Introduction

This book provides a quantitative, technical treatment of portfolio risk analysis with a focus on real-world applications. It is intended for both academic and practitioner audiences, and it draws its inspiration and ideas from both the academic and practitioner research literature. Quantitative modeling for portfolio risk management is an active research field. Virtually all institutional investment management firms use quantitative models as an integral part of their portfolio risk-management procedures. Research and development on these models takes place at investment management firms, brokerage houses, investment consultants, and risk-management software providers. Academic researchers have explored the econometric foundations of portfolio risk analysis models, the relative performance of the various types of models, the implications of the various models for the understanding of capital market behavior, and the models' implications for asset pricing equilibrium. This book attempts to synthesize the academic and practitioner research in this field. We argue that portfolio risk analysis requires a balanced, multidisciplinary perspective combining statistical modeling, finance theory, microeconomics, macroeconomics, and a behavioral–institutional understanding of modern capital markets.

Who Should Read This Book?

Among practitioners, an ideal reader of this book would be someone with a good background in finance and statistics working in the risk-management office of an institutional fund manager. He or she[1] may be relying entirely on vendor software for portfolio risk analysis, or entirely on routines developed in-house, or on a combination of in-house and vendor products. The book is not a how-to manual for building a portfolio risk analysis model, but it gives the reader a solid understanding about the many difficult issues and choices in model design and estimation and about current research frontiers in estimating and evaluating these models. Practitioners working in related functions, including

[1] To avoid unnecessary verbal clutter, in the remainder of the book we use the male pronoun for gender-neutral third-person singular.

portfolio managers, investment consultants, and risk analysis software providers, are also part of our target audience. As an important caveat, those working in the central risk-management offices of investment banks, with responsibility for a full range of trading desks, will find that this book is not comprehensive, since we do not cover risk analysis for financial engineering.

Among academics, an ideal reader of this book would be a graduate student or faculty member interested in portfolio-risk-related research. Many of the best new ideas in academic finance come from the interface with business practice. The book might be suitable as a main or secondary text in an advanced master's level course on portfolio management or financial risk management.

Topics Covered

In this section we discuss some broad topics that are included or excluded from the book and then briefly describe the content of each chapter.

Managerial, Accounting, and Regulatory Issues

The main objective of the modeling techniques described in this book is to provide accurate, reliable portfolio risk analysis and forecasts. Portfolio risk analysis serves as a crucial input into a range of managerial and regulatory decisions, but we do not address, except peripherally, the broader policy problems of financial risk management and regulation. We also do not address accounting and disclosure issues.

Portfolio Risk Analysis versus Portfolio Management

The book is addressed to the problems of portfolio risk analysis rather than the broader problems of portfolio management. There are two main reasons for this. First, the methods and techniques for portfolio risk analysis are distinct from those for portfolio management. Second, almost all institutional managers separate the portfolio management function from the risk analysis function; there are very strong business and regulatory reasons for this separation. Many of the worst scandals in the investment management industry occurred when this functional separation was breached or inadequately monitored. We touch upon portfolio management issues outside the risk analysis domain, such as asset valuation, security selection, and trading strategies, whenever it is illuminating for our main focus.

Portfolio versus Asset Risk Analysis

Consider a set of n asset returns $\boldsymbol{r}$ and a set of portfolio weights $\boldsymbol{w}$. Let r_w denote the portfolio return associated with this particular set of portfolio weights. If we know that the portfolio weights are permanently fixed, then r_w can be treated as if it were a single asset return. We can analyze the risk of r_w without analyzing the numerous risk relationships among all the constituent asset returns. The focus of the book is on *portfolio* risk analysis, by which we mean either that we do not know the portfolio weights $\boldsymbol{w}$ when we build the risk analysis model or we are allowing $\boldsymbol{w}$ to vary. We need to create a risk analysis model for potential portfolios r_w that will be accurate across a wide range of portfolio weight vectors $\boldsymbol{w}$. This requires a risk model of all n asset returns $\boldsymbol{r}$ and their interdependencies, rather than just a model of a particular portfolio return r_w.

Financial Engineering

Financial engineering and derivatives-trading strategies pose especially difficult problems for risk analysis. Financial engineers are willing and able to trade, at an acceptable price, virtually any nonlinear function of any future asset price or price path. This transforms the risk analysis problem into a subdiscipline of financial engineering; the portfolio aspect diminishes in importance. There are many high-quality texts on risk analysis for financial engineering and derivatives trading;[2] we do not attempt to cover this material. Risk managers oriented toward derivatives will find that this book covers only one part of their problem: risk analysis for portfolios of primary assets not including any financially engineered securities.[3] This is an important component of risk analysis for financial engineering, but it is not the whole story.

Limits to Coverage of Research Literature

Although the general perspective in this book is strongly empirical, we generally do not attempt to replicate the vast array of empirical findings on portfolio risk analysis. Instead, we provide references to empirical findings in the research literature and incorporate them into our analysis without attempting to re-derive them from raw data. We provide empirical illustrations when it seems particularly illuminating to do so, to underscore major conceptual points in the analysis or to highlight counterintuitive empirical findings.

[2] Including, for example, Christoffersen (2003), Crouhy et al. (2001), Dowd (2005), and Jorion (2007).

[3] We allow for the simplest types of derivatives overlays such as equity index futures.

Chapter Summaries

Chapter 1 sets out basic measures of return and risk. It compares arithmetic and logarithmic returns and discusses the relative advantages of each. It introduces the key measures of portfolio risk including variance, value-at-risk, and expected shortfall. It discusses the objectives of portfolio risk management and its limitations.

Chapter 2 examines the estimation and use of unstructured return covariance matrices. It discusses the problem of estimation error in covariance matrices and the implications for their use in portfolio management. Chapter 3 examines industry–country component models. In these simple models the cross section of returns is divided into industry-related returns, country-related returns, and asset-specific (neither country- nor industry-related) returns. This simple decomposition is surprisingly powerful in explaining the common components in the cross section of equity returns, and also has growing relevance for corporate bond markets as they broaden and deepen internationally.

Factor models of security returns are typically categorized as statistical, economic, or characteristic-based. Chapter 4 describes factor models of security returns and discusses statistical factor analysis. Chapter 5 deals with macroeconomic factor models in which the pervasive factors in returns are observable economic time series, such as output, inflation, and interest rate changes. Chapter 6 treats characteristic-based factor models, in which the factor sensitivities of assets are tied to the corporate characteristics and/or cash flow characteristics of assets. Not all factor models fit neatly into one of these three categories. We also discuss "hybrid" factor models, which have features from more than one of these pure types. Chapter 7 analyzes foreign exchange risk.

An important design choice in risk model construction is the approach to integration of risk analysis across asset classes and across national borders. This problem of risk model integration is considered in chapter 8. This chapter also surveys research on the level and trend of cross-border capital market integration, which has obvious relevance for the choice of integrated versus segregated risk modeling.

Due to its analytical convenience, the first eight chapters of the book have an emphasis on unconditional portfolio return variance as a risk measure. Chapters 9–14 broaden the perspective, using many other risk metrics besides variance, and explicitly account for risk dynamics. Chapter 9 explores models of dynamic volatility and dynamic correlations, and the choice of forecast horizon. Chapter 10 considers density estimation and the related problem of tail estimation and value-at-risk measures. Chapter 11 discusses credit risk, and chapter 12 liquidity risk.

Chapter 13 looks at risk analysis for alternative asset classes such as hedge funds, venture capital, and commodities. Chapter 14 deals with the performance evaluation of portfolio risk–return realizations and also the performance evaluation of portfolio risk-forecasting models. Chapter 15 provides a brief conclusion.

Useful Background

Understanding the material in the book requires at least an intermediate-level background in statistics, linear algebra, and finance theory. Some sections require more advanced knowledge in statistics or finance theory. For those who wish to refresh their knowledge or study these topics independently, we suggest some appropriate finance and statistics texts.

Greene (2008) provides a solid foundation in statistics and econometrics appropriate for understanding the material covered in this book. Readers wanting a finance-focused treatment at a more advanced level may benefit from reading Campbell et al. (1997). For general finance background, Bodie et al. (2009) and Elton et al. (2010) are two possible references.

Approximations Used in the Book

Statistical approximations are very important in portfolio risk analysis. Diversification is a key principle of portfolio management, and by its nature it relies on statistical approximations. There are also useful approximations as the chosen return interval becomes short or as the sample used for risk model estimation becomes large.

We use a simple common notation for the different types of approximations used in the book. Most of the approximations rely on one of three limiting variables: either the number of assets n, the number of time periods T, or the return measurement interval Δ (monthly, weekly, daily, hourly, etc.). We take the limiting approximation for large n, or large T, or for small Δ, holding all other variables constant. Which of these three limiting variables is being used in the approximation is indicated by a superscript on the approximately equals symbol:

$$\overset{n}{\approx}, \quad \overset{T}{\approx}, \quad \overset{\Delta}{\approx}.$$

If f and g go to zero with Δ, but the difference goes to zero more quickly, in the sense that

$$\frac{f-g}{\Delta} \overset{\Delta}{\approx} 0,$$

we write

$$f \overset{o(\Delta)}{\approx} g,$$

meaning that $f - g$ goes to zero relative to the magnitude, or "order," of Δ. This is useful if we are looking at two returns over a short time interval and want to say that the returns are approximately the same, even though both are approximately zero. The symbol "$\approx$" has the standard definition from introductory calculus; those who received at least a "B" in their introductory calculus course do not need to read this long footnote, while those who received a "C" or worse will not want to, but we include it anyway for completeness.[4]

We also rely on the two basic statistical approximations: limit in probability and limit in distribution. We add the superscript "pr" for limit in probability and "di" for limit in distribution. So, for example, let $\hat{m}$ denote the sample mean from T independent observations of a random variable with a true mean of zero, a variance of one, and finite higher moments. Then

$$\hat{m} \overset{\text{pr},T}{\approx} 0$$

is our notation for the law of large numbers (the sample mean approaches zero in probability) and

$$(T^{1/2})\hat{m} \overset{\text{di},T}{\approx} N(0,1)$$

is our notation for the central limit theorem (the sample mean scaled by the square root of the number of observations is approximately normal in its distribution). Readers not familiar with these standard statistical approximations are referred to Greene (2008, appendix D) or any standard statistics textbook. In a few isolated places we use approximations that do not fit neatly into these simple categories; we use the notation $\overset{*}{\approx}$ in these cases, and give references outside the text.

---

[4] Recall from introductory calculus that

$$f \overset{T}{\approx} a$$

means that for any $\epsilon > 0$ there exists a T^* such that $|f(T) - a| < \epsilon$ for all $T > T^*$;

$$f \overset{n}{\approx} a$$

has the analogous definition with n replacing T as the limiting variable. Approximations based on the return interval differ in that the limiting approximation relies on a suitably small (rather than suitably large) value of the limiting variable:

$$f \overset{\Delta}{\approx} a$$

means that for any $\epsilon > 0$ there exists a δ such that $|f(\Delta) - a| < \epsilon$ for all $\Delta < \delta$.