

Husmann, Sven; Shivarova, Antoniya; Steinert, Rick

Article — Published Version

Sparsity and stability for minimum-variance portfolios

Risk Management

Provided in Cooperation with:

Springer Nature

Suggested Citation: Husmann, Sven; Shivarova, Antoniya; Steinert, Rick (2022) : Sparsity and stability for minimum-variance portfolios, Risk Management, ISSN 1743-4637, Palgrave Macmillan, London, Vol. 24, Iss. 3, pp. 214-235,
<https://doi.org/10.1057/s41283-022-00091-0>

This Version is available at:

<https://hdl.handle.net/10419/311074>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



Sparsity and stability for minimum-variance portfolios

Sven Husmann¹ · Antoniya Shivarova¹ · Rick Steinert¹ 

Accepted: 14 February 2022 / Published online: 11 March 2022
© The Author(s) 2022

Abstract

The popularity of modern portfolio theory has decreased among practitioners because of its unfavorable out-of-sample performance. Estimation risk tends to affect the optimal weight calculation noticeably, especially when a large number of assets are considered. To overcome these issues, many methods have been proposed in recent years, but only a few address practically relevant questions related to portfolio allocation. This study therefore uses different covariance estimation techniques, combines them with sparse model approaches, and includes a turnover constraint that induces stability. We use two datasets of the S&P 500 to create a realistic data foundation for our empirical study. We discover that it is possible to maintain the low-risk profile of efficient estimation methods while automatically selecting only a subset of assets and further inducing low portfolio turnover. Moreover, we find that simply using LASSO is insufficient to lower turnover when the model's tuning parameter can change over time.

Keywords Minimum-variance portfolio · LASSO · Turnover constraint · Out-of-sample variance · Asset selection · Short-sale budget

Introduction and main idea

The mean-variance portfolio optimization of Markowitz (1952) is still one of the most widely used approaches for selecting an optimal portfolio of assets with uncertain returns. To implement this approach in practice, one needs to estimate two sets of parameters—expected asset returns and the covariances of asset returns—which are traditionally estimated using the sample means and sample covariances of past

✉ Rick Steinert
steinert@europa-uni.de

Sven Husmann
husmann@europa-uni.de

Antoniya Shivarova
shivarova@europa-uni.de

¹ Europa-Universität Viadrina, Große Scharrnstraße 59, 15230 Frankfurt (Oder), Germany



returns, respectively. Unfortunately, the literature has extensively found that portfolios based on these estimates exhibit extremely poor out-of-sample performance, as the errors in the parameter estimates are carried over to the portfolio weights.¹

However, one portfolio on the efficient frontier that does not require researchers to estimate the mean decreases the estimation error: the global minimum-variance (GMV) portfolio. Interestingly, DeMiguel et al. (2009b) show that the mean-variance portfolio is outperformed out-of-sample not only by the naive portfolio but also by the GMV portfolio in terms of the Sharpe ratio as well as the certainty-equivalent for most of their investigated datasets. This is a remarkable finding, as the GMV portfolio specifically aims to reduce the variance and not to increase the Sharpe ratio.

In this study, we therefore introduce a sparse and stable model approach that focuses on the minimum-variance portfolio. We regard a portfolio as sparse if it selects a small number of assets out of a large investment space and as stable if its weights exhibit only small changes for each rebalancing step over time.

We show that recent advances in estimating covariance matrices have improved the risk profile of this portfolio type vastly. Nonetheless, we find that even the most accurate and recent covariance estimation techniques need to be updated to meet the common requirements of investors, especially in terms of transaction costs. For that, we focus on two specifications: a sparse selection of a subset of all assets and low turnover. This study is the first to show that combining highly efficient covariance estimators with penalty terms and turnover constraints can lead to portfolios that have the same low risk as their unconstrained counterparts, while simultaneously keeping a lower number of assets as well as lower turnover. We further provide new evidence on the usage of a mere penalized optimization for the weights toward achieving low turnover.

The remainder of this paper is organized as follows. After this section, we review the subject to provide insights from the scientific literature as well as from a practitioner's perspective. We then introduce three model setups. In the empirical study section, we describe our methodology, provide empirical evidence, and draw conclusions. A summary of our work and final remarks are provided in the last section.

Review of the literature

In recent decades, research on estimating covariance matrices under specific premises has gained increasing popularity. Researchers from different fields have adopted various strategies to tackle the many issues arising because of high dimensionality in the data and ill-conditioned covariance estimation. For instance, Bouchaud and Potters (2009), Fan et al. (2013), and Ledoit and Wolf (2017) develop methods based on random matrix theory to estimate covariance matrices. Shrinking approaches have also been established, which have in common that the covariance

¹ See, for example, Michaud (1989), Best and Grauer (1991), Chopra and Ziemba (1993), Broadie (1993), Litterman (2003), and DeMiguel et al. (2009a, 2009b).



matrix estimated with maximum likelihood (ML) is mixed with one or many target matrices, as shown, for example, by Ledoit and Wolf (2004). Other researchers have focused more on the time dependency of returns; see, for instance, Engle et al. (2017) for a recent study.

Owing to the increase in the transparency of research and advancement of computational power, implementing these models has become gradually easier for practitioners. Nonetheless, with the rise in data availability, investors seek diversification in large markets, but are limited by organizational as well as legal restrictions. The aforementioned estimation procedures, however, do not always make it straightforward to include these requirements. For example, the option to choose from a large set of stocks creates the problem of selection, as it is unfavorable for investors to hold a large number of stocks with a small relative weight (Lobo et al. 2007; Takeda et al. 2013). This is mostly because of the fixed costs associated with including each asset. Further, the matter of transaction costs also plays a crucial role in portfolio choice. If, over time, the weights of the portfolio change too much and thus require frequent rebalancing, the investor faces unnecessarily high costs. To tackle these issues, researchers including Konno and Wijayanayake (2002) and Lobo et al. (2007) have independently developed methods to include costs in the portfolio optimization itself. Another important, commonly applied restriction is the exclusion of short-sale positions, which is usually induced by law.

These real-world constraints are not only relevant because of legal and other regulatory circumstances. Jagannathan and Ma (2003), for example, show that self-imposed short-sale constraints can significantly improve the out-of-sample performance of portfolios. In particular, so-called norm constraints, or penalizing constraints, which help induce sparsity or shrinkage, can both reduce the amount of necessary assets and improve the estimation accuracy. They are implemented by penalizing the p -vector norm of the asset weights with an additional factor, usually called δ . Since Brodie et al. (2009) and DeMiguel et al. (2009a) introduced the ℓ_1 -norm and squared ℓ_2 -norm to portfolio optimization, these have gained increasing attention in this field of research. The ℓ_1 -norm is primarily applied to create sparse portfolios (i.e., portfolios with only a few active positions and, thus, lower overall estimation risk). The squared ℓ_2 -norm controls the balance of a portfolio, which can be measured by the deviation of its weights from the weights of an equally weighted portfolio. Another norm tailor-made for subset selection is the ℓ_0 -norm, which overcomes the issue of the ℓ_1 -norm of deselecting potentially relevant assets (Takeda et al. 2013). Advances in this field have included, for instance, Fastrich et al. (2014), who compare standard cardinality constraints with different ℓ_p -norms. Despite its high quality, however, the ℓ_0 -norm suffers from an absence of feasible solutions for estimating portfolios.

Model setup

To introduce our modeling approach, we start with the standard approach for constructing a GMV portfolio with n assets and extend it further with restrictions to induce sparsity as well as stability. We assume that the investor uses a one-time optimization in the



current period and readjusts his or her investment decision in subsequent periods by repeatedly executing the one-time portfolio optimization. The true covariance matrix Σ , unknown to the investor, needs to be estimated using the return data of τ historic time points.

Standard minimum-variance portfolios

To find the portfolio exhibiting the lowest variance among all assets, an investor faces the following optimization problem:

$$\hat{w} = \arg \min_w w' \hat{\Sigma} w, \quad (1)$$

$$\text{s.t. } Aw = a, \quad (2)$$

where \hat{w} is the estimated vector of portfolio weights, $\hat{\Sigma}$ is the estimated covariance matrix, and (2) represents the sum constraint. The latter means that all weights must sum to 1, as we choose $A = \mathbf{1}_n$ and $a = 1$, where $\mathbf{1}_n$ is the n -dimensional vector of ones. We refer to the GMV model as the standard model.

Sparse minimum-variance portfolios

If a smaller number of assets are selected from the whole investment space, the optimization problem can be adjusted by adding an ℓ_1 -norm constraint, often referred to as the least absolute shrinkage and selection operator (LASSO), as follows:

$$\hat{w} = \arg \min_w w' \hat{\Sigma} w, \quad (3)$$

$$\text{s.t. } Aw = a, \quad (4)$$

$$\|w\|_1 \leq \delta. \quad (5)$$

Equation (5) introduces a sparsity parameter δ , which controls the shrinkage of the portfolio weights toward zero. Choosing a high δ value will lead to the same result as the standard optimization problem, whereas a sufficiently small δ will restrict the parameter space to a few assets. The solution to the optimization problem can still easily be found using standard quadratic programming with linear constraints, as it is possible to reallocate w into its positive part $w^+ = \max(w, 0)$ and $w^- = \max(-w, 0)$. The left-hand side of constraint (5) can then be rewritten as $\|w\|_1 = \mathbf{1}_n w^+ + \mathbf{1}_n w^-$. The whole optimization problem then becomes



$$\begin{aligned}
\begin{bmatrix} \hat{w}^+ \\ \hat{w}^- \end{bmatrix} &= \arg \min_{w^+, w^-} \begin{bmatrix} w^+ \\ w^- \end{bmatrix}^T \begin{bmatrix} \hat{\Sigma} & -\hat{\Sigma} \\ -\hat{\Sigma} & \hat{\Sigma} \end{bmatrix} \begin{bmatrix} w^+ \\ w^- \end{bmatrix} + \begin{bmatrix} \lambda 1_n \\ \lambda 1_n \end{bmatrix}^T \begin{bmatrix} w^+ \\ w^- \end{bmatrix} \\
\text{s.t. } [A, -A] \begin{bmatrix} w^+ \\ w^- \end{bmatrix} &= a \text{ and } \begin{bmatrix} 0_n \\ 0_n \end{bmatrix} \leq \begin{bmatrix} w^+ \\ w^- \end{bmatrix},
\end{aligned} \tag{6}$$

which is a quadratic optimization with linear constraints and a Lagrange parameter λ . Owing to the combination of (4) and (5), the parameter space cannot be fully restricted (e.g., to up to zero assets). Due to its composition, we refer to this model as the LASSO model.

Sparse and stable minimum-variance portfolios

In a real-world application of a minimum-variance portfolio, investors are prone to costs, which are related to the turnover of the portfolio (e.g., transaction costs). The introduction of a shrinkage-type constraint such as (5) can, to some degree, account for transaction costs, as the parameter δ will penalize high asset weights w_i and therefore indirectly reduce the possibility of vast changes between the subsequent rebalancing time points. However, we argue that LASSO alone cannot sufficiently decrease turnover, as it has no information on past weights. Although the weights overall are relatively small, if many of these change at a time or many weights that were earlier deselected (i.e., set to zero) are now selected, turnover might still be reasonably high. Hence, we include a specific turnover constraint that works as a proxy for transaction costs. The optimization problem now changes to

$$\hat{w} = \arg \min_w w^T \hat{\Sigma} w, \tag{7}$$

$$\text{s.t. } Aw = 1, \tag{8}$$

$$||w||_1 \leq \delta, \tag{9}$$

$$w \leq k + w_o, \tag{10}$$

$$-w \leq k - w_o, \tag{11}$$

where the stability-inducing constraints (10) and (11), respectively, form the turnover constraints dependent on the weights of the previous optimization step w_o and a tuning parameter k , which controls the allowed change in the positions for one rebalancing step in both directions. The usual turnover constraint $|w| \leq w_o + k$ is rewritten as (10) and (11) to avoid using non-linear constraints. With respect to the adjustment of w to w^+ and w^- , as explained before, the whole optimization problem remains a quadratic problem with linear constraints and is therefore efficiently

solvable with standard optimization software.² Due to its design, we refer to this model as the LASSO with turnover (TO) constraint, in short, LASSO + TO model.

Model discussion

In theory, estimating the GMV portfolio, as introduced in Sect. 3.1, should be sufficient to obtain the portfolio with the lowest variance and, thus, the lowest risk. The restriction introduced in Sect. 3.2 would then become obsolete, as no matter how large the investment space is, including assets will always reduce or keep the variance at the same level, but never increase it. The additional constraints (10) and (11) would then function as merely a practitioners' constraint by reducing turnover and thus, transaction costs.

However, as pointed out earlier, the covariance matrix needs to be accurately estimated to obtain optimal results. As researchers have pointed out, even small estimation differences can lead to vast deviations from the true efficient frontier; see, for instance, Jobson and Korkie (1980, 1981), and Frost and Savarino (1986, 1988) for some of the earliest studies of this topic. Moreover, Kan and Zhou (2007) argue that the unbiased ML estimator for Σ has unwanted properties for specific ratios $\frac{\alpha}{\tau}$, even when the underlying return data follow a normal distribution.

Several methods have been proposed to reduce the estimation error for the covariance matrix, most applying some form of shrinkage (e.g., Ledoit and Wolf (2004)). Nonetheless, as any estimation error in the covariance matrix directly influences the estimation of the weights w of the portfolio, some authors shrink the weights w directly by, for instance, combining it with another portfolio. For instance, Tu and Zhou (2011) apply a mean-variance portfolio combined with an equally weighted portfolio. Jagannathan and Ma (2003) even argue that any constraint on the optimization procedure might help reduce the estimation error.

All the model approaches introduced in Sects. 3.1–3.3 can therefore reduce the estimation error in their own way. To choose a suitable standard GMV portfolio, we use highly efficient and recent estimators for the covariance matrix Σ . The constraint (5) of our sparse model will not only create sparsity but also directly reduce the estimation error in the weights w , as these will be shrunk toward zero. In our sparse and stable model presented in Sect. 3.3, we further stabilize the portfolio estimations by introducing the turnover constraint. This directly forces the weights to change only in a small window of length $2k$ and hence indirectly applies another way of prohibiting estimation errors due to potential misspecifications in the data.

Overall, this framework allows us to study the behavior of LASSO, one of the most common sparsity-inducing methods, when the covariance estimate has already adjusted to potential problems. We further check whether the common assumption that LASSO can reduce turnover significantly on its own continues to hold true when the covariance matrix is already estimated sufficiently (e.g., Brodie et al.

² Setting $k = \infty$, models (7) and (3) become the same and when $k = \infty$ and $\delta = \infty$, models (7) and (1) yield the same results.



(2009)) By doing so, we can gain insights into how the introduction of a common turnover constraint changes the risk profile of these portfolios.

Empirical study

As our results are solely based on an empirical analysis, it is crucial to employ a suitable empirical setup to ensure the validity and reproducibility of our findings. As an investor who uses a minimum-variance portfolio seeks, by definition, the portfolio exhibiting the lowest variance, our empirical study must reflect this objective. Indeed, all the investigated portfolios include so-called tuning parameters—parameters important for the optimization procedure but not yet optimized by theoretical analysis. Instead, one of these parameters is identified by cross-validation in combination with machine power, whereas the other will be set to a constant value.

In our case, we have two tuning parameters: δ , resulting from the LASSO constraint, and k , emerging from the turnover constraint. Owing to computational restrictions, we use δ as the tuning parameter to achieve the lowest variance and therefore optimize its value with cross-validation. The parameter k of the turnover constraint is kept constant throughout the dataset, as we only use it to reduce turnover compared with the unconstrained benchmark.

Hence, we do not set the value of δ so that it meets specific well-known constraints such as the short-sale constraint (see DeMiguel et al. (2009a)). In contrast to other authors such as Zhao et al. (2019), we do not want to achieve any practitioner's rule for investment and therefore do not keep δ as a constant, independent of the present market situation. By contrast, we allow the δ value to change in every period, as we always want to achieve the optimization goal, that is, finding the portfolio with the lowest variance. This, in our opinion, a more realistic approach leads to a more likely change in the chosen assets and higher turnover. This in turn provides another reason for imposing an additional turnover constraint, as in model (7).

Data

For our empirical study, we use S&P 500 stock price data from the Thomson Reuters EIKON database. This covers daily data from January 1998 to the end of December 2018, with $T = 5282$ observations overall. Our analysis is based on discrete returns, calculated as $r_t = \frac{P_t - P_{t-1}}{P_{t-1}}$. To avoid transforming the data and therefore potentially distorting valuable information, we only focus on those stocks present throughout the data period (319 stocks). To check whether dimensionality influences our results, we analyze the surviving 319 companies as well as a randomly generated subset of 100 stocks of the original 319. All our models are estimated by taking into account the returns of approximately the past two years of trading (i.e., $\tau = 2 * 252 = 504$ observations). This results in 4778 trading days of out-of-sample returns from our different model approaches.

To illustrate the structure of our two datasets, Fig. 1 presents their correlation plots. This figure shows the sample correlation of each possible pair of stock returns.






Fig. 1 Correlation plots for the stocks of our two datasets. The blue correlation indicates a strong positive correlation; white, slight to no correlation; and red, a strong negative correlation. (Color figure online)

As all the correlations are close to zero or positive for both datasets, we order the stocks according to their first principal component.³ Friendly (2002) and Wei et al. (2017) provide examples of an R-implementation. Both figures show that a large number of stocks exhibit strong correlations with each other. However, from the top left to bottom right of the figures, the overall correlation diminishes to a slight positive correlation, in some cases even close to zero. Further, the random selection of 100 out of the 319 stocks does not visually break the underlying correlation structure of the data, as both plots seem to have a similar appearance.

Variance estimators

To thoroughly analyze whether the introduced LASSO and turnover constraints decrease the number of assets as well as turnover, while maintaining a low-variance profile, we use some recent and efficient variance estimators. Starting with one of the most commonly used estimators among practitioners and researchers, we calculate the sample covariance estimator, defined as

$$\hat{\Sigma}_S = \frac{1}{\tau - 1} (R - \hat{\mu}1') (R - \hat{\mu}1')',$$

where $R \in \mathbb{R}^{n \times \tau}$ is the matrix of past returns and $\hat{\mu} \in \mathbb{R}^n$ the vector of expected returns (here, estimated as average returns). At high concentration ratios, $q = n/\tau \rightarrow 1$, the empirical variance, although unbiased, exhibits high estimation variance and, therefore, a high out-of-sample estimation error. To minimize this estimation error, a linear shrinkage procedure can be applied to the unbiased sample

³ In this way, we can easily display the magnitude of the correlation structure without checking, for instance, 50,721 correlations unit by unit, as would be needed in the case of 319 stocks.



estimator by combining it with a target covariance matrix. Following Ledoit and Wolf (2003), the variance estimator becomes

$$\hat{\Sigma}_{LW_L} = s\hat{\Sigma}_T + (1-s)\hat{\Sigma}_S, \quad (12)$$

where $\hat{\Sigma}_T$ is the estimate of a specific target covariance matrix and s is a shrinkage constant with $s \in [0, 1]$. Assuming identical pairwise correlations between all n assets, the target matrix is substituted with the constant covariance matrix as in Ledoit and Wolf (2004).

A more sophisticated shrinking method is non-linear shrinkage, as suggested by Ledoit and Wolf (2017). As this estimator shrinks the eigenvalues individually; small, potentially underestimated eigenvalues are pushed up, while large, potentially overestimated eigenvalues are pulled down. Without going into further detail, we write the non-linear shrinkage estimator as

$$\hat{\Sigma}_{LW_{NL}} = V\hat{E}_{LW_{NL}}V', \quad (13)$$

where V is the matrix of the orthogonal eigenvectors and $\hat{E}_{LW_{NL}}$ is the diagonal matrix of the shrunk eigenvalues, as shown by Ledoit and Wolf (2012, 2015). Because $\hat{\Sigma}_{LW_{NL}}$ is proven to be asymptotically optimal within the class of rotationally equivariant estimators, we might expect it to perform better than any of the aforementioned estimators, especially in cases of large concentration ratios.

Hence, we further extend our analysis to factor-based covariance estimation methods, which assume a specific structure in the covariances of asset returns. One promising example of that family of variance estimators is the principal orthogonal complement thresholding (POET) estimator provided by Fan et al. (2013). Here, the principal components of the sample covariance matrix $\hat{\Sigma}_S$ are used as factors. Moreover, subsequent adaptive thresholding with a threshold parameter θ is applied to the covariance of the residuals of the estimated factor model (see, e.g., Cai and Liu 2011).⁴ Therefore, the POET estimator has the form:

$$\hat{\Sigma}_{POET} = \sum_{i=1}^K \hat{\xi}_i v_i v_i' + \hat{\Sigma}_{u,K}^\theta, \quad (14)$$

where v_i is the eigenvector to asset return i , $\hat{\xi}_i$ is the corresponding eigenvalue, and $\hat{\Sigma}_{u,K}^\theta$ is the idiosyncratic covariance matrix after the applied thresholding procedure with threshold level θ .

In particular, the estimators (13) and (14) estimate the GMV portfolios well, and thus, they can be considered to be the state-of-the-art among homoscedastic variance estimators for return data.

⁴ For the application in our study, we use the R-code provided by Fan et al. (2013) within the R-package *POET* and repeatedly apply a separate cross-validation to obtain the number of factors K , as suggested by the authors.



Performance measures

To evaluate the out-of-sample performance of each portfolio, we report various performance measures, starting with the out-of-sample portfolio standard deviation σ_p and Sharpe ratio SR_p , defined as

$$\hat{\sigma}_p = \frac{1}{T - \tau} \sum_{t=\tau}^{T-1} (w_t' r_{t+1} - \hat{\mu}_p)^2, \quad (15)$$

$$\widehat{SR}_p = \frac{\hat{\mu}_p - r_f}{\hat{\sigma}_p}, \quad (16)$$

where w_t are the portfolio weights chosen at time t , $w_t' r_{t+1}$ is the out-of-sample portfolio return, $\hat{\mu}_p = \frac{1}{T - \tau} \sum_{t=\tau}^{T-1} w_t' r_{t+1}$ is the out-of-sample portfolio expected return. For the computation of the Sharpe ratio, we assume a risk-free interest rate $r_f = 0$.

Since we consider a variance minimization problem, daily out-of-sample portfolio variance is of utmost importance. Hence, we check whether the calculated out-of-sample variance of the LASSO-based method in (3) as well as the LASSO and turnover-based method in (7) have significantly different standard deviations than their standard counterpart in (1). Therefore, we perform the two-sided HAC test with the Parzen kernel for the differences in variances, as described by Ledoit and Wolf (2008), and report the corresponding p -values.

Furthermore, in accordance with the literature on portfolio optimization and estimation risk reduction, to approximate the arising transaction costs (e.g., DeMiguel et al. 2009a; Dai and Wen 2018), we use the average daily turnover

$$\text{Turnover} = \frac{1}{T - \tau - 1} \sum_{t=\tau+1}^{T-1} \sum_{j=1}^n \left(|w_{j,t+1} - w_{j,t}| \right), \quad (17)$$

where $w_{j,t+}$ denotes the portfolio weight in asset j before rebalancing at $t + 1$ but scaled back to sum to 1 and $w_{j,t+1}$ is the portfolio weight in asset j after rebalancing at $t + 1$.

We next evaluate the portfolio composition with respect to the number of non-zero investments and short sales as well as the development of the short-sale budget over time, defined as

$$\text{Average assets} = \frac{1}{T - \tau} \sum_{t=\tau+1}^T \sum_{j=1}^n \mathbb{1}_{\{w_{j,t} \neq 0\}}, \quad (18)$$

$$\text{Average short sales} = \frac{1}{T - \tau} \sum_{t=\tau+1}^T \sum_{j=1}^n \mathbb{1}_{\{w_{j,t} < 0\}}. \quad (19)$$

To shed more light the introduced models from the perspective of risk exposure, we include two different portfolio concentration measures. First, the concentration ratio



determines the distribution of assets exposures within a portfolio and is defined as the aggregate share of the n_b -largest weights within a portfolio.

$$\text{Concentration ratio} = \frac{1}{T - \tau} \sum_{t=\tau}^T \sum_{j=1}^{n_b} |w_{j,t}|, \quad (20)$$

where we set $n_b = 5$ throughout our empirical study. Naturally, a lower concentration ratio implies better portfolio exposure and diversification.

Second, following Choueifaty and Coignard (2008), we compute the diversification ratio as the ratio of the weighted average of asset volatilities divided by the portfolio volatility.

$$\text{Diversification ratio} = \frac{1}{T - \tau} \sum_{t=\tau}^T \frac{w'_t \sigma_t}{\sqrt{w'_t \Sigma_t w_t}}, \quad (21)$$

where Σ_t is calculated as in Eq. (13) and $\sigma_t = \sqrt{\text{diag}(\Sigma_t)}$. Due to its definition, the diversification ratio takes up values ≥ 1 and is higher when the portfolio exhibits higher (better) diversification levels.

Finally, to gain more insight into the model structure, we analyze the final δ values of model types (3) and (7). All the values are reported on a daily basis.

Course of action

For our empirical work, we use a non-expanding rolling window study that incorporates cross-validation for our tuning parameter. As mentioned earlier, we evaluate the tuning parameter δ for the LASSO constraint, so that it may change each day. The parameter k is left constant over time, set to 0.0005 for the 319 S&P dataset and 0.001 for the 100 S&P dataset. These values were found by checking different values k for each dataset in a small subsample. Changing k by a reasonably large number did not result in vastly different outcomes. In general, choosing a too large value for k leads to a portfolio that still has high turnover, whereas choosing it to be too small worsens its risk/return profile. The more the assets considered, the lower k should be.

To analyze the impact of both the sparsity (LASSO) and the stability (LASSO + TO) constraints, we implement a simple one-fold cross-validation for the tuning parameter δ . However, because of the described model representation of (6), which allows us to simplify the absolute value constraint, we apply our cross-validation toward the λ Lagrange parameter instead of δ and restore all δ values in a second step by simply calculating $\|w\|_1 = \delta$.

We start with $t = 1$, January 2, 1998, and use the following daily returns up to $t = 504$ to create an in-sample dataset covering approximately two years of daily returns. From that data sample, we take another smaller subsample for our cross-validation consisting of the first $504 - 20 = 484$ observations. We then calculate models (3) and (7) using 20 different λ values chosen from a linear sequence of numbers from $\lambda_{t+1} = \lambda_t + 0.00001$ to 0, whereas we initialize λ_1 with 0.00001. The resulting



weights of these 20 models are then applied to the first subsequent daily return of the cross-validation subsample (here, the 485th observation) to create an individual daily portfolio return for both models. The subsample is then shifted by one and the procedure carried out again. This is repeated until we reach the 20 observations we previously omitted. We then compare the standard deviations of the 20 out-of-sample cross-validation returns for each λ for both models individually. After receiving an optimal λ , chosen to be that corresponding to the lowest standard deviation, we set our final λ to be λ_t for each model individually. Next, we calculate the weights using models (3) and (7) for all the selected data on daily in-sample returns. The true out-of-sample returns are then constructed by multiplying the calculated weights by the returns of the following period (here, the 505th observation). As model (1) needs no cross-validation, we calculate it only at this point to receive its out-of-sample portfolio return as well. We then proceed by shifting the former in-sample data by one period (i.e., a day) and repeat the procedure 4778 times until the last out-of-sample daily return covers December 31, 2018.

Results

After applying the above mentioned techniques to the 319 and 100 S&P 500 stock returns, we can analyze the suggested portfolio strategies. The major results of our study for the first dataset with 319 stocks are provided by Table 1 by sorting according to the implemented variance estimation methods in the columns and the three modeling approaches in the rows. All the values are based on out-of-sample data. First, we report the standard deviation p.a., which is the most important measure for comparing minimum-variance portfolios. Second, we use the Sharpe ratio p.a. to evaluate the return-risk profile of the suggested models. Moreover, as a proxy for transaction costs, we calculate the average daily turnover. We further report average assets, namely the average number of assets with weights different than zero included in the daily final portfolio choice, and the average short sales per day as the average of all negative weights. Finally, we compute both the concentration and the diversification ratios of all implemented model combinations. For a better overview of the data, every number in bold corresponds to the model that performed best under a certain variance estimation technique and performance measure. The underlined numbers represent the overall best model in that performance category.

In terms of standard deviation, the standard approach leads to the highest risk levels of all the estimation techniques, with only one exception (i.e., the POET estimator). The difference between the standard approach and the LASSO approach with POET is, however, not significant, with a p -value of almost 90%. For less effective estimators such as ML and Ledoit and Wolf (2003), LASSO-based models can significantly decrease the variance, as indicated by a p -value of almost 0. This finding is in accordance with that of Dai and Wen (2018) who conduct a similar study combining the estimator of Ledoit and Wolf (2003) with LASSO and compare it with the short sale-constrained GMV portfolio. Here, we can further compare the effect of LASSO across various estimators. Interestingly, the LASSO model with the linear shrinkage of Ledoit and Wolf (2003) performs even better than their



Table 1 Standard deviation p.a., Sharpe ratio p.a., average turnover per day, average assets as the mean of all non-zero weights, average short sales as the mean of all weights greater 0, the concentration, and diversification ratios for the different models applied to the 319 S&P dataset

	ML	Ledoit and Wolf (2003)	Ledoit and Wolf (2017)	Fan et al. (2013)
Standard deviation p.a.				
Standard	0.1387	0.1064	0.1004	<u>0.0974</u>
LASSO	0.1025 (0.0000)	0.0990 (0.0000)	0.1000 (0.5392)	0.0975 (0.8979)
LASSO + TO	0.0998 (0.0000)	0.0991 (0.0000)	0.0997 (0.5106)	0.0981 (0.4153)
Sharpe ratio p.a.				
Standard	0.7268	0.6644	0.7363	0.6819
LASSO	0.7400	0.6373	0.7128	0.6880
LASSO + TO	0.7153	0.6426	0.7396	<u>0.7781</u>
Turnover				
Standard	0.7856	0.3279	0.1757	0.1472
LASSO	0.4125	0.3979	0.3062	0.2811
LASSO + TO	0.1439	0.1379	0.1392	<u>0.1215</u>
Average assets				
Standard	319.00	319.00	319.00	319.00
LASSO	<u>120.19</u>	144.49	188.78	193.97
% of full	37.68	45.29	59.18	60.80
LASSO + TO	265.22	255.53	267.06	233.57
% of full	83.14	80.10	83.72	73.22
Average short sales				
Standard	152.26	154.30	133.11	141.73
LASSO	<u>44.94</u>	60.87	67.29	77.81
% of full	14.09	19.08	21.09	24.39
LASSO + TO	125.97	122.37	117.45	103.73
% of full	39.49	38.36	36.82	32.52
Concentration ratio				
Standard	0.6578	0.4683	<u>0.2238</u>	0.3058
LASSO	0.4668	0.4577	0.2571	0.3415
LASSO + TO	0.4401	0.4363	0.2516	0.3373
Diversification ratio				
Standard	1.7398	1.9365	<u>3.0017</u>	2.2852
LASSO	2.5232	2.2680	2.8770	2.3243
LASSO + TO	2.6667	2.3746	2.9479	2.3932

The column headers represent the used variance estimation technique and rows report the results of the GMV calculated with (1) (Standard), (3) (LASSO), or (7) (LASSO+TO). The results in bold represent the best model of that column for a specific measure and the underlined number represents the best overall model for that measure. The numbers in brackets represent the p -values of the test $H_0 : \sigma_{LASSO} = \sigma_{Standard}$ and $H_0 : \sigma_{LASSO+TO} = \sigma_{Standard}$, respectively. All the results are calculated based on the whole out-of-sample period consisting 4778 daily observations




non-linear version (Ledoit and Wolf (2017)), independent of the model with which it is combined. Moreover, the POET estimator strictly dominates all the other variance estimators irrespective of whether a LASSO constraint is imposed. One possible explanation is that POET is the only covariance estimation method to incorporate an underlying factor model, a popular theory in finance for explaining the cross-section of returns. In general, introducing the stability constraint within the LASSO + TO model does not noticeably weaken the results for variance compared with the LASSO model. Hence, both the LASSO and LASSO + TO models retain the low-variance profile of the standard approach and can even at times significantly outperform it, providing suggestive evidence that incorporating the LASSO might improve variance overall as well.

To add more information on the overall return-risk profile of the suggested portfolio models, we include the Sharpe ratio p.a. in our empirical study. The LASSO+TO portfolio strategy with the POET estimator results in the highest Sharpe ratio across all model combinations. Nevertheless, the sparsity and stability constraints as well as the efficiency of the covariance estimators do not seem to influence the Sharpe ratio levels consistently. This is not surprising, since we optimize both the portfolios and the sparsity parameters solely according to the underlying risk level. Moreover, the Sharpe ratios are known to suffer under high standard errors (Ledoit and Wolf 2008). Accordingly, the included Sharpe ratios serve as an informative addition to our already extensive empirical results. The main purpose of our study remains to identify a portfolio that maintains a low-risk profile while investing in less assets (sparsity) with less turnover (stability). Therefore, in the following, we analyze in detail the structure of the various portfolios.


In terms of turnover, calculated in Eq. (17), the findings are surprising and novel. The table suggests that LASSO models without the turnover constraint have overall higher turnover than the standard model, except for the ML estimator. Merely in the case of the latter, LASSO reduces the turnover, a finding that coincides with those of, for example, Brodie et al. (2009). Moreover, the more efficient a variance estimation method, the larger is the relative increase of turnover due to the incorporated sparsity constraint as in (3). Only when introducing the stability constraint, as in model (7), does turnover vastly decrease to levels below that of the standard models. This holds true for all the variance estimators. To further support our findings and provide possible explanations, Fig. 2 provides an overview of the δ values of the portfolios over time, demonstrating the variance estimation technique of Ledoit and Wolf (2017) as a benchmark.

Figure 2a shows the calculated δ values for the standard as well as the LASSO model over time. The parameter δ in Eq. (5) corresponds to the sum of absolute portfolio weights. While this parameter can only be calculated afterward in the standard model, it is directly linked to the most influential tuning parameter in the LASSO model: the λ Lagrange parameter. δ also provides information on the short-sale budget as well as practitioners' investment rules, as shown by Zhao et al. (2019). Naturally, as represented by the black line of standard models' δ , this parameter is stable over a few months but can have high fluctuations over years. In that sense, δ varies between 2.5 and 5 throughout the whole out-of-sample period. The LASSO model, however, does not reveal such high levels of δ , and thus implies a





(a) Calculated δ values for the LASSO and Standard model.



(b) Calculated δ values for the LASSO with the turnover constraint and Standard model.

Fig. 2 Out-of-sample calculated δ values for the GMV portfolio types over time, where variance is estimated following Ledoit and Wolf (2017)

lower short-sale budget than its standard counterpart. This can be seen by the orange dots, but even better by the blue line, which represents the simple moving average of 30 days for these orange dots. The moving average is only applied for visualization purposes to show the variation in δ .

In general, two main findings come from Fig. 2a. First, the overall short-sale budget with LASSO is always lower than or equal to that of the standard portfolio approach, which coincides with the findings in the literature (e.g., Brodie et al. (2009)). Second, the optimal δ for LASSO portfolios seems to be much more volatile



than that for the standard portfolio. In particular, the second finding causes the turnover of these portfolios to be high. If the chosen optimal λ of each cross-validation out-of-sample step needs to be adjusted for every new step to minimize the standard deviation, the portfolio weights fluctuate more than usual. In studies focusing on LASSO-constrained portfolios, this is usually overseen, as most research assumes a stable sparsity parameter λ and, thus, a stable δ over time (e.g., Zhao et al. (2019)). Figure 2b provides more evidence in this regard. The image illustrates the same characteristics, but now for the LASSO + TO model. The orange dots seem to be much closer to each other than before. In addition, the blue line is far less volatile. This higher stability of δ over time leads to more stable weights and, thus, lower turnover.


Table 1 provides more information on the average number of assets over time, calculated as in (18). Naturally, the standard portfolio includes all 319 assets, as no restriction is imposed. However, the LASSO and LASSO + TO portfolios both reduce the number of stocks in the portfolio. Again, for the unconstrained LASSO model, the number of included assets increases when a more efficient covariance estimation technique is used. For the worst performing estimator in terms of standard deviation (i.e., ML), the LASSO model only selects 37.68% of all stocks on average, whereas this is 60.80% for the POET estimator of Fan et al. (2013). Imposing an additional turnover constraint changes these results slightly, meaning that the number of included assets increases compared with the regular LASSO model. However, compared with the standard model, the LASSO + TO model is still strictly dominating in terms of sparsity, as it reduces the number of included assets to less than 84% of the whole asset universe for all cases.

To get a fuller picture on the included assets, we investigate the average short sales, calculated as in (19). In accordance with the literature and as already seen in Fig. 2, the LASSO model tends to constrain short sales, leading to a small number of short sales overall compared with the standard approach. In most cases, LASSO can generally halve the amount of short sales of the standard approach, whereas including a turnover constraint induces a smaller reduction than LASSO alone. However, the LASSO + TO model still decreases the short-sale budget overall compared with the standard method. Figure 3 provides further details.

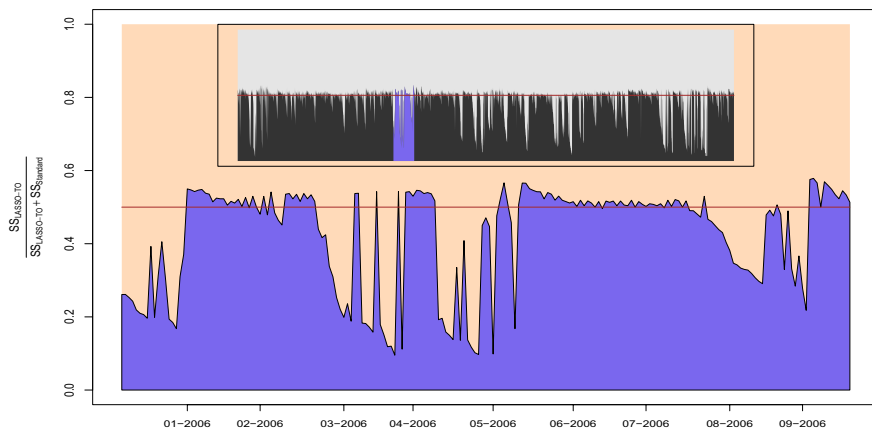
Here, we compare the amount of short sales of the LASSO models with that of the standard model over time. The black and gray areas of the two subfigures represent the full out-of-sample period, whereas the blue areas are an example of a chosen subset to provide an easier illustration. The figure specifically compares the proportion $\frac{SS_{LASSO(+TO)}}{SS_{LASSO(+TO)} + SS_{Standard}}$, where SS stands for the amount of short sales on a specific day for a specific portfolio. This number is beneficial to analyze, as it explains which model exhibits higher amounts of short sales for a given time point. The number 0.5 stands for the cut-off point, for which one or the other model will have higher short sales, assuming that lower short sales is advantageous.

Figure 3a, which compares the LASSO model with the standard model, shows that the LASSO model never exhibits more short sales than its standard version throughout the period. This is graphically illustrated by the black and blue areas never exceeding the red line. On some occasions (e.g., some days in August 2006),





(a) Comparison of short sales for the LASSO and Standard model.



(b) Comparison of short sales for the LASSO with the turnover constraint and Standard model.

Fig. 3 Development of short sales compared with the standard model over the whole out-of-sample period as the black area and a snapshot of 10 months as the blue area. Data are normalized, so that 0 stands for 0% of the combined short sales of standard and LASSO(+ TO) and 1 for 100%. The red line describes the breaking point of 0.5, where both methods show the same amount of short sales. (Color figure online)

the short-sale budget of the LASSO model becomes 0 (i.e., there were no short sales).

Figure 3b, which compares the LASSO + TO model with the standard model, does not share the same properties. Here, the red line is often exceeded by the black and blue areas, leading to a portfolio with higher short sales than under the standard method. Moreover, short sales were never absent from the portfolio. Nevertheless, as already shown in Table 1, the LASSO with the turnover constraint model still



reduces the short-sale budget, as the gray area below the red line is greater than the black and blue areas above the red line.

Finally, in Table 1, we report both the concentration (as in (20)) and the diversification (as in (21)) ratios across all model combinations. Here, first, we need to remind that a lower concentration ratio indicates a lower portfolio risk exposure. Reversely, a higher diversification ratio implies a better diversified portfolio. Overall, for the less efficient covariance estimation methods, we can observe strong effects of the introduced sparsity and stability constraints on both performance measures. The effects are most emphasized for the ill-conditioned ML estimator. In this case, the LASSO model decreases the portfolio concentration by approximately 30% and increases the portfolio diversification ratio by roughly 45%. Interestingly, both the LASSO and the LASSO + TO model with the ML estimator achieve higher diversification ratios than their counterparts with the POET estimator. Similar, although less pronounced, are the results considering the linear shrinkage estimator by Ledoit and Wolf (2003). For exactly these two estimation methods, both the LASSO and LASSO + TO models achieve the highest relative decrease in standard deviation. Intuitively, as LASSO models reduce the number of assets, it could be assumed that LASSO will increase the concentration ratio. On average, spreading an investment over 100 assets will result in lower absolute weights than spreading it over 10 assets. However, since the sparsity constraint is designed to remove assets with only a low absolute impact on the risk profile of the portfolio, the LASSO and LASSO + TO models manage to both reduce the number of assets and increase the concentration ratio for less efficient covariance estimators. Considering the more efficient, state-of-the-art covariance methods, it does not seem that the introduced LASSO and LASSO with turnover constraints can improve the concentration ratios. More importantly, the standard model with the non-linear shrinkage by Ledoit and Wolf (2017) achieves the best concentration and diversification levels across all model combinations. However, in this case, the differences to the LASSO and LASSO + TO models are negligible. We can therefore still conclude that the sparsity and stability constraints improve the overall portfolio profile, even in the case of highly efficient estimators. While we manage to keep a low-risk level and high diversification ratios, we strongly reduce the number of invested assets and the overall portfolio turnover rate.

To further strengthen our results, we examine 100 randomly selected stocks of the S&P 500. Table 2 shows the results. The findings in Table 1 are supported by those summarized in Table 2. Ledoit and Wolf (2017) and Fan et al. (2013) seem to be superior variance estimation techniques. Now, the standard deviation is almost always significantly better when using LASSO-based methods than standard methods, providing even more evidence that our proposed methods can reduce the out-of-sample variance of the minimum-variance portfolio. Considering the Sharpe ratios, we can observe some increase for the LASSO and LASSO + TO models, partly due to the respective decrease in the standard deviation of these particular portfolios. Turnover for the standard model with POET is still slightly lower than that for the LASSO + TO model simply because of our non-flexible choice of k for all the models, as described above. This result can be easily adjusted by imposing a tighter



Table 2 Standard deviation p.a., Sharpe ratio p.a., average turnover per day, average assets as the mean of all non-zero weights, average short sales as the mean of all weights greater 0, the concentration, and diversification ratios for the different models applied to the 100 S&P dataset

	ML	Ledoit and Wolf (2003)	Ledoit and Wolf (2017)	Fan et al. (2013)
Standard deviation p.a.				
Standard	0.1170	0.1137	0.1131	0.1136
LASSO	0.1123 (0.0014)	0.1116 (0.0703)	0.1116 (0.1314)	0.1112 (0.0051)
LASSO + TO	0.1122 (0.0000)	0.1119 (0.0407)	0.1116 (0.0643)	0.1113 (0.0143)
Sharpe ratio p.a.				
Standard	0.7414	0.6896	0.6776	0.6463
LASSO	0.6351	0.6376	0.6353	0.6730
LASSO + TO	0.7371	0.7186	0.7197	0.6710
Turnover				
Standard	0.1348	0.1041	0.0922	0.0504
LASSO	0.2111	0.1829	0.1633	0.1247
LASSO + TO	0.0878	0.0858	0.0864	0.0730
Average assets				
Standard	100.00	100.00	100.00	100.00
LASSO	56.61	59.13	63.60	62.28
% of full	56.61	59.13	63.60	62.28
LASSO + TO	83.43	81.71	85.35	73.99
% of full	83.43	81.71	85.35	73.99
Average short sales				
Standard	44.49	45.77	41.11	41.27
LASSO	20.17	22.62	20.97	21.12
% of full	20.17	22.62	20.97	21.12
LASSO + TO	37.90	37.42	36.20	30.06
% of full	37.90	37.42	36.20	30.06
Concentration ratio				
Standard	0.6800	0.6480	0.4946	0.4655
LASSO	0.6344	0.6355	0.5060	0.5273
LASSO + TO	0.6172	0.6227	0.4989	0.5392
Diversification ratio				
Standard	1.8016	1.6340	1.9574	1.5065
LASSO	1.9018	1.7500	1.9975	1.6764
LASSO + TO	1.9266	1.7690	2.0150	1.7107

The column headers represent the used variance estimation technique and rows report the results of the GMV calculated with (1) (standard), (3) (LASSO), or (7) (LASSO+TO). The results in bold represent the best model of that column for a specific measure and the underlined number represents the best overall model for that measure. The numbers in brackets represent the p -values of the test $H_0 : \sigma_{LASSO} = \sigma_{Standard}$ and $H_0 : \sigma_{LASSO+TO} = \sigma_{Standard}$, respectively. All the results are calculated based on the whole out-of-sample period consisting of 4778 daily observations



turnover constraint parameter k . All the other major findings remain the same: the LASSO alone does not reduce turnover because of the need to estimate λ for every period, but the LASSO with the turnover constraint model does. Furthermore, the LASSO models reduce the overall number of assets as well as the short-sale budget of the portfolios, while keeping a low-risk profile and high diversification level.

Summary and conclusion

In this study, we investigate different types of global minimum-variance portfolios in terms of their standard deviation and practically relevant features such as the number of included assets, a short-sale reduction, and a turnover constraint. We use realistic datasets with up to 319 stocks in one portfolio and find that highly efficient estimation techniques for minimum-variance portfolios can be combined with practitioners' requirements for such portfolios. Our proposed estimation setup is constructed to be easily implemented, as it is solvable with standard software for quadratic programming.

Adding common constraints found in the literature, we construct sparse and stable portfolios. Our detailed empirical analysis, covering almost 19 years of daily out-of-sample observations, show the distinct and novel features of using portfolio construction with efficient estimation techniques. Specifically, we make the following discoveries:

- LASSO-type models can retain the low-variance profile of highly efficient variance estimators or even lower it
- A standard LASSO constraint increases turnover when λ is allowed to change over time
- The LASSO with the turnover constraint model can reduce turnover drastically, while maintaining sparsity, keeping variance low, and reducing the short-sale budget

We therefore conclude that it is beneficial, especially to practitioners, to use the LASSO with the turnover constraint approach and combine it with a modern technique to estimate large covariances.

However, our results depend on the procedure used to obtain the lowest variance in the out-of-sample study, namely, the cross-validation for the different λ tuning parameters. While this is a common practice in many fields unrelated to portfolio optimization, some researchers tend to use preset parameter values (i.e., they do not allow them to change over time). Future research should aim to provide more evidence on whether one or the other method yields better results in terms of relevant performance measures such as variance.

Further, data with a daily frequency are an a priori assumption that latently influences the outcome of our study. Even if the investor decides to rebalance daily, it is still questionable whether using only end-of-day daily data is sufficient to estimate the necessary moments of the multivariate return distribution. Advancements



in intraday data analysis might therefore also be included in further research on this topic.

Finally, all our models inherit the assumption of homoscedasticity. This, however, conflicts with the inclusion of a time-varying λ to some extent. Variance models that can capture the time dependency in return data, such as DCC-GARCH models, might be able to overcome the need to allow λ to fluctuate over time. Recent advancements in combining non-linear variance estimation with time-dependent techniques might therefore also provide new insights into this matter.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Best, M.J., and R.R. Grauer. 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means some analytical and computational results. *Review of Financial Studies* 4 (2): 315–342.
- Bouchaud, J.-P. and M. Potters. (2009) Financial applications of random matrix theory: A short review.
- Broadie, M. 1993. Computing efficient frontiers using estimated parameters. *Annals of Operations Research* 45 (1): 21–58.
- Brodie, J., I. Daubechies, C. de Mol, D. Giannone, and I. Loris. 2009. Sparse and stable markowitz portfolios. *PNAS* 106 (30): 12267–12272.
- Cai, T.T., and W. Liu. 2011. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association* 106 (494): 672–684.
- Chopra, V.K., and W.T. Ziemba. 1993. The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management* 19 (2): 6–11.
- Choueifat, Y., and Y. Coignard. 2008. Toward maximum diversification. *The Journal of Portfolio Management* 35 (1): 40–51.
- Dai, Z., and F. Wen. 2018. Some improved sparse and stable portfolio optimization problems. *Finance Research Letters* 27: 46–52.
- DeMiguel, V., L. Garlappi, F.J. Nogales, and R. Uppal. 2009. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science* 55 (5): 798–812.
- DeMiguel, V., L. Garlappi, and R. Uppal. 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies* 22 (5): 1915–1953.
- Engle, R.F., O. Ledoit, and M. Wolf. 2017. Large dynamic covariance matrices. *Journal of Business & Economic Statistics* 9: 1–13.



- Fan, J., Y. Liao, and M. Mincheva. 2013. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society. Series B* 75 (4): 603–680.
- Fastrich, B., S. Paterlini, and P. Winker. 2014. Cardinality versus q-norm constraints for index tracking. *Quantitative Finance* 14 (11): 2019–2032.
- Friendly, M. 2002. Corrgrams: Exploratory displays for correlation matrices. *The American Statistician* 56 (4): 316–324.
- Frost, P.A., and J.E. Savarino. 1986. An empirical bayes approach to efficient portfolio selection. *Journal of Financial and Quantitative Analysis* 21 (3): 293–305.
- Frost, P.A., and J.E. Savarino. 1988. For better performance: Constrain portfolio weights. *Journal of Portfolio Management* 15 (1): 29–34.
- Jagannathan, R., and T. Ma. 2003. Risk reduction in large portfolios: why imposing the wrong constraints helps. *Journal of Finance* 58 (4): 1651–1683.
- Jobson, J.D., and R.M. Korkie. 1980. Estimation for markowitz efficient portfolios. *Journal of the American Statistical Association* 75 (371): 544–554.
- Jobson, J.D., and R.M. Korkie. 1981. Performance hypothesis testing with the sharpe and treynor measures. *Journal of Finance* 36 (4): 889–908.
- Kan, R., and G. Zhou. 2007. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42 (3): 621–656.
- Konno, H., and A. Wijayanayake. 2002. Portfolio optimization under dc transaction costs and minimal transaction unit constraints. *Journal of Global Optimization* 22 (1–4): 137–154.
- Ledoit, O., and M. Wolf. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10 (5): 603–621.
- Ledoit, O., and M. Wolf. 2004. Honey, i shrunk the sample covariance matrix. *Journal of Portfolio Management* 30 (4): 110–119.
- Ledoit, O., and M. Wolf. 2008. Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance* 15 (5): 850–859.
- Ledoit, O., and M. Wolf. 2012. Nonlinear shrinkage estimation of large-dimensional covariance matrices. *Annals of Statistics* 40 (2): 1024–1060.
- Ledoit, O., and M. Wolf. 2015. Spectrum estimation: a unified framework for covariance matrix estimation and pca in large dimensions. *Journal of Multivariate Analysis* 139: 360–384.
- Ledoit, O., and M. Wolf. 2017. Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks. *Review of Financial Studies* 30 (12): 4349–4388.
- Litterman, R.B., ed. 2003. *Modern investment management: An equilibrium approach*. Hoboken: Wiley.
- Lobo, M.S., M. Fazel, and S. Boyd. 2007. Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research* 152 (1): 341–365.
- Markowitz, H.M. 1952. Portfolio selection. *Journal of Finance* 7 (1): 77–91.
- Michaud, R.O. 1989. The markowitz optimization enigma: Is ‘optimized’ optimal? *Financial Analysts Journal* 45 (1): 31–42.
- Takeda, A., M. Niranjana, J.-Y. Gotoh, and Y. Kawahara. 2013. Simultaneous pursuit of out-of-sample performance and sparsity in index tracking portfolios. *Computational Management Science* 10 (1): 21–49.
- Tu, J., and G. Zhou. 2011. Markowitz meets talmud: A combination of sophisticated and naive diversification strategies. *Journal of Financial Economics* 99 (1): 204–215.
- Wei, T., V. Simko, M. Levy, Y. Xie, Y. Jin, and J. Zemla. 2017. Package ‘corrplot’. *Statistician* 56: 316–324.
- Zhao, Z., O. Ledoit, and H. Jiang. 2019. Risk reduction and efficiency increase in large portfolios: leverage and shrinkage. University of Zurich, Department of Economics, Working Paper, (328).

