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.

Pricing Insurance Risk: Theory and Practice, First Edition. Stephen J. Mildenhall and John A. Major.

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

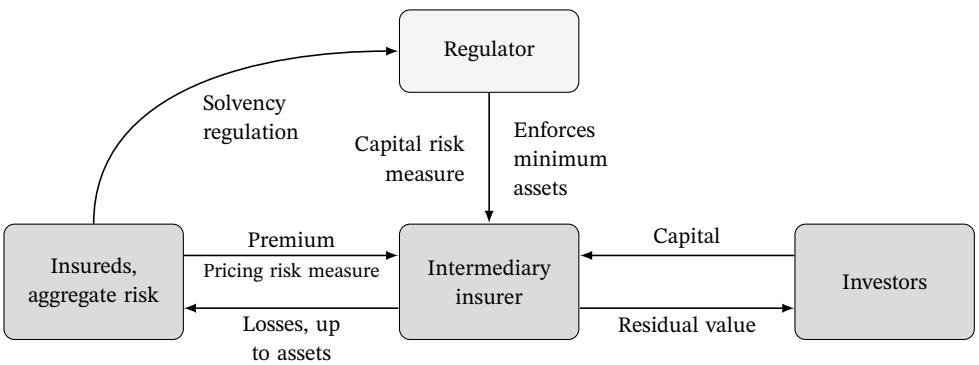

Figure 1.1 Players and their roles. The regulator applies a capital risk measure to determine required insurer assets. The pricing risk measure gives the cost of investors' capital. Assets in excess of losses are paid to investors as the residual value of the business.

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.

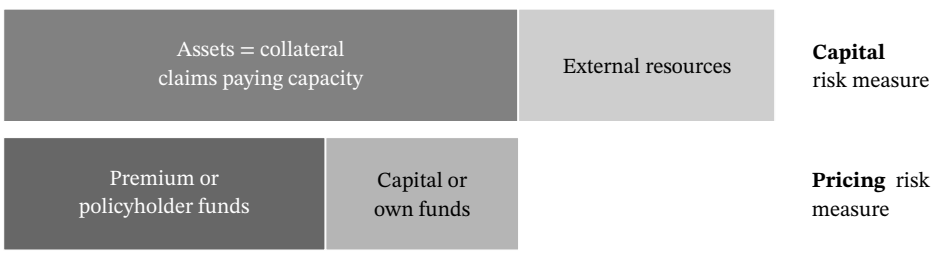

Figure 1.2 The different roles of capital and pricing risk measures.

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.