

Springer Finance

Attilio Meucci

Risk and Asset Allocation

风险和资产配置



Springer

世界图书出版公司
www.wpcbj.com.cn

Attilio Meucci

Risk and Asset Allocation

With 138 Figures



Attilio Meucci

Lehman Brothers, Inc.

745 Seventh Avenue

New York, NY 10019

USA

E-mail: attilio_meucci@symmys.com

Mathematics Subject Classification (2000): 15-xx, 46-xx, 62-xx, 65-xx, 90-xx

JEL Classification: C1, C3, C4, C5, C6, C8, G0, G1

Library of Congress Control Number: 2005922398

Corrected 3rd printing 2007

ISBN 978-3-540-22213-2 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2005

Scientific WorkPlace® is a trademark of MacKichan Software, Inc. and is used with permission. MATLAB® is a trademark of The Math Works, Inc. and is used with permission. The Math Works does not warrant the accuracy of the text or exercises in this book. This book's use or discussion of MATLAB® software or related products does not constitute endorsement or sponsorship by The Math Works of a particular pedagogical approach or particular use of the MATLAB® software.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

This reprint has been authorized by Springer-Verlag (Berlin/Heidelberg/New York) for sale in the Mainland China only and not for export therefrom.

图书在版编目 (CIP) 数据

风险和资产配置 = Risk and Asset Allocation;

英文/ (美) 梅乌奇 (Meucci, A.) 著. —影印本.

—北京: 世界图书出版公司北京公司, 2010. 2

ISBN 978-7-5100-0492-6

I. ①风… II. ①梅… III. ①风险管理—教材—英文

②资产管理—教材—英文 IV. ①F272. 3 ②F830. 593

中国版本图书馆 CIP 数据核字 (2010) 第 010621 号

书 名: Risk and Asset Allocation

作 者: Attilio Meucci

中 译 名: 风险和资产配置

责任编辑: 高蓉 刘慧

出 版 者: 世界图书出版公司北京公司

印 刷 者: 三河国英印务有限公司

发 行: 世界图书出版公司北京公司 (北京朝内大街 137 号 100010)

联系电话: 010-64021602, 010-64015659

电子信箱: kjb@wpchj.com.cn

开 本: 24 开

印 张: 23.5

版 次: 2010 年 01 月

版权登记: 图字: 01-2009-2045

书 号: 978-7-5100-0492-6/O · 707

定 价: 65.00 元

Preface

In an asset allocation problem the investor, who can be the trader, or the fund manager, or the private investor, seeks the combination of securities that best suit their needs in an uncertain environment. In order to determine the optimum allocation, the investor needs to model, estimate, assess and manage uncertainty.

The most popular approach to asset allocation is the mean-variance framework pioneered by Markowitz, where the investor aims at maximizing the portfolio's expected return for a given level of variance and a given set of investment constraints. Under a few assumptions it is possible to estimate the market parameters that feed the model and then solve the ensuing optimization problem.

More recently, measures of risk such as the value at risk or the expected shortfall have found supporters in the financial community. These measures emphasize the potential downside of an allocation more than its potential benefits. Therefore, they are better suited to handle asset allocation in modern, highly asymmetrical markets.

All of the above approaches are highly intuitive. Paradoxically, this can be a drawback, in that one is tempted to rush to conclusions or implementations, without pondering the underlying assumptions.

For instance, the term "mean-variance" hints at the identification of the expected value with its sample counterpart, the mean. Sample estimates make sense only if the quantities to estimate are market invariants, i.e. if they display the same statistical behavior independently across different periods. In equity-like securities the returns are approximately market invariants: this is why the mean-variance approach is usually set in terms of returns. Consider instead an investment in a zero-coupon bond that expires, say, in one month. The time series of the past monthly returns of this bond is not useful in estimating the expected value and the variance after one month, which are known with certainty: the returns are not market invariants.

Similarly, when an allocation decision is based on the value at risk or on the expected shortfall, the problem is typically set in terms of the portfolio's profit-and-loss, because the "P&L" is approximately an invariant.

In general, the investor focuses on a function of his portfolio's value at the end of the investment horizon. For instance, the portfolio's return or profit-and-loss are two such functions which, under very specific circumstances, also happen to be market invariants. In more general settings, the investor needs to separate the definition of his objectives, which depend on the portfolio value at a given future horizon, from the estimation of the distribution of these objectives, which relies on the identification and estimation of some underlying market invariants.

To summarize, in order to solve a generic asset allocation problem we need to go through the following steps.

Detecting invariance

In this phase we detect the market invariants, namely those quantities that display the same behavior through time, allowing us to learn from the past. For equities the invariants are the returns; for bonds the invariants are the changes in yield to maturity; for vanilla derivatives the invariants are changes in at-the-money-forward implied volatility; etc.

Estimating the market

In this step we estimate the distribution of the market invariants from a time series of observations by means of nonparametric estimators, parametric estimators, shrinkage estimators, robust estimators, etc.

Modeling the market

In this phase we map the distribution of the invariants into the distribution of the market at a generic time in the future, i.e. into the distribution of the prices of the securities for the given investment horizon. This is achieved by suitable generalizations of the "square-root-rule" of volatility propagation. The distribution of the prices at the horizon in turn determines the distribution of the investor's objective, such as final wealth, or profit and loss, etc.

Defining optimality

In this step we analyze the investor's profile. We ascertain the features of a potential allocation that are more valuable for a specific investor, such as the trade-off between the expected value and the variance of his objective, or the value at risk of his objective, etc.; and we determine the investor's constraints, such as budget constraints, reluctance to invest in certain assets, etc.

Only after performing separately the above steps can we proceed toward the final goal:

Computing the optimal allocation

At this stage we determine exactly or in good approximation the allocation that best suits the investor, namely the allocation that maximizes the valuable features of the investor's objective(s) given his constraints.

Nevertheless, the approach outlined above is sub-optimal: two additional steps are needed.

Accounting for estimation risk

It is not clear from the above that an allocation based on one month of data is less reliable than an allocation based on two years of data. Nevertheless, the effect of estimation errors on the allocation's performance is dramatic. Therefore we need to account for estimation risk in the optimization process.

Including experience

The most valuable tool for a successful investor is experience, or a-priori knowledge of the market. We need to include the investor's experience in the optimization process by means of a sound statistical framework.

Purpose of this book is to provide a comprehensive treatment of all the above steps. In order to discuss these steps in full generality and consistently from the first to the last one we focus on one-period asset allocation.

Audience and style

A few years ago I started teaching computer-based graduate courses on asset allocation and risk management with emphasis on estimation and modeling because I realized the utmost importance of these aspects in my experience as a practitioner in the financial industry. While teaching, I felt the need to provide the students with an accessible, yet detailed and self-contained, reference for the theory behind the above applications. Since I could not find such a reference in the literature, I set out to write lecture notes, which over the years and after many adjustments have turned into this book.

In an effort to make the reader capable of innovating rather than just following, I sought to analyze the first principles that lead to any given recipe, in addition to explaining how to implement that recipe in practice. Once those first principles have been isolated, the discussion is kept as general as possible: the many applications detailed throughout the text arise as specific instances of the general discussion.

I have tried wherever possible to support intuition with geometrical arguments and practical examples. Heuristic arguments are favored over mathematical rigor. The mathematical formalism is used only up to (and not beyond) the point where it eases the comprehension of the subject. The MATLAB[®] applications downloadable from symmys.com allow the reader to further visualize the theory and to understand the practical issues behind the applications.

A reader with basic notions of probability and univariate statistics could learn faster from the book, although this is not a prerequisite. Simple concepts of functional analysis are used heuristically throughout the text, but the reader is introduced to them from scratch and absolutely no previous knowledge of the subject is assumed. Nevertheless the reader must be familiar with multivariate calculus and linear algebra.

For the above reasons, this book targets graduate and advanced undergraduate students in economics and finance as well as the new breed of practitioners with a background in physics, mathematics, engineering, finance or economics who ever increasingly populate the financial districts worldwide. For the students this is a textbook that introduces the problems of the financial industry in a format that is more familiar to them. For the practitioners, this is a comprehensive reference for the theory and the principles underlying the recipes they implement on a daily basis.

Any feedback on the book is greatly appreciated. Please refer to the website symmys.com to contact me.

Structure of the work

This work consists of the printed text and of complementary online resources.

- **Printed text**

The printed text is divided in four parts.

Part I

In the first part we present the statistics of asset allocation, namely the tools necessary to model the market prices at the investment horizon. Chapters 1 and 2 introduce the reader to the formalism of financial risk, namely univariate and multivariate statistics respectively. In Chapter 3 we discuss how to detect the market invariants and how to map their distribution into the distribution of the market prices at the investment horizon.

Part II

In the second part we discuss the classical approach to asset allocation. In Chapter 4 we show how to estimate the distribution of the market invariants. In Chapter 5 we define optimality criteria to assess the advantages and disadvantages of a given allocation, once the distribution of the market is known. In Chapter 6 we set and solve allocation problems, by maximizing the advantages of an allocation given the investment constraints.

Part III

In the third part we present the modern approach to asset allocation, which accounts for estimation risk and includes the investor's experience in the decision process. In Chapter 7 we introduce the Bayesian approach to parameter estimation. In Chapter 8 we update the optimality criteria to assess the advantages and disadvantages of an allocation when the distribution of the market is only known with some approximation. In Chapter 9 we pursue optimal allocations in the presence of estimation risk, by maximizing their advantages according to the newly defined optimality criteria.

Part IV

The fourth part consists of two mathematical appendices. In Appendix A we review some results from linear algebra, geometry and matrix calculus.

In Appendix B we hinge on the analogies with linear algebra to introduce heuristically the simple tools of functional analysis that recur throughout the main text.

- **Online resources**

The online resources consist of software applications and ready-to-print material. They can be downloaded freely from the website symmys.com.

Software applications

The software applications are in the form of MATLAB programs. These programs were used to generate the case studies, simulations and figures in the printed text.

Exercise book

The exercise book documents the above MATLAB programs and discusses new applications.

Technical appendices

In order to make the book self-contained, the proofs to almost all the technical results that appear in the printed text are collected in the form of end-of-chapter appendices. These appendices are not essential to follow the discussion. However, they are fundamental to a true understanding of the subjects to which they refer. Nevertheless, if included in the printed text, these appendices would have made the size of the book unmanageable.

The notation in the printed text, say, "Appendix www.2.4" refers to the technical appendix to Chapter 2, Section 4, which is located on the internet. On the other hand the notation, say, "Appendix B.3" refers to the mathematical Appendix B, Section 3, at the end of the book.

A guided tour by means of a simplistic example

To better clarify the content of each chapter in the main text we present a more detailed overview, supported by an oversimplified example which, we stress, does not represent a real model.

Part I

A portfolio at a given future horizon is modeled as a random variable and is represented by a univariate distribution: in Chapter 1 we review univariate statistics. We introduce the representations of the distribution of a generic random variable X , i.e. the probability density function, the cumulative distribution function, the characteristic function and the quantile, and we discuss expected value, variance and other parameters of shape. We present a graphical interpretation of the location and dispersion properties of a univariate distribution and we discuss a few parametric distributions useful in applications.

For example, we learn what it means that a variable X is normally distributed:

$$X \sim N(\mu, \sigma^2), \quad (0.1)$$

where μ is the expected value and σ^2 is the variance.

The market consists of securities, whose prices at a given future horizon can be modeled as a multivariate random variable: in Chapter 2 we discuss multivariate statistics. We introduce the representations of the distribution of a multivariate random variable $\mathbf{X}$, namely the joint probability density function, the cumulative distribution function and the characteristic function. We analyze the relationships between the different entries of $\mathbf{X}$: the marginal-copula factorization, as well as the concepts of dependence and of conditional distribution.

We discuss expected value, mode and other multivariate parameters of location; and covariance, modal dispersion and other multivariate parameters of dispersion. We present a graphical interpretation of location and dispersion in terms of ellipsoids and the link between this interpretation and principal component analysis.

We discuss parameters that summarize the co-movements of one entry of $\mathbf{X}$ with another: we introduce the concept of correlation, as well as alternative measures of concordance. We analyze the multivariate generalization of the distributions presented in Chapter 1, including the Wishart and the matrix-variate Student t distributions, useful in Bayesian analysis, as well as very general log-distributions, useful to model prices. Finally we discuss special classes of distributions that play special roles in applications.

For example, we learn what it means that two variables $\mathbf{X} \equiv (X_1, X_2)'$ are normally distributed:

$$\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (0.2)$$

where $\boldsymbol{\mu} \equiv (\mu_1, \mu_2)'$ is the vector of the expected values and where the covariance matrix is the identity matrix, i.e. $\boldsymbol{\Sigma} \equiv \mathbf{I}$. We represent this variable as a unit circle centered in $\boldsymbol{\mu}$: the radius represents the two eigenvalues and the reference axes represent the two eigenvectors. As it turns out, the normal distribution (0.2) belongs to the special elliptical, stable and infinitely divisible classes.

In Chapter 3 we model the market. The market is represented by a set of securities that at time t trade at the price $\mathbf{P}_t$. The investment decision is made at the time T and the investor is interested in the distribution of the prices $\mathbf{P}_{T+\tau}$ at a determined future investment horizon τ . Modeling the market consists of three steps. First we need to identify the invariants hidden behind the market data, i.e. those random variables $\mathbf{X}$ that are distributed identically and independently across time.

For example suppose that we detect as invariants the changes in price:

$$\mathbf{X}_{t,\tilde{\tau}} \equiv \mathbf{P}_t - \mathbf{P}_{t-\tilde{\tau}}, \quad (0.3)$$

where the estimation horizon $\tilde{\tau}$ is one week.

Secondly, we have to associate a meaningful parametric distribution to these invariants

For example suppose that the normal distribution (0.2) with the identity as covariance is a suitable parametric model for the weekly changes in prices:

$$\mathbf{X}_{t,\tilde{\tau}} \sim N(\boldsymbol{\mu}, \mathbf{I}). \quad (0.4)$$

In this case the market parameters, still to be determined, are the entries of $\boldsymbol{\mu}$.

Finally, we have to work out the distribution of the market, i.e. the prices $\mathbf{P}_{T+\tau}$ at the generic horizon τ , given the distribution of the invariants $\mathbf{X}_{t,\tilde{\tau}}$ at the specific horizon $\tilde{\tau}$. This step is fundamental when we first estimate parameters at a given horizon and then solve allocation problems at a different horizon.

For example, suppose that the current market prices of all the securities are normalized to one unit of currency, i.e. $\mathbf{P}_T \equiv \mathbf{1}$, and that the investment horizon is one month, i.e. four weeks. Then, from (0.3) and (0.4) the distribution of the market is normal with the following parameters:

$$\mathbf{P}_{T+\tau} \sim N(\mathbf{m}, 4\mathbf{I}), \quad (0.5)$$

where

$$\mathbf{m} \equiv \mathbf{1} + 4\boldsymbol{\mu}. \quad (0.6)$$

In a market of many securities the actual dimension of risk in the market is often much lower than the number of securities: therefore we discuss dimension-reduction techniques such as regression analysis and principal component analysis and their geometrical interpretation in terms of the location-dispersion ellipsoid. We conclude with a detailed case study, which covers all the steps involved in modeling the swap market: the detection of the invariants; the "level-slope-hump" PCA approach to dimension reduction of the swap curve invariants, along with its continuum-limit interpretation in terms of frequencies; and the roll-down, duration and convexity approximation of the swap market.

Part II

In the first part of the book we set the statistical background necessary to formalize allocation problems. In the second part we discuss the classical approach to solve these problems, which consists of three steps: estimating the market distribution, evaluating potential portfolios of securities and optimizing those portfolios according to the previously introduced evaluation criteria.

In Chapter 4 we estimate from empirical observations the distribution of the market invariants. An estimator is a function that associates a number, the estimate, with the information i_T that is available when the investment decision is made. This information is typically represented by the time series of the past observations of the market invariants.

For example, we can estimate the value of the market parameter μ in (0.4) by means of the sample mean:

$$i_T \equiv \{\mathbf{x}_1, \dots, \mathbf{x}_T\} \mapsto \hat{\mu} \equiv \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t, \quad (0.7)$$

where we dropped the estimation interval from the notation.

We discuss general rules to evaluate the quality of an estimator. The most important feature of an estimator is its replicability, which guarantees that a successful estimation does not occur by chance. An estimator's replicability is measured by the distribution of its loss and is summarized by error, bias and inefficiency. Then we introduce different estimators for different situations: nonparametric estimators, suitable in the case of a very large number of observations; maximum likelihood estimators under quite general non-normal assumptions, suitable when the parametric shape of the invariants' distribution is known; shrinkage estimators, which perform better when the amount of data available is limited; robust estimators, which the statistician should use when he is not comfortable with a given parametric specification of the market invariants. Throughout the analysis we provide the geometrical interpretation of the above estimators. We conclude with practical tips to deal, among other problems, with outliers detection and missing values in the time series.

In Chapter 5 we show how to evaluate an allocation. The investor can allocate his money in the market to form a portfolio of securities. Therefore, the allocation decision is defined by a vector α whose entries determine the number of units (e.g. shares) of the respective security that are being purchased at the investment time T . The investor focuses on his primary objective, a random variable whose distribution depends on the allocation and the market parameters: different objectives corresponds to different investment priorities, such as benchmark allocation, daily trading (profits and losses), financial planning, etc.

For example, assume that the investor's objective is final wealth. If the market is distributed as in (0.5) the objective is normally distributed:

$$\Psi \equiv \alpha' \mathbf{P}_{T+\tau} \sim N(\alpha' \mathbf{m}, \sigma^2), \quad (0.8)$$

where $\mathbf{m}$ is given in (0.6) and σ^2 is a simple function of the allocation.

Evaluating an allocation corresponds to assessing the advantages and disadvantages of the distribution of the respective objective. We start considering stochastic dominance, a criterion to compare distributions globally: nevertheless stochastic dominance does not necessarily give rise to a ranking of the potential allocations. Therefore we define indices of satisfaction, i.e. functions of the allocation and the market parameters that measure the extent to which an investor appreciates the objective ensuing from a given allocation.

For example, satisfaction can be measured by the expected value of final wealth: a portfolio with high expected value elicits a high level of satisfaction. In this case from (0.6) and (0.8) the index of satisfaction is the following function of the allocation and of the market parameters:

$$(\alpha, \mu) \mapsto E\{\Psi\} = \alpha'(1 + 4\mu). \quad (0.9)$$

We discuss the general properties that indices of satisfaction can or should display. Then we focus on three broad classes of such indices: the certainty-equivalent, related to expected utility and prospect theory; the quantile of the objective, closely related to the concept of value at risk; and coherent and spectral measures of satisfaction, closely related to the concept of expected shortfall. We discuss how to build these indices and we analyze their dependence on the underlying allocation. We tackle a few computational issues, such as the Arrow-Pratt approximation, the gamma approximation, the Cornish-Fisher approximation, and the extreme value theory approximation.

In Chapter 6 we pursue the optimal allocation for a generic investor. Formally, this corresponds to maximizing the investor's satisfaction while keeping into account his constraints. We discuss the allocation problems that can be solved efficiently at least numerically, namely convex programming and in particular semidefinite and second-order cone programming problems.

For example, suppose that transaction costs are zero and that the investor has a budget constraint of one unit of currency and can purchase only positive amounts of any security. Assume that the market consists of only two securities. Given the current market prices, from (0.9) the investor's optimization problem reads:

$$\alpha^* \equiv \operatorname{argmax}_{\substack{\alpha_1 + \alpha_2 = 1 \\ \alpha \geq 0}} \alpha'(1 + 4\hat{\mu}), \quad (0.10)$$

where $\hat{\mu}$ are the estimated market parameters (0.7). This is a linear programming problem, a special case of cone programming. The solution is a 100% investment in the security with the largest estimated expected value. Assuming for instance that this is the first security, we obtain:

$$\hat{\mu}_1 > \hat{\mu}_2 \Rightarrow \alpha_1^* \equiv 1, \alpha_2^* \equiv 0. \quad (0.11)$$

In real problems it not possible to compute the exact solution to an allocation optimization. Nevertheless it is possible to obtain a good approximate solution by means of a two-step approach. The core of this approach is the mean-variance optimization, which we present in a general context in terms market prices, instead of the more common, yet more restrictive, representation in terms of returns. Under fairly standard hypotheses, the computation of the mean-variance frontier is a quadratic programming problem. In special cases we can even compute analytical solutions, which provide insight into the effect of the market on the allocation in more general contexts: for example, we prove wrong the common belief that uncorrelated markets provide better investment opportunities than highly correlated markets. We analyze thoroughly the problem of managing assets against a benchmark, which is the explicit task of a fund manager and, as it turns out, the implicit objective of all investors. We discuss the pitfalls of a superficial approach to the mean-variance problem, such as the confusion between compounded returns and linear returns which gives rise to distortions in the final allocation. Finally, we present a case study that reviews all the steps that lead to the optimal allocation.

Part III

In the classical approach to asset allocation discussed in the second part we implicitly assumed that the distribution of the market, once estimated, is known. Nevertheless, such distribution is estimated with some error. As a result, any allocation implemented cannot be truly optimal and the truly optimal allocation cannot be implemented. More importantly, since the optimization process is extremely sensitive to the input parameters, the sub-optimality due to estimation risk can be dramatic.

The parameter $\hat{\mu}$ in the optimization (0.10) is only an estimate of the true parameter that defines the distribution of the market (0.4). The true expected value of the second security could be larger than the first one, as opposed to what stated in (0.11). In this case the truly optimal allocation would read:

$$\mu_1 < \mu_2 \Rightarrow \bar{\alpha}_1 \equiv 0, \bar{\alpha}_2 \equiv 1. \quad (0.12)$$

This allocation is dramatically different from the allocation (0.11), which was implemented.

As a consequence, portfolio managers, traders and professional investors in a broader sense mistrust the "optimal" allocations ensuing from the classical approach and prefer to resort to their experience. In the third part of the book we present a systematic approach to tackle estimation risk, which also includes within a sound statistical framework the investor's experience or models.

Following the guidelines of the classical approach, in order to determine the optimal allocation in the presence of estimation risk we need to introduce a new approach to estimate the market distribution, update the evaluation

criteria for potential portfolios of securities and optimize those portfolios according to the newly introduced evaluation criteria.

In Chapter 7 we introduce the Bayesian approach to estimation. In this context, estimators are not numbers: instead, they are random variables modeled by their posterior distribution, which includes the investor's experience or prior beliefs. A Bayesian estimator defines naturally a classical-equivalent estimator and an uncertainty region.

For example, the Bayesian posterior, counterpart of the classical estimator (0.7), could be a normal random variable:

$$\hat{\mu}_B \sim N\left(\frac{1}{2}\hat{\mu} + \frac{1}{2}\mu_0, \mathbf{I}\right), \quad (0.13)$$

where μ_0 is the price change that the investor expects to take place. Then the classical-equivalent estimator is an average of the prior and the sample estimator:

$$\hat{\mu}_{ce} \equiv \frac{1}{2}\hat{\mu} + \frac{1}{2}\mu_0; \quad (0.14)$$

and the uncertainty region is a unit circle centered in $\hat{\mu}_{ce}$:

$$\mathcal{E} \equiv \{\mu \text{ such that } (\mu - \hat{\mu}_{ce})'(\mu - \hat{\mu}_{ce}) \leq 1\}. \quad (0.15)$$

Since it is difficult for the investor to input prior beliefs directly in the model, we discuss how to input them implicitly in terms of ideal allocations.

In Chapter 8 we introduce criteria to evaluate the sub-optimality of a generic allocation. This process parallels the evaluation of an estimator. The estimator's loss becomes in this context the given allocation's opportunity cost, i.e. a positive random variable which represents the difference between the satisfaction provided by the true, yet unattainable, optimal allocation and the satisfaction provided by the given allocation.

In our example, from (0.9) the opportunity cost of the sub-optimal allocation (0.10) reads:

$$OC(\alpha^*, \mu) = (1 + 4\mu)' \bar{\alpha} - (1 + 4\mu)' \alpha^*, \quad (0.16)$$

where $\bar{\alpha}$ is the truly optimal allocation (0.12).

We analyze the opportunity cost of two extreme approaches to allocation: at one extreme the prior allocation, which completely disregards any information from the market, relying only on prior beliefs; at the other extreme the sample-based allocation, where the unknown market parameters are replaced by naive estimates.

In Chapter 9 we pursue optimal allocations in the presence of estimation risk, namely allocations whose opportunity cost is minimal. We present allo-

cations based on Bayes' rule, such as the classical-equivalent allocation and the Black-Litterman approach. Next we present the resampling technique by Michaud. Then we discuss robust allocations, which aim at minimizing the maximum possible opportunity cost over a given set of potential market parameters. Finally, we present robust Bayesian allocations, where the set of potential market parameters is defined naturally in terms of the uncertainty set of the posterior distribution.

In our example, the sub-optimal allocation (0.10) is replaced by the following robust Bayesian allocation:

$$\alpha^* \equiv \underset{\substack{\alpha_1 + \alpha_2 = 1 \\ \alpha \geq 0}}{\operatorname{argmin}} \left\{ \max_{\mu \in \mathcal{E}} \operatorname{OC}(\alpha, \mu) \right\}, \quad (0.17)$$

where the opportunity cost is defined in (0.16) and the uncertainty set is defined in (0.15). The solution to this problem is a balanced allocation where, unlike in (0.11), both securities are present in positive amounts.

In general it is not possible to compute exactly the optimal allocations. Therefore, as in the classical approach to asset allocation, we resort to the two-step mean-variance setting to solve real problems.

Acknowledgments

I wish to thank Carlo Favero, Carlo Giannini, John Hulpke, Alfredo Pastor and Eduardo Rossi, who invited me to teach finance at their institutions, thereby motivating me to write lecture notes for my courses.

A few people provided precious feedback on different parts of the draft at different stages in its development, in particular Davide DiGennaro, Luca Dona', Alberto Elices, Silverio Foresi, Davide Guzzetti, Philip Stark, Dirk Tasche, Kostas Tryantafyllapoulos and an anonymous referee. Francesco Corielli and Gianluca Fusai furnished insightful comments and suggested new material for the book during many pleasant conversations throughout the last few years.

I am indebted to Catriona Byrne, Susanne Denskus and Stefanie Zoeller at Springer for their active support; to EDV-Beratung Frank Herweg for carefully correcting the proofs; and to George Pearson and John MacKendrick at MacKichan Software, Inc. for helping me discover the capabilities of Scientific Workplace[®], which I used to write this book.

A special thank is due to Jenifer Shiu, for her support during the last year of writing.

Greenwich, January 2005,

Attilio Meucci

Contents

Preface	XV
Audience and style	XVII
Structure of the work	XVIII
A guided tour by means of a simplistic example	XIX
Acknowledgments	XXVI

Part I The statistics of asset allocation

1	Univariate statistics	3
1.1	Building blocks	3
1.2	Summary statistics	9
1.2.1	Location	9
1.2.2	Dispersion	11
1.2.3	Higher-order statistics	14
1.2.4	Graphical representations	15
1.3	Taxonomy of distributions	16
1.3.1	Uniform distribution	16
1.3.2	Normal distribution	18
1.3.3	Cauchy distribution	20
1.3.4	Student t distribution	22
1.3.5	Lognormal distribution	24
1.3.6	Gamma distribution	26
1.3.7	Empirical distribution	28
1.T	Technical appendix	www
1.E	Exercises	www
2	Multivariate statistics	33
2.1	Building blocks	34
2.2	Factorization of a distribution	38
2.2.1	Marginal distribution	38
2.2.2	Copulas	40