A UNIFIED FRAMEWORK FOR FAST LARGE-SCALE PORTFOLIO OPTIMIZATION

Weichuan Deng[1], Ronakdilip Shah[1], Paweł Polak[1], Abolfazl Safikhani[2]

[1]Stony Brook University, [2]George Mason University.

Stony Brook University

Abstract

We develop a unified framework for fast large-scale portfolio optimization with shrinkage and regularization for different objectives such as minimum variance, mean-variance, and maximum Sharpe ratio with various constraints on the portfolio weights. For all of the optimization problems, we derive the corresponding quadratic programming problems and implement them in an open-source Python library. We use the proposed framework to evaluate the out-of-sample portfolio performance of popular covariance matrix estimators such as sample covariance matrix, linear and nonlinear shrinkage estimators, and the covariance matrix from the instrumented principal component analysis (IPCA). We use 65 years of monthly returns from (on average) 585 largest companies in the US market, and 94 monthly firm-specific characteristics for the IPCA model. We show that the regularization of the portfolio norms greatly benefits the performance of the IPCA model in portfolio optimization, resulting in outperformance linear and nonlinear shrinkage estimators even under realistic constraints on the portfolio weights. The corresponding code with the implementation in Python is available online at `https://github.com/PawPol/PyPortOpt`.

Portfolio Optimization Framework

The general theory of portfolio optimization, as introduced in a seminal work by [8], summarizes the trade-off between risk and investment return using the portfolio's mean and variance. The objective function is

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathcal{W}} \frac{1}{2}\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w},$$

where $\mathcal{W} := \left\{\mathbf{w}\in\mathbb{R}^N : \mathbf{w}'\boldsymbol{\mu} \geq \alpha_0 \text{ and } \mathbf{w}'\mathbf{1}_N = 1\right\}$, is a set of constraints on the portfolio weights which correspond to a fully invested portfolio with the expected return above the α_0 threshold. Get the close-form in-sample frontier,

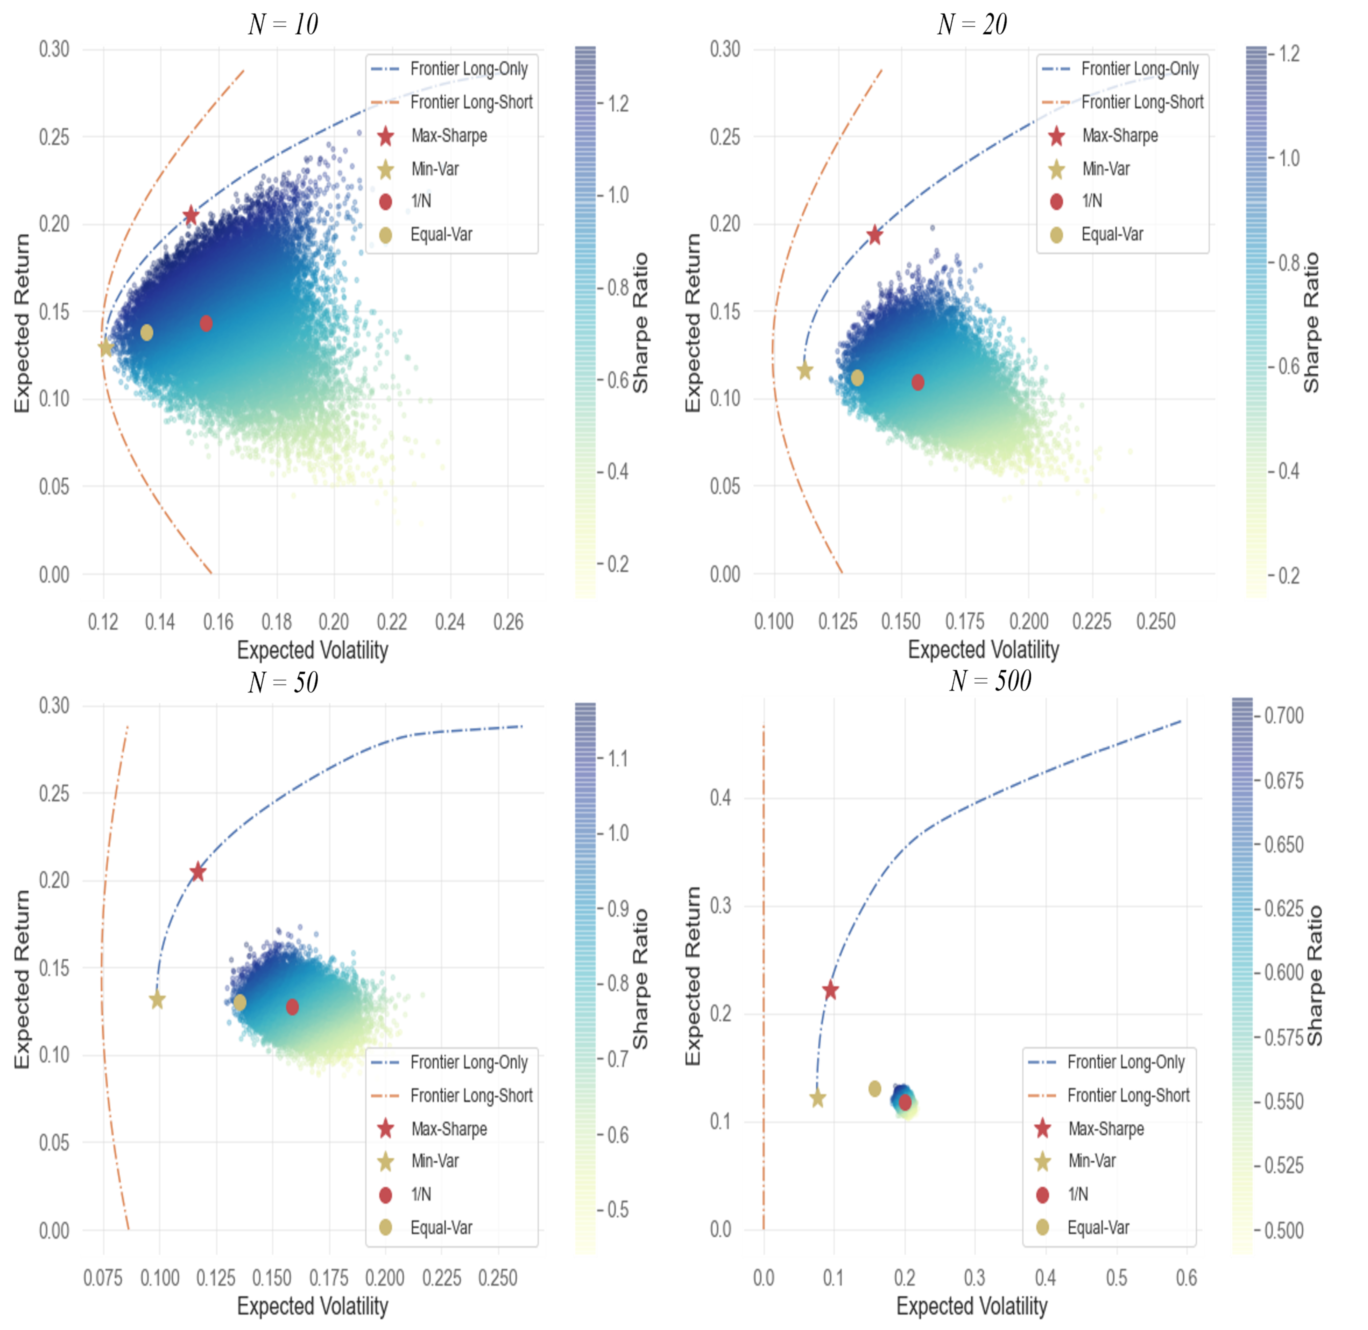

The above graphical illustration represents the in-sample efficient frontier. It underscores that in a high-dimensional setup, without appropriate portfolio settings, achieving an optimal risk-reward profile is far-fetched. Even the often-touted $1/N$ portfolio, celebrated for its naive diversification and performance, falls short in this context.

Portfolio Objective Function

Depending on the specific optimization goal, we formulate various objective functions.

- **Mean-Variance Optimization with Risk-Free Asset**

$$\mathbf{w}^* = \arg\min_{\mathbf{w}'\in\mathcal{W}} \frac{1}{2}\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w},$$

where $\mathcal{W} = \left\{\mathbf{w}\in\mathbb{R}^N : \mathbf{w}'\boldsymbol{\mu} + (1-\mathbf{w}'\mathbf{1}_N)r_f = \alpha_0\right\}$.

- **Maximum Sharpe Ratio Portfolio**

$$\arg\min_{\mathbf{w}\in\mathcal{W}} \frac{\sqrt{\gamma\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}\gamma}}{\gamma\mathbf{w}'(\boldsymbol{\mu}-r_f\mathbf{1}_N)} \iff \arg\min_{[\tilde{\mathbf{w}},\gamma]'\in\tilde{\mathcal{W}}} \tilde{\mathbf{w}}'\boldsymbol{\Sigma}\tilde{\mathbf{w}},$$

where $\tilde{\mathcal{W}} = \left\{[\tilde{\mathbf{w}},\gamma]'\in\mathbb{R}^{N+1} : 1 = \tilde{\mathbf{w}}'(\boldsymbol{\mu}-r_f\mathbf{1}_N), \mathbf{1}_N'\tilde{\mathbf{w}} = \gamma, \tilde{\mathbf{w}}\geq\mathbf{0}\right\}$.

In our portfolio optimization framework, we translate portfolio problems into standard quadratic programming (QP) problems. Subsequently, we employ the **OSQP solver** to tackle them.

Portfolio Constraints

Introducing specific constraints for each portfolio objective function, leading to various formulations:

- ℓ_2^2 **Penalized Portfolio Norms**: $\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathcal{W}} \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} + \lambda\|\mathbf{w}\|_2^2$
- ℓ_1 **Penalized Portfolio Norms**: $\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathcal{W}} \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} + \lambda\|\mathbf{w}\|_1$
- ℓ_1+ℓ_2^2 **Penalized Portfolio Norms**: $\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathcal{W}} \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} + \lambda_1\|\mathbf{w}\|_1 + \lambda_2\|\mathbf{w}\|_2^2$.

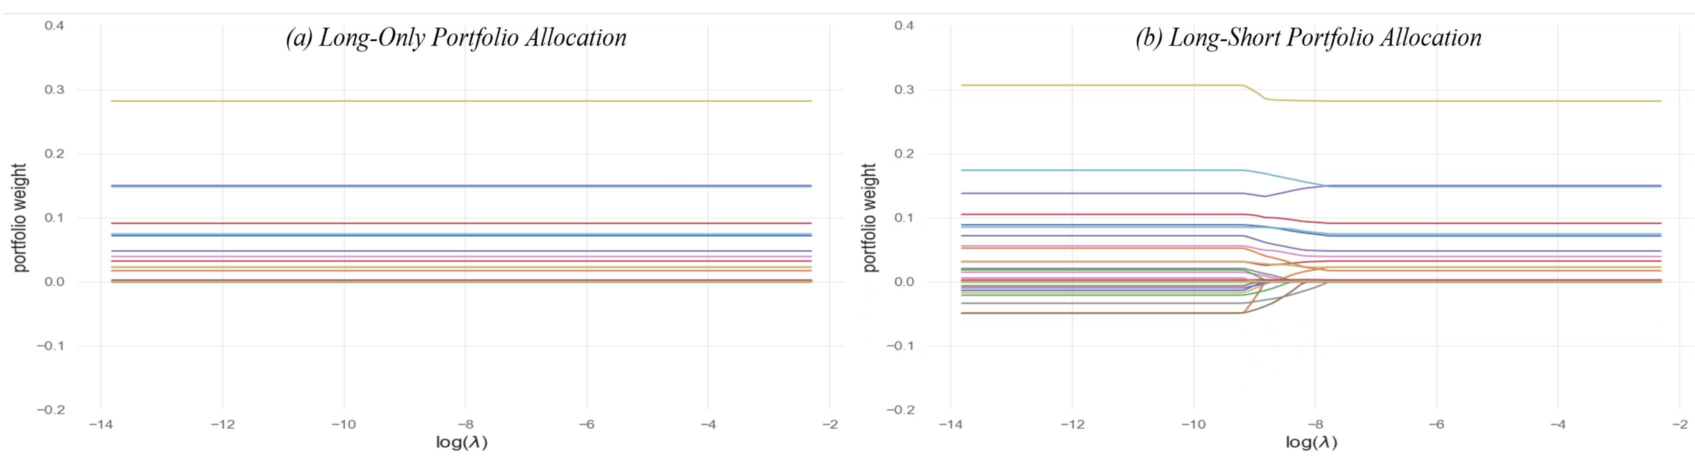

The feasible set of portfolio weights, $\mathcal{W}$, typically encompasses additional constraints. For example:

- **Long only**: $\mathcal{W} := \left\{\mathbf{w}\in\mathbb{R}^N : \mathbf{w}'\mathbf{1}_N = 1 \text{ and } w_i \geq 0, \forall i\right\}$.
- **Asset specific holding constraints**: $\mathcal{W} := \left\{\mathbf{w}\in\mathbb{R}^N : \mathbf{w}'\mathbf{1}_N = 1 \text{ and } L_i \leq w_i \leq U_i, \forall i\right\}$,

Additionally, our model incorporates various constraints for robust portfolio management, including benchmark exposure constraints, tracking error constraints, risk factor constraints, as well as linear and quadratic constraints.

Covariance Matrix Estimation

In Markowitz's portfolio theory, the mean vector ($\boldsymbol{\mu}$) and covariance matrix ($\boldsymbol{\Sigma}$) are presumed known. Practically, they require estimation from data, typically via historical sample mean and covariance under an iid assumption. Our portfolio optimization deploys three covariance matrix estimators:

- **Classical linear shrinkage covariance matrix estimator** [6] defined as $\hat{\boldsymbol{\Sigma}} = \hat{\delta}\hat{\mathbf{F}} + (1-\hat{\delta})\mathbf{S}$.
- **Nonlinear shrinkage covariance matrix estimator** from [5]. $\hat{\boldsymbol{\Sigma}}_t := \mathbf{U}_t\hat{\boldsymbol{\Delta}}_t\mathbf{U}_t'$.
- **IPCA shrinkage** $\mathbf{r}_{t+1} = \mathbf{Z}_t'\boldsymbol{\Gamma}_\beta\mathbf{f}_{t+1} + \boldsymbol{\epsilon}^*_{t+1}$, which leads to the shrinkage covariance matrix $\hat{\boldsymbol{\Sigma}}_{r_{t+1}} = \mathbf{Z}_t\hat{\boldsymbol{\Gamma}}_\beta\text{cov}(\hat{\mathbf{F}})\hat{\boldsymbol{\Gamma}}_\beta'\mathbf{Z}_t' + \hat{\mathbf{D}}$, where $\hat{\mathbf{D}} = \text{diag}(\text{cov}(\boldsymbol{r}_{t+1} - \mathbf{Z}_t'\hat{\boldsymbol{\Gamma}}_\beta\hat{\mathbf{f}}_{t+1}))$ is the diagonal matrix for the covariance of the residuals. We will compare $\hat{\boldsymbol{\Sigma}}_{r_{t+1}}$ with the two aforementioned shrinkage covariance matrix estimators and the sample covariance matrix estimator in our portfolio analysis below using our QP framework.

Empirical Analysis

- **Model Implementation**: Rolling window on maximum Sharpe ratio portfolio with the sample mean estimator and the IPCA covariance matrix estimator.

$$\arg\min_{\mathbf{w}\in\mathcal{W}} -\frac{\mathbf{w}'\boldsymbol{\mu}-r_f}{\sqrt{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}}} + \delta_1\frac{1}{\mathbf{w}'\boldsymbol{\mu}-r_f}\|\mathbf{w}\|_1 + \delta_2\frac{1}{(\mathbf{w}'\boldsymbol{\mu}-r_f)^2}\|\mathbf{w}\|_2^2.$$

- **CRSP Data**: From 1957 to 2021 monthly data. The number of assets ranges from 329 to 931 symbols every month, with an average number of 585 stocks.

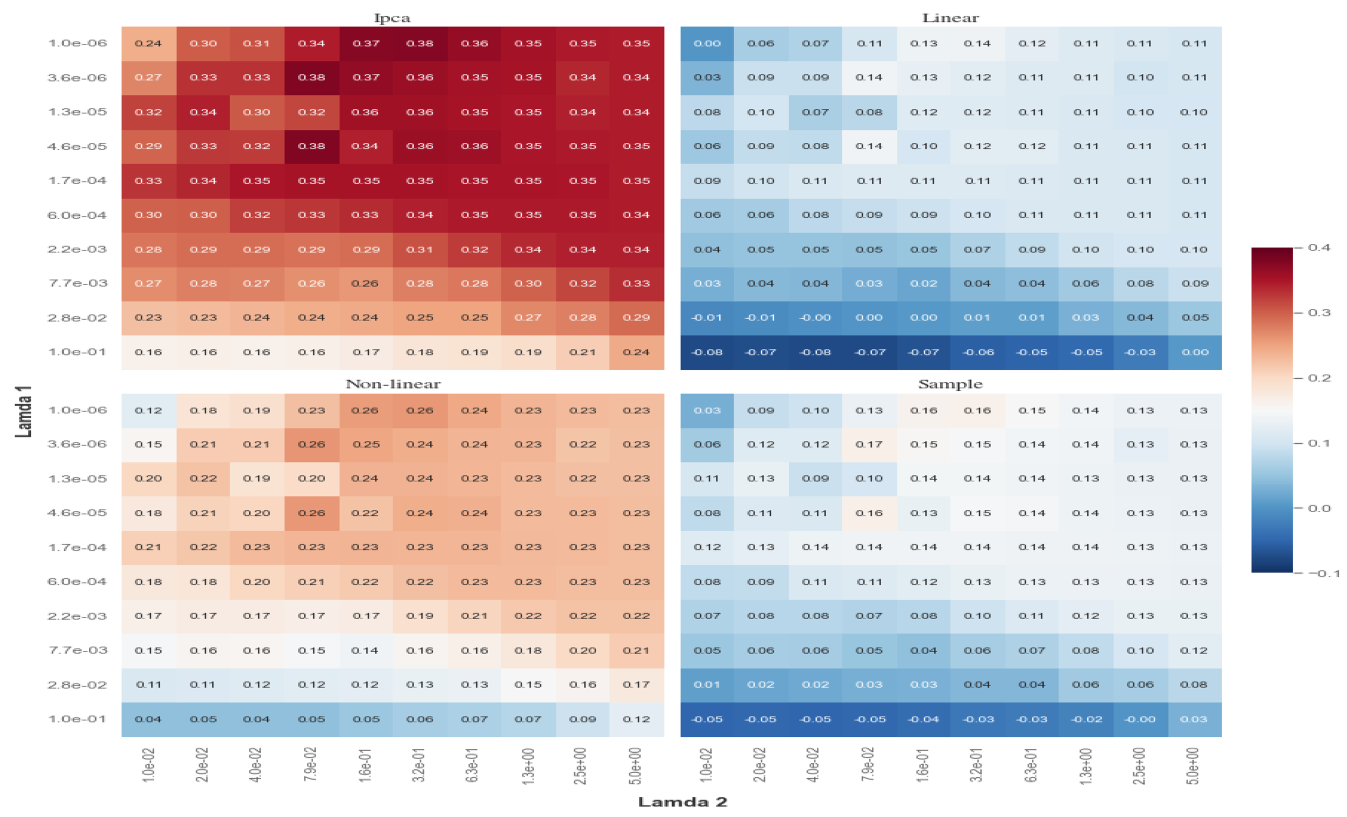

Sharpe ratio portfolio with IPCA estimates of the covariance matrix performs much better than all the benchmark models.

- **Result**: Apply 30 years rolling window with one month ahead for out-of-sample result.

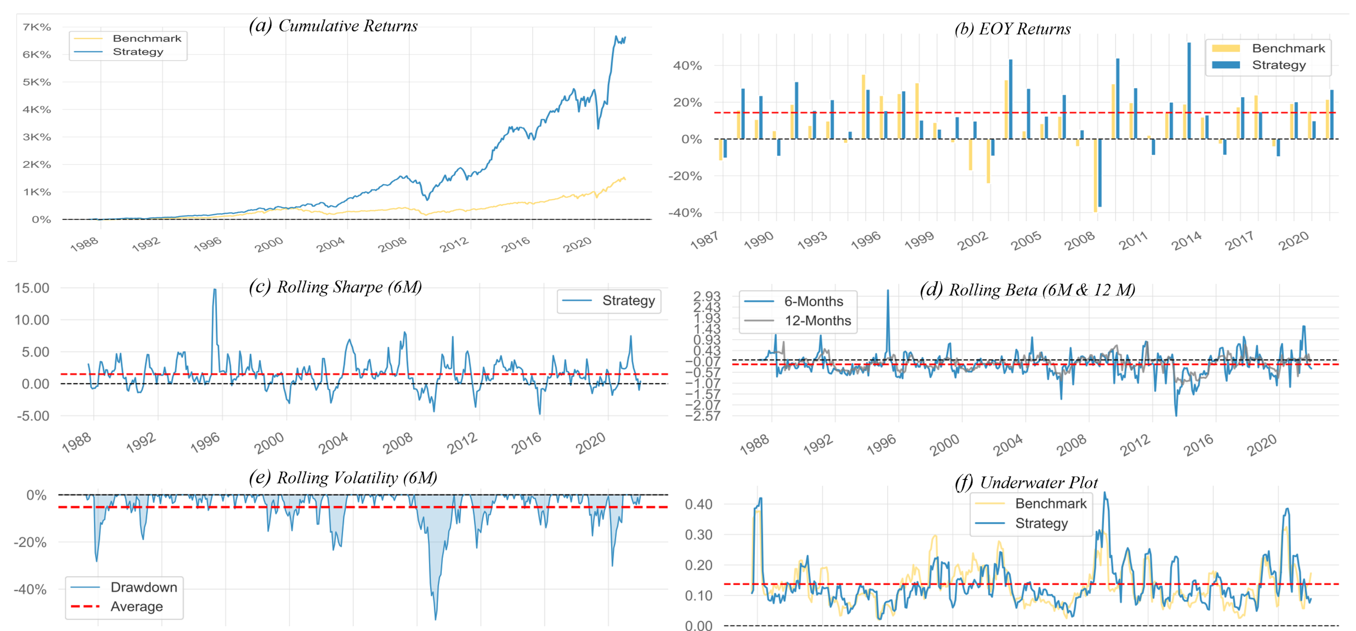

Conclusions and Extensions

- **Unified Framework Introduction**: This paper presents a unified framework for portfolio optimization, utilizing quadratic programming. This framework greatly enhances computational speed and accuracy, proving especially advantageous in large-scale portfolio problems.
- **IPCA Factor Model Performance**: When incorporated into a regularized portfolio optimization problem, the IPCA factor model demonstrates substantial improvements in portfolio performance, especially under long-short constraints.
- **Future Research Directions**: Future studies will apply our model to managed portfolios rather than individual stocks, with the aim of achieving a more stable mean estimator [7, 1, 3, 4, 2].

References:

[1] Svetlana Bryzgalova, Markus Pelger, and Jason Zhu. "Forest through the trees: Building cross-sections of stock returns". In: *Available at SSRN 3493458* (2020).

[2] Jianqing Fan, Yuan Liao, and Martina Mincheva. "Large covariance estimation by thresholding principal orthogonal complements". In: *Journal of the Royal Statistical Society. Series B, Statistical methodology* 75.4 (2013).

[3] Bryan T Kelly, Seth Pruitt, and Yinan Su. "Characteristics are covariances: A unified model of risk and return". In: *Journal of Financial Economics* 134.3 (2019), pp. 501–524.

[4] Serhiy Kozak, Stefan Nagel, and Shrihari Santosh. "Shrinking the cross-section". In: *Journal of Financial Economics* 135.2 (2020), pp. 271–292.

[5] Olivier Ledoit and Michael Wolf. "Analytical nonlinear shrinkage of large-dimensional covariance matrices". In: *The Annals of Statistics* 48.5 (2020), pp. 3043–3065.

[6] Olivier Ledoit and Michael Wolf. "Honey, I Shrunk the Sample Covariance Matrix". In: *The Journal of Portfolio Management* 30.4 (2004), pp. 110–119.

[7] Martin Lettau and Markus Pelger. "Factors that fit the time series and cross-section of stock returns". In: *The Review of Financial Studies* 33.5 (2020), pp. 2274–2325.

[8] H Markowitz. "Modern portfolio theory". In: *Journal of Finance* 7.11 (1952), pp. 77–91.