CHAPTER 1

Introduction

As the use of quantitative techniques has become more widespread in the financial industry, the issues of how to apply financial models most effectively and how to mitigate model and estimation errors have grown in importance. This book discusses some of the major trends and innovations in the management of financial portfolios today, focusing on state-of-the-art robust methodologies for portfolio risk and return estimation, optimization, trading, and general management.

In this chapter, we give an overview of the main topics in the book. We begin by providing a historical outlook of the adoption of quantitative techniques in the financial industry and the factors that have contributed to its growth. We then discuss the central themes of the book in more detail, and give a description of the structure and content of its remaining chapters.

QUANTITATIVE TECHNIQUES IN THE INVESTMENT MANAGEMENT INDUSTRY

Over the last 20 years there has been a tremendous increase in the use of quantitative techniques in the investment management industry. The first applications were in risk management, with models measuring the risk exposure to different sources of risk. Nowadays, quantitative models are considered to be invaluable in all the major areas of investment management, and the list of applications continues to grow: option pricing models for the valuation of complicated derivatives and structured products, econometric techniques for forecasting market returns, automated execution algorithms for efficient trading and transaction cost management, portfolio optimization for asset allocation and financial planning, and statistical techniques for performance measurement and attribution, to name a few.

Today, quantitative finance has evolved into its own discipline—an example thereof is the many university programs and courses being offered in the area in parallel to the "more traditional" finance and MBA programs. Naturally, many different factors have contributed to the tremendous development of the quantitative areas of finance, and it is impossible to list them all. However, the following influences and contributions are especially noteworthy:

- The development of modern financial economics, and the advances in the mathematical and physical sciences.
- The remarkable expansion in computer technology and the invention of the Internet.
- The maturing and growth of the capital markets.

Below, we highlight a few topics from each one of these areas and discuss their impact upon quantitative finance and investment management in general.

Modern Financial Economics and the Mathematical and Physical Sciences

The concepts of portfolio optimization and diversification have been instrumental in the development and understanding of financial markets and financial decision making. The major breakthrough came in 1952 with the publication of Harry Markowitz's theory of portfolio selection.[1] The theory, popularly referred to as *modern portfolio theory,* provided an answer to the fundamental question: How should an investor allocate funds among the possible investment choices? Markowitz suggested that investors should consider *risk and return* together and determine the allocation of funds among investment alternatives on the basis of the trade-off between them. Before Markowitz's seminal article, the finance literature had treated the interplay between risk and return in a casual manner.

The idea that sound financial decision making is a quantitative trade-off between risk and return was revolutionary for two reasons. First, it posited that one could make a quantitative evaluation of risk

[1] Harry M. Markowitz, "Portfolio Selection," *Journal of Finance* 7, no. 1 (March 1952), pp. 77–91. The principles in Markowitz's article were later expanded upon in his book *Portfolio Selection*, Cowles Foundation Monograph 16 (New York: John Wiley & Sons, 1959). Markowitz was awarded the Nobel Prize in Economic Sciences in 1990 for his work.

and return *jointly* by considering portfolio returns and their comovements. An important principle at work here is that of portfolio diversification. It is based on the idea that a portfolio's riskiness depends on the covariances of its constituents, not only on the average riskiness of its separate holdings. This concept was foreign to classical financial analysis, which revolved around the notion of the value of single investments, that is, the belief that investors should invest in those assets that offer the highest future value given their current price. Second, it formulated the financial decision-making process as an optimization problem. In particular, the so-called mean-variance principle formulated by Markowitz suggests that among the infinite number of portfolios that achieve a particular return objective, the investor should choose the portfolio that has the smallest variance. All other portfolios are "inefficient" because they have a higher variance and, therefore, higher risk.

Building on Markowitz's work, William Sharpe,[2] John Lintner,[3] and Jan Mossin[4] introduced the first asset pricing theory, the capital asset pricing model—CAPM in short—between 1962 and 1964. The CAPM became the foundation and the standard on which risk-adjusted performance of professional portfolio managers is measured.

Modern portfolio theory and diversification provide a theoretical justification for mutual funds and index funds, that have experienced a tremendous growth since the 1980s. A simple classification of fund management is into active and passive management, based upon the *efficient market hypotheses* introduced by Eugene Fama[5] and Paul Samuelson[6] in 1965. The efficient market hypothesis implies that it is not possible to outperform the market consistently on a risk-adjusted basis after accounting for transaction costs by using available information. In active management, it is assumed that markets are not fully efficient and that a fund manager can outperform a market index by using specific information, knowledge, and experience. Passive management, in con-

[2] William F. Sharpe, "Capital Asset Prices," *Journal of Finance* 19, no. 3 (September 1964), pp. 425–442. Sharpe received the Nobel Prize in Economic Sciences in 1990 for his work.

[3] John Lintner, "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolio and Capital Budgets," *Review of Economics and Statistics* 47 (February 1965), pp. 13–37.

[4] Jan Mossin, "Equilibrium in a Capital Asset Market," *Econometrica* 34, no. 4 (October 1966), pp. 768–783.

[5] Eugene F. Fama, "The Behavior of Stock Market Prices," *Journal of Business* 38 (January 1965), pp. 34–105.

[6] Paul A. Samuelson, "Proof that Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review* 6, no. 2 (Spring 1965), pp. 41–49. Samuelson was honored with the Nobel Prize in Economic Sciences in 1970.

trast, relies on the assumption that financial markets are efficient and that return and risk are fully reflected in asset prices. In this case, an investor should invest in a portfolio that mimics the market. John Bogle used this basic idea when he proposed to the board of directors of the newly formed Vanguard Group to create the first index fund in 1975. The goal was not to outperform the S&P 500 index, but instead to track the index as closely as possible by buying each of the stocks in the S&P 500 in amounts equal to the weights in the index itself.

Despite the great influence and theoretical impact of modern portfolio theory, today—more than 50 years after Markowitz's seminal work—full risk-return optimization at the asset level is primarily done only at the more quantitatively oriented firms. In the investment management business at large, portfolio management is frequently a purely judgmental process based on qualitative, not quantitative, assessments. The availability of quantitative tools is not the issue—today's optimization technology is mature and much more user-friendly than it was at the time Markowitz first proposed the theory of portfolio selection—yet many asset managers avoid using the quantitative portfolio allocation framework altogether.

A major reason for the reluctance of investment managers to apply quantitative risk-return optimization is that they have observed that it may be unreliable in practice. Specifically, risk-return optimization is very sensitive to changes in the inputs (in the case of mean-variance optimization, such inputs include the expected return of each asset and the asset covariances). While it can be difficult to make accurate estimates of these inputs, estimation errors in the forecasts significantly impact the resulting portfolio weights. It is well-known, for instance, that in practical applications equally weighted portfolios often outperform mean-variance portfolios, mean-variance portfolios are not necessarily well-diversified, and mean-variance optimization can produce extreme or non-intuitive weights for some of the assets in the portfolio. Such examples, however, are not necessarily a sign that the *theory* of risk-return optimization is flawed; rather, that when used *in practice*, the classical framework has to be modified in order to achieve reliability, stability, and robustness with respect to model and estimation errors.

It goes without saying that advances in the mathematical and physical sciences have had a major impact upon finance. In particular, mathematical areas such as probability theory, statistics, econometrics, operations research, and mathematical analysis have provided the necessary tools and discipline for the development of modern financial economics. Substantial advances in the areas of robust estimation and robust optimization were made during the 1990s, and have proven to be of great importance for the practical applicability and reliability of portfolio management and optimization.

Any statistical estimate is subject to error—estimation error. A robust estimator is a statistical estimation technique that is less sensitive to outliers in the data. For example, in practice, it is undesirable that one or a few extreme returns have a large impact on the estimation of the average return of a stock. Nowadays, Bayesian techniques and robust statistics are commonplace in financial applications. Taking it one step further, practitioners are starting to incorporate the uncertainty introduced by estimation errors directly into the optimization process. This is very different from the classical approach, where one solves the portfolio optimization problem as a problem with deterministic inputs, without taking the estimation errors into account. In particular, the statistical precision of individual estimates is explicitly incorporated in the portfolio allocation process. Providing this benefit is the underlying goal of *robust portfolio optimization*.

First introduced by El Ghaoui and Lebret[7] and by Ben-Tal and Nemirovski,[8] modern robust optimization techniques allow a portfolio manager to solve the robust version of the portfolio optimization problem in about the same time as needed for the classical portfolio optimization problem. The robust approach explicitly uses the distribution from the estimation process to find a robust portfolio in *one single* optimization, thereby directly incorporating uncertainty about inputs in the optimization process. As a result, robust portfolios are less sensitive to estimation errors than other portfolios, and often perform better than classical mean–variance portfolios. Moreover, the robust optimization framework offers great flexibility and many new interesting applications. For instance, robust portfolio optimization can exploit the notion of statistically equivalent portfolios. This concept is important in large-scale portfolio management involving many complex constraints such as transaction costs, turnover, or market impact. Specifically, with robust optimization, a manager can find the best portfolio that (1) minimizes trading costs with respect to the current holdings and (2) has an expected portfolio return and variance that are statistically equivalent to those of the classical mean-variance portfolio.

An important area of quantitative finance is that of modeling asset price behavior, and pricing options and other derivatives. This field can

[7] Laurent El Ghaoui, and Herve Lebret, "Robust Solutions to Least-Squares Problems with Uncertain Data," *SIAM Journal on Matrix Analysis and Applications* 18 (October 1997), pp. 1035–1064.

[8] Aharon Ben-Tal, and Arkadi S. Nemirovski, "Robust Convex Optimization," *Mathematics of Operations Research* 23, no. 4 (1998), pp. 769–805; and Aharon Ben-Tal, and Arkadi S. Nemirovski, "Robust Solutions to Uncertain Linear Programs," *Operations Research Letters* 25, no. 1 (1999), pp. 1–13.

be traced back to the early works of Thorvald Thiele[9] in 1880, Louis Bachelier[10] in 1900, and Albert Einstein[11] in 1905, who knew nothing about each other's research and independently developed the mathematics of Brownian motion. Interestingly, while the models by Thiele and Bachelier had little influence for a long time, Einstein's contribution had an immediate impact on the physical sciences. Historically, Bachelier's doctoral thesis is the first published work that uses advanced mathematics in the study of finance. Therefore, he is by many considered to be the pioneer of financial mathematics—the first "quant."[12]

The first listed options began trading in April 1973 on the Chicago Board Options Exchange (CBOE), only one and four months, respectively, before the papers by Black and Scholes[13] and by Merton[14] on option pricing were published. Although often criticized in the general press, and misunderstood by the public at large, options opened the door to a new era in investment and risk management, and influenced the introduction and popularization of a range of other financial products including interest rate swaptions, mortgage-backed securities, callable bonds, structured products, and credit derivatives. New derivative products were made possible as a solid pricing theory was available. Without the models developed by Black, Scholes, and Merton and many others following in their footsteps, it is likely that the rapid expansion

---

[9] Thorvald N. Theile, "Sur la Compensation de Quelques Erreurs Quasi-Systématiques par la Méthodes de Moindre Carrés [On the Compensation of Some Quasi-Systematic Errors by the Least Square Method]," *Vidensk. Selsk. Skr. 5* (1880), pp. 381–408.

[10] Louis Bachelier, "Théorie de la Speculation [Theory of Speculation]," *Annales Scientifiques de l'École Normale Supérieure Sér.*, 3, 17 (1900), pp. 21–86

[11] Albert Einstein, "On the Movement of Small Particles Suspended in Stationary Liquid Demanded by the Molecular-Kinetic Theory of Heat," in R. Fürth (ed.), *Investigations of the Theory of Brownian Movement* (New York: Dover Publications, 1956).

[12] The term "quant" which is short for *quantitative analyst* (someone who works in the financial markets developing mathematical models) was popularized, among other things, by Emanuel Derman in his book *My Life as a Quant* (Hoboken, NJ: John Wiley & Sons, 2004). On a lighter note, a T-shirt with the words "Quants Do It with Models" circulated among some quantitative analysts on Wall Street a few years ago.

[13] Fischer S. Black and Myron S. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, no. 3 (1973), pp. 637–659. Scholes received the Nobel Prize of Economic Science in 1997 for his work on option pricing theory. At that time, sadly, Fischer Black had passed away, but he received an honorable mention in the award.

[14] Robert C. Merton, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science* 4, no. 1 (Spring 1973), pp. 141–183. Merton received the Nobel Prize of Economic Science in 1997 for his work on option pricing theory.

of derivative products would never have happened. These modern instruments and the concepts of portfolio theory, CAPM, arbitrage and equilibrium pricing, and market predictability form the foundation not only for modern financial economics but for the general understanding and development of today's financial markets. As Peter Bernstein so adequately puts it in his book *Capital Ideas*: "Every time an institution uses these instruments, a corporation issues them, or a homeowner takes out a mortgage, they are paying their respects, not just to Black, Scholes, and Merton, but to Bachelier, Samuelson, Fama, Markowitz, Tobin, Treynor, and Sharpe as well."[15]

Computer Technology and the Internet

The appearance of the first personal computers in the late 1970s and early 1980s forever changed the world of computing. It put computational resources within the reach of most people. In a few years every trading desk on Wall Street was equipped with a PC. From that point on, computing costs have declined at the significant pace of about a factor of 2 every year. For example, the cost per gigaflops[16] is about $1 today, to be compared to about $50,000 about 10 years ago.[17] At the same time, computer speed increased in a similar fashion: today's fastest computers are able to perform an amazing 300 trillion calculations per second.[18] This remarkable development of computing technology has allowed finance professionals to deploy more sophisticated algorithms used, for instance, for derivative and asset pricing, market forecasting, portfolio allocation, and computerized execution and trading. With state-of-the-art optimization software, a portfolio manager is able to calculate the optimal allocation for a portfolio of thousands of assets in no more than a few seconds—on the manager's desktop computer!

[15] Peter L. Bernstein, *Capital Ideas* (New York: Free Press,1993).

[16] Flops is an abbreviation for *floating point operations per second* and is used as a measure of a computer's performance. 1 gigaflops = 10^9 flops.

[17] See Michael S. Warren, John K. Salmon, Donald J. Becker, M. Patrick Goda, Thomas Sterling, and Grégoire S. Winckelmans, "Pentium Pro Inside: I. A Treecode at 430 Gigaflops on ASCI Red. II. Price/Performance of $50/Mflop on Loki and Hyglac," *Supercomputing '97,* Los Alamitos, 1997, IEEE Computer Society; and Wikipedia contributors, "FLOPS," *Wikipedia, The Free Encyclopedia*, http://en.wikipedia.org/w/index.php?title=FLOPS&oldid=90585825 (accessed December 1, 2006).

[18] As of November 2006, the IBM BlueGene/L system with 131072 processor units held the so-called Linpack record with a remarkable performance of 280.6 teraflops (that is, 280.6 trillions of floating-point operations per second). See TOP500, www.top500.org.

But computational power alone is not sufficient for financial applications. It is crucial to obtain market data and other financial information efficiently and expediently, often in real time. The Internet and the World Wide Web have proven invaluable for this purpose. The World Wide Web, or simply the "Web," first created by Tim Berners-Lee working at CERN in Geneva, Switzerland around 1990, is an arrangement of interlinked, hypertext documents available over the Internet. With a simple browser, anybody can view webpages that may contain anything from text and pictures, to other multimedia based information, and jump from page to page by a simple mouse click.[19] Berners-Lee's major contribution was to combine the concept of hypertext with the Internet, born out of the NSFNet developed by the National Science Foundation in the early 1980s. The Web as we know it today allows for expedient exchange of financial information. Many market participants—from individuals to investment houses and hedge funds—use the Internet to follow financial markets as they move tick by tick and to trade many different kinds of assets such as stocks, bonds, futures, and other derivatives simultaneously across the globe. In today's world, gathering, processing, and analyzing the vast amount of information is only possible through the use of computer algorithms and sophisticated quantitative techniques.

Capital Markets

The development of the capital markets has of course had a significant impact on quantitative finance and the investment management industry as a whole. Investors today have a vast number of assets available in the capital markets, from more traditional assets such as stocks, bonds, commodities (precious metals, etc.) and real estate to derivative instruments such as options, futures, swaps, credit linked securities, mortgage-backed securities and other structured products, and specialized financial insurance products. These securities and products allow market participants to get exposure to, or to hedge risks—sometimes very specific risks. For example, a corporate bond portfolio manager may decide to hedge specific credit risks in his portfolio using a credit default swap, or a proprietary trader can short equity volatility by selling a volatility swap.

However, the number of assets available alone is not enough to guarantee success, if the assets are only traded infrequently and in small volumes. Successful capital markets have to be liquid, allowing market participants to trade their positions quickly and at low cost. An asset is

[19] A recent study concluded that as of January 2005 there are over 11.5 billion public webpages available on the Internet, see Antonio Gulli and Alessio Signorini, "The Indexable Web is More than 11.5 billion pages," 2005, Dipartimento di Informatica at Universita' di Pisa and Department of Computer Science at University of Iowa.

said to be liquid if it can be converted to cash quickly at a price close to fair market value. The U.S. capital markets are the most liquid in the world with Japan and the United Kingdom following. Cash, being the basic liquid asset, does not fluctuate in value—it itself defines price. All other assets can change in value and have an uncertain future price, making them risky assets. Naturally, informed investors will only hold less liquid and risky assets if they can expect to earn a premium, a risk premium.

With the tremendous increase in the number of assets—and with it, the amount of investment opportunities—it is hard, even for larger investment houses, to track and evaluate the different markets. Quantitative techniques lend themselves for automatic monitoring and analysis of the full multitude of securities. These tools give quantitative analysts, portfolio managers, and other decision makers the opportunity to summarize the vast amount of information available, and to present it in a cohesive manner. Modern financial and the econometric models rely on the access to accurate data, often with as long history as possible. It is typically much easier to obtain clean and trustworthy financial data from mature and liquid markets. In fact, the lack of reliable data is one of the inherent problems in applying sophisticated quantitative models to more illiquid markets. In these cases, practitioners are forced to rely on simulated data, make stronger assumptions in their models, or use less data-intensive models.

CENTRAL THEMES OF THIS BOOK

The purpose of this book is to provide a comprehensive introduction and overview of the state-of-the-art of portfolio management and optimization for practitioners, academics, and students alike. We attempt to bridge the gap from classical portfolio theory, as developed in the early 1950s, to modern portfolio optimization applications used in practice today. In particular, we provide an up-to-date review of robust estimation and optimization methods deployed in modern portfolio management, and discuss different techniques used in order to overcome the common pitfalls associated with classical mean-variance optimization. We discuss recent developments in quantitative trading strategies, trade execution, and operations research. While we focus on real world practical usability, and emphasize intuition and fundamental understanding, we try not to sacrifice mathematical rigor whenever possible.

We note that the concept of robustness in investment science extends beyond statistical and modeling methods. It suggests a new approach to financial forecasting, asset allocation, portfolio management, and trading. As a matter of fact, the concept of a *robust quantitative investment framework* seems to be gaining ground in the quantitative investment community, and is loosely defined by the following four stages:

1. Estimate reliable asset forecasts along with a measure of their confidence.
2. Deploy a robust model for portfolio allocation and risk management.
3. Manage portfolio rebalancing and trading costs efficiently as market conditions change.
4. Monitor and review the entire investment process on a regular basis.

The last stage includes the ability to evaluate past performance, as well as to measure and analyze portfolio risk. The role of quantitative models for econometric forecasting and optimization at each of these stages is very important, especially in large-scale investment management applications that require allocating, rebalancing, and monitoring of thousands of assets and portfolios.

From a broad perspective, the topics in this book can be categorized in the following four main areas: robust estimation, robust portfolio allocation, portfolio rebalancing, and management of model risk.

Robust Estimation

Models to predict expected returns of assets are routinely used by major asset management firms. In most cases, these models are straightforward and based on factors or other forecasting variables. Since parameter estimation in these financial models is data-driven, they are inevitably subject to estimation error. What makes matters worse, however, is that different estimation errors are accumulated across the different stages in the portfolio management process. As a result, the compounding of small errors from the different stages may result in large aggregate errors at the final stage. It is therefore important that parameters estimated at the different stages are reliable and robust so that the aggregate impact of estimation errors is minimized.

Given the existing plethora of financial forecasting models, the entire topic of robust statistical estimation is too extensive to cover in this book.[20] We will, however, touch upon several major topics. In particular, we review some fundamental statistical techniques for forecasting returns, show how robust statistical estimators for important inputs in the portfolio optimization process can be obtained, and how a robust

[20] For an overview of equity forecasting models, see Frank J. Fabozzi, Sergio M. Focardi, and Petter N. Kolm, *Financial Modeling of the Equity Market: From CAPM to Cointegration* (Hoboken, NJ: John Wiley & Sons, 2006).

portfolio allocation framework minimizes the impact of estimation and model errors. We describe robust frameworks for incorporating the investor's views such as shrinkage techniques and the Black-Litterman model to produce informed forecasts about the behavior of asset returns.

Robust Portfolio Allocation

Robust asset allocation is one of the most important parts of the investment management process, and the decision making is frequently based on the recommendations of risk-return optimization routines. Several major themes deserve attention. First, it is important to carefully consider how portfolio risk and return are defined, and whether these definitions are appropriate given observed or forecasted asset return distributions and underlying investor preferences. These concerns give rise to alternative theories of risk measures and asset allocation frameworks beyond classical mean-variance optimization. Second, the issue of how the optimization problem is formulated and solved in practice is crucial, especially for larger portfolios. A working knowledge of the state-of-the-art capabilities of quantitative software for portfolio management is critical. Third, it is imperative to evaluate the sensitivity of portfolio optimization models to inaccuracies in input estimates. We cover the major approaches for optimization under uncertainty in input parameters, including a recently developed area in optimization—*robust optimization*—that has shown a great potential and usability for portfolio management and optimization applications.

Portfolio Rebalancing

While asset allocation is one of the major strategic decisions, the decision of how to achieve this allocation in a cost-effective manner is no less important in obtaining good and consistent performance. Furthermore, given existing holdings, portfolio managers need to decide how to rebalance their portfolios efficiently to incorporate new views on expected returns and risk as the economic environment and the asset mix change. There are two basic aspects of the problem of optimal portfolio rebalancing. The first one is the robust management of the trading and transaction costs in the rebalancing process. The second is successfully combining both long-term and short-term views on the future direction and changes in the markets. The latter aspect is particularly important when taxes or liabilities have to be taken into account. The two aspects are not distinct, and in practice have to be considered simultaneously. By incorporating long-term views on asset behavior, portfolio managers may be able to reduce their overall transaction costs, as their portfolios do not have to be rebalanced as often. Although the interplay between the different aspects is complex to evaluate and model, disciplined portfolio rebalancing using an optimizer provides portfolio managers with new opportunities.

Managing Model Risk

Quantitative approaches to portfolio management introduce a new source of risk—model risk—and an inescapable dependence on historical data as their raw material. Financial models are typically predictive—they are used to forecast unknown or future values on the basis of current or known values using specified equations or sets of rules. Their predictive or forecasting power, however, is limited by the appropriateness of the inputs and basic model assumptions. Incorrect assumptions, model identification and specification errors, or inappropriate estimation procedures inevitably lead to model risk, as does using models without sufficient out-of-sample testing. It is important to be cautious in how we use models, and to make sure that we fully understand their weaknesses and limitations. In order to identify the various sources of model risk, we need to take a critical look at our models, review them on a regular basis, and avoid their use beyond the purpose or application for which they were originally designed.

OVERVIEW OF THIS BOOK

We have organized the book as follows. Part I (Chapters 2, 3, and 4) introduces the underpinnings of modern portfolio theory. Part II (Chapters 5, 6, 7, and 8) summarizes important developments in the estimation of parameters such as expected asset returns and their covariances that serve as inputs to the classical portfolio optimization framework. Part III (Chapters 9, 10, and 11) describes the tools necessary to handle the optimization step of the process. Part IV (Chapters 12, 13, and 14) focuses on applications of the methods described in the previous parts, and outlines new directions in robust portfolio optimization and investment management as a whole.

We start out by describing the classical portfolio theory and the concepts of diversification in Chapter 2. We introduce the concepts of efficient sets and efficient frontiers, and discuss the effect of long-only constraints. We also present an alternative framework for optimal decision making in investment—expected utility optimization—and explain its relationship to classical mean-variance optimization.

Chapter 3 extends classical portfolio theory to a more general mean-risk setting. We cover the most common alternative measures of risk that, in some cases, are better suited than variance in describing investor preferences when it comes to skewed and/or fat-tailed asset return distributions. We also show how to incorporate investor preferences for higher moments in the expected utility maximization framework, and discuss polynomial goal programming. Finally, we introduce a new approach to portfolio selection with higher moments proposed by Malevergne and Sornette, and illustrate the approach with examples.

Chapter 4 provides an overview of practical considerations in implementing portfolio optimization. We review constraints that are most commonly faced by portfolio managers, and show how to formulate them as part of the optimization problem. We also show how the classical framework for portfolio allocation can be extended to include transaction costs, and discuss the issue of optimizing trading impact costs across multiple client accounts simultaneously.

Chapter 5 introduces a number of price and return models that are used in portfolio management. We examine different types of random walks, present their key properties, and compare them to other trend-stationary processes. We also discuss standard financial models for explaining and modeling asset returns that are widely used in practice—the Capital Asset Pricing Model (CAPM), Arbitrage Pricing Theory (APT), and factor models,

The estimation of asset expected returns and covariances is essential for classical portfolio management. Chapter 6 covers the standard approaches for estimating parameters in portfolio optimization models. We discuss methods for estimating expected returns and covariance matrices, introduce dimensionality reduction techniques such as factor models, and use random matrix theory to illustrate how noisy the sample covariance matrix can be. In Chapter 7, we provide an introduction to the theory of robust statistical estimation.

Chapter 8 presents recent developments in asset return forecasting models, focusing on new frameworks for robust estimation of important parameters. In particular, we discuss shrinkage methods and the Black-Litterman approach for expected return estimation. Such methods allow for combining statistical estimates with investors' views of the market.

The subject of Chapter 9 is practical numerical optimization, our goal being to introduce readers to the concept of "difficult" versus "easy" optimization problems. We discuss the types of optimization techniques encountered in portfolio management problems—linear and quadratic programming, as well as the more advanced areas of convex programming, conic optimization, and integer programming. We explain the concept of optimization duality and describe intuitively how optimization algorithms work. Illustrations of the various techniques are provided, from the classical simplex method for solving linear programming problems to state-of-the-art barrier- and interior-point methods.

Classical optimization methods treat the parameters in optimization problems as deterministic and fully accurate. In practice, however, these parameters are typically estimated from error-prone statistical procedures or based on subjective evaluation, resulting in estimates with significant estimation errors. The output of optimization routines based on poorly estimated inputs can be seriously misleading and often useless. This is a reason why optimizers are sometimes cynically referred to as "error maximizers." It is important to know how to treat uncertainty in the estimates of input parameters in optimization problems. Chapter 10 provides a taxonomy of methods for optimization under uncertainty. We review the main ideas behind stochastic programming, dynamic programming, and robust optimization, and illustrate the methods with examples.

Chapter 11 contains practical suggestions for formulating and solving optimization problems in real-world applications. We review publicly and commercially available software for different types of optimization problems and portfolio optimization in particular, and provide examples of implementation of portfolio optimization problems in AMPL (an optimization modeling language) and MATLAB (a popular modeling environment).

Chapter 12 focuses on the application of robust optimization and resampling techniques for treating uncertainty in the parameters of classical mean-variance portfolio optimization. We present robust counterparts of the classical portfolio optimization problem under a variety of assumptions on the asset return distributions and different forms of estimation errors in expected returns and risk.

In Chapter 13, we describe recent trends and new directions in the area of robust portfolio management, and elaborate on extensions and refinements of some of the techniques described elsewhere in this book. In particular, we provide an overview of more advanced topics such as handling the underestimation of risk in factor models, robust applications of alternative risk measures, portfolio rebalancing with transaction and trading costs, and multiperiod portfolio optimization.

The last chapter of the book, Chapter 14, provides an outlook of some important aspects of quantitative investment management. We address the use of derivatives in portfolio management, currency management in international portfolios, and benchmark selection. We examine the most widespread quantitative and model-based trading strategies used in quantitative trading today, and discuss model risk including data snooping and overfitting. The chapter closes with an introduction to optimal execution and algorithmic trading.

The appendix at the end of the book contains a description of the data used in illustrations in several of the chapters.