Journal of Banking & Finance 37 (2013) 1232–1242

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

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

On portfolio optimization: Imposing the right constraints

Patrick Behr [a], Andre Guettler [b,*], Felix Miebs [b]

[a] Brazilian School of Public and Business Administration, Getulio Vargas Foundation, Rio de Janeiro, Brazil
[b] Department of Finance, Accounting and Real Estate, EBS Business School, Wiesbaden, Germany

ARTICLE INFO

Article history:
Received 6 December 2011
Accepted 26 November 2012
Available online 28 December 2012

JEL classification:
G11

Keywords:
Portfolio optimization
Shrinkage
Mean squared error
Bootstrap

ABSTRACT

We reassess the recent finding that no established portfolio strategy outperforms the naively diversified portfolio, $1/N$, by developing a constrained minimum-variance portfolio strategy on a shrinkage theory based framework. Our results show that our constrained minimum-variance portfolio yields significantly lower out-of-sample variances than many established minimum-variance portfolio strategies. Further, we observe that our portfolio strategy achieves higher Sharpe ratios than $1/N$, amounting to an average Sharpe ratio increase of 32.5% across our six empirical datasets. We find that our constrained minimum-variance strategy is the only strategy that achieves the goal of improving the Sharpe ratio of $1/N$ consistently and significantly. At the same time, our developed portfolio strategy achieves a comparatively low turnover and exhibits no excessive short interest.

1. Introduction

Since the foundation of modern portfolio theory by Markowitz (1952), the development of new portfolio strategies has become a horserace-like challenge among researchers. The sobering finding that theoretically optimal, utility maximizing portfolios perform poorly out-of-sample[1] can be attributed to the error prone estimation of expected returns, leading to unbalanced optimization results. This result directed researchers' attention to the minimum-variance portfolio, the only portfolio on the efficient frontier that simply requires the variance–covariance matrix as input parameter for the optimization. For instance, Merton (1980), Jorion (1985), and Nelson (1992) remark that variance–covariance estimates are relatively stable over time and can, hence, be predicted more reliably than expected returns. Nevertheless, DeMiguel et al. (2009b) have argued that no single portfolio strategy from the existing portfolio selection literature outperforms the naively diversified portfolio, $1/N$, consistently in terms of out-of-sample Sharpe ratio. Similar to Fletcher (2009), we evaluate in this paper a broader range of minimum-variance portfolios to challenge the findings of DeMiguel et al. (2009b). Additionally, we develop a constrained minimum-variance portfolio strategy that outperforms $1/N$ in terms of a lower out-of-sample variance and a higher Sharpe ratio while, at the same time, yielding a turnover and short interest that do not hamper the practical implementation of this portfolio strategy.

We propose a minimum-variance portfolio strategy with flexible upper and lower portfolio weight constraints. Incorporating these constraints into the portfolio optimization process trades off the reduction of sampling error and loss of sample information (Jagannathan and Ma, 2003). On the one hand, weight constraints ensure that portfolio weights are not too heavily driven by the sampling error inherent in parameter estimates based on historical data, which often leads to highly concentrated portfolios.[2] On the other hand, portfolio weight constraints cause a misspecification of the optimization problem because the resulting portfolio weights are less driven by potentially useful sample information (Green and Hollifield, 1992). Consequently, incorporating portfolio weight constraints into the portfolio optimization problem is promising if the input parameters are error prone.

We calibrate portfolio weight constraints such that the desired reduction of sampling error and the concomitant loss of sample information is traded-off. Using this shrinkage theory based framework, we introduce portfolio weight constraints which depend on the error inherent in the empirical variance–covariance matrix estimate. In particular, we impose the set of lower and upper portfolio weight constraints that minimizes the sum of the mean squared errors (MSE) of the covariance matrix entries. The latter serves as a loss function, quantifying the trade-off between the reduction of sampling error and loss of sample information.

* Corresponding author. Tel.: +49 611 7102 1236; fax: +49 611 7102 101236.
E-mail addresses: patrick.behr@fgv.br (P. Behr), andre.guettler@ebs.edu (A. Guettler), felix.miebs@ebs.edu (F. Miebs).

[1] See Frost and Savarino (1986), Jorion (1986), Michaud (1989), and Black and Litterman (1992), among others.

[2] See Green and Hollifield (1992), Chopra (1993), and Chopra and Ziemba (1993) for evidence concerning this point.

0378-4266/$ - see front matter
http://dx.doi.org/10.1016/j.jbankfin.2012.11.020

Our empirical results show that our constrained minimum-variance portfolio with lower and upper portfolio weight constraints achieves substantial out-of-sample variance reductions in comparison to various minimum-variance portfolios. We observe that the variance of our portfolio strategy achieves the lowest variances among all twelve considered portfolio strategies. In terms of risk adjusted performance, we observe that our portfolio strategy generates a 32.5% higher Sharpe ratio than $1/N$. This Sharpe ratio increase is statistically significant on five out of six datasets. Further, we observe that our portfolio strategy achieves on average a higher Sharpe ratio than every other benchmark strategy. This finding is robust with respect to the estimation window period, which we vary from 120 to 240 months.

Concerning the importance of weight constraints for the out-of-sample portfolio performance, we observe that imposing solitary lower or lower and upper weight constraints results in an equally effective risk reduction. The risk adjusted performance of both portfolios is on average equally good. However, we find that the constrained minimum-variance portfolio with lower and upper portfolio weight constraints achieves a less volatile Sharpe ratio over the various datasets than the constrained minimum-variance portfolio with solitary lower portfolio weight constraints. This is reflected in the statistical significance of the outperformance over $1/N$. While imposing lower and upper weight constraints yields on five out of six datasets a significantly higher Sharpe ratio than $1/N$, imposing solitary lower weight constraints yields only on two datasets a significantly higher Sharpe ratio. Further, we observe that the constrained minimum-variance portfolio with lower and upper portfolio weight constraints yields a lower turnover and short interest than the constrained minimum-variance portfolio with solitary lower weight constraints. Hence, we find that imposing lower and upper portfolio weight constraints is beneficial with respect to the resulting out-of-sample performance and the practical implementation of the portfolio strategy.

The impact of our ex ante calibrated weight constraints varies with respect to the size of the investment universe. While the imposed constraints are loose for small portfolios, they are comparatively tight for larger investment universes. Specifically, we observe for portfolios comprising 30 or more assets, that the ex ante calibrated lower portfolio weight constraints are close to zero, i.e. a short-sale constraint. While the lower portfolio weight constraints of our minimum-variance strategies with solitary lower, respectively lower and upper weight constraints are similar across all investment universes, we find that the tightness of the additional upper constraint is particularly pronounced for larger universes.

Our paper contributes to three lines of literature. First, we add to the prevalent discussion whether optimized portfolios represent a preferable investment vehicle over $1/N$ by amending the empirical evidence of DeMiguel et al. (2009b) and Fletcher (2009). Contrary to recent contributions to this ongoing discussion by Pflug and Pichler (2012) and Kritzman et al. (2010) that assess the conditions rendering $1/N$ an optimal strategy, i.e. the reasons for the unsatisfying performance of optimal portfolios, we develop a portfolio strategy that outperforms $1/N$ consistently and significantly. Thus, our paper relates to Chevrier and McCulloch (2008), DeMiguel et al. (2009a), and Tu and Zhou (2011), who claim to develop portfolio strategies achieving consistently higher Sharpe ratios than $1/N$. However, neither Chevrier and McCulloch (2008) nor Tu and Zhou (2011) provide statistical inference for their results.

Second, we extend previous work on portfolio optimization in presence of constraints. Alexander and Baptista (2006) and Alexander et al. (2007) evaluate the imposition of drawdown constraints, while DeMiguel et al. (2009a) and Gotoh and Takeda (2011) assess the impact of norm constraints on the portfolio optimization process. Our paper relates more closely to Frost and Savarino (1988), Grauer and Shen (2000), and Jagannathan and Ma (2003), evaluating the role of portfolio weight constraints. While the aforementioned work on weight constraints is concerned with arbitrarily chosen or ex post determined upper and/or lower constraints, we postulate a framework to determine these constraints ex ante. Our new approach should thus perform better out-of-sample given that flexible ex ante constraints are better able to suit the data at hand.

Third, our framework represents a new approach to the estimation of the variance–covariance matrix for portfolio optimization. Similar to the Ledoit and Wolf (2003, 2004a,b) shrinkage strategies, our approach imposes a data-dependent structure on the variance–covariance matrix. Our approach, however, requires fewer assumptions than the aforementioned shrinkage estimators. In particular, our framework requires neither any distributional assumptions, such as iid returns, nor the identification of a shrinkage target, which may have a significant impact on the out-of-sample portfolio performance.

The remainder of the paper is organized as follows. Section 2 outlines our methodology and data, while Section 3 contains the empirical results. Section 4 reports the robustness checks, Section 5 concludes.

2. Methodology and data

2.1. Calibrating portfolio weight constraints

We consider the standard myopic constrained minimum-variance portfolio optimization problem. In particular, the objective is the minimization of the portfolio variance, $w'\Sigma w$, where w denotes the column vector of optimal portfolio weights and Σ the population variance–covariance matrix. Since Σ is not observable, an estimate based on the available sample information has to be used instead. For the purpose of our constrained minimum-variance portfolio strategy, we use the sample variance–covariance matrix, $S = \frac{1}{\tau-1}R'R$, as an estimator of Σ, where R denotes the $\tau \times N$ matrix of de-meaned (in-sample) returns, τ the number of in-sample returns, and N the number of assets. Accordingly, the sample based estimate of the portfolio variance is given by $\hat{w}'S\hat{w}$, where $\hat{w}$ represents the sample estimate of w. Formally, we may express the constrained minimum-variance portfolio optimization problem as follows[3]:

$$\min_{\hat{w}} \quad \hat{w}'S\hat{w} \tag{1}$$
$$\text{s.t.} \quad \hat{w}'1_N = 1, \tag{2}$$
$$\hat{w}_i \geqslant w_{\min}, \quad \text{for } i = 1, 2, \ldots, N, \tag{3}$$
$$\hat{w}_i \leqslant w_{max}, \quad \text{for } i = 1, 2, \ldots, N. \tag{4}$$

The Kuhn–Tucker conditions (necessary and sufficient) are accordingly:

$$\sum_j \widehat{S}_{i,j}\hat{w}_j - \lambda_i + \delta_i = \lambda_0 \geqslant 0, \quad \text{for } i = 1, 2, \ldots, N,$$
$$\lambda_i \geqslant 0 \quad \text{and } \lambda_i = 0 \quad \text{if } \hat{w}_i > w_{\min}, \quad \text{for } i = 1, 2, \ldots, N,$$
$$\delta_i \geqslant 0 \quad \text{and} \delta_i = 0 \quad \text{if } \hat{w}_i < w_{max}, \quad \text{for } i = 1, 2 \ldots, N.$$

The notation is as follows: 1_N denotes the column vector of ones of appropriate size, λ the column vector of Lagrange multipliers for the lower portfolio weight constraint, δ the column vector of multipliers for the upper portfolio weight constraints, λ_0 the multiplier for the portfolio weights to sum up to unity, and $w_{\max}$ as well as $w_{\min}$ denote column vectors with uniform elements such that every asset has the same upper and lower weight constraint. Let $\hat{w}^{++}$ denote the solution to the constrained minimum-variance portfolio optimization problem in (1)–(4) and 1_N a conformable column vector of ones. We may then state the following proposition:

[3] We solve the constrained minimum-variance portfolio optimization problem using the Mosek quadratic programming solver quadprog for MATLAB.