European Journal of Operational Research 234 (2014) 356–371

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

60 Years of portfolio optimization: Practical challenges and current trends

Petter N. Kolm [a,*], Reha Tütüncü [b], Frank J. Fabozzi [c]

[a] Courant Institute, New York University, 251 Mercer St., New York, NY 10012, United States
[b] Goldman Sachs Asset Management, 200 West St., New York, NY 10282, United States
[c] EDHEC Business School, 393, Promenade des Anglais BP3116, 06202 Nice Cedex 3, France

ARTICLE INFO

Article history:
Available online 1 November 2013

Keywords:
Black–Litterman
Estimation errors
Mean–variance optimization
Multi-period optimization
Portfolio constraints
Portfolio optimization

ABSTRACT

The concepts of portfolio optimization and diversification have been instrumental in the development and understanding of financial markets and financial decision making. In light of the 60 year anniversary of Harry Markowitz's paper "Portfolio Selection," we review some of the approaches developed to address the challenges encountered when using portfolio optimization in practice, including the inclusion of transaction costs, portfolio management constraints, and the sensitivity to the estimates of expected returns and covariances. In addition, we selectively highlight some of the new trends and developments in the area such as diversification methods, risk-parity portfolios, the mixing of different sources of alpha, and practical multi-period portfolio optimization.

1. Introduction

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 (Markowitz, 1952). 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? First, Markowitz quantified return and risk of a security, using the statistical measures of its expected return and standard deviation. Second, Markowitz suggested that investors should consider return and risk together, and determine the allocation of funds among investment alternatives on the basis of their return-risk trade-off. Before Markowitz's seminal article, the finance literature had treated the interplay between return and risk in an ad hoc fashion.

The idea that sound financial decision-making is a quantitative trade-off between return and risk was revolutionary for two reasons. First, it posited that one could make a quantitative evaluation of portfolio return and risk jointly by considering security returns and their co-movements. 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 correlations 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 optimization (MVO) problem 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.

Markowitz's work has had a major impact on academic research and the financial industry as a whole. Some internet searches we did as of the writing of this paper revealed the following numbers:

- 19,016 articles in Google Scholar cite Markowitz's original paper "Portfolio Selection".
- When searching for "modern portfolio theory" we obtained:
  About 590,000 hits in Google.
  531 YouTube videos.
  217 books on Amazon.
  Many thousands of tweets on Twitter.

MVO is used both for constructing portfolios of individual assets (*asset level*) and for asset allocation (*asset class level*). While in this paper we focus on the former application, the majority of the techniques we discuss are applicable to asset allocation, optimization on the asset class level is often considered easier than on the asset level, primarily because of the small number of asset classes. Interestingly, today more than 60 years later, risk-return optimization

* Corresponding author. Tel.: +1 (212) 998 4855.
E-mail addresses: petter.kolm@nyu.edu (P.N. Kolm), reha.tutuncu@gs.com (R. Tütüncü), frank.fabozzi@edhec.edu (F.J. Fabozzi).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.10.060

at the asset level is still primarily done only at the larger and/or more quantitatively oriented firms. However, with the availability of optimization tools customized for portfolio and risk management more and more investment managers are using some form of risk-return optimization as part of their portfolio construction process. A major reason for the surprisingly slow adaptation by investment managers to apply quantitative risk-return optimization is that they have observed that directly "out-of-the-box" portfolio optimization tends to be unreliable in practice. Specifically, risk-return optimization can be very sensitive to changes in the inputs, especially when the return and risk estimates are not well aligned or when the problem formulation uses multiple, interacting constraints. As a result, many practitioners consider the output of risk-return optimization to be opaque, unstable, and/or unintuitive.

Estimation errors in the forecasts significantly impact the resulting portfolio weights. For example, it is well-known that in practical applications equally weighted portfolios often outperform mean–variance portfolios (DeMiguel, Garlappi, & Uppal, 2009; Jobson & Korkie, 1981; Jorion, 1985), mean–variance portfolios are not necessarily well-diversified (Green & Hollifield, 1992), portfolio optimizers are often "error maximizers" (Michaud, 1998), and mean–variance optimization can produce extreme or non-intuitive weights for some of the assets in the portfolio (Black & Litterman, 1991, 1992). Such examples, however, are not necessarily a sign that the theory of risk-return optimization is flawed. Rather, it means that the classical framework has to be modified *when used in practice* in order to achieve reliability, stability, and robustness with respect to model and estimation errors. We will review some of the common approaches for this purpose in this paper.

Our intention with this article is not to provide a survey of MVO, its extensions and related areas. Some surveys include Steinbach (2001), Rubinstein (2002), Fabozzi, Kolm, Pachamanova, and Focardi (2007), and Markowitz (2014). Admittedly, there are many important contributions and works that we do not cover due to space constraints. The main goal with this article is twofold.

First, we address some of the key aspects related to using portfolio optimization in practice. The inclusion of transaction costs in the portfolio selection problem may present a challenge to the portfolio manager, but is an important practical consideration. We discuss a standard approach on how to extend traditional asset allocation models to incorporate transaction costs. In practice, it is common to amend the mean–variance framework with various types of constraints that take specific investment guidelines and institutional features into account. We discuss the use of various categories of constraints in portfolio construction and methods that quantify their impact on the portfolios generated. One of the main criticisms of the MVO approach focuses on its dependence on estimated parameters; specifically, expected returns and covariances, and its sensitivity to errors in these estimates. We outline various approaches that exist in the literature to mitigate the impact of estimation errors, including Bayesian methods, the Black–Litterman approach, and robust optimization techniques.

Second, we selectively highlight some of the new trends and developments in MVO and its related areas. Due to space constraints, we cannot survey all new trends in this area. While admittedly our choice is subjective, it is based on what we believe are as some of the important developments in this area for the use of MVO and its extensions in practice. In particular, we discuss the recent focus on diversification methods and provide a summary of the developments related to risk-parity portfolios. We also provide a formalization of the problem of mixing different sources of alpha and address some of the challenges that arise from these formulations. Finally, we outline some of the recent literature on practical multi-period portfolio optimization.

The paper is organized as follows. In Section 2, we review classical MVO. In Section 3, we discuss some of the most common ways on how to address the challenges one encounters when implementing MVO in practice. In Section 4, we highlight some interesting new directions and trends in MVO and related areas.

2. Mean–variance optimization

We consider an investment universe of n assets $S_1, S_2, \ldots, S_n$ with uncertain future returns $r_1, r_2, \ldots, r_n$. We denote by $r = [r_1, \ldots, r_n]^\top$ the vector of these returns. A portfolio is represented by the n-dimensional vector $\omega = [\omega_1, \ldots, \omega_n]^\top$ where ω_i denotes the proportion of the total funds invested in security i. The (uncertain) return of the portfolio, r_P, depends linearly on the weights

$$r_P(\omega) = \omega_1 r_1 + \cdots + \omega_n r_n = \omega^\top r$$

We denote by σ_i the standard deviation of r_i, ρ_{ij} denote the correlation coefficient of the returns of assets S_i and S_j (for $i \neq j$), and Σ the (symmetric) $n \times n$ covariance matrix of the returns of all the assets, i.e.

$$\Sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1n} \\ \sigma_{21} & \sigma_{22} & \cdots & \sigma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{n1} & \sigma_{n2} & \cdots & \sigma_{nn} \end{bmatrix}$$

where $\sigma_{ii} = \sigma_i^2$ and $\sigma_{ij} = \sigma_{ji} = \rho_{ij}\sigma_i\sigma_j$ (for $i \neq j$). All valid covariance matrices are positive semidefinite matrices (i.e. $\omega^\top \Sigma \omega \geqslant 0$ for all ω), or equivalently, all of their eigenvalues are nonnegative. In this paper, we assume that Σ satisfies the stronger property of positive definiteness, namely that $\omega^\top \Sigma \omega > 0$ for all $\omega \neq 0$. This is equivalent to assuming that none of the assets $S_1, S_2, \ldots, S_n$ can be perfectly replicated by a combination of the remaining assets. Positive definiteness assumption ensures that Σ is an invertible matrix. For a given portfolio ω, we can compute the variance and the standard deviation of the portfolio return as

$$V(\omega) = \omega^\top \Sigma \omega$$
$$\sigma(\omega) = \sqrt{\omega^\top \Sigma \omega}$$

The standard deviation of the portfolio return $\sigma(\omega)$, also referred to as portfolio volatility, is frequently used as a measure of risk of the portfolio ω.

We let Ω, a subset of $\mathbb{R}^n$, denote the set of permissible portfolios. In particular, $\omega \in \Omega$ means that the portfolio weights have to satisfy the constraints we impose upon our portfolio.

We represent the expected returns of the securities by

$$\mu = \begin{bmatrix} \mu_1 \\ \vdots \\ \mu_n \end{bmatrix}$$

where $\mu_i = E(r_i)$ for all $i = 1, \ldots, n$.

Using this notation, the MVO problem takes the form

$$\max_{\omega \in \Omega} \mu^\top \omega - \lambda \cdot \omega^\top \Sigma \omega$$

where λ is an investor specific risk aversion parameter that determines the trade-off between expected portfolio return and portfolio risk.

Alternative formulations of the MVO problem are obtained by either maximizing the expected return subject to an upper limit on the portfolio variance, or by minimizing the portfolio variance subject to a lower limit on the expected return, i.e.