1 Introduction

Financial markets let people trade promises of future payments. These payment promises are called financial *assets*. Prominent examples are bonds or company stocks. A bond is a way to borrow money, and it promises its buyer a future debt repayment with interest. Stock is used to raise capital and promises its holder future dividend payments. In addition to those, many other types of assets exist, but at their core they all are tradable contractual claims on future cash flows. Supply and demand determine the prices at which to buy or sell them—the prices at which to *enter positions*.

Now, asset prices and thus the values of the positions in them change over time. This can have fundamental causes, e.g., a company discovering a new drug or a country recovering from a recession, but it can also be due to the market activity itself, as witnessed in the hefty fluctuations of stock prices even on slow-news days. These latter, seemingly random market movements in asset prices constitute *market risk* to positions.

One way to assess this kind of risk is to determine the potential impact of specific market movements on positions' values. Such what-if inquiries often mimic or simulate market shocks and are then called *stress tests*. If the shocks are tiny and standardized, the results are called *sensitivities*; they serve to track and compare the positions' asset exposures.

In the same vein, you can try to estimate the plausibility and impact of future asset price changes from historically observed ones and condense the effect on the positions into some aggregate measure of risk. A prominent measure is the so-called *value-at-risk* or *VaR*, which is a hypothetical daily loss expected to be breached once every hundred days—in other words: *the probability that tomorrow there will be a loss larger than the VaR is* 1%. A sibling of the VaR is the *expected shortfall* or *ES*, which yields the estimated average loss over several worst-case scenarios.

These and similar measures are used to monitor the positions' risk profile, to signal critical market conditions, to limit exposures, and to otherwise meet requirements as prescribed by law and banking regulation authorities. They may

M. Auer, *Hands-On Value-at-Risk and Expected Shortfall*, Management for Professionals, https://doi.org/10.1007/978-3-319-72320-4_1

also determine a bank's crucial *capital requirements*, the amount of capital a bank has to hold given its exposure to the markets, and thus the cost of doing business. The risk numbers and their workings therefore matter to a large audience besides the risk manager and his IT guy: to the trader, the accountant, the compliance officer, the board member, and the regulator.

With the stakeholders many and varied, a risk model should not merely accurately capture risks but also be transparent and easy to explain, avoid recurring modifications and the resulting scrutiny, and operate reliably under real-world duress like imperfect data. We now set out to describe a system designed to achieve these goals.

Part I of this book outlines the basic risk measures and their relations, and it proposes a simple VaR approach (a *filtered historical* one).

Part II describes how to apply the risk measures to common questions about a risk profile's bearings, and it details risk measure properties, time series behavior, and model sanity-checks.

Part III illustrates a possible overall design of a risk system and presents ways to implement this system into software.

Finally, the appendix collects some mathematical foundations, a brief digression on risk-neutral pricing, and links to further reading material.

Before all that, however, the next chapter aims to give an intuitive introduction to how risk can be thought of and compressed into one single number.