

Wiley Series in Probability and Statistics

PRICING INSURANCE RISK

THEORY AND PRACTICE

STEPHEN J. MILDENHALL

JOHN A. MAJOR

$$\begin{aligned}\bar{P}(a) &= \int_0^a g(S(x)) dx \\ &= xg(S(x)) \Big|_0^a + \int_0^a xg'(S(x)) dF(x) \\ &= \int_0^a xg'(S(x)) dF(x) + ag(S(a)) \\ &= E[(X \wedge a)g'(S(X))].\end{aligned}$$

$$\begin{aligned}D^n \rho_X(X_i) &:= E[X_i \tilde{Z}_X] \\ D^j \rho_{X, \tilde{X}}(X_i) &:= E[X_i Z_{\tilde{X}}]\end{aligned}$$

WILEY

Pricing Insurance Risk

Pricing Insurance Risk

Theory and Practice

Stephen J. Mildenhall and John A. Major

WILEY

This edition first published 2022
© 2022 John Wiley & Sons, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by law. Advice on how to obtain permission to reuse material from this title is available at <http://www.wiley.com/go/permissions>.

The right of Stephen J. Mildenhall and John A. Major to be identified as the authors of this work has been asserted in accordance with law.

Registered Office

John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, USA

Editorial Office

111 River Street, Hoboken, NJ 07030, USA

For details of our global editorial offices, customer services, and more information about Wiley products visit us at www.wiley.com.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats.

Limit of Liability/Disclaimer of Warranty

While the publisher and authors have used their best efforts in preparing this work, they make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Library of Congress Cataloging-in-Publication Data

Names: Mildenhall, Stephen J., author. | Major, John A., author.

Title: Pricing insurance risk : theory and practice / Stephen J Mildenhall, John A Major.

Description: Hoboken, NJ : John Wiley and Sons, 2022. | Includes bibliographical references and index.

Identifiers: LCCN 2021062792 (print) | LCCN 2021062793 (ebook) | ISBN 9781119755678 (hardback) |

ISBN 9781119755692 (pdf) | ISBN 9781119756521 (epub) | ISBN 9781119756538 (ebook)

Subjects: LCSH: Risk (Insurance) | Risk management.

Classification: LCC HG8054.5 .M55 2022 (print) | LCC HG8054.5 (ebook) | DDC 368–dc23/eng/20220111

LC record available at <https://lcn.loc.gov/2021062792>

LC ebook record available at <https://lcn.loc.gov/2021062793>

Cover Design: Wiley

Cover Image: © zef art/Shutterstock

Set in 9.5/12.5pt STIXTwoText by Integra Software Services Pvt. Ltd, Pondicherry, India

10 9 8 7 6 5 4 3 2 1

Dedicated to our wives, Helen and Diane

Contents

Preface *xii*

1 Introduction 1

- 1.1 Our Subject and Why It Matters 1
- 1.2 Players, Roles, and Risk Measures 2
- 1.3 Book Contents and Structure 4
- 1.4 What's in It for the Practitioner? 7
- 1.5 Where to Start 9

2 The Insurance Market and Our Case Studies 13

- 2.1 The Insurance Market 13
- 2.2 Ins Co.: A One-Period Insurer 15
- 2.3 Model vs. Reality 16
- 2.4 Examples and Case Studies 17
- 2.5 Learning Objectives 25

Part I Risk 27

3 Risk and Risk Measures 29

- 3.1 Risk in Everyday Life 29
- 3.2 Defining Risk 30
- 3.3 Taxonomies of Risk 31
- 3.4 Representing Risk Outcomes 36
- 3.5 The Lee Diagram and Expected Losses 40
- 3.6 Risk Measures 54
- 3.7 Learning Objectives 60

4 Measuring Risk with Quantiles, VaR, and TVaR 63

- 4.1 Quantiles 63
- 4.2 Value at Risk 70
- 4.3 Tail VaR and Related Risk Measures 85

4.4	Differentiating Quantiles, VaR, and TVaR	102
4.5	Learning Objectives	102
5	Properties of Risk Measures and Advanced Topics	105
5.1	Probability Scenarios	105
5.2	Mathematical Properties of Risk Measures	110
5.3	Risk Preferences	124
5.4	The Representation Theorem for Coherent Risk Measures	130
5.5	Delbaen's Differentiation Theorem	137
5.6	Learning Objectives	141
5.A	Lloyd's Realistic Disaster Scenarios	142
5.B	Convergence Assumptions for Random Variables	143
6	Risk Measures in Practice	147
6.1	Selecting a Risk Measure Using the Characterization Method	147
6.2	Risk Measures and Risk Margins	148
6.3	Assessing Tail Risk in a Univariate Distribution	149
6.4	The <i>Intended Purpose</i> : Applications of Risk Measures	150
6.5	Compendium of Risk Measures	153
6.6	Learning Objectives	156
7	Guide to the Practice Chapters	157
Part II Portfolio Pricing 161		
8	Classical Portfolio Pricing Theory	163
8.1	Insurance Demand, Supply, and Contracts	163
8.2	Insurer Risk Capital	168
8.3	Accounting Valuation Standards	178
8.4	Actuarial Premium Calculation Principles and Classical Risk Theory	182
8.5	Investment Income in Pricing	186
8.6	Financial Valuation and Perfect Market Models	189
8.7	The Discounted Cash Flow Model	192
8.8	Insurance Option Pricing Models	200
8.9	Insurance Market Imperfections	210
8.10	Learning Objectives	213
8.A	Short- and Long-Duration Contracts	215
8.B	The Equivalence Principle	216
9	Classical Portfolio Pricing Practice	217
9.1	Stand-Alone Classical PCPs	217
9.2	Portfolio CCoC Pricing	223
9.3	Applications of Classical Risk Theory	224

9.4	Option Pricing Examples	227
9.5	Learning Objectives	231
10	Modern Portfolio Pricing Theory	233
10.1	Classical vs. Modern Pricing and Layer Pricing	233
10.2	Pricing with Varying Assets	235
10.3	Pricing by Layer and the Layer Premium Density	238
10.4	The Layer Premium Density as a Distortion Function	239
10.5	From Distortion Functions to the Insurance Market	245
10.6	Concave Distortion Functions	252
10.7	Spectral Risk Measures	255
10.8	Properties of an SRM and Its Associated Distortion Function	259
10.9	Six Representations of Spectral Risk Measures	261
10.10	Simulation Interpretation of Distortion Functions	263
10.11	Learning Objectives	264
10.A	Technical Details	265
11	Modern Portfolio Pricing Practice	271
11.1	Applying SRMs to Discrete Random Variables	271
11.2	Building-Block Distortions and SRMs	275
11.3	Parametric Families of Distortions	280
11.4	SRM Pricing	285
11.5	Selecting a Distortion	292
11.6	Fitting Distortions to Cat Bond Data	298
11.7	Resolving an Apparent Pricing Paradox	304
11.8	Learning Objectives	306
Part III	Price Allocation	307
12	Classical Price Allocation Theory	309
12.1	The Allocation of Portfolio Constant CoC Pricing	309
12.2	Allocation of Non-Additive Functionals	312
12.3	Loss Payments in Default	324
12.4	The Historical Development of Insurance Pricing Models	326
12.5	Learning Objectives	337
13	Classical Price Allocation Practice	339
13.1	Allocated CCoC Pricing	339
13.2	Allocation of Classical PCP Pricing	347
13.3	Learning Objectives	348

14	Modern Price Allocation Theory	349
14.1	The Natural Allocation of a Coherent Risk Measure	349
14.2	Computing the Natural Allocations	365
14.3	A Closer Look at Unit Funding	369
14.4	An Axiomatic Approach to Allocation	385
14.5	Axiomatic Characterizations of Allocations	392
14.6	Learning Objectives	394
15	Modern Price Allocation Practice	397
15.1	Applying the Natural Allocations to Discrete Random Variables	397
15.2	Unit Funding Analysis	404
15.3	Bodoff's Percentile Layer of Capital Method	413
15.4	Case Study Exhibits	421
15.5	Learning Objectives	439
Part IV Advanced Topics		441
16	Asset Risk	443
16.1	Background	443
16.2	Adding Asset Risk to Ins Co.	444
16.3	Learning Objectives	447
17	Reserves	449
17.1	Time Periods and Notation	449
17.2	Liability for Ultimate Losses	450
17.3	The Solvency II Risk Margin	461
17.4	Learning Objectives	468
18	Going Concern Franchise Value	469
18.1	Optimal Dividends	469
18.2	The Firm Life Annuity	472
18.3	Learning Objectives	476
19	Reinsurance Optimization	477
19.1	Background	477
19.2	Evaluating Ceded Reinsurance	477
19.3	Learning Objectives	481
20	Portfolio Optimization	483
20.1	Strategic Framework	483
20.2	Market Regulation	484
20.3	Dynamic Capital Allocation and Marginal Cost	485
20.4	Marginal Cost and Marginal Revenue	487

- 20.5 Performance Management and Regulatory Rigidities 488
- 20.6 Practical Implications 490
- 20.7 Learning Objectives 491

A Background Material 493

- A.1 Interest Rate, Discount Rate, and Discount Factor 493
- A.2 Actuarial vs. Accounting Sign Conventions 493
- A.3 Probability Theory 494
- A.4 Additional Mathematical Terminology 500

B Notation 503

References 507

Index 523

Preface

Pricing Insurance Risk is a topic of great concern to actuaries, especially property-casualty actuaries, our primary audience. But it is also relevant to those working in other fields, including risk management and Enterprise Risk Management, capital modeling, catastrophe modeling, financial regulation, and solvency assessment. Insurance risk is managed through pooling, unlike financial risk that is managed through hedging. The title could have been Pricing Non-Hedgeable Risk.

The book came about through a confluence of supporting factors. We had worked independently for many years on the problem of defining the value of risk management and risk transfer (especially in the context of property catastrophe risk) and “escaping the efficient frontier.” Don Mango brought us together to work with him and Jesse Nickerson to present a multipart tutorial on spectral risk measures at the Casualty Actuarial Society Spring 2018 meeting. The tutorial was so successful that we felt it deserved a wider audience and set about developing a monograph: “Spectral Risk Measures for the Working Actuary.” As we proceeded to refine our thinking and presentation, we realized there was so much more to be explained. Three and a half years and 1200 git commits later, we had this book.

The literature is rich with good answers to many fundamental questions about insurance risk that are consistent with finance theory and relatively easy to apply. Much is known, in the sense of being out there in the literature, but too much is not *widely* known by people who would benefit from that knowledge. Actuarial education and practice in this area lags the state of the art. We have encountered actuaries struggling with ill-defined terminology and concepts with multiple names. We have seen confusion wrought by inappropriate application of finance theories (remember the underwriting beta?). Our newly minted US Fellows are often ignorant of the latest developments because they are not on the exam syllabus and there has not been an easy way to incorporate them.

This book presents these good answers in one systematic and comprehensive source for the first time, making them much more accessible to actuaries and other practitioners. With this book we intend to raise the bar in actuarial education, enable clear communications, and improve the efficiency of actuaries everywhere by delivering a fresh map of the conceptual territory. We wish we could take credit for the theory we present, but most of it is around twenty years old. We are simply reporting the work of others.

Insurance pricing is multidisciplinary, combining actuarial science and risk theory, probability and statistics, finance and economics, accounting and law. As we organized and

synthesized a body of literature as nearly as old as the industry itself to tell the story of insurance pricing, we tried to be sensitive to its historical development—a play some of which we watched unfold in real time. It is a story we both found fascinating. From defining underwriting profit and a reasonable target profit in the 1920s to arguments about investment income in the 1960s. From systematic risk and option pricing theory applications in the 1980s to a more insurance-specific model in the 1990s. And most recently to the introduction of coherent, convex, and even star-shaped risk measures. We hope the reader has time to appreciate the giants on whose shoulders we are lucky enough to stand and can join us in taking in the spectacular vistas of the meaning, quantification, and management of risk they have revealed.

In putting together this book we tried to stay reasonably rigorous without getting lost in a theorem-proof wilderness. We feel strongly that knowing how to use a technique is not helpful if you are unsure that it is valid to use in the first place! We include technical remarks and provide pointers into the literature (about 300 bibliographic references) for those who want a more thorough understanding of “why.” For the practitioner, we included nearly 100 examples and 150 exercises. The Learning Objectives at the end of each chapter summarize what we hope the reader will take away from it. **Bold** words and phrases introduce terminology that is used throughout the book.

We aimed this book primarily at property-casualty actuaries, at minimum two years of experience as a student actuary with basic knowledge of insurance coverage and structuring, and having passed the beginning mathematics exams. We expect readers with different backgrounds will still be able to get something from the book. A lot of the insurance and finance terminology is only an internet search away. Mathematics background should include calculus and basic probability—sample spaces, discrete vs. continuous random variables, normal and lognormal distributions, integration by parts, etc. Of course, for an in-depth understanding, more background, especially in probability theory, is better.

The manuscript was prepared using free software. It was written in Markdown and converted to TeX using Pandoc. TikZ was used for the figures and diagrams, and all the graphs and plots were made using Python, Pandas, and Matplotlib. We used R for the statistical analysis and to double check Python (they always agreed). Spreadsheets were used for the discrete examples. We both remember when computers booted from (genuinely) floppy disks. The existence of so much free software, of such a high quality, is an unexpected joy.

We owe a debt of gratitude to many people. In academia, keeping us accurate, we thank Dani Bauer, Stuart Klugman, Andreas Tsanakas, Ruodu Wang, Shaun Wang, and George Zanjani. In business, keeping us real, we thank Avi Adler, John Aquino, Neil Bodoff, Julia Chu, Andrew Cox (1978–2021), Dan Dick, Paul Eaton, Bryon Ehrhart, Kent Ellingson, Stephen Fiete, Bob Fox, Jonathan Hayes, Greg Heerde, Wouter Heynderickx, Rodney Kreps, Morton Lane, Mike McClane, Tessa Moulton, Parr Schoolman, Paul Schultz, Jason Trock, Gary Venter, Steve White, and Rebecca Wilkinson.

Special thanks go to Don Mango for starting this all; to Jesse Nickerson for his early involvement in the research and his comments on drafts; and to Yuriy Krvavych and Lawrence McTaggart for their comments on drafts. Richard Goldfarb stands out for particular thanks, having provided very detailed and pertinent feedback that resulted in numerous improvements. Stephen: I would like to recognize the influence of Glenn Meyers and Richard

Derrig (1941–2018) early in my career—they taught me how to think about pricing insurance risk. I am enormously grateful to my wife, Helen, who started proofreading the manuscript at a late stage and found herself learning the material in a crash course. Her fresh perspective and unyielding commitment to clarity helped improve the presentation in uncountably many ways. John: I would like to thank Jack Caron, Bernie Shorr, and Aaron Stern for opening doors.

1

Introduction

In order to make insurance a trade at all, the common premium must be sufficient to compensate the common losses, to pay the expense of management, and to afford such a profit as might have been drawn from an equal capital employed in any common trade.

Adam Smith, *The Wealth of Nations* (Book 1, Ch X, Part I, 5th Edition, 1789)

1.1 Our Subject and Why It Matters

Pricing insurance risk is the last mile of underwriting. It determines which risks are accepted onto the balance sheet and makes an insurer's risk appetite operational. It is critical to successful insurance company management.

As the last mile, pricing depends on all that has come before. Actuaries and underwriters have analyzed and classified the risk, trended and developed losses, and on-leveled premiums to pick a best-estimate prospective loss cost. Accountants have allocated fixed and variable expenses. Simulation models place the new risk within the context of the company's existing portfolio. The mechanics of all this work is the subject of much of the actuarial education syllabus: experience and exposure rating, predictive analytics, and advanced statistical methods. That is not the subject of this book! All of that prior effort determines the expected loss, and we take it as a given. Pricing adds the *risk margin*—to afford capital a reasonable return. The risk margin is our subject.

Since risk margins are often small, how is it they deserve a whole book? Because risk considerations have an outsized market impact. True, personal property may only earn a single-digit margin. But that business often relies on reinsurance priced with margins of 50% or more. When the reinsurance markets fail or become stressed—as seen after Hurricane Andrew and the Northridge earthquake, for example—the tail of high-risk-margin business wags the dog of much larger property lines. Risk margins are critical to the functioning of the insurance market. Even for lines with thin margins, the collective risk and return decisions of firms have profound macro impacts over time such as the secular increases in homeowners pricing over the last twenty years.

We emphasize *insurance risk*. We do not discuss credit risk nor operational risk. We have only a little to say about asset risk and nothing about interest rate risk. Market risk, underwriting cycles, competitive threats? Sorry, all off-topic. We are focused on the risk of losses arising from insurance contracts. We lean heavily towards a property-casualty perspective and, within that, towards catastrophe risk; however, the principles we lay out apply to any insurance risk. This is not a book about Enterprise Risk Management (ERM) although we do have a few words to say about optimization and portfolio management.

The goal of this book is to demonstrate how to

1. compute a reservation price (technical premium, required premium) for the **portfolio**, and
2. allocate it to portfolio **units** (policies, lines of business, etc.) in a defensible manner

starting from a model of the insured risks. These pricing techniques have powerful applications. They allow us to assess the performance of different units, evaluate needed reinsurance, and optimize overall strategy.

1.2 Players, Roles, and Risk Measures

Figure 1.1 shows the participants in the insurance pricing problem. Insureds, left, pay premiums to the insurer and in turn receive loss payments. The regulator, on top, observing the risk that the insurer is taking on, imposes asset requirements. Investors, right, provide capital and in turn receive the residual value (remaining assets) after losses are paid.

Insureds buy insurance because of their aversion to risk and because they are required to do so to drive a car, buy a house with a mortgage, etc. Regulators play a social policy role, addressing three principal concerns. First, to ensure mandated third-party insurance provides effective protection. Second, to manage the externality of losses exceeding assets. And third, to prevent insureds being fleeced by excessive premiums. The first concern is present in any tort-based system. We loosely identify the second as European and the third as

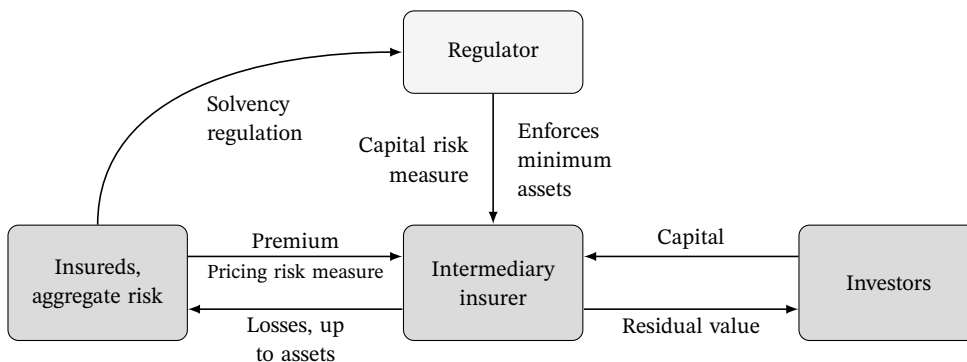


Figure 1.1 Players and their roles. The regulator applies a capital risk measure to determine required insurer assets. The pricing risk measure gives the cost of investors' capital. Assets in excess of losses are paid to investors as the residual value of the business.

American. We focus on the second concern, asset adequacy. Our development of technical premiums naturally aligns with the third fairness consideration if we assume that capital markets require fair returns.

Investors indirectly determine premiums because premiums plus capital add up to and fund assets, Figure 1.2. Investors' willingness to provide capital to insurers translates into a pricing risk measure, which the insurer applies to the covered risks. *Premium and asset levels are separate problems and need separate tools.*

Two important questions arise from insurance company promises to pay certain sums of money contingent on random events.

1. Are there sufficient assets to honor those promises?
2. Are investors being adequately compensated for taking on those risks?

Crucially, we need to talk about not one but *two* different risk measures to answer these questions.

Question 1 concerns risk tolerance and is usually answered by an economic capital model. It determines the assets necessary to back an existing or hypothetical portfolio at a given level of confidence. This exercise is also reverse engineered: given existing or hypothetical assets, what are the constraints on business that can be written?

We can imagine a regulator—interpreted broadly as an external authority—considering a portfolio of risks that the insurer proposes to cover. The regulator specifies the amount of assets the insurer must hold to cover the risk. If there is a shortfall after losses are realized, it will be made up by parties external to the insurer, e.g., a guarantee fund or other government entity, or the insureds themselves insofar as they are not reimbursed for claims. The regulator seeks to minimize the nonpayment externality, balanced with a desire for economical insurance.

A **capital risk measure** is applied to economic capital model output to quantify the level of assets the insurer must hold. Value at Risk (VaR) or Tail Value at Risk (TVaR) at some high confidence level, such as 99.5% or 1 in 200 years, are both popular, but other possible measures exist.

Question 2 concerns risk pricing or risk appetite. We must determine the expected profit insureds need to pay in total to make it worthwhile for investors to bear the portfolio's risk. Regulated insurers are invariably required to hold capital on a regulated balance sheet. We generally assume a funding constraint where premium and investor supplied capital are the only sources of funds. Then, the **pricing risk measure** determines the split of their asset funding between premium and capital.




Figure 1.2 The different roles of capital and pricing risk measures.

The capital and pricing risk measures should not be confused. Historically, capital risk measures have been studied in the context of finance and regulation, e.g., Artzner *et al.* (1999). In contrast, actuaries have studied pricing risk measures as premium calculation principles (Goovaerts, De Vylder, and Haezendonck 1984). The recent popularity and focus on coherent risk measures has overshadowed actuarial premium calculation principles and led to some confusion about the two distinct purposes of risk measures. Much of the recent literature implicitly or explicitly refers to the capital domain only. However, practitioners dealing with issues such as business unit performance, premium adequacy, and shareholder value are operating in the pricing domain. Taking a risk measure suitable for one use and applying it to the other invites unexpected and confusing results. Instead, we must understand how the capital and pricing measures work together in a complex, nonlinear manner to determine technical prices.

The top-down pricing process we have described may not seem commonplace, although those working in catastrophe reinsurance should find our process familiar. Most individual risk pricing actuaries can legitimately claim to use a bottom-up approach. Nevertheless, deep within almost every company lies a corporate financial model functioning exactly as we describe. It asks: How much capital is needed? What is the cost of that capital? What overall margin is necessary? And, how should it be allocated to each unit?

1.3 Book Contents and Structure

The book has four main parts: one on measuring risk, one about portfolio pricing, one about pricing units within a portfolio, and one addressing advanced topics. The high level overview we provide here supplements the introductory paragraphs in each chapter.

1.3.1 Part I: Measuring Risk

Part I is about **risk**. What is risk, and how can it be measured and compared? We discuss the mathematical formalism and practical application of representing an insured risk by a random variable. We define a risk measure as a functional taking a random variable to a real number representing the magnitude of its risk. We give numerous examples of risk measures and the different properties they exhibit.

Some properties are more or less mandatory for a useful risk measure, and they lead us to coherent risk measures. Coherent risk measures have an intuitive representation, providing us with some guidance on forming and comparing them. Spectral risk measures (SRMs)—also known as distortion risk measures—are a subset of coherent measures. They have additional properties and a particularly straightforward representation via a distortion function. Spectral risk measures can be viewed in four equivalent ways:

1. as expected values with varying distorted probabilities,
2. as a weighted sum of TVaRs at different thresholds,
3. as a weighted sum of VaRs at different thresholds, where the weights have specific properties, and
4. as the worst expected value across a set of different probability scenarios.

Spectral risk measures alter or distort the underlying pattern of probabilities and compute expected values based on the new probabilities, analogous to the effect of stochastic discount factors in modern finance. The distorted probability treats large losses as more likely, creating a positive pricing margin. TVaR is the archetypal SRM. It is simple yet powerful and has many desirable properties. We gain analytical insights into the nature of SRMs because they are all weighted averages of TVaRs. For example, we can allocate any SRM-derived quantity by bootstrapping a TVaR allocation.

1.3.2 Part II: Portfolio Pricing

Part II is about **portfolio pricing**, where the entire portfolio is treated as a single risk. Risk is related to return, suggesting we should apply a risk measure to portfolio losses and use the result to indicate a price. Our principal goal is to determine what price is sufficient for assuming the portfolio risk. Secondary goals include understanding, making inferences about, and calibrating to, market prices.

Insurance is characterized by risk transfer through risk pooling. Figure 1.1 combines all insureds into one portfolio. It shows how the capital and pricing risk measures interact to determine the insurer's risk pool premium. Part II of the book treats the cash flows on the lower right, between the insurer and the investors.

We are aware that pricing actuaries and underwriters do not set premiums; markets do. However, the aggregate effect of individual company risk-return decisions drives quotes and acceptances in the market. When we talk about setting premiums, we understand it in the framework of evaluating market pricing or offering a quote.

How are the parameters of a pricing model determined? This is a difficult question that must be answered to put theory into practice. We provide examples showing how different parametrization methods perform, link pricing to capital structure, and calibrate an SRM to catastrophe bond pricing.

1.3.3 Part III: Price Allocation

Insurers must allocate a portfolio price and margin to its constituent units to sell policies and run their business. **Price allocation** is the topic of Part III.

We examine how units contribute to portfolio risk. For example, the models may show several outcomes that lead to insolvency. Which units are the drivers of losses in those scenarios? Parallel questions can be asked of other, less catastrophic, levels of loss.

Having computed technical premiums as distorted expected values, we can then apply the same distorted probabilities to unit losses, based on their co-occurrence with the total portfolio losses, to allocate the premiums to units. This technique provides a high degree of consistency and synchronization in calculating technical premiums by unit.

There is a particular approach to handling business unit performance assessment and reinsurance decision making that makes use of a capital measure, typically VaR or TVaR, but appears not to make use of a pricing metric. This approach is rooted in return on risk-adjusted capital style financial logic. It takes two steps: allocate capital and then assign a cost

of capital *hurdle rate* that every unit must meet on its allocated capital. All decisions, such as reinsurance purchases, use the same cost of capital as a benchmark.

Practitioners recognize that this approach tends to place uncomfortably large burdens on catastrophe exposed units relative to units that do not participate much in solvency threatening events. In addition, when applied to reinsurance purchasing, it tends to favor, almost without exception, deals that operate at high levels of loss with low probabilities. We show that the problem here stems from the implicit use of what we call the Constant Cost of Capital (CCoC) SRM. The overall hurdle rate for the entire portfolio may not be appropriate for every unit. What is needed is a pricing risk measure—*different* from CCoC—that responds to varying levels of riskiness with different required rates of return. Whereas Part II discusses the construction of such measures, Part III discusses how to deploy them.

1.3.4 Part IV: Advanced Topics

The last part of the book touches on five topics that go beyond the coverage of previous chapters: asset risk, reserves, going concern and franchise value, reinsurance optimization, and portfolio optimization.

1.3.5 Further Structure

Parts II and III divide portfolio pricing from allocation considerations. Within each part, we further distinguish **classical** from **modern** approaches, and **theory** from **practice**. The hierarchy, reflected in the sequence of eight core chapters, is:

- Chapters 8 and 9: classical portfolio pricing theory and practice,
- Chapters 10 and 11: modern portfolio pricing theory and practice,
- Chapters 12 and 13: classical price allocation theory and practice, and
- Chapters 14 and 15: modern price allocation theory and practice.

Our dividing line between classical and modern is 1997, the average publication date of three highly influential papers. Relating to Part I: Artzner *et al.* (1997) introduced coherent risk measures and revolutionized thinking about measuring risk. Relating to Part II: Wang (1996) introduced the premium density and developed the theory of pricing by layer. And, relating to Part III: Phillips, Cummins, and Allen (1998) rigorously derived financial prices in a multiline company accounting for default. The classical versus modern bifurcation serves a convenient organizational purpose but should not be taken too seriously.

Classical pricing is predominantly actuarial and risk theoretic. A stability requirement, often linked to the probability of eventual ruin, determines required assets. There is no direct consideration of the cost of capital. On the other hand, modern approaches combine risk theory with financial and mathematical economics and decision science. They relate risk to the investors' return and the cost of capital and pay attention to uncertainty and risk under pooling. They leverage powerful mathematics to understand intuitively reasonable risk measures.

1.4 What's in It for the Practitioner?

This book is intended to be a practitioner-friendly reference as well as providing a theoretical framework. Our methods must have a firm theoretical foundation to justify their real world application. Many topics are inherently technical and require a mathematical background to understand fully. At the same time, the methods we describe can and should be implemented in practice. We have structured the book so readers eager to get their hands dirty can do that more easily. Throughout Parts II and III, we alternate theoretical and practical chapters. The Practice chapters present a range of simple numerical examples and apply all the methods we propose to three realistic Case Studies.

We also pay more attention to institutional arrangements—the way things get done—than the typical theoretical presentation. Different forms of capital, capital compared to equity, accounting, and the mechanism of default, especially equal priority, all play essential roles. Furthermore, we address certain standard practices in the industry and subsume them within our analytical framework so the reader can better appreciate their properties and behavior.

Over the years we have found that putting these tools and techniques into practice raises the following questions.

Which risk measure should I use? A common followup question asks if the risk measure should be sensitive to the tail of the distribution or volatility in the body, which translates into a concern about solvency vs. quarterly earnings. As always in modeling, the measure must be appropriate to the intended purpose. Our framework separates the amount of capital from its cost: the capital risk measure is necessarily tail focused, whereas the pricing risk measure captures investor return expectations. Additionally, the connection between capital *structure* and the *pricing* risk measure is a fundamental insight.

How do I reconcile and manage different economic and regulatory views of risk? Often followed by, “Who cares? We manage to [Rating Agency’s] capital model”. We agree: you don’t care. The rating agency model is a binding constraint on the *amount* of capital for many insurers. You control the form and influence the *cost* of capital. Again, two risk measures. It isn’t a question of reconciliation; it is a question of understanding each measure’s distinct purposes.

Should pricing target a return on all capital, or can there be *excess capital*? Genuinely excess capital is exceptionally rare. Our model produces a cost of capital specific to each company, which varies with the amount of capital. A better capitalized company has a lower percentage cost of capital, other things being equal, because higher layer capital is less stressed and exposed to risk. As a result, the problem of applying a uniformly high cost of capital, producing uneconomic premiums, should not occur. The *frictional costs* of capital are, however, constant for all layers of capital. Indeed, they could be increasing if the management of overcapitalized companies has an incentive to engage in frivolous, self-aggrandizing activities.

How do I determine the cost of capital? Does it vary by unit? The risk cost of capital is the weighted average cost over the actual capital used. Debt and reinsurance have known

costs. The cost of equity capital is normally estimated using a peer study. The cost of capital varies by unit according to which capital layers each unit consumes. The frictional cost of capital typically does not vary by unit.

Can risk margins ever be negative? Classical and modern approaches to pricing are unanimous that the risk margin must be positive for the portfolio. However, negative margins are appropriate for some units within the portfolio. They occur for units that are hedges, with losses arising more in situations where the portfolio has lower losses and less when the portfolio has more significant losses. The common practice of paying a positive margin for ceded reinsurance proves the point: the outward cash flow (premium) is greater in expectation than the inward cash flow (recovery). Looking at expectations makes reinsurance seem inappropriate, but the key to the value of reinsurance (or any hedge) is *when* those cash flows occur.

How do I use a risk measure to determine reservation prices? Chapter 10 and Chapter 14 show how pricing and capital risk measures combine to determine premiums. Chapter 20 offers some more advanced considerations.

The reader will recognize a gap between our simplified models of insurance operations and the complexity of the real world. The practitioner who has mastered Parts I, II, and III and is starting to think seriously about implementing risk measures will likely come up with numerous “What about...?” questions. The following more advanced questions commonly arise for insurers with functioning and integrated risk pricing systems. They are addressed in Part IV.

How do I handle asset risk? How do I incorporate risky assets in the model? How much capital does asset risk consume? Should I treat asset risk in a fundamentally different way from insurance risk? We conclude that an additional degree of freedom emerges, but not to any good use. Section 8.8 discusses the impact of asset risk on pricing and the market value of equity in an option pricing model. Chapter 16 shows that investing in a risky asset typically *lowers* the fair price (and quality) of insurance being sold.

How do I price for reserve risk? I write business that takes years to settle. It is unrealistic to assume all losses are paid in one year. How do I incorporate reserves into the model? Reserve volatility consumes underwriting capacity. However, our model shows that the allocated margins are small when reserves are stable. In a sense, reserves can provide ballast for the prospective portfolio. IFRS and other accounting conventions have begun to require a risk margin for reserves for better earnings recognition. We discuss the Solvency II Cost of Capital Risk Margin and a real option approach to reserves in Chapter 17.

How do I manage a going concern? I don’t manage for just one year and then dissolve the business; I manage a going concern with brand recognition and franchise value. How does that change the model? Chapter 18 outlines the theory of *optimal dividends* and a simple model of franchise value.

How can I optimize ceded reinsurance purchases? I can see how assumed reinsurance can be treated as selling another line of business, but how do I think about ceded reinsurance?

More specifically, how should I go about optimizing it? Chapter 19 discusses how to evaluate and optimize a ceded reinsurance program.

How can I optimize my insurance portfolio? I used to think about optimizing my capital usage against capital constraints. Now I think I should be optimizing my cost of capital, but that doesn't seem to be what you are recommending. Is there a disconnect here? Chapter 20 explores the complex interaction of cost allocation, benefit allocation, and premium regulation. It uncovers some unavoidable market distortions.

1.5 Where to Start

If you have read this far, you likely have a pricing problem. It may be embedded in a broader effort—business unit assessment or portfolio optimization or strategic planning—but it comes down to a pricing problem at its core. At a high level, our recommendations sound simple:

1. Establish your asset requirement.
2. Establish your portfolio cost of capital.
3. Select and calibrate a consistent spectral risk measure.
4. Use what we call the natural allocation to allocate the margin to each unit.

These recommendations presume a lot of work has already been done: gathering and organizing relevant data, developing a mathematical model or numerical tabulation (simulated sample) of the portfolio risks, establishing loss cost estimates for the units, etc. As we said, pricing is the last mile.

The asset requirement should be easy to determine since an external authority usually promulgates it. However, it may require some work to compute, using a standard (e.g., regulatory) capital risk measure. If you find no obvious binding capital constraint, remember that *management's* risk tolerance is irrelevant; only the *owner's* risk tolerance matters. Try to divine it. This step can be incredibly challenging for mutual companies. If you are engaged in an optimization project, then a capital risk measure is necessary because you will have to what-if the capital requirement. If the problem involves the current portfolio only, say a business unit profitability assessment or reinsurance purchase decision, you need only calculate current required assets.

The portfolio cost of capital may similarly be handed down from on high. It can be expressed as a rate of return or a monetary margin amount; these are interchangeable representations. In the unlikely case you get to set your portfolio profitability target, you need to examine your firm's balance sheet—fortunately, this is required in the next task.

Selecting and calibrating a pricing risk measure—specifically a spectral risk measure—is the biggest challenge. We have evolved away from our early fondness for particular parametric SRMs (especially the ones we invented). We now recommend using bespoke nonparametric or semiparametric distortion functions to more closely mirror actual funding costs. It may be that you are not modeling the entire firm's portfolio but only a part of it.

If you do not have access to the whole company risk profile, fear not. You should treat the task as if the parent company is the investor and the portfolio is the company—even though this is a case of suboptimizing. The point here is that the SRM gives *shape* to how the overall required margin is distributed across layers of assets at risk. More specific advice on selecting a distortion function is given in Section 11.5.

With these inputs in hand, allocating margin via the natural allocation is almost a trivial numerical exercise.

Of course, we hope you will read the whole book eventually, but we are not so naïve as to assume you have the time to sit down and read it cover to cover. It takes a lot more to explain and understand *why* than *how*. Why spectral risk measures make sense and do not violate the received wisdom of finance theory, and why the natural allocation is justified in being treated as canonical and not merely one of many equally acceptable alternatives—these issues take many more pages than explaining the mechanics of computation. We hope you will appreciate the why and read the whole book. But if you want to jump ahead to a quick grasp of the how, we recommend the following. Make sure to do enough of the exercises as you go along to feel secure that you “get it.”

- Read about the insurance market and Ins Co., our model company, in Chapter 2.
- Review the introductory material on risk measures in Chapter 3. This should be material you already know. But do pay special attention to the Lee diagram in Section 3.5.
- Some of the material on VaR and TVaR in Chapter 4 may be new to you, so make sure you are comfortable with the basics.
- Section 5.1 lays out the big picture of how Ins Co. approaches the task of analyzing its capital and pricing needs.
- Read the practical applications in Chapter 6 and Guide to the Practice Chapters, Chapter 7.
- Read about classical risk theory in Section 8.4 and the DCF model in Section 8.7. This will help tie the later material back to material you likely have already seen.
- See how classical premium calculation principles work out on our case studies in Section 9.1.
- Read the sections in Chapter 10, Modern Portfolio Pricing Theory, down through Section 10.8. This is core theory about SRMs.
- Read and work examples in Chapter 11, Modern Portfolio Pricing Practice, down through Selecting a Distortion in Section 11.5. Read this last section twice and bookmark it for the day you need to select a distortion for your own purposes.
- Browse Classical Price Allocation Theory, Chapter 12, down through Loss Payments in Default, Section 12.3. This is material that should be more or less familiar to actuaries.
- See how classical price allocation works out on our Case Studies in Chapter 13.
- Read the first two sections in Chapter 14, Modern Price Allocation Theory. This covers the natural allocation of a coherent risk measure, properties and characterization of allocations, computational algorithms, and comments on selecting an allocation. This is the core theory about allocating SRMs. If you are looking for ways to visualize multidimensional risk, read Section 14.3, especially Section 14.3.7, as well.

- Read Modern Price Allocation Practice, Chapter 15. This is essential “how-to” material.
- If reserves feature prominently in your project, you may want to read Chapter 17 in Part IV. This also covers the Solvency II risk margin in Section 17.3.
- If reinsurance purchasing features prominently in your project, you may want to read Chapter 19.
- If you are working with portfolio optimization, you may want to read Chapter 20.

2

The Insurance Market and Our Case Studies

In this chapter, we outline the operation of the insurance market as we model it and describe the hypothetical Ins Co. used in our examples. We then introduce a Simple Discrete Example and three more realistic Case Studies which are used throughout the book to illustrate the methods presented in the theory chapters.

2.1 The Insurance Market

Insurers are **one period, limited liability** entities with no existing liabilities. We consider multi-period insurers in Part IV only. The insurer is called **Ins Co.** It sells insurance **policies** to **insureds** or **policyholders**. Policyholder and insured are treated as synonyms, and both include **claimants**. Ins Co.'s **portfolio** is the collection of policies it writes. The length of the period is usually one year. Its length is relevant only because of the time value of money, since interest is a rate per year.

Policyholder liabilities are any amounts Ins Co. owes to policyholders. The two biggest are loss reserves and unearned premium reserves. We incorporate reserves in Part IV only. In property-casualty insurance, loss reserves cover claims that have been incurred but not paid, whether reported or not. In life insurance, liabilities include policy values for long duration contracts.

Assets are the total financial resources owned by Ins Co. that it can use to meet its policyholder liabilities. A regulator usually stipulates **required assets**, a minimum amount of assets that Ins Co. must hold; see Section 1.2.

Ins Co., like all firms, finances its assets by issuing **liabilities**. It sells policies, creating policyholder liabilities, in exchange for premiums, and it raises capital from **investors** by selling them its residual value (equity) or other promises (debt, reinsurance).

Investors can be shareholders when Ins Co. is a stock company or insureds when it is a mutual company or debt holders or reinsurers.

Insurers are **intermediaries** between insureds and investors. Intermediary always means an insurance company intermediary, and never an agent or broker.

Portfolio components are referred to as **units**. A unit can be a single policy or a group of policies or be defined by line, geography, branch, business unit, or other characteristics. Unit can also represent the segmentation between reinsurance ceded losses and retained losses.

Ins Co.'s **aggregate loss** is the sum of losses from its portfolio over one period.

Ins Co.'s operations are controlled by eight variables: **(expected) loss, premium, assets, margin, capital, loss ratio, cost of capital,** and **leverage**. The first five are monetary quantities, and the last three are unitless ratios. They obey five relationships:

1. premium equals loss plus margin,
2. assets equal premium plus capital, which we call the **asset funding constraint**,
3. loss ratio equals loss divided by premium,
4. cost of capital equals margin divided by capital, and
5. asset leverage equals assets divided by premium.

Figure 2.1 lays out these variables and relationships. Monetary quantities are the vertices, the bold diagonal lines correspond to the ratios, and the two shaded triangles signify the asset and premium sum conditions.

Premium is the amount charged for providing insurance. Premium is net of (i.e., excludes) underwriting expenses but includes an allowance for risk called the **margin**. Profit, profit load, profit margin, risk margin, and risk load are all synonyms for margin.

Premium is the critical variable; it is the foundation of the schematic. It is the bridge between investor cash flows on the left and insurance cash flows on the right. At the expected outcome, premium is shared, with margin flowing to investors and expected loss to the insured.

Policyholders are liable for their expected loss—as Adam Smith pointed out in 1789; by “common loss” he means expected loss. Financing the remaining assets is the **shared liability** of policyholders and investors. The shared liability equals assets minus expected loss, or equivalently capital plus expected margin. **Pricing apportions the shared liability to policyholders and investors.**

The **loss ratio** is the ratio of loss to premium. Because premiums exclude expenses, a 90% loss ratio includes a healthy margin. The premium **markup** is the inverse expected loss ratio,

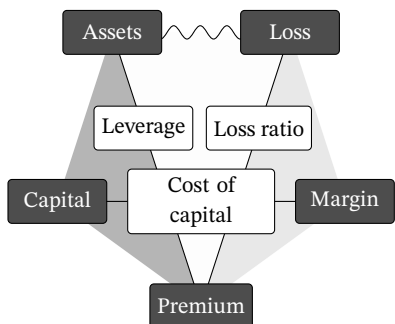


Figure 2.1 The eight variables that control insurance operations and five relationships between them.

the ratio of premium to expected loss. Catastrophe bond pricing often quotes markups rather than loss ratios. **Premium leverage** refers to the ratio of premium to capital.

The margin is distinct from the **contingency provision**, which the Actuarial Standards Board (2011) defines as a correction for persistent biases in ratemaking. It says the “contingency provision is not intended to measure the variability of results and, as such, is not expected to be earned as profit.”

A **catastrophe** or **catastrophe event** refers to an single event causing loss to multiple units, such as a hurricane, typhoon, earthquake, winter storm, terrorist attack, or pandemic. A **catastrophe loss** is the total loss across all units from a catastrophe event. A **catastrophe unit** means a unit prone to catastrophe losses. A **catastrophe risk** is a peril likely to result in catastrophe losses. Catastrophe risks tend to attract large margins, making them particularly interesting.

At various points we mention **catastrophe models**. These are computer simulation tools used to estimate potential catastrophe losses from an insurance portfolio. Mitchell-Wallace *et al.* (2017) provides helpful background about the operation and use of catastrophe models.

Losses in a **thick-tailed** unit have a high coefficient of variation, are right-skewed and leptokurtic (high kurtosis), and have a significant probability of assuming a very substantial value. Catastrophe losses are usually thick-tailed.

A **long-tailed** unit has a slow payout pattern, meaning claims are not settled until many years after they occur.

Reinsurance is a type of insurance, so we say insurance to cover both, and reinsurance if that is all we mean. Cedents cede business to reinsurers.

The accounting distinction between capital and equity causes unnecessary confusion.

Capital refers to funds intended to assure the payment of obligations from insurance contracts, over and above reserves for policyholder liabilities. Capital is also referred to as **net assets**. The book value of capital depends on accounting conventions. Capital is usually regulated by statute. **Surplus** is a synonym for capital used in US statutory accounts.

Equity is the value of the owner’s residual interest. In a stock company, it is called shareholder’s equity. Accounting equity is typically lower than capital since debt can be included in capital but not equity. Equity also has a market value for public stock companies, based on the value of shares outstanding. Equity levels are not regulated. Accounting equity can be negative. The market value of equity is always non-negative because of limited liability.

2.2 **Ins Co.: A One-Period Insurer**

In this section, we introduce the hypothetical insurer called **Ins Co.**, that we use as the base for our theory and examples. Ins Co. is a limited liability company that intermediates between insureds and investors.

Ins Co.’s customers are **insureds** who are subject to risks they wish to insure, for the three reasons explained in Section 8.1.1. Insureds who use insurance for risk transfer or financing