distributions Known on a Subset	130
	Part III Additional Information on the Structure	143
	8 Additional Information on Functionals of the Risk Vector	146
	8.1 Higher-Dimensional Marginals	146
	8.2 Functionals of the Risk Vector	151
	9 Partially Specified Risk Factor Models	161
	9.1 Partially Specified Factor Models	162
	9.2 Risk Bounds	164
	9.3 Bounds for Convex Risk Measures	170
	9.4 Relation between VaR and TVaR Bounds	175
	9.5 Application to Credit Risk Portfolios	184
	10 Models with a Specified Subgroup Structure	186
	10.1 Subgroup Models with Positive and Negative Internal Dependence	187
	10.2 Models with Independent Subgroups	194
	10.3 VaR Bounds in Subgroup Models with Dependence Information within and between the Subgroups	202
	Part IV Risk Bounds under Moment Information	219
	11 Bounds on VaR, TVaR, and RVaR under Moment Information	222
	11.1 VaR Bounds: Bounded Domain	222
	11.2 VaR Bounds: Unbounded Domain	236
	11.3 TVaR and RVaR Bounds	237
	11.4 Proofs	239

12	Bounds for Distortion Risk Measures under Moment Information	243
	12.1 Bounds for Distortion Risk Measures: Bounded Domain	243
	12.2 Bounds for Distortion Risk Measures: Unbounded Domain	252
	12.3 Proofs	254
13	Bounds for VaR, TVaR, and RVaR under Unimodality Constraints	266
	13.1 VaR Bounds	267
	13.2 RVaR and TVaR Bounds	274
	13.3 Application to Model Risk Assessment of a Credit Risk Portfolio	278
	13.4 Proofs	279
14	Moment Bounds in Neighborhood Models	283
	14.1 Problem Formulation	283
	14.2 Bounds on Distortion Risk Measures	285
	14.3 Bounds on Range VaR	291
	14.4 Application to an Insurance Portfolio	294
	14.5 Proofs	299
	14.6 Appendix: Isotonic Projection	305
	14.7 Appendix: Models for Insurance Portfolio Losses	307
	References	309
	Index	320

Preface

What is this book about? This book deals with the ubiquitous problem of model uncertainty, which is one of the most relevant topics in quantitative risk management. It identifies a series of relevant instances of model uncertainty, such as departures from assumed independence, incomplete dependence information, factor models that are only partially specified, or portfolio information that is only available on an aggregate level (e.g., mean and variance of the portfolio loss). It provides the necessary tools to quantify this model uncertainty and in particular to determine the best upper and lower risk bounds for various risk aggregation functionals of interest.

Why did we write this book? While there are good textbooks available dealing with basic methods, concepts, and models in quantitative risk management, this book is the first systematic treatment of the topic of model risk. Not only does it elaborate on the necessary theoretical results for the determination of risk bounds, but it also provides numerical procedures for the effective evaluation of these bounds.

For whom is this book? The book is a relevant reference text on the topic of model risk assessment for actuaries, risk managers, and regulators. It also serves as a textbook for graduate courses on the topic of risk bounds and model uncertainty within the general subject area of risk management in quantitative finance and insurance. The methodology draws on diverse quantitative disciplines ranging from mathematical finance, probability, and statistics to actuarial mathematics.

Acknowledgments: The book is a culmination of a long series of research papers on the topic of model risk. We are indebted to numerous colleagues and former Ph.D. students who either coauthored some of these research papers or who have helped us in our understanding of model risk and the mathematics underlying it. In particular, we would like to mention the substantial contributions of Paul Embrechts and Giovanni Puccetti, who in fact initiated this type of research, as well as the early basic contributions of Etienne de Vylder. Our cooperation includes joint work with Jonathan Ansari, Valeria Bignozzi, Kris Boudt, Andrew Chernih, Ka Chun Cheung, Dries Cornilly, Michel Denuit, Corrado De Vecchi, Paul Embrechts, Luc Henrard, Edgars Jakobsons, Rodrigue Kazzi, Thibaut Lux, Dennis Manko, Silvana Pesenti, Giovanni Puccetti, Daniel Small, Jan Tuitman, Ruodu Wang, Julian Witting, and Jing Yao. We are also grateful to the many anonymous referees whose comments helped us to improve the quality of these research papers and thus of the book. We thank

David Tranah and the team at Cambridge University Press for all their help in the production of this book. Special thanks go to Monika Hattenbach for her dedicated work in preparing the manuscript. Finally, we are indebted to our partners, Gabi, Sonja, and Christophe, for their patience and for forgiving our absentmindedness.

Introduction

This book deals with the problem of model uncertainty, which is a most relevant topic in quantitative risk management. It identifies a series of relevant instances of model uncertainty, such as departures from assumed independence, incomplete dependence information, factor models that are only partially specified, or portfolio information that is only available on an aggregate level (e.g., mean and variance of the portfolio loss). It provides necessary tools to quantify this model uncertainty and in particular to determine the best upper and lower risk bounds for various risk aggregation functionals of interest.

Making sound decisions under uncertainty generally requires quantitative analysis and the use of models. However, a “perfect” model does not exist since some divergence between the model and the reality it attempts to describe cannot be avoided. In a broad sense, model risk is about the extent to which the quality of model-based decisions is sensitive to underlying model deviations and data issues. Quantifying model risk is a key problem in nearly all applied disciplines, including epidemiology, engineering, finance, and insurance. For instance, sums of variables (portfolios) are at the core of the insurance business, as the insurer counts on diversification effects to control the risk of the entire portfolio. For an insurance portfolio, the assumption of independence between the policies is sometimes realistic, in which case the insurer can, for instance, resort to the central limit theorem or to Monte Carlo methods to quantify the maximum loss in a given period of time at a certain probability level (i.e., the VaR). In the majority of cases, however, the individual risks are influenced by one or more common factors, such as geography or economic environment, and it is difficult to specify the joint distribution. Another example concerns the establishment of the capital buffers banks need to put aside to absorb unforeseen losses for a portfolio of risky loans. Doing so requires accurate estimates of the likelihood that various obligors default together, which is very difficult due to a scarcity of data. In this book we focus on model risk in a financial and insurance context.

Model risk may have a real impact on society, such as damage to an institution’s reputation or even systemic risk implications. For example, Long-Term Capital Management (LTCM) was a hedge fund that used quantitative models based on normality assumptions but neglected the importance of stress-testing. In 1998 it lost 4.5 billion and required the financial intervention by the Federal Reserve (Lowenstein, 2008). In fact, one of the drivers of the worldwide 2008–2009 financial crisis was the almost blind reliance on certain model assumptions (Salmon, 2009).

While the problem of model uncertainty in statistics has a long history – being developed there under the notion of robust statistics – at the time of writing this book, measuring and managing model risk in a financial context is a relatively new activity. Model risk has, however, already taken a prominent position in the agenda of regulators and supervisors. For instance, in February 2017, the European Central Bank published a guide to the targeted review of internal models (TRIM) in which it is declared that every institution “should have a model risk management framework in place that allows it to identify, understand, and manage its model risk” (ECB, 2017). However, in a status update on TRIM published in June 2018, more than a quarter of the companies supervised had no model risk management framework in place. In their discussion paper on the review of specific items in the Solvency II Regulation, the Actuarial Association of Europe is insisting on focusing more on model risk assessment (AAE, 2017). In the UK, the prudential regulation authority of the Bank of England published some notes on stress test model management principles in which they insist on the necessity of understanding and accounting for the assessment of model uncertainties (PRA, 2018).

In addition to the inherent necessity for insurers and banks to continuously monitor and challenge the models they use in their operations (pricing, product design, risk management), the need for model risk management is further strengthened by the fact that the insurance and banking market is innovating at a very fast pace. New products emerge with unique characteristics, such as driverless car liability insurance, artificial intelligence/robotics liability insurance, and nanotechnology liability insurance. The advent of these new products calls for new pricing and reserving models. In response to this, and in particular driven by the lessons learnt from the 2008–2009 financial crisis, regulators and rating agencies are thus increasing pressure on the financial industry to measure the risk they run and to demonstrate that enough capital is available for absorbing adverse shocks. In order to mitigate this model risk efficiently without restraining too much the arrival of new products and models, the model risk management function (MRM) is emerging in the financial industry. The well-established Professional Risk Managers’ International Association (PRMIA) follows this evolution in that it is giving more and more attention to the MRM function in its seminars and other activities.

In the remainder of this introductory chapter, we explain some basic notions of risk assessment and risk models. Furthermore, we describe some basic tools of how to measure aggregate risk, and we conclude by providing an overview of the various subjects dealt with in this book.

A. Risk Assessment and Risk Models

The risk assessment of a multidimensional portfolio (X_1, X_2, \dots, X_n) is a core issue in risk management of financial institutions. In particular, this problem appears naturally for an insurance company. Any insurer is exposed to different risk factors (e.g., non-life risk, longevity risk, credit risk, market risk, operational risk), has different business lines, or has an exposure to several portfolios of clients. In this regard, one typically attempts to measure the risk of a sum, $S = \sum_{i=1}^n X_i$ or of another aggregation function of

the risk vector X , in which the individual risks X_i depict losses (claims of the different customers, changes in the different market risk factors, . . .) using a risk measure such as the variance, the VaR, or the tail value-at-risk¹ (TVaR). It is clear that solving this problem is mainly a numerical issue once the joint distribution of (X_1, X_2, \dots, X_n) is completely specified. Estimating a multivariate distribution or testing its adequacy is in general a difficult task. In many cases, the actuary will be able to use mathematical and statistical techniques to describe the marginal risks X_i , but the dependence among the risks is not available, or only partially known. In other words, the assessment of portfolio risk based on specific models is prone to model misspecification (model risk).

From a mathematical point of view, it is often convenient to assume that the random variables X_i are mutually independent, because powerful and accurate computation methods such as Panjer's recursion and the technique of convolution can then be applied. In this case, one can also take advantage of the central limit theorem, which states that the sum of risks, S , is approximately normally distributed if the number of risks is sufficiently high. In fact, the mere existence of insurance is based on the assumption of mutual independence between the insured risks, and sometimes this complies, approximately, with reality. In the majority of cases, however, the different risks will be interrelated to a certain extent. For example, a sum S of dependent risks occurs when considering the aggregate claims amount of a non-life insurance portfolio because the insured risks are subject to some common factors such as geography, climate, or economic environment. The cumulative distribution function of S can no longer be easily specified.

Standard approaches to estimating a multivariate distribution of a portfolio (X_1, X_2, \dots, X_n) consist in using a multivariate Gaussian distribution or a multivariate Student t distribution, but there is ample evidence that these models are not always adequate. More precisely, while the multivariate Gaussian distribution can be suitable as a fit to a dataset "on the whole," it is usually a poor choice if one wants to use it to obtain accurate estimates of the probability of simultaneous extreme ("tail") events or, equivalently, if one wants to estimate the VaR of the aggregate portfolio $S = \sum_{i=1}^n X_i$ at a given high confidence interval; see McNeil et al. (2015). The use of the multivariate Gaussian model is also based on the (wrong) intuition that correlations are enough to model dependence (Embrechts et al., 1999, 2002). This fallacy also underpins the variance-covariance standard approach that is used for capital aggregation regulatory frameworks such as Basel III and Solvency II, and which also appears in many risk management frameworks in the industry. Furthermore, in practice, there are not enough observations that can be considered as tail events. In fact, there is always a level beyond which there is no observation. Therefore, if one makes a choice for modeling tail dependence, it has to be somewhat arbitrary.

In the literature, one can find flexible multivariate models that allow a much better fit to the data, for example using pair copula constructions and vines (see, e.g., Aas et al., 2009 or Czado, 2010 for an overview). While these models have theoretical and intuitive appeal, their successful use in practice requires a dataset that is sufficiently

¹ In the literature it is also called the expected shortfall and the conditional value-at-risk, among others.

rich. However, no model is perfect, and while such developments are clearly needed for an accurate assessment of portfolio risk, they are only useful to regulators and risk managers if they are able to significantly reduce the model risk that is inherent in risk assessments.

B. Measuring Aggregate Risk

Insurance companies essentially exchange premiums against (future) random claims. Consider a portfolio containing n policies, and let X_i , $i = 1, 2, \dots, n$, denote the loss, defined as the random claim net of the premium, of the i th policy. In order to protect policyholders and other debtholders against insolvency, the regulator will require the portfolio loss $S = X_1 + X_2 + \dots + X_n$ to be “low enough” as compared to the available resources, say a capital requirement K , which means that the available capital K has to be such that $S - K$ is a “safe bet” for the debtholders, i.e., one is “reasonably sure” that the event “ $S > K$ ” is of minor importance (Tsanakas and Desli, 2005; Dhaene et al., 2012).

It is clear that measuring the riskiness of $S = X_1 + X_2 + \dots + X_n$ is of key importance for setting capital requirements. However, there are several other reasons for studying the properties of the aggregate loss S . Indeed, an important task of an enterprise risk management (ERM) framework concerns capital (risk) allocation, i.e., the allocation of total capital held by the insurer across its various constituents (subgroups), such as business lines, risk types, and geographical areas, among others. Indeed, doing so makes it possible to redistribute the cost of holding capital across the various constituents so that it can be transferred back to the depositors or policyholders in the form of charges (premiums). Risk allocation also makes it possible to assess the performance of the different business lines by determining the return on allocated capital for each line. Finally, the exercise of risk aggregation and allocation may help to identify areas of risk consumption within a given organization and thus to support the decision-making concerning business expansions, reductions, or even eliminations; see Panjer (2001) and Tsanakas (2009).

When measuring the aggregate risk S , it is also important to consider the context at hand. In particular, different stakeholders may have different perceptions of riskiness. For example, depositors and policyholders mainly care only about the probability that the company will meet its obligations. Regulators primarily share the interest of depositors and policyholders and establish rules to determine the required capital to be held by the company. However, they also care about the magnitude of the loss given that it exceeds the capital held, as this is the amount that needs to be funded by society when a bailout is needed. Formally, they care about the *shortfall* of the portfolio loss S with solvency capital requirement $\varrho(S)$; that is,

$$(S - \varrho(S))_+ := \max(0, S - \varrho(S)). \quad (0.1)$$

The shortfall is thus part of the total loss that cannot be covered by the insurer. It is also referred to as the *loss to society* or the *policyholders' deficit*. In view of their limited liability, shareholders do not really have to care about the residual risk but

rather focus on the properties of the variable $S - (S - \varrho(S))_+ = \min(S, \varrho(S))$. In summary, various stakeholders may have different perceptions and sensitivities with respect to the meaning of the risk they run, and they may employ different paradigms for defining and measuring it.

As for measuring the risk, the two most influential risk measures are the *value-at-risk* (VaR) and the *tail value-at-risk* (TVaR).² For a given confidence level α , they are denoted by VaR_α and TVaR_α , respectively, and are defined as

$$\text{VaR}_\alpha(S) = \min\{x \mid P(S \leq x) \geq \alpha\}, \quad 0 < \alpha < 1 \quad (0.2)$$

and

$$\text{TVaR}_\alpha(S) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_q(S) \, dq, \quad 0 < \alpha < 1. \quad (0.3)$$

So, VaR_α is the minimum loss one observes with probability $1 - \alpha$, whereas TVaR_α is the average of all upper VaRs. A general class of risk measures that includes VaR and TVaR is the distortion risk measures, which are essentially weighted averages of VaRs.

Another interesting risk measure that is not in this class is the “tail risk,” which is the probability that the aggregate risk exceeds some level K , i.e., $P(S \geq K)$.

C. Regulatory Frameworks

The report of the Basel Committee on Banking Supervision (2010) describes the modeling methods used by financial firms and regulators in various countries to aggregate risk. It also aims to identify the conditions under which these aggregation techniques perform as anticipated in the model and to suggest potential improvements. The report expresses doubts about the reliability of internal risk aggregation results that incorporate diversification benefits: “Model results should be reviewed carefully and treated with caution, to determine whether claimed diversification benefits are reliable and robust.” In this subsection, we very briefly summarize how the main “regulatory frameworks” deal with risk aggregation and diversification.

Basel III Regulation for Banks

One calculates a bank’s overall minimum capital requirement as the sum of capital requirements for the credit risk, operational risk, and market risk, without recognizing possible diversification benefits between the three risk types. As diversification is not acknowledged, the practice of adding up capital is seen as a conservative approach. However, this also depends on the risk measure that is used for establishing capital. If VaR is used, it might indeed occur that the VaR of the aggregate risk is higher than the sums of the individual VaRs. By contrast, if a coherent risk measure ϱ is used, then it always holds that $\varrho(X_1 + \cdots + X_n) \leq \varrho(X_1) + \cdots + \varrho(X_n)$.

² VaR appears to be the most popular risk measure in practice, among both regulators and risk managers; see, for example, Jorion (2006).

As for market risk, banks have the choice between two methods. They may benefit from diversification if they use an internal model approach (IMA). With the standardized measurement method (SMM), the minimum capital requirement for market risk is the sum of the capital charges calculated for each individual risk type (interest rate risk, equity risk, foreign exchange risk, commodities risk, and price risk in options).

Solvency II

The Solvency Capital Requirement (SCR) under Solvency II is defined as the VaR at 99.5 % over a time horizon of one year. When aggregating risks, insurers may benefit from diversification: They have the option to use an internal model (without any particular method prescribed) or a standard formula. The standard formula aggregates risks using a correlation matrix (variance-covariance approach) to take dependencies into account.

Comparison and Comments on Regulatory Frameworks

Generally, regulatory rules incorporate diversification by taking into account some correlation effect to reduce the total capital (at least in some subcategories). Overall, we observe that regulators all implicitly assume that the sum of the risk numbers is the worst possible situation. “No diversification benefits” is then synonymous to “adding up risk numbers (VaRs).”

The easiest method to aggregate risks is the variance-covariance approach (which is explicitly mentioned in Solvency II above and is also used by the Australian regulator OSFI (2014)). It builds on the assumption that the correlation matrix is enough to describe the dependence and that it is possible to aggregate risks based on this correlation matrix. Its strength is in its being a simple approach, but it is only a correct approach for elliptical multivariate distributions, such as the Gaussian multivariate distribution. Furthermore, correlation is a linear measure of dependence and does not capture tail dependence adequately. Using such a method to aggregate risk may perhaps be fine for having some idea on the distribution “globally” but fails when it comes to assess the risk in the tail; also note that capital requirements are typically based on tail risk measures, such as VaR at 99.5 %, which essentially reflects the outcome of a 1-in-200-year scenario.

Instead of using the variance-covariance approach, one may use copulas to aggregate the individual risks. This approach is rather flexible and allows one to separate the risk assessment of the marginal distribution of individual risks and their dependence. By specifying a given copula to model some dependence, it is then possible to recognize tail dependence among some risks. However, determining the “right” copula to use is a very hard task that is prone to significant model risk, as we will see later in this book. Statistical methods to fit a multivariate model involve large numbers of parameters and copula families. In addition, understanding the outputs of the model will then require good expertise in the copula approach to understand the impact of each assumption made on the dependence. This is a concern and a challenge among institutions.

Another way to capture tail risks and tail dependence is to understand “where the dependence comes from” and to model the real risk drivers of the dependence between individual risks of the portfolio and understand their interactions. The report of the Basel Committee on Banking Supervision (2010) suggests using “scenario-based aggregation.” Aggregation through scenarios boils down to determining the state of the firm under specific events and summing profits and losses for the various positions under the specific event. In other words, it means that one needs to incorporate information that one knows about the dependence in some specific states.

As observed in the report of the Basel Committee on Banking Supervision (2010), the results of scenario-based aggregation are easier to interpret, with more meaningful economic and financial implications, but use of this method again requires deep expertise to identify risk drivers, derive meaningful sets of scenarios with relevant statistical properties, and then use them to obtain a full loss distribution. A lot of the inputs in these kinds of model come from experts’ judgements. Overall, there is no clear unique solution to the problem of risk aggregation. Each method has its pros and cons and may be helpful in some situations and useless in others.

D. Overview of the Contents

As described in the introduction, the main topic of the book is to establish risk bounds for aggregate risk functionals under various forms of model uncertainty for the underlying model. The model uncertainty is described by different types of (partial) knowledge on the underlying model: knowledge of the marginals, various forms of partial knowledge of the dependence structure, or knowledge of the structure of the model, as given for example by (partially specified) factor models or by (partially specified) subgroups of the model. An important part of the book addresses the possible improvement of risk bounds due to information on various functionals of the risk vector, in particular given by moment bounds for the aggregate portfolio. Combination with a neighborhood model assumption yields further improvement of the risk bounds.

The book consists of four main parts. The basic assumption in Parts I–III in order to derive risk bounds is that the marginal distributions of the risk vector $X = (X_1, \dots, X_n)$ are known, say $X_i \sim F_i$, $1 \leq i \leq n$. This assumption can be realistically made in many applications since it is much easier to model and test simple hypotheses compared to the task of modeling and testing the joint distribution of X . Throughout the book, we present focus on risk bounds for the aggregated portfolio $S = \sum_{i=1}^{i=n}$. Nevertheless in Parts I–III the assumed knowledge of the marginals makes it possible to also derive various risk bounds for other risk functionals, for example, for the maximal risk. In Part IV, we do not assume knowledge of the marginals but only use information that pertains to a portfolio sum S .

In **Part I**, an introduction to the problem of risk bounds with information on the marginal distributions is given. Also, a basic algorithm to determine these bounds – the rearrangement algorithm – and a basic solution method – the dual bounds – are introduced.

In Chapter 1, we introduce some basic notions of risk measures (like VaR, TVaR, and convex risk measures) and describe some corresponding worst case VaR and TVaR portfolios. We also give a rearrangement formulation of various forms of determining worst or best case risk bounds. We describe the connection of upper and lower risk bounds to convex ordering properties, discuss comonotonicity and countermonotonicity (antimonotonicity), and give some basic results to obtain worst case VaR portfolios resp. portfolios with maximal tail risk.

The “standard bounds” for tail risks go back to classical sources, like Sklar (1973) or Moynihan et al. (1978). For $n = 2$ they are shown in Makarov (1981) and Rüschendorf (1982) to be sharp, but for $n > 2$ they typically only deliver rough (i.e., not necessarily sharp) bounds.

The conditional moment method gives an upper bound on the tail risk of the portfolio sum in terms of conditional moments of the marginals. This type of upper risk bound produces sharp bounds under a mixing condition on the upper tail. In the final subsection of this part, we discuss in more detail the notion of mixability, describe some basic results due to Wang and Wang (2011, 2016) on mixability, and explain its role in the determination of convex minima of portfolio sums and similarly for best and worst case portfolios.

Chapter 2 is devoted to the motivation and introduction of the rearrangement algorithm (RA) as introduced in Puccetti and Rüschendorf (2012a), which is a fundamental tool to determine sharp upper and lower bounds for the tail risks resp. for the VaR. It is basically motivated by the formulation of the problem of determining risk bounds as a rearrangement problem and in a second step by a further reduction to an assignment problem. Also, a variation of the RA, the block rearrangement algorithm (BRA), is introduced, which improves some aspects of the RA. Several interesting applications of these algorithms are indicated.

Chapter 3 gives an introduction to Hoeffding–Fréchet functionals, which describe the largest resp. the smallest value of a risk functional over all possible dependence structures, when fixing the marginals. The main result is a duality theorem for these functionals that is a far reaching extension of the dual representation in the Monge–Kantorovich mass transportation problem. A reduction of the admissible dual functions to a simple class of admissible dual functions introduced in Embrechts and Puccetti (2006b) leads to good dual bounds for the tail risk. Sharpness of these bounds is established under a mixing condition in the homogeneous case in Puccetti and Rüschendorf (2013).

An easy-to-calculate upper bound for the worst case VaR of an aggregated portfolio is given by the TVaR of a corresponding comonotonic sum. In Chapter 4 we derive asymptotic sharpness of this upper TVaR bound under an asymptotic mixing condition. A version of this result also holds in the infinite mean case and in the inhomogeneous case.

In contrast to Part I, we assume in **Part II** not only that we know the marginal distributions of the risk vector $X = (X_1, \dots, X_n)$, say $X_i \sim F_i$, $1 \leq i \leq n$, but that we have additional information on the dependence among these n variables.

In Chapter 5 we present the method of improved standard bounds. We review bounds on the distribution function that improve upon the Hoeffding–Fréchet bounds

by including, in addition to the marginal information, some positive or negative dependence information on the distribution functions of the form $F \leq G$ or $F \geq G$. These improved Hoeffding–Fréchet bounds lead to improved upper and lower VaR bounds. To do so, we build on results of Williamson and Downs (1990), Denuit et al. (1999), Embrechts et al. (2003), Rüschendorf (2005), Embrechts and Puccetti (2006a,b), and Puccetti et al. (2016), among others. We present some examples in which the information significantly reduces the standard bounds. Further examples of this type are discussed in Rüschendorf (2017a,b). The last section of this chapter presents an extension of the method of improved standard bounds by Lux and Rüschendorf (2018) to include two-sided dependence information of the form $G \leq F \leq H$ or related bounds for the copula C or for the survival functions \bar{F} .

Chapter 6 deals with bounds on VaR under the additional constraint that a bound on the variance of the sum is known (as in Bernard et al., 2017c) or that information on higher moments on the sum (as in Bernard et al., 2017a) is available. As compared to assuming only knowledge of the marginal means or variances (so-called moment bounds), knowledge of the marginal distributions may make it possible to improve the VaR bounds. However, when supplementing the information on the marginal distributions with information on the portfolio variance, it turns out that, when the variance constraint is “low” enough, the analytical VaR-bounds coincide with moment bounds (studied in full detail in Part IV). We then propose a numerical method based on the RA-algorithm to approximate sharp VaR bounds given the marginal constraints and additionally bounds on the variance of the sum. A corresponding rearrangement algorithm is no longer available if only moment bounds are assumed. We show that for large portfolios, our analytical bounds are nearly sharp. However, in the context of smaller portfolios (or when the portfolio depicts significant concentration), one cannot expect the bounds to be sharp. The RA-algorithm can accommodate this situation and makes it possible to approximate sharp bounds. As a by-product, this algorithm also sheds light on the composition of extreme portfolios.

Chapter 7 is devoted to studying the case when information on the joint distribution of the risk vector $X = (X_1, \dots, X_n)$ is partially available, by knowing the conditional distribution of X on a subset S in R^n the marginal distributions as well as the probability of the subset, or alternatively, knowing bounds for the distribution function F_X resp. the copula C_X on the corresponding subsets. For example, it may be the case that C_X of X is known on a central domain due to availability of sufficient statistical data, allowing for a precise model within this domain. Alternatively, some positive or negative dependence information may be available in the upper or lower tail area.

We derive VaR bounds in several situations of interest in which the conditional distribution of X is known on a central domain. The corresponding improvements of the Hoeffding–Fréchet bounds are derived in Rachev and Rüschendorf (1994), Nelsen et al. (2004), Tankov (2011) and Bernard and Vanduffel (2014) in the two-dimensional case and by Puccetti et al. (2016) and Lux and Papapantoleon (2019) in the n -dimensional case.

We solve several cases exactly or numerically when we know the distribution function on a subset or we know the conditional distribution on a subset (Bernard and

Vanduffel, 2015; Puccetti et al., 2016). While knowing the distribution function F on an open domain A implies knowledge of P^X on A , the converse direction is not valid, even for rectangle domains A . So both types of information are typically different. Based on a numerical approach using the RA-algorithm, the value of this type of information for improving VaR bounds for the sum is determined

In **Part III** the risk bounds and range of dependence uncertainty obtained in the unconstrained case is essentially improved by including relevant structural information on the underlying class of models. This additional structural information may lead to a positive or a negative dependence restriction and thus induce improvements of the upper resp. the lower risk bounds, as shown in Part II.

We deal in this part with several types of structural information, investigate their impact, and show in several applications their potential in reducing dependence uncertainty. In some applications it is possible to combine these structural assumptions with dependence assumptions, as in Part II, leading to particular useful reduction results.

In Chapter 8 we consider first in Section 8.1 the case that besides the one-dimensional marginals, some higher dimensional marginals are also known. This assumption allows us to derive improvements of the classical Fréchet bounds by using a duality theorem. For general higher-dimensional marginal systems, the “reduction method” allows us to derive good upper and lower bounds by an associated reduced risk model with simple marginals. This reduced problem can be solved by the RA. If the higher-dimensional marginals exhibit strong positive dependence, this leads to improvements of the lower bounds in comparison to the unconstrained case; if they exhibit negative dependence, this leads to improvements of the upper bounds.

In Section 8.2 we consider a general case of additional constraints given by the distribution or the expectations of a class of functionals. This includes in principle the constraints due to higher-dimensional marginals (infinite set of restrictions) but also the case of variance constraints as in Part II. We give several improved lower and upper dual bounds for such classes of constraints. Using martingale constraints (infinite set of restrictions), this method can also be used to derive improved price bounds for options.

Chapter 9 gives a detailed discussion on partially specified factor models (PSFM). This model assumption is based on an underlying factor model $X_i = f_i(Z, \varepsilon_i)$ with a systematic risk factor Z and individual risks $\varepsilon_1, \dots, \varepsilon_n$. In comparison to the usual assumption of completely specified factor models, which is in general a hard to verify assumption, in PSFMs only the joint distributions of (ε_i, Z) are specified, which is a much simpler to verify model assumption. We show that the assumption of PSFM may lead to strongly reduced risk bounds and that it can be combined in a particular effective and flexible way with other dependence assumptions, like variance bounds for the aggregated portfolio.

Chapter 10 deals with a systematic investigation of the assumption that the risk vector X is split into k subgroups. For a comparison vector Y , conditions are given for the comparison of the subgroup sums of X and Y and further for the comparison of the copulas of the vectors of subgroup sums to imply a relevant comparison theorem between the aggregated portfolios $\sum_{i=1}^n X_i$ and $\sum_{i=1}^n Y_i$. Also, this criterion allows flexible and effective applications and can be combined in a useful

way with further constraints, for example, with the assumption of PSFMs within the subgroups.

In **Part IV** risk bounds are studied under the assumption that moment bounds of the portfolio sums are given. It is shown in Chapter 6 that adding a variance constraint on the portfolio loss $S = \sum X_i$ in addition to knowledge of the marginal distribution of the portfolio components X_i may lead to an improved VaR upper bound that corresponds to the classic Cantelli moment bound (involving mean and variance). This observation motivates us to study risk bounds for a portfolio sum $S = \sum X_i$ when the mean of S (but *not* the marginal distributions of its components) is given as well as some of its higher order moments. This approach to assessing model risk is, for instance, useful when loss statistics are only available at an aggregate level.

In Chapter 11 we study moment bounds for the risk measures VaR, TVaR, and the range value-at-risk (RVaR), which can be seen as a generalization of VaR and TVaR. Most of our results assume that only the first two moments are known (but no other higher-order moment). For the important case of VaR, however, we propose under an additional domain restriction a method that can in principle deal with any number of known moments. As a main tool for deriving the bounds, we use convex ordering and its generalization, s -convex order.

In Chapter 12 we provide a detailed study on moment bounds for distortion risk measures. In Section 12.1 we build on results in Cornilly et al. (2018) to derive moment bounds under an additional domain restriction. In Section 12.2 we dispense with domain restrictions and derive bounds using the tool of isotonic projections combined with the use of the Cauchy–Schwarz inequality.

In Chapters 13 and 14 we study the influence of structural information on the moment bounds for VaR, TVaR, and RVaR. First, in Chapter 13, we study bounds when, in addition to moment information, it is also known that the distribution is unimodal. Second, in Chapter 14 we study bounds when the distribution is assumed to stay in the neighborhood of a reference distribution. We use the Wasserstein distance as a metric to determine the neighborhood and use the tool of isotonic projections to obtain sharp bounds on VaR, TVaR, and RVaR.