Chapter 1
Introducing Economic Capital

Just as a cautious businessman avoids investing all his capital in one concern, so wisdom would probably admonish us also not to anticipate all our happiness from one quarter alone.

(Sigmund Freud)

This is a book with a central focus on the notion of economic capital and its applications. It thus makes good sense to begin with a clear description of this surprisingly complex concept. Our story begins with the banking business, although the following ideas will apply equally well—with some perhaps slight adjustments—to a broader range of financial institutions including insurance companies, investment dealers, or brokerages. Banks, as we all know, lend money. Such a mandate entails a number of ancillary financial activities. Loan origination and counterparty evaluation is a central preoccupation, cost-effective funds must be raised in local and international markets to finance these loans, short-term liquidity needs to be earmarked and invested to manage daily flows and unexpected outcomes, and key portfolio asset-and-liability mismatches need to be hedged.[1] In a few words, a modern bank needs to lend, fund, invest, and hedge. All of these efforts, even hedging, give rise to various financial (and non-financial) risks that need to be identified, measured, and managed. A bank's success, as a venture, depends importantly on its effectiveness in navigating these risks.

The first challenge facing a bank arises from the diverse nature of these risks. A movement in an interest or exchange rate, for example, is somehow different from the failure of a counterparty to meet their payment obligations. Measuring the former risk is quite likely to involve conceptual and practical differences from the latter one. This is only a simple example; the true range of risks is, in fact, much broader. Although they might arise and behave in different ways, it is useful and important to manage one's risks in a global sense. This idea is loosely referred to as enterprise risk management. A powerful tool for the joint management of these

[1] Many banks, of course, also raise funds via customer deposits.

D. Bolder, *Modelling Economic Capital*, Contributions to Finance and Accounting,
https://doi.org/10.1007/978-3-030-95096-5_1

risks is referred to as economic capital. Economic capital, to be criminally brief, is a firm-wide, model- and assumption-based measure of business risk.

Calculating economic capital requires not just one, but typically multiple financial models. The multiplicity of models arises precisely from the variety of risks that we seek to incorporate into the final measure. In a perfect world, a single, unified model would describe the totality of an institution's risks. The diversity of the underlying risks and, quite frankly, our shortcomings in measuring them makes such an approach practically unfeasible. The following chapters will focus principally on the credit-risk dimension, but we would be remiss if we did not touch upon—at least, conceptually—the other important sources of risk faced by financial institutions.

Although financial modelling is a crucial element in the computation of economic capital, jumping straight into technical considerations is probably not a great idea. Instead, to somewhat mangle a popular adage, we should *think before we model*.[2] With a clear view of what we wish to accomplish—and some of the potential stumbling blocks along the way—digestion of mathematical and statistical details can occur much more smoothly and naturally. This chapter, therefore, seeks to provide a gentle, but comprehensive, introduction to the key ideas, motivations, and concepts associated with the notion of economic capital.

1.1 Presenting the Nordic Investment Bank

Although our intention is to attack economic capital from a broad perspective, the ideas and methodologies presented in the following chapters are employed by a specific institution: the Nordic Investment Bank (NIB). NIB is an international financial institution focused—to put it very briefly—on environmental and productivity related lending activity in the Nordic and Baltic regions. Established in 1975 by the five Nordic countries, it makes loans to both private and public-sector entities, but not individuals.[3] It also finances its activities through market borrowing, not customer deposits. NIB's loans are offered on competitive market terms.

The NIB has been estimating, internally reporting, and using economic capital since 2004. Although not a regulated entity, it has nonetheless followed the Internal Capital Adequacy Assessment Process (ICAAP) since 2016.[4] ICAAP, as a concept and process, relates to the Basel Committee for Banking Supervision's Pillar 2 (i.e.,

[2] The typical advice is to *think before you speak*, but experience indicates that thinking is profitable before engaging in a rather wider set of activities from bungee jumping to financial modelling.

[3] The five Nordic countries include—moving from west to east—Iceland, Norway, Denmark, Sweden, and Finland. The three Baltic countries—from south to north, Lithuania, Latvia, and Estonia—joined the NIB in 2005.

[4] See ECB [15] for a useful guide to the occasionally byzantine world of ICAAP.

the supervisory review) aspect introduced in Basel II.[5] In 2020, through a desire to modernize its risk-governance framework, and precipitated by practical challenges associated with antiquated statutory gearing limits, the NIB introduced a number of changes to its statutes. A new set of economic-capital, leverage, and liquidity management principles—along with some associated changes in the responsibilities of its Control Committee—were introduced.

These statutory adjustments were, to put it mildly, a game changer. Prior to the revised statutes, the principal constraint on NIB's financial activity was a volume-based gearing ratio. While broadly effective, such an approach did not differentiate between the riskiness of each unit of lending or treasury activity to our overall activities. Volume certainly matters, but so does the underlying risk profile of what an institution adds to its balance sheet. Economic-capital offers an approach to capture both dimensions. It does so, however, at a cost. Volume-based ratios are typically easily defined and computed; economic-capital, by contrast, is a relatively complex and nuanced object. Its complexity, of course, is necessary to adequately reflect the intricacy of the risks facing a modern financial institution.

This book is, in many ways, a child of NIB's recent statutory change. While most of the elements described in the forthcoming chapters have been performed for years—if not decades—prior to the recent revision of its statutes, it was a natural moment to carefully rethink and revise the overall framework. Such an effort naturally involved careful documentation of the attendant choices; the result was this book. Putting these ideas into the public domain serves, at least, *two* purposes. First, and perhaps most centrally, it provides clarity about our modelling practices and permits qualified and interested external parties to provide useful constructive feedback on our choices. Placing these ideas into circulation is thus an important step towards mitigating the inherent model risk associated with construction and implementation of an economic-capital framework. Transparency is also a fundamental aspect in public discourse among NIB's member countries. Modelling practices, their mathematical complexity notwithstanding, are no exception.[6] Secondly, many practical lessons have been learned during this process that, in our view, will be of interest to a broader audience. Openness, sharing, and solicitation of external feedback are thus the key objectives in the production of this publication.

The following chapters, therefore, are not based upon theoretical examples of how one might possibly use economic capital. Instead, we will address these issues within the context of a real-life entity. Naturally, there are important limits to the amount of proprietary information that can be shared in such a publication. No customer information will be provided and most figures will be related to high-level values already found in other publications such as financial reports. In many cases,

[5] Basel II's introduction and discussion ranged from roughly 2004 to 2006.

[6] This relates to the idea—see, for example, Elmgren [16]—of open government. While open government is challenging, and not always uncontroversial, transparency has a long history in the Nordic region.

we will use fictitious (although illustrative) data to demonstrate our choices. Specific figures and details regarding credit obligors are, however, rather beside the point and *not* pertinent to a meaningful discussion of the topics. This is a book about choosing, parametrizing, implementing, and applying economic-capital methodologies. As in any large-scale financial modelling venture, both objective and subjective decisions must be taken. It will be useful for the reader to understand and experience how these methodological choices have been taken within the NIB context across the entirety of the credit-risk dimension. Great pains have nonetheless been taken to ensure that, despite the usefulness of a specific perspective, the notions considered in the coming chapters are of general interest and utility.

Colour and Commentary 1 (SOME READING INSTRUCTIONS)*: This work has a lot of ground to cover. This makes it inevitably long and detailed. We can loosely think of reading this book as visiting a large touristic destination such as London, Paris, New York, or Beijing.*[a] *There is much to see, learn, and experience, but it can be difficult to get started and decide where to allocate one's time. Conceptualizing this book as a city visit, its consumption requires a plan. The preface and this first chapter set the stage: these are roughly equivalent to a short tourist guide. To complement the technical discussion, extensive mathematics, and examples found in each chapter, numerous distinct grey text boxes (such as this one) are provided.*[b] *Each box seeks to highlight various key points, takeaways, challenges, or observations. Their purpose is to help the reader to navigate, digest, and organize the material. To extend our tourist-attraction analogy, we can view these grey text boxes as the various destinations found on popular jump-on and jump-off bus tours offered to tourists in these major cities. Skipping through the boxes, in any given chapter, can assist the reader in deciding when (and where) to* jump *into the details.*

[a] No promise is made that exploring the coming pages will be anywhere near as exciting as your next vacation.

[b] There are, in fact, roughly 150 of these text boxes scattered over the forthcoming 12 chapters.

1.2 Defining Capital

Before there was economic capital, there was capital. As the introductory quote suggests—stemming from famed psychoanalyst, Sigmund Freud, whose contributions mostly occurred a century ago—the idea of capital is *not* new. A definition was provided by Adam Smith in his *Wealth of Nations* published in the late eighteenth

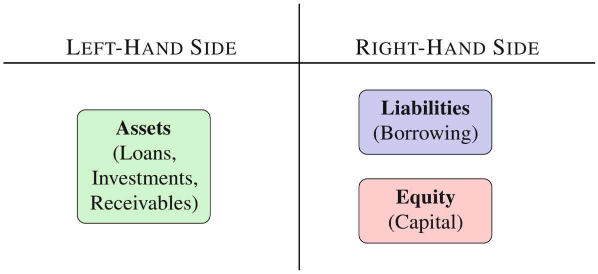

Fig. 1.1 *A stylized balance sheet*: This schematic represents a (very) stylized firm balance sheet. It illustrates the conceptual interplay between assets, liabilities, and equity.

century.[7] The idea of capital, albeit with a strong political-economy dimension, is also treated extensively in the works of Karl Marx written during the 1880s.[8] There are many different definitions that one might consider, but they all basically boil down to an owner's investment or stake in an enterprise. Although it can be broken down more finely, this (slightly simplified) working definition amply serves our purposes.

It is, of course, useful to be slightly more concrete. Figure 1.1 provides a, certainly quite familiar, stylized view of a firm's balance sheet. The left-hand side illustrates firm assets such as loans, investments, and receivables. These are, more or less, the elements of the business that are used to generate revenues. The right-hand side basically deals with sources of financing associated with the revenue-generating left-hand side. It includes liabilities and the overall equity, or capital, position.

Capital (or equity) is, from double-entry book-keeping, a residual figure. It is the excess of a firm's assets over its liabilities. For an accounting scholar, using the terms capital and equity interchangeably is likely considered to be a grave act of negligence.[9] This not being an accounting treatise, a bit of simplification and cheating is permitted. Our objective is, in this regard, not precision, but rather illumination.

Accounting niceties aside, the important point from Fig. 1.1 is that the difference between a firm's assets and liabilities provides invaluable insight into a business' state of affairs. A large equity position indicates a comfortable distance between a firm's assets and its obligations. This is generally a good thing. If it is too large, however, then the efficiency and profitability of a firm can suffer. When this equity

[7] Although not the most famous event occurring in this year, this work—see Smith and Krueger [43]—was originally published in 1776. It defines capital as "that part of a man's stock, which expects to afford him revenue."

[8] The curious, and perhaps masochistic, reader is referred to the original source: Marx et al. [29].

[9] Specifically, the difference between assets and liabilities is referred to as shareholder's equity. Capital, in an accounting sense, is generally considered to be a technical subset of this equity position. Capital is thus viewed as the owner's ongoing investment in the firm, whereas equity is the owner's share of the business.

or capital position becomes too small, conversely, then a small bump in the road might put the firm's continued existence into question. This rough logic falls under the broader category of capital analysis; it is a critical pillar—along with solvency, liquidity, and competitive position—in the determination of a firm's overall credit quality. It reduces to the following simple question: does a firm have sufficient equity?

This balance-sheet view is not only very interesting to accountants and credit analysts, but it is also the jumping-off point for much of corporate finance. The appropriate mix of liabilities and equity, for example, is referred to as the capital-structure decision and is a rich area of research.[10] Dividend policy, the market for corporate control, corporate governance, and topics related to shareholder activism also all look at the firm through this lens.

Our viewpoint, with an eye on the notion of economic capital, is slightly different. We are looking for inherent limitations with this view of the firm. The missing ingredient, as one might expect, is *risk*. It turns out that—particularly for financial institutions—the risk dimension manifests itself primarily in a firm's assets. This makes eminent sense, since a corporation's revenues are typically generated from its assets. For most firms, assets include receivables, physical stock, cash, and inventory. Liabilities, conversely, include payables and bank and/or market debt. The riskiness of the assets of a generic firm—other than perhaps receivables—is typically not dramatic. These assets are, from an accounting perspective, usually held variously at historical or amortized cost.

The story is somewhat different for financial institutions. Some investments—including marketable securities—are held at fair value. This implies that, as market conditions change, valuation changes flow through profit-and-loss. This is useful, because it creates a link between asset riskiness and the firm's equity (or capital) position. Loan obligations, however, are not typically fair valued; the reasons are various, but the lack of liquid secondary markets for their resale is one important explanation. Changes in the (often difficult to assess) value of loan assets are thus not directly incorporated into the balance sheet. These changes do, however, make an indirect appearance. Loan impairments—often referred to as expected credit losses or ECL—attempt to capture average credit losses and impending credit deterioration. These impairment values also flow through the profit-and-loss statement and ultimately impact the firm's equity position.

1.2.1 The Risk Perspective

Fair value and impairment calculations represent valiant efforts, on the behalf of the accounting community, to incorporate asset riskiness into a firm's financial statements. These are particularly helpful for financial institutions such as the

[10] The most sensible starting point in this area is Myers [33, 34].

NIB. There is, however, still something missing. Risk is, practically, a combination of outcomes and likelihood. Risk managers are typically worried about very negative outcomes with low likelihoods (or probability). To really assess the comfortableness—or, in more popular jargon, adequacy—of a firm's capital position, this extreme perspective needs to be incorporated.

What do we really mean by extreme? Let's consider some concrete examples of highly negative outcomes for a financial institution:

- one, or more, large, structurally important loan or bond obligors could simultaneously move into default and fail to pay the firm back;
- interest rates or exchange rates or credit spreads—or all of the above—could dramatically move against the firm;
- one or more large swap counterparties could go out of business forcing the firm to novate (i.e., liquidate and replace) an important part of their swap exposure; or
- key systems or controls could fail and/or large-scale fraudulent activity could occur.

These are all (rather disheartening) examples of credit, market, counterparty, and operational risk, respectively. Each depict adverse events that a firm might face, but which do not make any appearance in our schematic definition of capital from Fig. 1.1.

The balance-sheet equity or capital amount is thus, rather unfortunately, not particularly informative about such extreme down-side asset risks. Is such a pessimistic perspective simply a fiction of the minds of paranoid risk managers? Risk managers may indeed by overly suspicious by nature and training, but the last 25 years have offered no shortage of dramatic financial events: an incomplete (but still depressing) list includes the 1997 Asian crisis, the dot-com meltdown of 2000–2001, the global financial crisis of 2007–2009 and its decade long aftermath, at least two waves of European sovereign-debt crises, and the financial turmoil associated with the COVID-19 pandemic in 2020.[11] Really understanding and assessing the adequacy of firm's capital position thus requires us to systematically go beyond average, or expected, losses.

We have established the need, for the appropriate assessment of a firm's capital adequacy, to incorporate worst-case risks associated with the asset-side of the balance sheet. We have further established how this viewpoint is particularly important for financial institutions given the composition of their assets. There is, however, another important characteristic associated with financial institutions. Deposit-taking banks, investment banks, and insurance companies play a central role in any economy: credit intermediation. In other words, they facilitate the financing of firms and individuals either directly by extension of credit or insurance

[11] An excellent, if somewhat discouraging, broader view of the unnerving regularity of financial turbulence throughout history is Reinhart and Rogoff [38].

or indirectly through enabling financial transactions.[12] If such financial institutions fail or run into trouble, the effect on the economy is disproportionate. There is a knock-on effect to other firms through a reduced ability to efficiently, or cost-effectively, finance their continued operations. In really extreme situations, otherwise well-functioning firms can face dire solvency problems simply due to problems with financial institutions.

Financial firms thus have a high degree of systemic importance. This explains their preferential access to liquidity arrangements through central banks, deposit insurance, and, in extreme cases, direct government bailouts. While this is certainly a contentious area involving use of taxpayer money, these measures all arise from the unique centrality and importance of financial institutions in any given economy. As Peter Parker learned the hard way, however, *with great power, comes great responsibility.*[13] Systemically important financial institutions may have access to a broad range of (tax-payer funded) support functions, but there is a flip side to the coin. They also face disproportionately high levels of oversight and regulation. The regulatory community has thus been, virtually since its inception, at the forefront of the movement to incorporate this extreme-risk perspective into the assessment of a firm's asset riskiness and, by extension, its capital position. Our simple question, *"does a firm have enough capital?"*, takes on a new degree of importance in this setting.

1.2.2 Capital Supply and Demand

This detour through the origins of capital, as a financial concept and its potential shortcomings, has finally led us to the rationale for economic capital. Economic capital is basically an *economically* motivated alternative representation—working through the firm's assets—of the balance-sheet capital position. Abstracting (for the moment) from the details of its computation, imagine that we can actually calculate such an adjusted capital amount. What would be the point? We could directly compare this risk-adjusted capital figure to the actual balance-sheet capital (or equity) position. If our approximation was bigger than the balance-sheet value, it would be somewhat disappointing. It would mean that the firm would not have enough capital to weather really extreme worst-case risk outcomes. We would probably, by contrast, be encouraged if our asset-risk adjusted capital position exceeded the balance-sheet value. This is the heart of the economic-capital concept. It is a measure used by regulators—and indeed any interested internal or external

[12] For a much more complete review of the banking sector, and its economic centrality, see Siklos [42].

[13] This is a humorous, but still entirely pertinent, reference to Spider-Man. If you really desire some additional background on this, very tangential topic, it's easy to find. Bendis [2] is not a bad place to start.

stakeholder—to assess a firm's capital adequacy. Again, we return to our simple question: does the firm have enough capital?

Regulators have, in particular, a somewhat unique perspective on the idea of economic capital. Although they are certainly interested in understanding the sufficiency of a financial institution's capital position, they are also intimately concerned with the *minimum* amount of capital each firm should hold. This prescriptive way of looking at the problem makes sense; regulators are trying to create a level playing field among financial institutions within their jurisdiction. While minimum capital requirements and economic capital are intimately related, they are not precisely the same thing. The former is a lower bound, while the latter is a general concept.

To use economic capital to assess a firm's capital adequacy, it is helpful to introduce the notions of capital supply and demand. Capital supply is the actual amount of a firm's available capital; rather simply, it is the difference between balance-sheet assets and liabilities.[14] The capital demand is one's economic capital computation; it is an estimate of how much capital is potentially required, or demanded under adverse conditions, by the riskiness of the firm's assets. Capital adequacy analysis thus essentially reduces to the comparison of these two quantities. Capital supply less capital demand is often referred to as (capital or equity) headroom. In the unfortunate case of capital demand exceeding supply, we can speak of negative headroom or capital shortfall. As a general rule—which paves over much complexity and nuance—we would like a firm's headroom to be reasonably large and positive. Figure 1.2 provides a schematic illustration of the case when capital supply exceeds capital demand (and vice versa).

While hopefully helpful from a conceptual perspective, to this point we have only addressed why one might be interested in something like economic capital. The actual quantity, and how it is produced, probably remains frustratingly vague. Numerous questions arise. How do we, for example, actually adjust the specific assets? More structure is definitely needed in this area. What do we actually mean by worst-case extreme risk? This could clearly mean different things to different firms. We might also ask: how do we accommodate disparate types of risk? There is also the need to incorporate the interactions between individual assets in our portfolios. We are, after all, modern financial professionals concerned with the portfolio perspective. These are all important questions that we will seek to address in the following discussion.

[14] There are often slight technical adjustments to this figure, which are not to be ignored, but simple balance-sheet arithmetic gets us to the right order of magnitude.

Colour and Commentary 2 (DEFINING CAPITAL): *Equity, in a balance-sheet sense, refers to an owner's share of the firm. Practically, it is the residual asset value after subtracting the firm's liabilities. Colloquially, this is referred to as the enterprise's capital position.*[a] *Accounting practices—such as fair valuation and loan-impairment calculations—take into account the average, or expected, riskiness in firm assets. They do not, however, address extreme worst-case risk outcomes. This perspective, given the depressingly regular incidence of financial crises, is important and cannot be ignored. Assessing a firm's capital adequacy—which is a fancy way of asking if it has enough capital—requires this extreme-risk perspective. Here enters the concept of economic capital. Economic capital is an* economically *motivated alternative representation—focused on the riskiness of the firm's assets—of the balance-sheet capital position. This reduces questions of capital adequacy to a comparison of actual balance-sheet equity values (i.e., capital supply) to one's economic-capital estimate (i.e., capital demand).*

[a] There are, from an accounting perspective, subtle and important differences between capital and equity. This fact notwithstanding, these two terms are treated as broadly synonymous in everyday financial parlance.

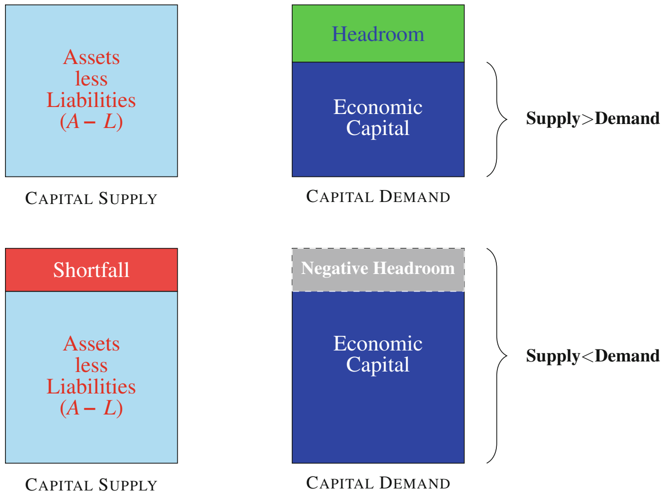

Fig. 1.2 *Capital-adequacy mechanics*: The preceding schematic illustrates the interaction between the notions of capital supply, capital demand, and headroom. All else equal, we prefer that a firm's supply of capital exceeds the demands upon it.

1.3 An Enormous Simplification

Let us begin with a statement of full transparency: there does *not* exist a unique, correct way to compute economic capital. Any individual who claims otherwise is either being less than truthful, or trying to sell you something, or both. Economic capital is a model- and assumption-based metric. There is instead a range of possible choices with varying degrees of defensibility and potentially differing objectives. Depending on your perspective, this might be either refreshing or discouraging. The variability in economic-capital computations should nonetheless not be overstated. There do, of course, exist general principles and accepted practices for the measurement of economic capital. The Basel Committee on Banking Supervision, for example, has been a driving force, from the regulatory side, in finding a common language for economic capital and standardizing comparison across firms and jurisdictions.[15] This offers a useful backdrop for making decisions on one's economic-capital calculation, but it still represents a few options among many possible choices.

Despite its non-uniqueness and the potential complexity behind one's economic-capital risk model, computing economic capital is practically fairly easy to understand. To see this we introduce a simplification; this is not how economic capital is computed, but rather a conceptual tool. Imagine the following procedure. For each asset on a firm's balance sheet—be it a loan, a deposit, a corporate bond, or an interest-swap contract—we assign it a *number*. Not just any number, of course, but we follow some rather structured rules. In particular, each of our asset numbers takes a value between 0 and 1. More specifically,

1. a value of zero implies that the asset has no risk; and
2. a value of one implies that it has no value.

The obvious corollary of this property is that the greater the asset's risk, the larger our assigned number. The extremes, while informative, are likely not terribly interesting. A value of zero would probably only apply to local-currency cash position with a reputed central bank (and even this is debatable). Assigning a value of one, conversely, is basically equivalent to writing off the entire value of the asset. This would presumably only occur in the event of default with absolutely no prospect of any recovery through bankruptcy proceedings. In the vast majority of cases, our number will be comfortably in the unit interval with, in high-quality institutions, a fairly strong tendency to lie at the lower end of the scale.[16]

The reader is probably thinking that such a number assignment might be an interesting exercise, but how is it related to the idea of economic capital? The answer is surprisingly simple. If one multiplies this *number* by the asset value,

[15] Much has been written on the Basel accords; BIS [3, 4, 5] are excellent starting points for some of the practical elements of minimum regulatory capital requirement computations. Chapter 11 will delve much more deeply into the regulatory world.

[16] Different lending institutions, of course, specialize in different parts of the credit spectrum.

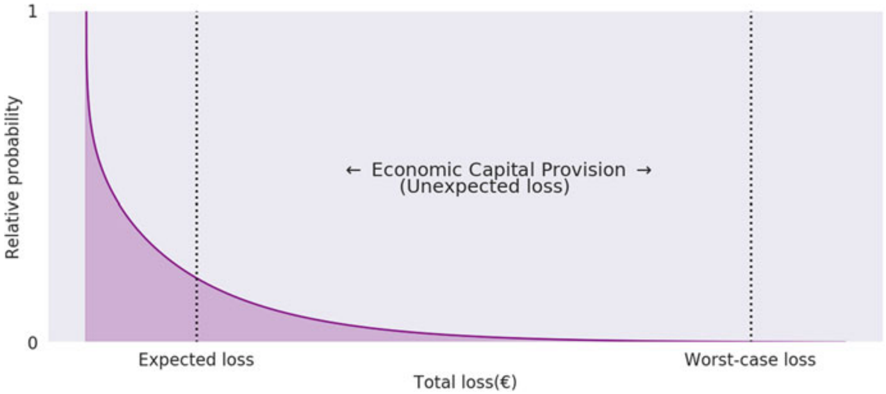

Fig. 1.3 *Economic capital at a glance*: This figure illustrates a stylized portfolio loss distribution. Economic capital is defined, in a technical sense, as the difference between the worst-case outcome and the expected, or average, loss. This difference is often referred to as unexpected loss.

one arrives at its economic-capital figure. The total economic capital figure for a given firm is thus the sum of the product of these numbers—each between zero and one—and their asset values.[17] Consider a simple (aggregate) example: a firm's total assets are EUR 1 billion and each and every number is 8%. In this case, a firm's economic capital figure is EUR 80 million. If the firm's actual equity position is EUR 100 million, then it can be considered to have adequate (or even excess) capital. An equity position of EUR 60 million, however, would likely be considered problematic.

Our asset numbers are clearly related to the notion of risk. Bigger numbers, after all, are assigned to riskier assets. While useful to understand and interpret economic capital, these numbers are not used to compute it. Estimating economic capital requires wrestling with asset risk. Risk, when you get down to it, depends on the interplay between value outcomes and their associated probabilities and likelihoods. The collection of all possible outcomes (or events) and likelihood is referred to as a statistical distribution or probability law. We require such an object to assess the risk of our asset portfolio. Figure 1.3 illustrates a stylized, firm-wide, asset-portfolio loss distribution: for the moment, let's not worry about where it might come from.[18] This distribution will help us identify the actual practical definition of economic capital. It specifically outlines the various firm asset-loss outcomes along with an assessment of their relative probability. Low levels of loss occur with a high likelihood, while extreme (worst-case) losses are rare. The expected, or average, portfolio loss lies, as it should, at the centre of the overall probability mass.

[17] Those familiar with regulatory risk weights will recognize this idea.

[18] This is a job for future chapters.

As discussed previously, the main point of economic capital is to understand what happens to a firm's assets when things go really wrong. In the context of our asset-loss distribution in Fig. 1.3, it means considering events out in its far right-hand side. This is referred to as the tail of the distribution. It is in this unsavoury neighbourhood that we will meet extreme, worst-case losses. Inspection of Fig. 1.3 reveals that the further we move into the distribution's tail, the lower the probability of the loss outcome. This makes logical sense, but it strongly suggests that we cannot really talk about *worst-case* losses without also including some assessment of their probability. We could, for example, talk about a 0.1% worst-case loss. This basically suggests that portfolio losses of this magnitude would occur once out of every 1000 random draws from this firm-asset probability law. This worst-case probability is referred to as a confidence level and it is typically denoted by the first letter in the Greek alphabet, α. Moreover, by convention, we talk about it as $1-\alpha$.[19] So, to return to our example, a 0.1% of a worst-case loss translates into a $1 - 0.1\% = 99.9\%$ confidence level.

It is tempting to characterize economic capital as this $1 - \alpha$ worst-case loss. This would not be quite correct. As we saw previously, the accounting community has already done a good job—through adjustments such as fair valuation and loan impairments—of incorporating expected loss. Economic capital, therefore, is defined as the $1 - \alpha$ worst-case outcome minus the expected loss. This difference—in an unfortunate choice of terminology—is often referred to as the unexpected loss. Conceptually, it is useful to think of economic capital as the additional reduction of asset value—over and above the expected loss already addressed by one's accountants—that might occur if things really get ugly.

Figure 1.3 and our recent definitions are practically quite helpful. The economic capital is thus formally defined as the $1 - \alpha$ confidence-level unexpected asset loss. To compute our number between zero and one, at the overall portfolio level, we simply evaluate the following ratio:

$$\text{Firm's Asset } \textit{Number} = \frac{\text{Asset Portfolio Unexpected Loss}}{\text{Asset Portfolio Value}}. \tag{1.1}$$

The larger this number, of course, the more capital a firm would require to insure, or cushion, itself against large losses associated with the inherent riskiness of its assets. This ratio also depends directly on the value of α. As α gets smaller—and by extension, the larger our confidence level, $1 - \alpha$—the bigger the numerator in Eq. 1.1, and the more economic capital required.

The role of our confidence level, $1 - \alpha$, warrants further investigation. How, for example, does one choose α? If left to anxious and distrustful risk managers, α might become infinitesimally small leading to astronomical levels of economic capital. Delegating this task to the marketing team, conversely, might very well

[19] Confidence level is admittedly not a great name, but the idea is to describe the degree of confidence that a firm's asset losses will not exceed $1 - \alpha$.

create the opposite result. Happily, in practice, α is *not* set by either party. The actual choice of α will typically depend on the current, or desired, credit quality of the firm. A firm with a AAA rating—or one that seeks to be considered among the ranks of AAA entities—will typically have a very low level of α. Such firms, to warrant the strength of their rating, are expected to have a combination of asset quality and capital adequacy to weather a very severe financial storm. At a one-year horizon for AAA entities, for example, it is common to set $\alpha = 0.01\%$ leading to an intimidating 99.99% confidence interval. Indeed, the level of α can be considered to roughly correspond to the firm's probability of default over the time horizon under examination. This intuitive approach thus suggests very different levels of α for AAA or BBB or CCC firms.[20]

We have already discussed the importance of comparing economic capital to a firm's actual equity position, let's now make it a bit more precise. We define a firm's capital headroom as,

$$\text{Firm's Capital Headroom} = \underbrace{\text{Asset Portfolio Unexpected Loss}}_{\text{Economic Capital}} - \text{Firm's Equity Position.} \tag{1.2}$$

The economic capital is the unexpected loss associated with a $1-\alpha$ confidence level; it is an alternative, economically motivated representation of a firm's equity position informed by asset riskiness. The difference between a firm's actual equity position and its economic-capital estimate is highly informative. Headroom tells us a number of things. Its sign describes whether a firm is adequately capitalized for its asset risks or not. The magnitude of an institution's headroom speaks to its capacity to take new risks. Overly large or small levels of headroom will raise questions. In the former case, one might worry about the efficiency—and also potentially profitability—of the firm. In the latter, fears will arise about risk-taking capacity and a firm's ability to maintain their current credit rating.[21]

Our notion of economic capital as assigning a number, between zero and one, to each individual asset is represented in Eq. 1.1. That said, it isn't quite complete; Eq. 1.1 computes this number at the portfolio level. The portfolio level is useful and interesting for high-level oversight by regulators and credit analysts, but day-to-day management of economic capital requires more granularity. What we seek is

[20] Capital adequacy is a key—although certainly not the only—element in a credit agency's decision process. All else equal, for example, a downgrade is likely imminent (but not certain) for a AAA-rated firm that finds itself with AA levels of equity.

[21] If the headroom is large in either direction, one may be able adjust the confidence level, $1 - \alpha$, to generate a manageable, positive headroom level. When possible, and abstracting from other important determinants, one could impute the firm's implied credit rating associated with their equity position.

something more like:

$$\text{An Individual Asset's } \textit{Number} \approx \frac{\text{Individual Asset's Unexpected Loss}}{\text{Individual Asset's Value}}. \quad (1.3)$$

A closer look reveals, of course, that to get these numbers, we need some assessment of the individual asset's unexpected loss. Our economic-capital number and unexpected loss are thus two ways of describing the same underlying concept; the reader is invited to employ the perspective that feels most intuitive to them.

Equation 1.3 falls into the category of things that are relatively easy to write down, but rather harder to accomplish. There are, at least, *two* complicating factors:

1. identifying and measuring the unexpected losses associated with each individual asset; and
2. taking into account this asset-level risk within the context of the overall portfolio.

Both of these aspects are absolutely central. Different assets are naturally subject to varying types of risk. Our economic-capital measure would be incomplete if it did not incorporate these main risks. The second point is a bit more subtle, but equally important. Assets cannot generally be considered in isolation. A swap contract, for example, may be used to offset the interest- or exchange-rate risk associated with a bond. A loan in one sector or region may actually diversify a firm's overall lending portfolio.

The good news is that this chapter does not immediately attempt to provide specific technical answers to these thorny practical challenges. Instead, we will try to introduce a number of concepts that should help us motivate, evaluate, and understand the nature of these technical choices investigated in later chapters.

Colour and Commentary 3 (ECONOMIC CAPITAL IS JUST A NUMBER)*: Classic, and technically correct, definitions of economic capital will focus on the unexpected element—defined at a specific level of confidence—of the firm's asset-loss distribution. This can be difficult to digest; an alternative conceptual notion might help. Extracting all of the complexity associated with models and assumptions from the discussion, it can be useful and illuminating to think about economic capital as a number between zero and one. Zero corresponds to an absence of risk, whereas one implies a complete lack of asset value. As a general rule, the larger the number, the riskier the assets. Whether this number is 0.1, 0.12, or 0.20, when it is multiplied by the firm's assets, the economic-capital value is revealed. Analytically, we can take this a step further and imagine that each individual asset is assigned its own economic-capital number. Practically, such a figure can be constructed by dividing the asset's worst-case (unexpected) loss—with a* $1 - \alpha$ *confidence*

(continued)

Colour and Commentary 3 (continued)
level—by the current asset value. This collection of numbers provides not only a recipe for computing overall economic capital, it also allows management to slice-and-dice *the economic capital by different dimensions. Moreover, one can compare the relative riskiness of different assets within one's portfolio simply by using their assigned number. The individual asset-level description of economic capital is referred to as risk attribution (or allocation) and it is an invaluable tool in the management of economic capital.*

1.4 Categorizing Risk

Computation of economic capital is thus, in its essence, a risk-measurement problem. There are, sadly, many different types of risk. Indeed, classifying and constructing lists of different types of risk is a popular activity among risk managers. Borrowing the notion from the physical sciences—most particularly the study of biology—such lists are often referred to as risk taxonomies. Figure 1.4 provides a visualization of such a risk taxonomy for the individual risks embedded in a typical financial institution's assets.

The usual suspects make an appearance in our taxonomy: market, default and migration, counterparty, and operational risks.[22] We will have rather more to say about the similarities and differences between these risks in the coming pages. Capital buffers are, by contrast, rather specific to the calculation of economic capital. These are, as the name suggests, extra allocations towards the economic-capital calculation to account for uncertainty in its measurement. The core intention, which originated with financial regulators, is to ensure an abundance of caution into one's assessment of a firm's capital position.

There is, of course, more to the capital-buffer story than simply a regulatory recommendation to wear both a belt and suspenders. The rationale and reason for the presence of capital buffers is tightly linked to *two* extreme viewpoints to be considered in the parametrization of one's risk models. As highlighted in Fig. 1.3, risk requires working with asset-loss distributions. In other words, it is necessary to formulate forecasts of the range of possible asset values one's portfolio might take. The nature of our asset-loss distributions, or our range of estimates regarding the future, are based on parameters. These parameter choices depend importantly on the information available to us at the time of the calculation.[23]

---

[22] McNeil et al. [30] provide a nice description of financial risks for the reader seeking more colour and detail.

[23] We have much more to say on the question of parameter selection in Chap. 3.

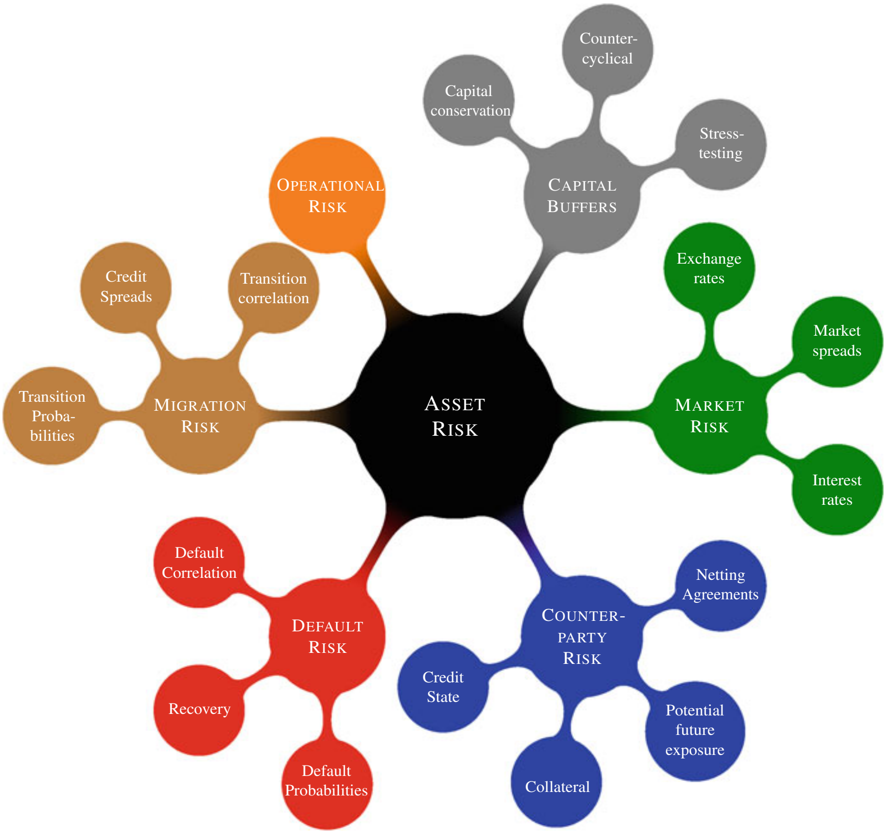

Fig. 1.4 *A simple risk taxonomy*: The preceding figure describes, for a typical financial institution, a set of major risks that typically make an appearance in economic-capital calculations. Attached to each type of risk, where appropriate, are some key risk drivers.

Information and conditionality are two closely related topics. Imagine you were told that you needed to estimate the height of a 16-year-old girl. Absent any additional information, you would likely provide the average height for young women, of age 16, across the entire population. This is referred to as an unconditional estimate; you have limited information. Imagine further, however, that you were told this young woman played competitive basketball. You would likely *condition* upon (or take into account) this new information, and might revise your estimate upwards. This is called, unsurprisingly, a *conditional* estimate.[24] The quality of one's estimate is a separate question. The distinction, in this case, is information.

[24] If you also knew her position—an example of further conditioning information—you might further adjust your guess.

In a risk-management setting, we may use the same ideas to forecast our asset-portfolio losses. It is a bit more artificial, but the basic viewpoint remains the same. In the first case, we can construct an unconditional estimate by using model parameters, estimated over long time periods, to determine our risk metrics. If we have used a multi-decade period to determine our parameters, we are thus essentially acting as if the risk measure could have occurred at any time over this period. No specific time information, to belabour the point, is assumed in this case. In practical circles, this is referred to as a *through-the-cycle* approach. The reason is that the parameter estimates are assumed to incorporate information across multiple economic cycles (i.e., a long stretch of time). Such a forecast is probably a sensible average estimate, over a long time period, but could be quite poor at any specific moment.[25]

A conditional estimate operates at the opposite end of the spectrum. We would construct our forecast as of today, incorporating the most recent possible information. In other words, we actively decide to incorporate timing information. Our parameter estimates will thus use only a relatively short window of time—say, for example, one year of daily data—and we are likely to place even higher weight on the latest values. This represents a sensible, and generally quite accurate, forecast under current circumstances. It might, however, be quite poor in two years time—or, conversely, when looking three years into the past. Such a forecast is, for fairly obvious reasons, termed a *point-in-time* estimate.

Time is thus the critical piece of conditioning information in the economic-capital setting. It is a relatively well-accepted principle that economic-capital computations should be performed using the through-the-cycle approach. The reasoning is threefold. First of all, economic capital is intended to capture extreme events. Longer datasets are, literally by construction, more likely to capture some extreme observations. Second, point-in-time risk estimates tend to be procyclical. This means that in quiet economic conditions they tend to be low, while they increase rapidly under crisis conditions. The impact of risk-measure procyclicality is rapid increases in economic capital—and consequent model-driven deterioration of a firm's capital adequacy—in crisis conditions. This does not seem to be a particularly positive characteristic in an economic-capital estimator. Finally, there is a general desire to be able to compare economic-capital estimates across time. With the point-in-time approach, changes in economic capital might arise from either changes in the portfolio composition or market conditions. It is difficult to isolate these effects. Overall, therefore, a through-the-cycle estimator is better positioned to capture extremes, avoid procyclicality, and permit intertemporal comparison of a firm's economic-capital position.

Against this backdrop, we may now revisit the idea of a capital buffer. Although the through-the-cycle approach attempts to capture extreme outcomes, it does

[25] Just like using the average population height could fail rather seriously when forecasting the height of our 16-year old basketball player.

have an important long-term average element.[26] Capital buffers attempt to address this potential shortcoming by adding something extra to the economic-capital computation for the possibility of actual outcomes being somewhat worse than what the unconditional, through-the-cycle perspective might suggest. Indeed, in some regulatory environments, the level of these buffers is partially determined in a countercyclical manner. Such a mechanism, in the hands of regulators driving minimum regulatory capital requirements, is an example of a macroprudential tool.[27]

Capital buffers have the benefit of being based on easy-to-compute formulaic representations provided by the regulatory community. The other members of our risk taxonomy highlighted in Fig. 1.4 are not so lucky. They must be computed through the blood, sweat, and tears of mathematical and statistical assumption as well as often complex computational implementation. This might sound discouraging, but for some risk types, the situation is even worse. In the real world, risks can be usefully decomposed into *two* disparate groups:

- those that we can measure with some degree of confidence; and
- those that we cannot.

Our usual suspects—market, migration, default, and counterparty risks—are, even if it can be at times technical and difficult, bravely estimated with varying degrees of accuracy. Capital buffers, given their macroprudential perspective, are rather hard to handle. Since responsibility for this component has been co-opted by our friendly neighbourhood regulator, this is happily *not* our problem. This leaves us with operational risk; it is the canonical example of a risk that we cannot reliably estimate. Vast in scope, operational risk simultaneously encompasses loss-generating events such as innocent data-entry errors, holes in one's IT-security firewall, fraudulent activity, and so-called *acts of God* such as hurricanes, earthquakes, or pandemics. Each of these elements are difficult to understand and measure on a standalone basis; their joint examination quickly becomes a quantitative analyst's nightmare. The consequence is that operational-risk economic capital is typically computed in a highly simplified formulaic manner.[28]

The remainder of this chapter, and indeed the entire book, will focus on those risks that we can actually measure. Credit risks—migration, default, and counterparty—will take centre stage. Our discussion of the important market-risk dimension will remain at the conceptual level.[29] This is not to say that difficult-to-

---

[26] The ability to capture extremes is affected by parameter selection, but certainly goes beyond this choice. The mathematical and statistical structure of the model also impact this dimension in important ways.

[27] See ECB [14] for more background on this key aspect of financial-stability policy.

[28] The typical approach is to use some notion of the firm's size as a proxy for its exposure to operational risks.

[29] The reader is referred to Bolder [7] or Jorion [20] for much more discussion on the measurement of market risk.

measure risks—such of those of the operational flavour—are not important. On the contrary, they are central.[30] They nonetheless involve rather less internal modelling and computation, placing them somewhat outside of the scope of this discussion.

Colour and Commentary 4 (RISK CONDITIONALITY)*: The nature of the data used to parametrize one's risk models has, rather unsurprisingly, an important influence on the result. This notion can be made more precise. Long-horizon, relatively low frequency datasets yield unconditional—or long-term average—estimates. Risk measures based on such data are referred to as* through-the-cycle *estimators. Short-horizon, higher frequency datasets lead to conditional (or local) estimates; the resulting risk measures are termed* point-in-time *metrics. The treatment of time is thus the critical difference between these alternative approaches. Both perspectives—and the many possible choices between these* two *extremes—are perfectly valid. Which approach one chooses will naturally depend on one's objectives. Economic capital tends to use the through-the-cycle lens, because it better captures extremes, avoids procyclicality, and permit intertemporal comparison of a firm's risk position. NIB, following this general practice, also adopts the through-the-cycle perspective. The through-the-cycle viewpoint—also widely employed by the regulatory community—is a key reason for the presence of the capital buffers arising in economic-capital calculations. These buffers thus attempt, to a certain extent, to carefully bring the point-of-time perspective into the analysis.*[a]

[a] The point-in-time angle will not be ignored. It plays a central role in loan-impairment and stress-testing analysis.

1.5 Risk Fundamentals

With the conceptual frame hopefully moderately clear in the reader's mind, it is time to turn our attention to the risk-measurement dimension. As should be abundantly clear, approximating the $1 - \alpha$ unexpected loss for each individual asset—and thereby the entire portfolio—will require some heavy machinery. Our immediate objective is not to describe the machine in detail, but rather outline some of its main design features.

Although certainly not exhaustive, *three* key risk questions will be addressed in this section:

[30] Chapelle [12] is an excellent starting point for the various strategies used by operational risk managers to overcome these challenges.

1. identifying a fundamental characterization of risk;
2. examining the inherent tension between portfolio diversification and concentration; and
3. reviewing *two* alternative approaches for the construction of financial models (or indeed, any type of mathematical model).

The disparate risks in our taxonomy behave in different ways. The first order of business is to gain an understanding of how precisely these risks might differ and what this implies for their measurement. Equally important, we will try to understand the commonalities—through a fundamental characterization of risk—between our disparate risks. The portfolio perspective is another area of correspondence between different risk types. The twin concepts of diversification and concentration provide important insights into the measurement of economic capital.

The final topic in this section relates to some foundational concepts in mathematical modelling. Financial modelling is *not*—to be brutally honest—everyone's favourite topic. Some view modelling with suspicion and distrust, others with open hostility, while many consider it to be a necessary evil.[31] The simple reality, however, is that mathematical models are tools. Their role is to illuminate and inform the decision-taking process, but certainly not to take the decisions themselves. Against this backdrop, discussion of possible modelling choices and their implications for economic-capital estimation should hopefully be welcome.

1.5.1 Two Silly Games

The following exercise will require some patience from the reader. It has—at first blush, at least—no obvious link to finance or economics and is quite certainly not the classic approach towards introducing economic capital. There is, of course, some method to this madness. Hopefully the reader's patience will be rewarded and the process might even prove (moderately) entertaining.

We'd like to try playing *two* rather silly, and entirely fictitious, games of chance. Both games involve simultaneously rolling two (fair) dice as one might when playing a board game or visiting a casino.[32] The result of each roll of the dice has a financial consequence for the participant. The actual outcome depends on the sum of our dice values (or pips), which we will refer to as $\mathcal{S}$.

The rules—and specific results—of each game are summarized in Table 1.1. In the first game, one gains €0.05 if $\mathcal{S}$ is even and experiences a loss of equal and opposite sign when $\mathcal{S}$ is odd. In brief, even you win, odds you lose. The second game is only slightly more nuanced. If $\mathcal{S} = 2$, you lose a whopping €1.75, whereas

[31] A small minority, conversely, view mathematical models as the answer to all of our problems. For the quantitative analyst, this is more disturbing than scepticism and wariness, because it suggests a critical misunderstanding of the role of these tools.

[32] Each die is of the six-sided traditional variety, where the dots on each side are referred to as *pips*.

Table 1.1 *The rules*: This table highlights the simple rules associated with our two silly games. The final outcome of each game is determined by the sum of the sides of two rolled dice, which is denoted as $\mathcal{S}$.

Game #1	Game #2
If $\mathcal{S}$ is even, you win €0.05	If $\mathcal{S} > 2$, you win €0.05
If $\mathcal{S}$ is odd, you lose €0.05	If $\mathcal{S} = 2$, you lose €1.75

Table 1.2 *Game results*: The adjacent table provides a complete description of the full set of outcomes and likelihoods associated with a single realization of our two silly games. It also provides profit-and-loss results along with the expected financial receipt.

	Dice outcomes		Game P&L		Expected P&L	
$\mathcal{S}$	Count	Prob.	#1	#2	#1	#2
2	1	0.03	0.05	**−1.75**	0.001	**−0.049**
3	2	0.06	**−0.05**	0.05	**−0.003**	0.003
4	3	0.08	0.05	0.05	0.004	0.004
5	4	0.11	**−0.05**	0.05	**−0.006**	0.006
6	5	0.14	0.05	0.05	0.007	0.007
7	6	0.17	**−0.05**	0.05	**−0.008**	0.008
8	5	0.14	0.05	0.05	0.007	0.007
9	4	0.11	**−0.05**	0.05	**−0.006**	0.006
10	3	0.08	0.05	0.05	0.004	0.004
11	2	0.06	**−0.05**	0.05	**−0.003**	0.003
12	1	0.03	0.05	0.05	0.001	0.001
Total	36	1.00	–	–	0.000	0.000

in all other dice outcomes, a small gain of €0.05 is achieved. In short, with *snake eyes* you lose big, in all other cases you win small.[33]

Playing each game a single time is probably fairly uninteresting, but imagine playing either game 25 times in a row. The final financial result will then simply be the sum of these 25 repetitions. Which game would you prefer to play? The profit-and-loss profiles of each game are quite different. One's choice will certainly depend on individual preferences and assessment of the relative odds of success.

There are, it turns out, no secrets about the throwing of two dice. Table 1.2 unpacks the full details of our two silly games. With two six-sided dice, there are 36 different possible combinations that might be obtained. There are only, however, 11 possible outcomes for the sum of both sides (i.e., $\mathcal{S}$). Some values of $\mathcal{S}$—such as two or 12—can only occur in a single way. Others—such as six, seven, and eight—arise through multiple variations of dice throws. Six possible values of $\mathcal{S}$ are even, while only five are odd. Adding up the probability of the odd and even outcomes, however, we find that they are equally likely.

Table 1.2 also helpfully displays the expected profit-and-loss associated with a single repetition of our silly games. This is simply, following from basic probability

[33] Rolling a pair of dice with a sum of two is colourfully referred to as snake eyes in the game of *craps*. This rather odd gambling activity incidentally seems to appear in a surprising number of Hollywood movies.

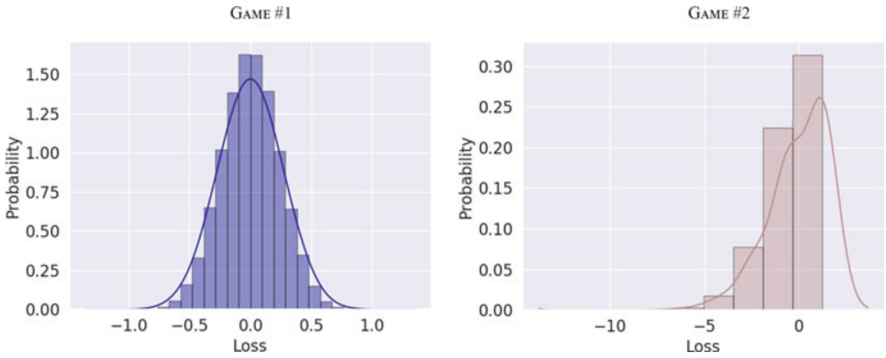

Fig. 1.5 *Silly game pay-off profiles*: The two preceding graphics outline—across tens of thousands of simulated realizations—the distribution of the pay-off profiles associated with our two silly games.

and the discreteness of our sample space, the sum of the probability weighted outcomes.[34] To be honest, the game pay-offs were specifically selected so that expected profit-and-loss value of each game is equivalent; in short, these games were cooked. On average, there is *no* differences between these games.

Perhaps the reader, in light of this new information, has revised her choice of which game to play. After all, it seems that, on average, one should be indifferent between playing these two silly games. There is, nevertheless, ample scope *not* to be entirely detached. Even if the expected outcome agrees between these two games, they do not have the same *risk* characteristics. A single roll of the dice can create a loss of €1.75 in the second game. This admittedly occurs with a small (i.e., 3%) probability. To obtain such a loss in the first game, however, one would need 35 consecutive odd rolls of the dice! While not physically impossible, the probability of such an event is vanishingly small.[35] There is thus a dramatic difference between our two silly games.

Figure 1.5 helps us visualize the precise nature of these differences by displaying the pay-off distribution for our silly games. These graphics were produced by simulating, on a computer, hundreds of thousands of realizations of 25 consecutive repetitions of each game. Both are centred around zero; this is where the similarities end. The first game's profit-and-loss profile is symmetric—that is, there is an equally probability of gain and loss. Moreover, the most extreme gains or losses do not appear to exceed about €1. Our second game's loss distribution is, by contrast,

[34] It suffices to show the expected outcome for a single game, because each repetition of our silly games is independent. The pay-off for n identical repetitions of our silly games is, therefore, merely the sum of their n expected pay-outs. See Casella and Berger [11] for more background on the notion of statistical expectation.

[35] Specifically, it is $0.5^{35} = \frac{1}{34{,}359{,}738{,}368}$, which is effectively zero. Many popular lotteries, to provide some context, actually offer significantly higher probabilities of success. See Zabrocki [49] for a Canadian example.

highly asymmetric. The negative side of the distribution extends much further than the profit side. Worst-case losses appear—albeit with very low probability—to exceed €10. The best-case profit outcome is less than half of the worst-case loss. Simply put, the second game offers more potential for reward and, of course, risk. After examining Fig. 1.5, the reader might be forgiven for deciding to change her choice of game.

Why might we care about these silly games? You certainly did not decide to read this chapter for a lengthy diatribe about games of chance. The reason these games matter is that their pay-offs correspond—at least, in a stylistic sense—to *two* central characters in our risk taxonomy: market and credit risk. Market risk is, at least approximately, symmetric. Interest and exchanges rates or market-spread movements tend to have roughly equal probability of rising or falling. These market-risk factors can, and do, change in large, sudden swings. Most of the time, however, their movements are the sum of many, small cumulative up and down steps. Mechanically, this is the pay-off profile of the first game.

Credit risk is a different animal. The bond investor or lender earns a typically small credit spread or lending margin. Very occasionally, a credit obligor defaults and fails to meet part, or all, of its contractual obligation. In these rare cases, the loss outcome is very large. As a result, credit risk is occasionally described as picking up nickels in the dirt in front of a steamroller. One earns, with apparently low risk, a steady stream of small returns. In some rare cases, one slips and dismayingly meets the business end of the steamroller. This is, again approximately, a description of our second game's pay-off distribution. It also, fairly graphically, illustrates the inherent motivation for banks to have a firm handle on the credit risks embedded in their lending portfolios.

Figure 1.6 displays—to permit comparison with Fig. 1.5—illustrative market- and credit-loss distributions. Loss-centric risk managers have a tendency to write losses—unlike in the analysis of our two games—as a positive number. This

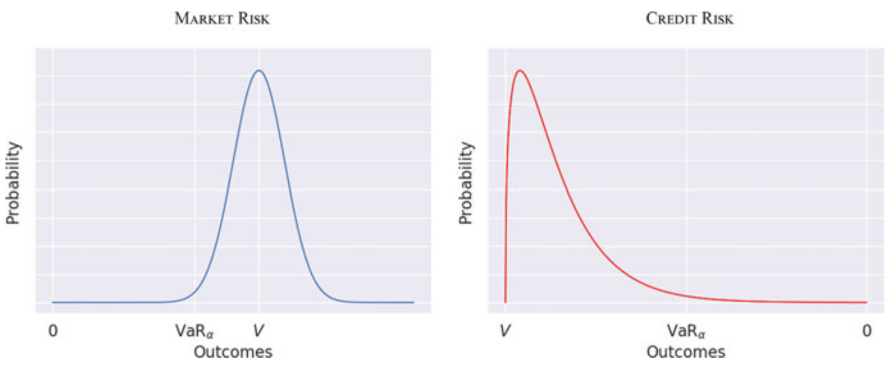

Fig. 1.6 *Market- and credit-loss distributions*: These figures display illustrative market- and credit-loss distributions. Loss, following common practice, is displayed as a positive number. These distributions are, in a stylized way, roughly consistent with our silly game pay-off profiles from Fig. 1.5.

convention makes no practical difference; it simply leads to the classical right-skewed credit-loss distribution.

While our two dice-rolling games are certainly silly, the results are surprisingly consistent with the differences between credit and market risk. The relative outcomes and magnitudes are naturally somewhat exaggerated in the context of our silly games. The negative financial outcome in the second game has a probability of $\frac{1}{36} \approx 0.03$. In our current portfolio, on average, the default probability is about $\frac{60}{10{,}000} = 0.006$; for the highest quality credit countparties, it is only a fraction of this average amount. At the same time, the probability of market-risk gains and losses are *not*—at any given point in time—precisely equal. Furthermore, large jumps in market-risk factors can, and do, occur with some regularity. These practical differences notwithstanding, our two silly games can be considered as a useful analogy for these two risk archetypes. Credit and market-risk arise from *two* rather different investment strategies or, if you prefer, games.

As useful as the analogy posed by our simple games might be, the reality of financial-market risk is dramatically more complicated and nuanced. The set of possible financial outcomes, in our two silly games, is known and discrete.[36] We may count and assign a probability to every possible dice throw. Actual asset values, sadly, do not work this way. A virtual infinity of outcomes are possible, and in most cases, the assignment of probabilities is a subjective exercise. Even worse, we do *not* even know what it is that we do not know. In other words, in the real world there are possible, but perhaps very improbable, outcomes that we have not even thought to consider. This is Taleb [44]'s infamous black swan.[37] This unsettling situation is increasingly referred to as Knightian uncertainty after the influential publication, Knight [25]. This does not imply that our silly game analogies are without use, it merely counsels caution and humbleness in the face of the task ahead of us.

Understanding this distinction highlighted by our silly games, nonetheless has a number of benefits. The fundamental structure of market and credit risk is different. Market risk is comprised of large number of small positive, and negative, outcomes. Credit risk is dominated by a few large, but inherently rare, events. It thus makes complete sense that we, as risk managers, need to model and measure these divergent risks using distinct tools. With that important point in mind, characterization of both risks starts from the same place. It is necessary to identify the underlying economic factors that drive risk. This brings us to a fundamental decomposition of risk that impacts, you will not be surprised to hear, market and credit risk in alternative ways.

[36] Our ability to perfectly understand the associated sample space is a principal reason why games of chance are so beloved by statisticians and probabilists.

[37] Or, if you prefer another colour of swan, see Bolton et al. [10].

Colour and Commentary 5 (RISK ARCHETYPES): *We introduced—to motivate how varying risks in our taxonomy might arise—two silly games involving the throwing of a pair of dice. The financial results of playing either game, in expectation at least, are equal. Their risk profiles differ dramatically. The first game—an allegory for market risk—involves a symmetric loss distribution constructed from the sum of many small positive and negative underlying drivers. Credit risk is encapsulated in the second game. It involves a modest positive return in most states of the world, but with the potential for a large, low-probability loss lurking in the background. This is intended to reflect the mechanics of default and the associated highly skewed credit-loss distribution. The morale of these silly games is that key asset-value risks arise in very different ways and, as a consequence, we require distinct strategies to estimate their contribution to our economic capital figures.*

1.5.2 A Fundamental Characterization

The value of asset portfolios, for a variety of reasons, rise and fall over time. The underlying drivers of these changes in asset valuation are legion. One of the main objectives of a risk taxonomy is to classify and organize these so-called risk drivers. Our two silly games take this a step further by trying to stylistically describe the characteristics of different kinds of risk. These efforts are still lacking something concrete. They do *not* help us understand *why* certain asset values change differently from others. To address this point, we may decompose these asset-value movements into two (very) broad categories:

1. general, or *systemic*, market movements that impact the value of all (or most) assets; and
2. *idiosyncratic*, or specific, movements relating to a single (or only a few) assets.

This amounts to a categorization of changes in asset valuations into common and specific risk poles. This idea is far from new. It lies at the heart of the capital asset pricing model (CAPM) and arbitrage pricing theory (APT).[38]

Some concrete examples are useful. When the general level of interest rates rises in an economy, it affects many assets. Some might even argue that it impacts *all* assets. This is a case of a systemic risk. Similar situations occur with exchange rates,

[38] See, for more background on these classical although still very relevant financial theories, Treynor [45], Sharpe [40, 41], Lintner [26], and Ross [39].

commodity prices, or key macroeconomic variables such as inflation and output.[39] These quantities are correspondingly referred to as systemic risk factors.

On the other hand, some asset-value effects are associated with a specific entity. A given company—due to poor governance structures, failure of an important project, or simply bad luck—might run into financial difficulty. The impact of such financial distress, and associated asset-value movements, is typically restricted to the company in question. We refer to these firm-specific elements as idiosyncratic risk factors.

How might we make use this risk decomposition? It is very tempting to place market risk into the systemic category and assign credit risk to the idiosyncratic side. Market risk is, after all, largely about examining the effect of systemic factors on one's portfolio: interest rate, exchange rate, credit-spread, basis-swap, or market-volatility movements. Moreover, finance theory tells us that idiosyncratic market risk is, with a moderate amount of effort and organization, readily diversified away. As such, it is not compensated, or priced, in financial markets. Credit risk—often linked to a firm's liquidity, solvency, capital, and competitive positions—matches fairly closely to our firm-specific notion of idiosyncratic risk.

There is a significant amount of truth to this unequivocal, binary allocation of market and credit risk into these systemic and idiosyncratic buckets. Unfortunately, it is not quite true. Some element of what is typically classified as market risk can be considered to be quite idiosyncratic. Credit spreads do move in concert with other similar credits, but some (often variable and difficult-to-estimate) part of that spread remains firm specific. Moreover, bond yields are importantly influenced by security specific factors surrounding market-microstructure and liquidity.[40]

Credit risk is, of course, a bit more complicated. Much of what we think of as credit risk is clearly idiosyncratic. Firms have a strong tendency to run into trouble as a consequence of internal firm-specific choices. Nevertheless, firms do not operate in a vacuum. They have competitors who exert an important influence on their financial situation. They are also tied into a network of suppliers and clients. Competitive actions or the financial health of key members in their operating network can, and do, impact their fate. These important latter effects cannot be easily categorized as firm specific or idiosyncratic. Defaults are thus logically correlated through general economic conditions and industry interdependencies. Credit risk has—like its cousin, market risk—elements of both dimensions.

This latter point is the key message of Fig. 1.7. Neither systemic nor idiosyncratic are—despite their important conceptual differences—completely distinct. The corollary is that market and credit risk cannot be allocated, in a wholesale

[39] One can, and probably should, argue that it is actually the macroeconomic variables who are driving the financial outcomes.

[40] See Amihud et al. [1] for thoughts on liquidity and O'Hara [36] for an excellent entry point into the theory of market microstructure.

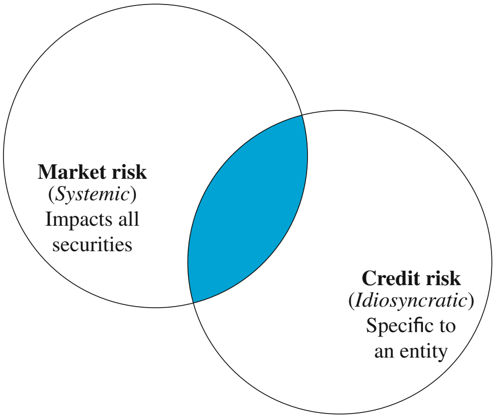

Fig. 1.7 *Independent risks?*: This schematic is designed to help the reader appreciate the interplay between systemic and idiosyncratic risks. While a useful distinction, these two notions of risk are not truly independent. Market and credit risk accordingly inherit this overlap.

manner, to these broad categories. As we will see in the following section, this has important implications. The fundamental distinction relates back to a centrally important concept in both finance and economic-capital estimation: concentration and diversification.

Colour and Commentary 6 (A BROAD RISK CATEGORIZATION)*: It is informative and useful to break the risk associated with asset-value movements into two broad categories: systemic and idiosyncratic. Systemic risk stems from underlying factors whose movements—naturally, in different ways—affect all assets. Idiosyncratic risk is specific and relates only to the individual firm. Conventional wisdom often seeks to categorize market risk as systemic and credit risk as idiosyncratic. There is some truth in this arrangement, but it is not quite correct. What is more, it is wrong in important ways. Idiosyncratic risk, in financial theory, is not rewarded. This is because, with sufficient effort, it can be diversified away. Following this logic, were credit risk to be entirely idiosyncratic, it could be safely ignored. This latter statement is patently false, implying the incorrectness of treating credit risk as an entirely idiosyncratic phenomena.*[a] *Indeed, investigation of the systemic nature of credit risk yields important insight into a key element of economic-capital calculations: portfolio concentration.*

[a] The implications of treating market risk as entirely systemic, however, are rather less severe.

1.5.3 Introducing Concentration

Almost everyone who has ever taken an introductory course in finance has performed a diversification exercise. It can take a variety of forms, but it essentially amounts to indicating how the variance of portfolio returns (i.e., risk) can be reduced by adding non perfectly correlated assets. This concept, first introduced by Markowitz [28], is the central idea behind diversification. There is, of course, a limit to the benefits of diversification. Systemic risk, which impacts all assets, simply cannot be fully diversified away.

Idiosyncratic risk is another matter. It can, and should, be eliminated through diversification. Credit and market risk, however, vary importantly along this perspective. The lion's share of diversification benefits in an equity portfolio appear to be achievable with a few dozen stocks. Replicating a sovereign-bond benchmark—that might include hundreds of individual instruments—can be achieved with only a handful of bond holdings. In short, the market-risk component of a fixed-income portfolio is readily diversified. Statements suggesting that idiosyncratic risk is not compensated are made, almost invariably, in a market-risk context.

Default- or credit-risk diversification turns out to be a bit more slippery. Why might this be the case? Our silly games can help provide us some insight. A market-risk loss represents, in most cases, only a small proportion of the instrument's value. Remember, market-risk movements are typically the sum of numerous, relatively small market movements. Multiple securities, being buffeted in different directions by the vagaries of market-risk factors, have many more opportunities to offset each others' gains and losses. Credit-risk losses, while rare, are also more sizable; in the worst case, you can lose the entire investment. Such rare events are not readily offset by other asset-price movements.

The consequence of this difference is that credit risk is significantly more sensitive—than its market-risk counterpart—to the composition of one's asset portfolio. In contrast to the market-risk setting, true diversification in credit-risk portfolios can require hundreds or even thousands of individual securities. Even then, it is necessary that no one single exposure dominates the portfolio.

This latter claim might be hard for some readers to accept. Even those who buy this idea might benefit from a demonstration. Let us, therefore, consider a concrete, motivational example. Imagine that we construct *two* portfolios where everything—with one important exception—is identical. The total values of the portfolios are equal and set to €1000.[41] Every instrument further shares the same probability of default at $\frac{50}{10{,}000} \approx 0.005$. Moreover, in the event of default, we assume that only 60% of the exposure's value will be recovered.

Despite these striking similarities, there is a key distinction. One portfolio has only 100 individual risk exposures of varying size; some are quite large, while others are relatively small. The other portfolio has 10,000 individual exposures.

[41] You may multiply this value by whatever value you wish to make it a bit more realistic.

Table 1.3 *Competing portfolios*: What are the differences, in terms of credit risk, associated with diversified and concentrated portfolios? One way to find out is to construct two portfolios that differ solely along the concentration dimension. The underlying table summarizes two such example portfolios.

Measure	Diversified	Concentrated
Total portfolio size	€1000	
Number of positions	10,000	100
Mean position size	€0.10	€10.00
Minimum position size	€0.10	€0.08
Maximum position size	€0.10	€47.80
Default probability	50 bps.	
Loss-given default	40.00%	
Confidence level	99.97%	

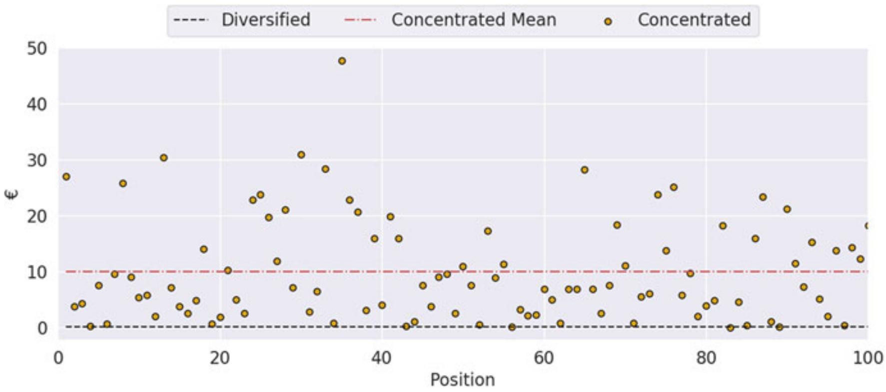

Fig. 1.8 *Sample exposure profiles*: This graphic shows the distribution of exposure sizes associated with the diversified and concentrated portfolios introduced in Table 1.3. The distinction is visually striking.

Each exposure takes a value of $\frac{1}{10{,}000}$. To put it more directly, one set of exposures is diversified, while the other is rather concentrated. The key summary statistics associated with these two portfolios are summarized, for convenience, in Table 1.3.

Figure 1.8 helps out by illustrating the contrasting distributions of exposure sizes associated with our two sample portfolios. The diversified portfolio exposures are not terribly interesting; they all share the same value. The same is not true for our concentrated portfolio. One very large exposure accounts for almost 5% of the total portfolio. There is another handful of exposures in the 2 to 3% range. Others appear to be very close to zero.

One might legitimately ask if this is another cooked example. The simple answer is yes. At the same time, however, it is not so different from some real-world portfolios. The concentrated example might represent a corporate-bond book or

Table 1.4 *Concentration results*: Here we observe the differences, in terms of idiosyncratic and systemic risk written in percentage terms of our portfolio value, for our competing portfolios. Systemic risks differ very little, whereas the idiosyncratic dimension diverges dramatically.

	Diversified	Concentrated
Idiosyncratic	0.1%	3.1%
+ Systemic	7.6%	9.0%

a commercial lending portfolio.[42] Such portfolios typically include a relatively modest number of heterogeneous exposures; note, however, that these 100 positions are already substantially larger than what is typically required for a comfortable degree of diversification in the market-risk environment. The diversified case, conversely, might be a large credit-card or student-loan portfolio. It is comprised of many relatively homogeneous positions. It is extremely unlikely, of course, that either of these cases would involve identical default probabilities and loss-given-default assumptions. This aspect was definitely, and unapologetically, manipulated to focus our attention entirely on the concentration dimension.

The interesting question after having introduced these two portfolios, of course, is what is the impact on each portfolio's credit-risk economic capital calculation? To answer this query, we naturally need to compute their economic capital. Abstracting from the technical details, we use two alternative models. One, fairly unrealistic, model assumes that all credit risk is idiosyncratic in nature. The other includes both idiosyncratic and systemic elements in a sensible manner.[43] The utility of these two distinct models is to permit us to assess the impact of both systemic and idiosyncratic risks on our economic-capital estimates.

Table 1.4 summarizes the results of our experiment. The individual values are provided in percentage terms. As one would have expected, the idiosyncratic risk for the diversified portfolio approaches zero. The 10,000 position portfolio has thus succeeded in diversifying away idiosyncratic risk. The concentrated portfolio has not. Roughly one third of the total risk, for this 100 exposure portfolio, remains idiosyncratic. The full benefits of diversification thus remain at some distance. When we move to consider the full credit-risk model in the second line of Table 1.4, we see that the total risk of these two portfolios are surprisingly similar. The diversified portfolio is naturally less risky, but not dramatically so. Figure 1.9 strikingly visualizes this point.

There are, at least, three takeaways from this exercise. First, systemic risk simply cannot be diversified. This is hardly new—it is, in fact, a central lesson from the CAPM model—but it is reassuring to see it resurface in this context. It is

[42] In actuality, these would involve more positions, but stylistically they are similar.

[43] For the interested reader, the idiosyncratic choice is the binomial independent-default model, while the second (systemic and idiosyncratic) approach is a one-factor Gaussian threshold model. See Bolder [8, Chapters 2 and 4] for more background on these models. We will also revisit these general approaches, within the NIB setting, in Chap. 2.

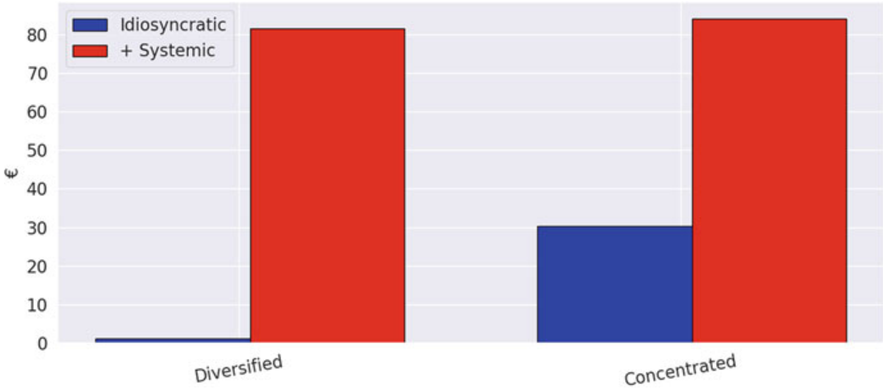

Fig. 1.9 *Achieving diversification*: This figure provides, in currency terms, the idiosyncratic and systemic risk results for the two distinct portfolios introduced in Table 1.3. Total risk is not dramatically dissimilar, but the idiosyncratic elements differ. Concentration clearly matters.

furthermore hard work, as a second point, to diversify away idiosyncratic risk in the case of credit risk. 100 heterogeneous exposures does not even bring us close. Finally, the systemic aspect is definitely *not* restricted to the world of market risk. Indeed, our example and Table 1.4 clearly illustrate the fundamental importance of systemic factors in the estimation of credit-risk economic capital.

In this fabricated example, we have only examined the simplest type of concentration: that related to exposure size. Credit obligations also differ along a host of other dimensions such as their geographic region, industrial sector, credit category, default-recovery assumptions, or even firm size. The more similarities, within a portfolio, among these aspects, the greater the potential for concentration.[44] Diversification in a credit-risk setting is thus even more difficult than what is suggested in this example.

Concentration and diversification are opposite side of the same coin. Both are intimately related to the fundamental notions of idiosyncratic and systemic risks. Appreciating this fact is an important step towards understanding models of economic capital. A key responsibility of such models is to appropriately combine these two dimensions to describe our asset-loss distributions. For the measurement of market risk, this involves leaning strongly (or even completely) in the systemic direction. Credit risk, conversely, requires a careful hand to manage the interactions between idiosyncrasy and diversification on the one hand and systemic issues and concentration on the other.

[44] Concentration is consequently a core theme in the measurement of credit risk. See Lütkebohmert [27] for much more on strategies and techniques to manage it.

Colour and Commentary 7 (ONE COIN, TWO SIDES)*: A cornerstone of finance theory holds that idiosyncratic risk, which is readily diversified, is not compensated. While true in a market risk sense, the narrative is somewhat more complicated in the credit-risk situation. A simple, stylized example of a concentrated and diversified portfolio illustrates that elimination of idiosyncratic risk is not trivial task for credit-risk portfolios. Hundreds, if not thousands, of positions are necessary. Moreover, an eye must be also be trained on other possible sources of concentration such as geographical location and industry. Asset riskiness, and thus estimation of economic capital, depends importantly on portfolio composition. This analysis indicates that the structure of one's portfolio also has different implications depending on the type of risk under examination.*

1.5.4 Modelling 101

We now, in our quest to better understand economic-capital calculations, turn our attention to financial modelling. Let us begin with a wise statement from Holland [19]:

> Model building is the art of selecting those aspects of a process that are relevant to the question being asked.

To those less accustomed to working with mathematical models, it might be a bit startling to read about *selection* in this context. A bit of reflection, however, reveals that it is inevitable. Models are a simplified representation of a complex, indeed unknowable, reality. In the physical sciences, our understanding may be superior to that in the social realm, but it still remains far from perfect.[45] We simply cannot create models that encapsulate all of the complexity of the real world; choices must be made. Holland [19]'s basic point is that it's more important, in principle, to appreciate the problem and the conceptual logic behind one's modelling choice than the fine print of the mathematical details.[46]

This reasoning will help to explain why so much attention has been placed on the nature of the risks associated with a firm's assets. It also suggests that this (first) conversation about economic-capital models will focus, as intimated by Holland [19], on the selection of those relevant aspects. The technical details will be thoroughly covered in subsequent chapters.

---

[45] Compare this to our two silly games, where all possibilities and probabilities can be characterized to perfection.

[46] Naturally, the details are also important, but they quickly become irrelevant if high-level decisions are poorly taken.

There are many types and flavours of mathematical models, but one overriding classification can be helpful. The key distinction surrounds *why* certain events, which one is attempting to model, occur. One class of models—which we will refer to as *structural*—attempts to capture some aspects of the underlying mechanism driving key events. Imagine, for example, building a model of monetary policy. A structural approach would describe the logical linkages between central-bank rate decisions, the monetary transmission mechanism, and inflation outcomes. Structural models describe, in mathematical terms, how key variables *should* interact. This requires some theoretical foundations upon which the model can build. The alternative approach is data driven. Placing less weight on theoretical considerations, one uses statistical techniques to extract empirical relationships from the data. These are termed *reduced-form* models. Image and voice recognition software are excellent examples of reduced-form models; they use massive amounts of data to train algorithms to recognize certain combinations of events. The *why* aspect is not central.

Ultimately, all models have parameters. These parameters are, almost invariably, estimated from real-world data. One might justifiably enquire, in this context, as to the real difference between structural and reduced-form models. The key distinction is *causality*. A structural model attempts to describe causal relationships between variables within the real-world data-generating process. Reduced-form models are more agnostic about causality, they attempt to identify empirical relationships embedded in the data.

The difference may appear subtle and perhaps abstract, but it is a useful magnifying glass through which to examine and evaluate models. Neither approach, for example, is dominant. Both have their own strengths and weaknesses. Reduced-form models lean heavily on the data; they are, as a consequence, fairly sensitive to non-representative or erroneous inputs. They are also particularly vulnerable when predicting outcomes outside of one's data range. Structural models place a greater importance on theory. Should the theoretical relationships be flawed, these shortcomings will be inherited by the model. In most estimations of economic capital, however, both modelling techniques make an appearance.

Estimation of market-risk, to be more concrete, is typically performed with a reduced-form model. The value of financial instruments—be they loans, deposits, bonds, or swaps—practically depends on a host of underlying market-risk factors. Different financial instruments, however, have varying degrees of responsiveness to change in these factors. Some instruments might be very sensitive to exchange rates, whereas others are invariant to them. A common, and clever, approach to characterizing market risk thus involves writing

$$\text{Asset-value movements} = \sum_{k=1}^{K} \text{Factor Sensitivity}_k \cdot \text{Factor change}_k, \tag{1.4}$$

where we have K distinct market-risk factors. In words, the contribution of various market-risk factors are assumed to be additive. Each individual market-risk factor

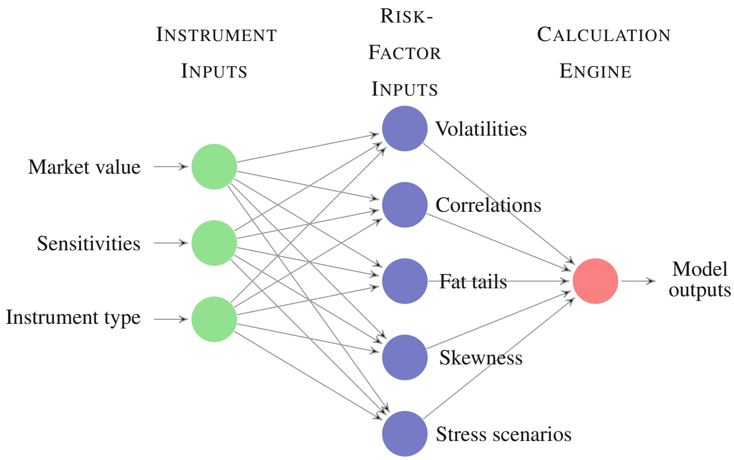

Fig. 1.10 *A market-risk model*: The preceding schematic organizes the instrument- and risk-factor-level inputs into the typical market-risk calculation. The market-risk factor system is normally described using a data-driven reduced-form model.

contribution is simply the product of the instrument's factor sensitivity and the factor's movement.[47] A critical aspect of measuring market risk, therefore, involves the construction and management of these sensitivities.

Equation 1.4 has a convenient form, but it says nothing about the sign, magnitude, and interaction of these market-risk factor changes. Getting a handle on market risk will inevitably involve tackling this aspect. The reduced-form assumption enters at this point. Market-risk models rarely attempt to describe why interest rates, exchange rates, and key market spreads interact. Instead, they collect large amounts of market data including daily, weekly, or monthly observations of all important market-risk factors. Risk estimates are then based on estimation of the joint statistical properties of this data, or alternatively by extracting risk-factor movements from a period of extreme market turmoil.[48] In either case, the role of this complex system of market-risk factors is informed by observed, empirical data relationships.

Figure 1.10 provides a high-level schematic of the main ingredients involved in a typical market-risk calculation. Key inputs are profitably separated into *two* groups: those related to the financial instruments and those linked to the market-risk factors. Collecting and organizing the instrument sensitivities is hard work, but most of these

[47] See Bolder [7, Part I] for a more detailed discussion of these ideas.

[48] This is an almost criminally brief description. See Bolder [7, Part IV], and subsequent chapters, for more specifics.

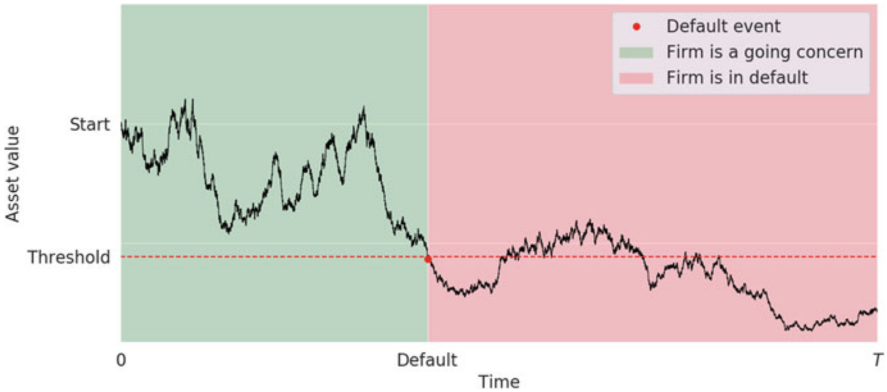

Fig. 1.11 *A structural relationship*: Merton [31]'s key insight was to identify default as occurring when a firm's asset values fall below some predefined—typically liability related—threshold. This is a classical example of a theoretically motivated structural model.

values are readily available through internal systems. The art and science of market-risk modelling relate to the risk-factor system. By and large, this aspect is captured through a data-driven reduced-form model.

Market-risk measurement is thus predominately reduced-form and focused on systemic market risk factors.[49] What about credit risk? Measuring credit risk, the reader may not be surprised to learn, usually involves a different strategy relative to the market-risk case. NIB's credit-risk economic-capital model—similar to many other institutions—follows a structural approach.[50] The key theoretical insight dates back to Merton [31], who offered a clever, and intuitive, description of the underlying *cause* of firm default. It begins, as one would hope, with the firm's balance sheet. The insight is that default occurs when the firm's equity position hits zero; in other words, the asset value is equal to, or larger, than the firm's commitments. Using the language of the previous sections, the firm's supply of capital is exhausted. This observation yields a road map for modelling credit risk.

Figure 1.11 provides a useful visualization of Merton [31]'s theoretical conclusion. The firm's assets move over time in a random fashion. If, at any point, they fall below some predefined threshold, then default is assumed to have occurred. This threshold value needs to be determined, but it basically reduces to some function of the firm's liabilities and the firm's overall financial position. This idea, as is the case in all structural models, effectively endogenizes a key aspect of the phenomena one is trying to model. That is, default is determined inside the model through the interaction of other key variables.

[49] The canonical references in this area are Morgan/Reuters [32] and Jorion [21].

[50] There is, however, an entire class of reduced-form credit risk models. See, for example, Bolder [8, Chapter 3] for several concrete examples.

Merton [31]'s insight, while incredibly useful, still requires some polishing to make it a model. The classical structure continuous-time in Merton [31] is now rarely used, but has been supplanted with a number of practical points introduced by Vasicek [46, 47, 48] and Gupton et al. [17, 18]. Without getting to deeply into the details—covered amply in subsequent chapters—one asset-related quantity is particularly helpful in understanding the inner workings of this model. Each credit obligor is assigned a latent state variable closely related to its asset value. It is comprised of the following *two* familiar components:

$$\text{Asset-value index} = \text{Systemic element} + \text{Idiosyncratic element}. \tag{1.5}$$

This requires some unpacking. Each firm's asset-value index is the sum of a systemic and idiosyncratic component. The systemic factor impacts all firm assets; this is the mechanism driving default dependence in the model. The idiosyncratic piece is specific to the individual firm. The genius of Eq. 1.5 is that it explicitly incorporates our fundamental risk characterization.

The individual components of Eq. 1.5 are random and are assumed to be independent.[51] Equation 1.5 gives us what we need to build a proper model. Mechanically, we can think of the credit event as being determined by a number of random draws. We pick a common systemic value out of a hat; it is shared by all credit obligors. We then pull a separate idiosyncratic factor for each credit obligor out of a long line of hats and then proceed to construct our asset-value indices. There is a credit event for each credit counterpart with the following generic form:

$$\text{Credit Event} = \begin{cases} \text{Default} : \text{Asset-value index} \leq \text{Threshold} \\ \text{Survival} : \underbrace{\text{Asset-value index} > \text{Threshold}}_{\text{Otherwise}} \end{cases}. \tag{1.6}$$

Armed with a practical definition of each firm's credit event, the rest is (copious) mathematical and statistical detail. Unlike the market-risk setting, however, realistic implementations of Eqs. 1.5 and 1.6 involve some fairly heavy computations. Drawing random numbers from hats may be relatively easy on a computer, but extremely large numbers must be drawn from an almost breathtaking array of hats to obtain a reasonable degree of accuracy.[52] We will touch on the implementation challenges in Chap. 4.

Figure 1.12—constructed analogously to Fig. 1.11—organizes the instrument- and portfolio-level inputs into a visualization of the credit-risk calculation. The simulation engine, which is responsible for all of the random-number generation, requires many inputs. The first group relates to the individual instruments while the second involves instrument attributes central to systemic-risk considerations. As previously discussed, these instrument features help the model characterize the

[51] This basically means that the randomness of the systemic component is determined completely separately from the idiosyncratic element.

[52] There are probably enough hats involved to make even Britain's Queen Elizabeth II blush.

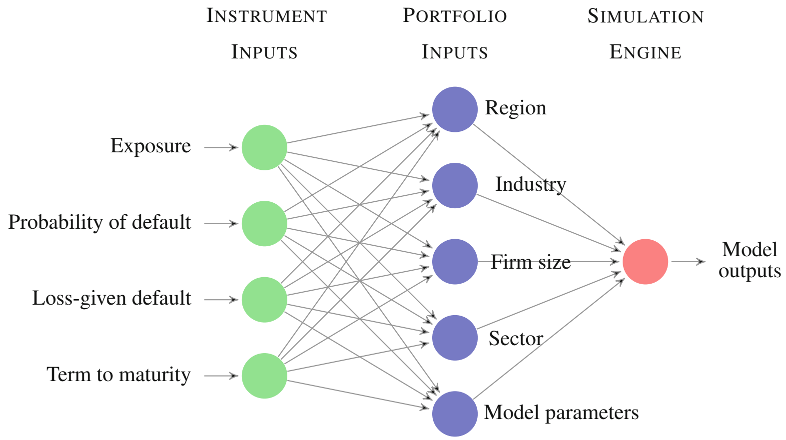

Fig. 1.12 *A credit-risk model*: This schematic organizes the instrument- and portfolio-level inputs into a visualization of the credit-risk calculation. The instrument and portfolio data fields relate principally to the idiosyncratic and systemic dimensions, respectively. The simulation engine combines these inputs using a structural modelling approach.

degree of portfolio concentration. The simulation engine then collects and combines these inputs using a structural modelling approach to generate economic-capital estimates.

Colour and Commentary 8 (HANDLING CAUSALITY)*: In this motivational introductory chapter, detailed technical model explanations and derivations are misplaced. A serious discussion of economic-capital estimation would be nonetheless incomplete without addressing the modelling dimension. At a high level, models are profitably characterized by how they handle causality. Structural models rely on theory to endogenize key aspects of the model; in this sense, they attempt to assign causality. Reduced-form models, conversely, rely on empirical data to exogenously describe these key relationships. No approach is, in a general sense, superior to the other and many practical implementations make comfortable use of both philosophies. Market-risk economic-capital model, in general, falls into the reduced-form camp; complex interactions of systemic market-risk factors are typically informed exogenously by empirical data. Credit-risk calculations, by contrast, are often structurally motivated. The structural, or endogenous, default element is theoretically determined through an (indirect) examination of each firm's balance sheet thereby linking back to the idea of capital supply and demand. This approach also easily incorporates both the systemic and idiosyncratic elements so important to credit risk.*

1.6 Managing Models

We have, in this discussion, only scratched the surface of the intricacies involved in computing credit-risk economic capital. We have followed Holland [19]'s advice and highlighted some of the relevant aspects. No responsible modelling discussion, however, would be complete without a few words on model governance. Computation of economic capital, in many organizations, is based on internal models. In the literature, these are typically referred to as decision-support models.[53] Decision-support models can be found in medical, agricultural, logistical, and business applications. Interesting and relevant examples might include a medical diagnosis tool, an early-warning tornado system, mergers-and-acquisition decision analysis, or selecting the optimal product mix in a pulp-and-paper mill.

Such models are typically developed internally when, as in the case of economic capital, their development and understanding has strategic importance. With the explicit introduction of economic capital into NIB's statutes, this hurdle would seem to have been comfortably cleared. Simply because a model is strategically important and merits internal development, however, does not imply a free pass. On the contrary, its criticality warrants sustained and careful oversight. The possibility that models are themselves a source of risk—as introduced by Derman [13] and Rebonato [37]—should not be taken lightly. Such model risk (perhaps) need not be directly incorporated into our economic-capital computations, but it definitely needs to be managed.

Managing model risk requires a governance structure. Any internal model—whether for decision support, regulatory oversight, or for policy making—is an ongoing work in progress. It consequently involves a continuous cycle of validation, analysis and improvement, verification, implementation, and usage. It never really ends; Fig. 1.13 reflects this fact with a circular schematic.

Managing model risk essentially amounts to an exercise in quality control.[54] As a consequence, much focus in model risk is assigned to the model-validation process. While the importance of independent, external validation cannot be denied, internal model developers and owners also have a number of additional responsibilities. These include the use of benchmark models, careful computer-code and system management, production of detailed diagnostics and sensitivity analysis, and ongoing stress testing. These are not part-time activities, but central aspects of a well-functioning internal modelling environment.[55] As a consequence, these ideas will show up repeatedly in the following chapters.

One additional model-owner responsibility towers over all others: documentation and transparency. Quality control is virtually impossible without clarity surrounding model choices and assumptions. Scientific discourse, which has proven rather effec-

[53] See Keen [22, 23] and Keen and Morton [24] for the origins of this interesting academic literature.

[54] See Bolder [9] for a more detailed argumentation of this point and its implications.

[55] This, in turn, ties back to the quantitative analysis axioms introduced in the preface.

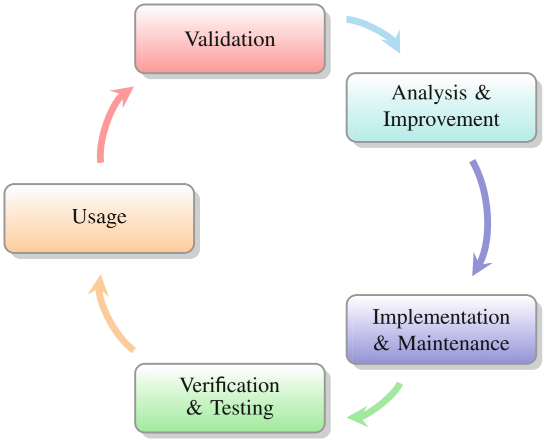

Fig. 1.13 *Model-improvement cycle*: Any model, to be of practical usefulness, requires ongoing review, validation, and improvement. It also needs to be used. This cycle implies that no working model is a static entity, but rather a dynamic one.

tive in recent centuries, slavishly follows this principle. Scientific ideas are, literally by construction, forced to run the gauntlet of peer review and publication. Modelling ideas are also strengthened and improved through their explicit documentation and exposure to external consideration and critique. This reasoning is the driving force behind the production of this book and argues for its dissemination to the widest possible audience.[56]

Colour and Commentary 9 (MODEL RISK): *Economic-capital computations depend critically on modelling assumptions and choices. Given their strategic importance, these models are typically internally developed in the form of decision-support systems. Perhaps somewhat ironically, these models are themselves a source of risk. Classified as model risk in our taxonomy, there are no current requirements to add it to economic capital. Common sense and sound business practice, however, argue strongly for extensive governance measures surrounding these models. This can take a range of forms. Model validation receives the most attention—and is certainly very important—but model owners also have numerous additional responsibilities. The most important, and perhaps least well appreciated, element relates to documentation and transparency. Carefully writing down one's methodologies is a valuable discipline. Model quality is moreover directly proportional to its degree of exposure to external critique. These reflections were the catalyst for the production of this book and the justification for its broad dissemination to an external audience.*

[56] Again, our axioms *"write it down"* and *"seek external criticism"* are clearly reflected in these actions.

Table 1.5 *Economic-capital results*: This table illustrates NIB's economic-capital figures—from the 2020 annual report in NIB [35, pp. 21–23]—by their principal risk dimensions. All figures are denominated in EUR millions.

Dimension	Subtotal	Total
Default risk	1534	
Migration risk	487	
Credit risk		2021
Market risk		567
Operational risk		101
Conservation buffer	306	
Countercyclical buffer	–	
Stress-test buffer	120	
Capital buffers		426
Economic capital (i.e., capital demand)		3115
Headroom		697
Adjusted common equity (i.e., capital supply)		3812

1.7 NIB's Portfolio

Remaining at the conceptual level can only take us so far, because economic capital is a concrete quantity. To complete this introductory chapter, therefore, it is interesting and useful to examine various perspectives on actual NIB economic-capital computations. Table 1.5 thus provides a bird's eye view illustrating the entire economic-capital computation of roughly EUR 3.1 billion as the end of 2020.[57] It also distinctly includes the various members of our risk taxonomy.[58]

The values from Table 1.5 are drawn from the 2020 NIB annual report. Other institutions provide similar disclosures both in their financial statements and Pillar 3 regulatory reports.[59] Although not a regulated entity, the spirit of the computations from Table 1.5 is consistent with fundamental regulatory principles and best practice.

A few useful conclusions can be drawn. Credit risk, through either default or migration, is the dominant risk in NIB's portfolio. This is the case for literally all lending institutions. It accounts, through NIB's lending and investment activities, for almost two thirds of the total. Moreover, the credit-risk component is multifaceted including both a default and a credit-migration dimension. As will be made much more precise in Chap. 2, default risk relates to an obligor not repaying their obligations whereas migration touches on the capital impact associated with a deterioration of borrower credit quality. Both elements are important, although the

[57] These values are found in NIB [35, pp. 21–23].

[58] Counterparty risk is embedded in the default and credit-migration figures.

[59] BIS [6] rather carefully outlines a publication template for the contents and format of key economic-capital metrics.

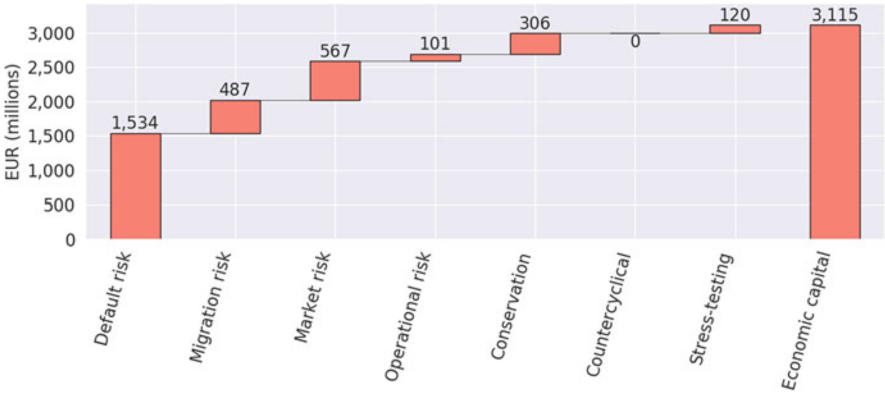

Fig. 1.14 *The waterfall graphic*: The preceding figure illustrates the main results from Table 1.5 in graphic format; it is often referred to as a waterfall graphic. It underscores the central importance of the credit-risk dimension.

default aspect dominates NIB's portfolio. The size and complexity of the credit-risk aspect is the principal justification for the disproportionate amount of time and energy invested, in the following chapters, into the credit-risk dimension. There is simply much to discuss and analyze; this fact is further underscored visually by Fig. 1.14.

A second point is that Fig. 1.14 provides an opportunity to perform a rapid analysis of NIB's capital adequacy position. In short, it looks to be fairly healthy. The headroom is between 15 to 20% of the equity position providing scope to further grow lending and treasury activities. From NIB [35, pp. 25], we may also read the total NIB asset position as about EUR 35.4 billion. Using this fact and Table 1.5, we may also proceed to compute our economic-capital numbers. For the entire portfolio, it is about $\frac{3.1}{35.4} \approx 0.088$ or 8.8% if we represent it as a percentage. This suggests that, on average, about 9% of each asset's value needs to be set aside for their riskiness. A separate number can also usefully be computed, depending on how far one wishes to go, for taxonomy type, sub-portfolio, or risk factor. The credit and market-risk numbers, for example, amount to about 5.7% and 1.6%, respectively.

NIB's credit-risk economic-capital model is, as is common practice, based on a long-term, unconditional, through-the-cycle approach. To correct for this aspect, among other things, the regulatory community has introduced the notion of capital buffers. These represent an additional cushion to account for variations in current economic conditions. NIB, although not a regulated entity, also follows this regulatory practice. The buffers themselves come in three flavours: conservation, countercyclical, and stress-test. The conservation buffer is the closest thing to a pure through-the-cycle adjustment and is simply a linear multiple of the minimum regulatory capital requirements.[60] The countercyclical buffer, conversely, is an

[60] Again, we will turn our attention to the regulatory perspective in Chap. 11.

example of a macro-prudential tool. It is set by various jurisdictions—also as a percentage of regulatory capital—against the flow of current macroeconomic conditions. In other words, regulatory authorities build it up during good times, only to release it during period of economic downturn or crisis. This explains why it takes a value of zero in Table 1.5; the vast majority of authorities released the countercyclical buffer during March and April of 2020 at the inception of the global COVID-19 pandemic.

The final piece of the puzzle in Table 1.5 is the so-called stress-testing buffer. This is an additional contribution to overall economic capital, which is intended to capture the impact of severe adverse macro-financial conditions. It is basically another layer of prudence. It can feel somewhat excessive to add yet another safeguard to a collection of worst-case risk measures and regulator-specified buffers. While there is truth to this reflection, this buffer stems from a separate stress-testing exercise. There is value, and the potential for significant insight into the strengths and weaknesses of one's asset portfolio, in such an analysis. For regulated entities, the level of the stress-testing buffer is determined via assessment by one's supervisor. At the NIB, the level is specified by Board decision supported by quantitative analysis. A much more detailed discussion of this aspect is found in Chap. 12.[61]

Colour and Commentary 10 (ASSESSING CAPITAL ADEQUACY)*: Financial institutions are constantly asked by key stakeholders—such as regulators, credit agencies, investors, borrowers, and indeed, themselves—a very simple question: do they have enough capital? The fancy term for attempts to answer this question is capital-adequacy analysis. It is not a simple question to answer, because it depends on the multidimensional risks faced by the firm's assets. Economic capital, as a concept, was developed to help by providing a lens through which one can holistically examine the riskiness of these assets. The preceding analysis of NIB's capital adequacy as of December 2020 is a good example of the various elements at play. Similar analyses are performed across literally thousands of financial institutions across the world. In the coming chapters, we will delve much more deeply into the mathematical structures and assumptions necessary to generate these high-level figures. Given that there is no shortage of details, a rather high amount of care needs to be given towards the organization of this discussion.*

[61] Treatment of stress testing in the final chapter of this work should not be seen as a commentary on its importance. It shows up last simply because, to adequately discuss it, one needs to have a firm global understanding of all the various elements.

1.8 Looking Forward

Having supplied some conceptual insight into the idea of economic capital and looked more concretely at the NIB situation, the final task in this chapter is to lay out the structure of the remaining discussion. Our stated objective is to provide a reasonably complete description of a framework for the measurement and practical application of credit-risk economic capital. A framework suggests multiple interdependencies between various components; this needs to be reflected in the organization of our exposition. Writing such a book is thus like building a conceptual house. Each chapter represents an important element of the overall construction; moreover, the chapters depend upon on another in a sequential fashion. It is difficult, for example, to discuss loan-pricing issues or parameter selection without a good understanding of the economic-capital model. Figure 1.15 correspondingly displays the forthcoming 11 chapters from the bottom upwards as building blocks.

Part I is the foundation of our conceptual house. It lays out the methodological details associated with our credit-risk economic-capital model. This covers *three* separate chapters: the model itself, the determination of model parameters, and details associated with the implementation. Following from the axioms introduced in the preface, although NIB has a single production model, two additional challenger models are computed on a daily basis.[62] This is to permit an ongoing point of comparison (or sanity check) for interpretation, trouble-shooting, and communication. The order of these three chapters is quite important. The base methodology is the natural starting point in Chap. 2. Since models are not particularly useful without parameters, Chap. 3 immediately addresses this question following the specification of the modelling details. Finally, model estimation involves stochastic simulation and, as a consequence, it is computationally intensive. Strategies for managing this complexity form a central part of Chap. 4.

Part II addresses the important economic-capital application of loan pricing. Loan pricing needs to be performed in a quick and flexible manner to permit loan originators to consider multiple financing alternatives with their clients. Any sensible assessment of loan pricing requires economic-capital inputs, which are the result of a reasonably slow and complex simulation procedure. This presents a fundamental problem. The classic solution—adopted in many institutions—involves use of an approximation for the economic-capital consumption associated with a specific loan. Chapter 5 is, therefore, dedicated to the construction of a (semi-)closed-form approximation of default and migration risk. This turns out to be both a worthwhile investment and a helpful tool, since we will also make liberal use of it in our loan-impairment and stress-testing calculations. Chapter 6 then turns our attention to the conceptual and practical details of loan pricing. This touches not

[62] This a direct application of our quantitative analysis axiom *"if you can help it, never do anything just one way."*

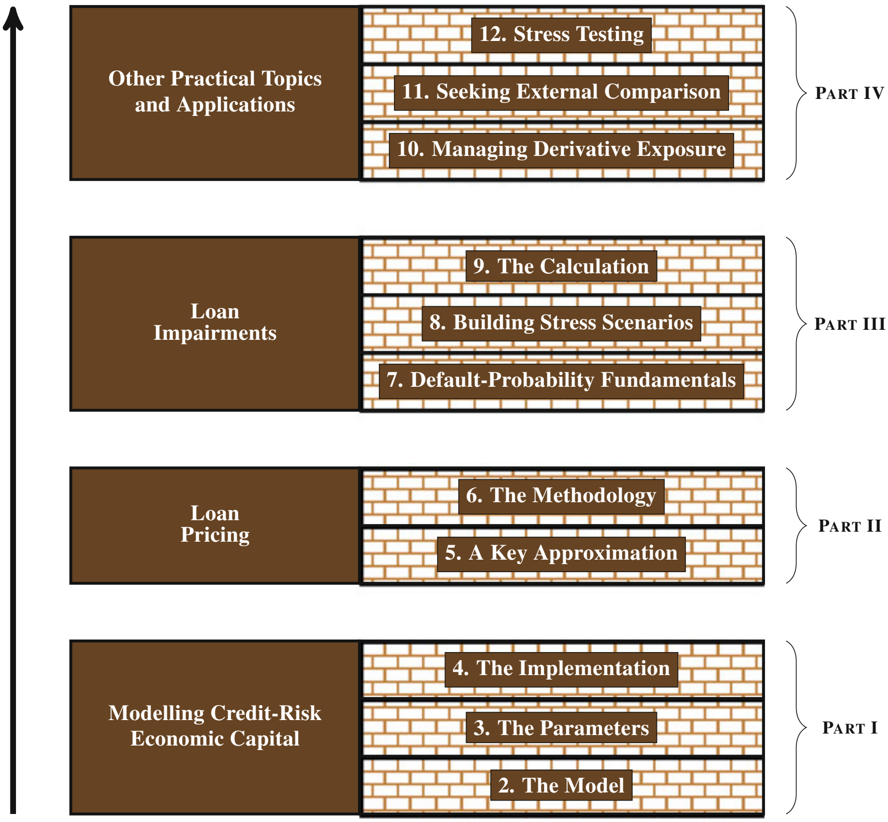

Fig. 1.15 *Our chapter plan*: This schematic describes the organization of the forthcoming 11 chapters in this book. The overall credit-risk economic-capital framework—including models and applications—are organized into four logical sections. We have represented them from the bottom upwards as building blocks; each chapter thus represents a key step towards the subsequent discussion.

only on many key aspects of economic capital, but also extends into important ideas in corporate finance.

Part III of this book focuses on another application of economic capital: loan impairments. At first glance, the link to economic capital might seem tenuous. Loan impairments are an accounting measure that typically ignore the tail of the loss distribution, but instead focus on the expectation. Indeed, this is why loan impairments are also widely referred to as *expected* credit losses. There are, however, at least *three* conceptual linkages between economic capital and loan impairments. The first relates to the structure of the International Financial Reporting Standard #9 (IFRS 9). Launched in January 2018, IFRS 9 introduced an important forward-looking dimension into the computation of expected credit losses. Future default-probability outcomes need to be predicated on forecasts of key macro-financial variables. The machinery constructed to manage this difficult task is also employed in economic-capital stress-testing analysis. The heavy-lifting associated with the construction of these forward-looking scenarios is found in Chaps. 7 and 8.

The second link between economic capital and loan impairments stems from the role of loan impairments in capital supply. An increase in loan impairments will flow through a firm's profit-and-loss statement, which ultimately reduces its equity position. All else equal, this will lead to a decrease in economic-capital headroom. Often the reasons for an increase in loan impairments—greater risk in the loan portfolio—will also generate an increase in credit-risk economic capital further reducing the firm's headroom. This squeezing of headroom between capital supply and demand is a central feature of stress-testing analysis. Excluding loan impairments from this discussion would thus mean missing out on an important aspect of stress testing. The third, and final, link relates to an inherent risk of commercial lending stemming from portfolio composition. Its forward-looking aspect notwithstanding, the IFRS 9 expected-loss computation does not explicitly capture one's portfolio composition. For a reasonably concentrated and high-quality credit portfolio—as found in many commercial-lending institutions (including NIB)—the failure to capture the concentration element can lead to a systemic underestimation (i.e., downward bias) of expected loss. This suggests the use of a so-called concentration or portfolio-composition adjustment, which makes use of the credit-risk economic-capital model. The combination of these three conceptual linkages implies (in its most general form) that loan impairments and economic capital are, in fact, intertwined activities. The details of the expected credit loss calculation—including a proposed adjustment for portfolio composition—are covered in Chap. 9.

Part IV, the final section of this book, deals with a collection of (loosely related) practical topics and applications. Chapters 10 and 11, for example, both have their feet firmly anchored in the regulatory world. Chapter 10 addresses the fascinating and subtle realm of counterparty credit risk. Modern banks need to make use of derivative contracts—swaps, forwards, and in some cases, options—to manage their positions. Unlike loans, bonds, or deposits, the computation of credit exposures—a key input into the economic-capital calculation—is a challenge. The difficulty stems principally from the fact that, at any given point in time, a derivative contract may be either an asset or a liability to the firm. This status, of course, has rather important implications for credit risk. We thus dedicate the totality of Chap. 10 to a discussion of the (flexible and helpful) regulatory approach for the prudent, risk-based estimation of derivative exposures. Chapter 11 tackles a common challenge for an international institution. As a non-regulated entity, its actions are not prescribed by a supervisor. At the same time, regulatory guidance and rating agencies provide an essential external point of comparison for internal risk-management activities. Chapter 11 examines, therefore, multiple regulatory approaches and both Pillar I and II calculations as well as a well-known rating-agency methodology. It describes NIB's continuous effort to identify and evaluate external benchmarks. This is not, however, an NIB-specific undertaking; all institutions benefit from constructive external comparison. Chapter 12 concludes, building on much of the previous discussion, with a consideration of the stress-testing dimension. With the nominal goal of motivating the stress-testing buffer from Table 1.5, it also addresses the incorporation of severely adverse scenarios as well as the idea of reverse stress testing.

Colour and Commentary 11 (A CONCEPTUAL HOUSE)*: This work seeks to provide a reasonably complete description of a framework for the measurement and practical application of credit-risk economic capital. The very word, framework, suggests multiple pieces and interdependencies. The various elements of a framework cannot, unfortunately, be considered in just any order. Writing such a book is thus like building a conceptual house. This means that we need to deal with the basement before working on the roof. The organization of the remaining 11 chapters—organized in* four *distinct parts—attempt to reflect this fact. Part I is the foundation of our house; it deals with the credit-risk economic-capital model, its many parameters, and the implementation challenges stemming for the use of a simulation model. Parts II and III can be viewed as the walls and rooms of our conceptual structure. They focus on two important applications: loan pricing and impairments. In doing so, they also introduce some centrally important technical ideas—approximation of economic capital and macro-financial stress scenarios—for our practical toolbox. Part IV, let's call it the roof of our building, concludes with a collection of practical topics and applications. These are not least in importance—after all, they are keeping the rain off of our heads—but instead are best considered when equipped with a global understanding of the key concepts earned in the previous chapters.*

1.9 Wrapping Up

The principal objective of this chapter was to provide a comprehensive, although gentle, introduction to the computation of economic capital. A few key takeaways bear repeating. First of all, economic capital is basically a worst-case measure of asset riskiness used to complement the *average*, or expected, risk perspective provided in financial statements. The second key point is that economic capital is ultimately—for each individual asset, credit obligor, or sub-portfolio—simply a number between zero and one. An economic-capital number of zero implies a riskless position. Outside of this extreme case, the number increases proportionally with an asset's overall risk. Since asset risks are complex and multifaceted, we need to make use of fairly complex mathematical models to actually compute economic capital. The principal job of these models is to sensibly capture idiosyncratic and systemic effects, the actual portfolio composition, and notions of concentration and diversification. These models themselves, as a final point, also create their own risks. Model governance and oversight are key tools in managing this aspect. Two particularly important model-risk mitigants, documentation and model transparency, are key drivers behind the production of this book.

The forthcoming chapters consider *two* main perspectives: measurement and management of credit-risk economic capital. The economic-capital measurement chapters delve more deeply into methodological modelling questions as well practical parametrization and implementation issues. Other modelling chapters focus on the fast (semi-)closed-form approximation of economic capital and building stress scenarios consistent with macro-financial forecasts. This is the foundation of our credit-risk economic-capital framework. The management-focused chapters principally address key applications of economic capital: loan pricing, calculation of loan impairments, and stress testing. These aspects build on our modelling foundation and provide helpful tools to assist in guiding the lending organization.

References

1. Y. Amihud, H. Mendelsohn, and L. P. Pedersen. Liquidity and asset prices. *Foundations and Trends in Finance*, 1(4):269–364, 2005.
2. B. M. Bendis. Ultimate Spider-Man. (1):1–52, 2000.
3. BIS. The internal ratings-based approach. Technical report, Bank for International Settlements, January 2001.
4. BIS. International convergence of capital measurement and capital standards: A revised framework. Technical report, Bank for International Settlements, January 2004.
5. BIS. An explanatory note on the Basel II IRB risk weight functions. Technical report, Bank for International Settlements, July 2005.
6. BIS. Pillar 3 disclosure requirements–updated framework. Technical report, Bank for International Settlements, December 2018.
7. D. J. Bolder. *Fixed Income Portfolio Analytics: A Practical Guide to Implementing, Monitoring and Understanding Fixed-Income Portfolios*. Springer, Heidelberg, Germany, 2015.
8. D. J. Bolder. *Credit-Risk Modelling: Theoretical Foundations, Diagnostic Tools, Practical Examples, and Numerical Recipes in Python*. Springer-Nature, Heidelberg, Germany, 2018.
9. D. J. Bolder. Stealing Others' Strategies: Lessons on Model Risk. In *Evolving Practices in Public Investment Management: Proceedings of the Seventh Public Investors Conference*, 2020.
10. P. Bolton, M. Despres, L. A. P. da Silva, F. Samama, and R. Svartzman. *The Green Swan: Central Banking and Financial Stability in the Age of Climate Change*. Bank for International Settlements, January 2020.
11. G. Casella and R. L. Berger. *Statistical Inference*. Duxbury Press, Belmont, California, 1990.
12. A. Chapelle. *Operational Risk Management*. John Wiley & Sons, Ltd, 2018.
13. E. Derman. Model risk. Goldman Sachs Quantitative Strategies Research Notes, April 1996.
14. ECB. Macro-Prudential Policies and Tools. Technical report, European Central Bank, June 2010. Financial Stability Review: Section IV.A.
15. ECB. ECB Guide to the Internal Capital Adequacy Assessment Process (ICAAP). Technical report, European Central Bank, November 2018.
16. A. Elmgren. *Open Government in the Nordics*. Aarhus University, Denmark, February 2019.

17. G. M. Gupton, G. Coughlan, and M. Wilson. *CreditMetrics — Technical Document*. Morgan Guaranty Trust Company, New York, 1997.
18. G. M. Gupton, C. C. Finger, and M. Bhatia. *CreditMetrics — Technical Document*. Morgan Guaranty Trust Company, New York, 2007.
19. J. H. Holland. *Hidden Order: How Adaptation Builds Complexity*. Addison-Wesley, New York, NY., 1995.
20. P. Jorion. Risk Management Lessons from Long-Term Capital Management. *European Financial Management*, 6:277–300, 2000.
21. P. Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill Ryerson Limited, New York, NY, 2006.
22. P. G. W. Keen. *Decision Support Systems: A Research Perspective*. Center for Informations Systems Research. Afred P. Sloan School of Management, 1980.
23. P. G. W. Keen. Value Analysis: Justifying Decision Support Systems. *MIS Quarterly*, 5(1): 1–15, 1981.
24. P. G. W. Keen and M. S. S. Morton. *Decision support systems: an organizational perspective*. Addison-Wesley series on decision support. Addison-Wesley Pub. Co., 1978.
25. F. H. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin Co, Boston, MA, 1921.
26. J. Lintner. The valuation of risky asset and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, 1965.
27. E. Lütkebohmert. *Concentration Risk in Credit Portfolios*. Springer Verlag, Berlin, 2008.
28. H. Markowitz. Portfolio Selection. *The Journal of Finance*, 7(1):77–91, 1952.
29. K. Marx, R. Institut Marksizma-Leninizma (Moscow, I. für Marxismus-Leninismus beim ZK der SED., and I. Marx-Engels-Stiftung. *Karl Marx, Friedrich Engels Gesamtausgabe (MEGA): Abt. "Das Kapital" und Vorarbeiten*. Dietz, 2008. Original Publication Dates: 1867–1883.
30. A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Tools and Techniques*. Princeton University Press, New Jersey, 2015.
31. R. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470, 1974.
32. J. Morgan/Reuters. *RiskMetrics — Technical Document*. Morgan Guaranty Trust Company, New York, 1996.
33. S. C. Myers. Determinants of corporate borrowing. *Journal of Financial Economics*, (5): 147–175, 1977.
34. S. C. Myers. Still searching for optimal capital structure. *Journal of Applied Corporate Finance*, 1(6):4–14, 1993.
35. NIB. Nordic Investment Bank Financial Report 2020. URL: https://www.nib.int/filebank/a/1613377638/93a3a266fc787bb3e6dd86c3a2d563d3/10767-NIB_Financial_Report_2020.pdf, February 2020.
36. M. O'Hara. *Market Microstructure Theory*. Blackwell Publishers Cambridge, Massachusetts, 1995.
37. R. Rebonato. Theory and practice of model risk management. Quantitative Research Centre (QUARC) Technical Paper, 2003.
38. C. M. Reinhart and K. S. Rogoff. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, Princeton, New Jersey, 2009.
39. S. A. Ross. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 35(3): 341–360, 1976.
40. W. F. Sharpe. A simplified model for portfolio analysis. *Management Science*, 9(2):277–293, 1963.
41. W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, 1964.
42. P. Siklos. *Money, Banking and Financial Institutions: Canada in the Global Environment*. McGraw-Hill Ryerson Limited, Toronto, Ontario, 1994.
43. A. Smith and A. B. Krueger. *The Wealth of Nations*. Bantam Classics, 2003. Original Publication Date: 1776.

44. N. N. Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House, New York, NY, 2007.
45. J. L. Treynor. Market value, time, and risk. Unpublished Manuscript, August 1961.
46. O. A. Vasicek. Probability of loss on loan distribution. KMV Corporation, February 1987.
47. O. A. Vasicek. Limiting loan loss probability distribution. KMV Corporation, August 1991.
48. O. A. Vasicek. The distribution of loan portfolio value. *Risk*, (12):160–162, December 2002.
49. M. Zabrocki. Calculating the Probabilities of Winning Lotto 6/49. York University: Mathematics and Statistics, March 2003.