

Portfolio optimization

in the presence of estimation errors

on the expected asset returns

Gérard Cornuéjols¹, Özgün Elçi¹, and Matthias Köppe²

¹Carnegie Mellon University, Tepper School of Business, 5000 Forbes Ave, Pittsburgh, PA 15213, USA

²University of California, Davis, Department of Mathematics, One Shields Avenue, Davis, CA 95616, USA
gc0v@andrew.cmu.edu, elci@cmu.edu, mkoeppe@math.ucdavis.edu

February 10, 2022

ABSTRACT: It is well known that the classical Markowitz model for portfolio optimization is extremely sensitive to estimation errors on the expected asset returns. Robust optimization mitigates this issue. We focus on ellipsoidal uncertainty sets around the point estimates of the expected asset returns. We investigate the performance of diagonal estimation-error matrices in the description of the uncertainty set. We show that the class of diagonal estimation-error matrices can achieve an arbitrarily small loss in the expected portfolio return as compared to the optimum. We then formulate the problem of finding the best estimation error matrix as a bilevel program. The bilevel model allows us to numerically analyze the error when there are multiple estimates for the expected return and/or when there are additional restrictions on the structure of the estimation-error matrix. We extend our analysis to show that the diagonal estimation-error matrices can achieve an arbitrarily small loss even when there are multiple estimates for the expected return. We finally focus on analyzing the use of identity matrix as an estimation-error matrix. The results of our simulation show that robust portfolio models featuring an identity matrix as an estimation-error matrix outperform the classical Markowitz model when the size of the uncertainty set is chosen properly.

Keywords: portfolio optimization; robust optimization; estimation error; bilevel programming

1. Introduction Consider a portfolio optimization problem where we want to invest in n assets. If the return vector $\mathbf{r} \in \mathbb{R}^n$ is given, we formulate the problem as $\text{maximize}\{\mathbf{r}^\top \mathbf{x} : \mathbf{x} \in \Delta\}$ where $\mathbf{x} \in \mathbb{R}^n$ denotes the fraction of investment in each asset and Δ denotes the feasible region. In this paper we consider $\Delta := \{\mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1, \mathbf{x} \geq 0\}$. The constraints $\mathbf{x} \geq 0$ restrict the model to portfolios with long positions only. This problem has a trivial optimal solution: Only invest in the asset that has the highest return.

In practice, however, the assets are risky and the return vector \mathbf{r} is random. The classical mean-variance portfolio optimization problem introduced by [Markowitz \(1952\)](#) addresses this issue by maximizing the expected return of the portfolio subject to a constraint on the risk modeled as the variance of the portfolio return.

$$\underset{\mathbf{x}}{\text{maximize}} \quad \boldsymbol{\mu}^\top \mathbf{x} \quad (1a)$$

$$\text{subject to} \quad \mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} \leq v \quad (1b)$$

$$\mathbf{x} \in \Delta. \quad (1c)$$

Here, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ denote the expectation vector and covariance matrix of the asset returns, respectively. In practice, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are estimated, and it has been observed that the Markowitz model tends to amplify estimation errors. In particular, small errors in $\boldsymbol{\mu}$ may produce large changes in portfolio holdings (see, e.g., [Best and Grauer, 1991](#); [Chopra and Ziemba, 1993](#); [Michaud, 2008](#)). Approaches to mitigate the effect of estimation errors in $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ on portfolio construction have lead to a vast literature. Among these we name [Horst et al. \(2001\)](#), [Chopra \(1993\)](#), [Jagannathan and Ma \(2003\)](#), [Goldfarb and Iyengar \(2003\)](#), [Tutuncu and Koenig \(2004\)](#), [Ceria and Stubbs \(2006\)](#), [Scherer \(2007\)](#), [DeMiguel et al. \(2009\)](#), [Lim et al. \(2012\)](#), and [Ban et al. \(2018\)](#). In this paper, we pursue this line of work, focusing on the uncertainty in the $\boldsymbol{\mu}$ estimates.

We assume that the expected return vector $\boldsymbol{\mu}$ is unknown and belongs to an ellipsoidal uncertainty set given by

$$\mathcal{U} := \left\{ \boldsymbol{\mu} \in \mathbb{R}^n : (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})^\top \boldsymbol{\Xi}^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) \leq \kappa^2 \right\} \quad (2)$$

where $\hat{\boldsymbol{\mu}}$ is the estimated expected return and the positive definite matrix $\boldsymbol{\Xi}$ is referred to as the estimation-error matrix. Using the uncertainty set \mathcal{U} , we formulate the robust portfolio optimization problem as follows

$$\max_{\mathbf{x} \in \mathcal{X}} \min_{\boldsymbol{\mu} \in \mathcal{U}} \left\{ \boldsymbol{\mu}^\top \mathbf{x} \right\} \quad (3)$$

where \mathcal{X} denotes the feasible set (1b) – (1c). We assume that the covariance matrix $\boldsymbol{\Sigma}$ is known. Consequently, \mathcal{X} is the intersection of an ellipsoid with a simplex.

This problem can be reformulated following [Ben-Tal and Nemirovski \(1999\)](#) as

$$\max_{\mathbf{x} \in \mathcal{X}} \left\{ \hat{\boldsymbol{\mu}}^\top \mathbf{x} - \kappa \sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}} \right\}. \quad (4)$$

We refer the reader to [Goldfarb and Iyengar \(2003\)](#), [Tutuncu and Koenig \(2004\)](#) and [Fabozzi et al. \(2007\)](#) for more detailed discussions on robust portfolio optimization. The *risk-like* term $\kappa \sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}}$ in formulation (4) can be interpreted as an estimation risk that must be considered by risk-averse investors on top of the risk caused by the variance $\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}$ of the portfolio return ([Fabozzi et al., 2007](#), p 371).

Our contribution. In this paper, we provide a theoretical analysis on the choice for the estimation-error matrix $\boldsymbol{\Xi}$ in robust portfolio optimization. The literature on selecting/constructing an estimation-error matrix is scarce ([Gotoh and Takeda, 2011](#)). [Stubbs and Vance \(2005\)](#) provide a comprehensive overview on practical approaches for computing estimation-error matrices. They also recommend the use of diagonal estimation-error matrices to practitioners. On the other hand, there are several studies in which a scalar multiple of $\boldsymbol{\Sigma}$ is used as the estimation-error matrix (see, e.g., [Scherer, 2007](#); [Olivares-Nadal and DeMiguel, 2018](#)). Among these, [Scherer \(2007\)](#) has a skeptical take on robust optimization and shows that such a choice for $\boldsymbol{\Xi}$ is equivalent to some other well known shrinkage approaches.

In our work, we begin by discussing the difference between true and actual frontier - a well known approach that is used to quantify the impact of the estimation errors. We show that when $\boldsymbol{\Xi}$ is a multiple of $\boldsymbol{\Sigma}$, robust optimization cannot improve the actual performance of the Markowitz model. This analysis supports the recommendation of [Stubbs and Vance \(2005\)](#) where they claim that such a choice is not used in practice and is not recommended. We then focus on the use of diagonal estimation-error matrices and show that the class of diagonal estimation-error matrices can achieve an arbitrarily small loss in the expected portfolio return as compared to the optimum. We accomplish this by reformulating the optimality conditions of the robust portfolio problem when we have a single expected return vector.

We then investigate diagonal estimation-error matrices in the presence of multiple expected returns. To this end, we propose a bilevel model that computes the loss and we show that the diagonal

estimation-error matrices can achieve an arbitrarily small loss even when there are multiple estimates for the expected return. The proposed bilevel model also provides a principled way to construct estimation-error matrices using data, though it is a non-convex optimization problem and it is difficult to solve in practice.

We finally focus on the use of an identity matrix for Ξ . This choice requires calibrating a single parameter (κ) for the robust problem. We perform computational experiments to test whether robust optimization can perform better than the Markowitz model. Our results demonstrate that a good choice for κ can significantly improve the performance of the Markowitz model, especially when the expected return estimates are not reliable.

The rest of the paper is organized as follows. We introduce true, estimated and actual frontiers in Section 2. In Section 3, we show that robust optimization can improve on the actual performance. We present the analysis of the robust optimization problem in Section 4. Section 5 generalizes the analysis of the robust model to incorporate multiple expected returns. We analyze the use of an identity matrix as estimation-error matrix in Section 6. Section 7 concludes the paper.

2. True, estimated, and actual frontiers The sensitivity of mean-variance portfolio optimization models to estimation errors on the expected asset returns is well documented in the literature (see, e.g., Best and Grauer, 1991; Chopra, 1993; Michaud, 2008). It is sometimes referred to as the error maximization tendency of mean-variance optimization. In order to quantify the effect of estimation errors, Broadie (1993) made a distinction between true, estimated, and actual frontiers. A *frontier* plots the maximum expected return as a function of the risk threshold (Markowitz, 1952). The *true frontier* is computed by using the true expected returns of the assets, a quantity in fact unknown to the decision maker. The *estimated frontier* is computed by using the estimated expected returns. It describes what appears to be the maximum return based on the estimated parameters. The *actual frontier* shows what really occurs when we invest in the portfolio optimized using the estimated parameters. We next describe how we compute these three frontiers.

We solve the problem $\text{maximize}\{\boldsymbol{\mu}^\top \mathbf{x} : (1b) - (1c)\}$ to obtain the optimal solution \mathbf{x}^* where $\boldsymbol{\mu}$ is the true (but unknown to the investor) vector of expected asset returns. Using \mathbf{x}^* , we construct the true frontier $\boldsymbol{\mu}^\top \mathbf{x}^*$. Let $\hat{\boldsymbol{\mu}}$ be the vector of estimated expected asset returns. We solve $\text{maximize}\{\hat{\boldsymbol{\mu}}^\top \mathbf{x} :$

(1b) – (1c)} to obtain the *Markowitz solution estimate* $\hat{\mathbf{x}}^M$. Using $\hat{\mathbf{x}}^M$, we construct the *estimated Markowitz frontier* $\hat{\boldsymbol{\mu}}^\top \hat{\mathbf{x}}^M$ and the *actual Markowitz frontier* $\boldsymbol{\mu}^\top \hat{\mathbf{x}}^M$.

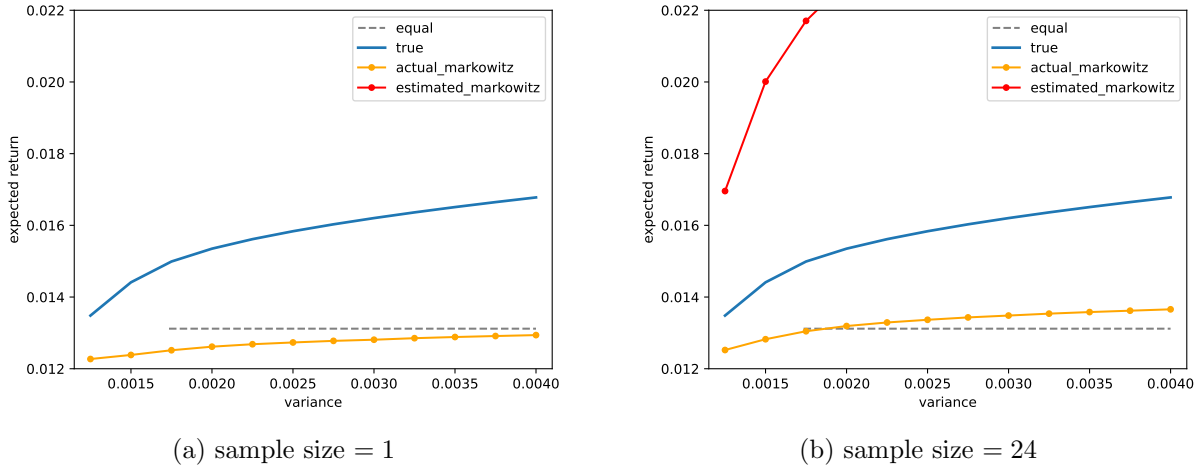


Figure 1: True, actual Markowitz, and estimated Markowitz frontiers.

In Figure 1, we illustrate *the gap* between the true and the actual frontiers on real-world data. Details on the data set can be found in Appendix A. We assume that the asset returns are normally distributed with distribution $\text{Normal}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The vector $\hat{\boldsymbol{\mu}}$ is the sample average of n random samples. Equivalently, we generate $\hat{\boldsymbol{\mu}}$ from $\text{Normal}(\boldsymbol{\mu}, \boldsymbol{\Sigma}/n)$. We illustrate the cases $n = 1$ and $n = 24$. Each point in the figures represents the average value over 10000 trials. The dashed-line represents the performance of the equal portfolio where we invest equally in all the assets.

We see in Figure 1a that the actual performance of the Markowitz solution estimate is poor when the sample size is equal to 1. In fact, the average expected return values are worse than the equal-weight portfolio. On the other hand, we see that the performance of the Markowitz solution estimate improves significantly in Figure 1b with more accurate return estimates with a sample size of 24. As n goes to $+\infty$, the actual Markowitz frontier converges to the true frontier. Furthermore, we plot the estimated frontier in Figures 1a and 1b. We see that the estimated frontier is far away from the true frontier so much that it is off the chart in Figure 1a.

Similarly, solving the robust optimization problem (4), we obtain the *robust solution estimate* $\hat{\mathbf{x}}^R$. Using $\hat{\mathbf{x}}^R$, we construct the *estimated robust frontier* $\hat{\boldsymbol{\mu}}^\top \hat{\mathbf{x}}^R$ and the *actual robust frontier* $\boldsymbol{\mu}^\top \hat{\mathbf{x}}^R$. A key aspect here is the modeler’s choice of the estimation-error matrix Ξ .

3. Robust optimization can improve the actual performance A possible choice for Ξ is to make this matrix proportional to Σ , namely $\Xi = \rho\Sigma$. Such an idea has been proposed by Horst et al. (2001). Notice however that, for solutions \mathbf{x} that satisfy the risk constraint $\mathbf{x}^\top \Sigma \mathbf{x} \leq v$ at equality (which is the most interesting case), the objective $\hat{\boldsymbol{\mu}}^\top \mathbf{x} - \kappa \sqrt{\mathbf{x}^\top \Xi \mathbf{x}}$ of the robust problem (4) becomes

$$\hat{\boldsymbol{\mu}}^\top \mathbf{x} - \kappa \sqrt{\mathbf{x}^\top \rho \Sigma \mathbf{x}} = \hat{\boldsymbol{\mu}}^\top \mathbf{x} - \kappa \sqrt{\rho v}.$$

The last term is just a constant, and therefore the robust problem reduces to the Markowitz model (1). The estimation errors do not affect the optimal portfolio! In the remainder of this paper, we consider more promising choices for Ξ .

We next show that portfolios obtained through robust optimization can achieve a better actual performance when the matrix Ξ is chosen appropriately. To illustrate the potential of robust optimization,

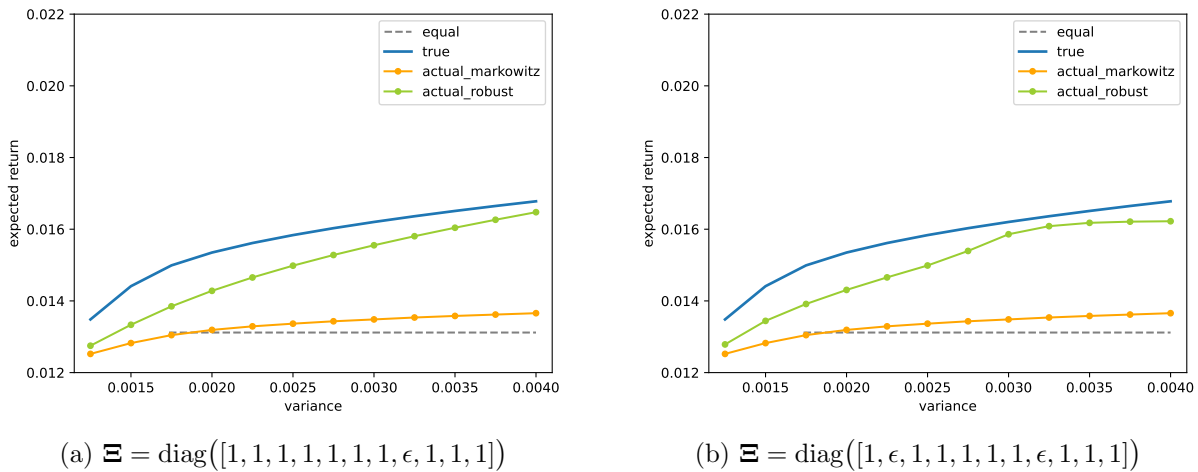


Figure 2: Robust frontiers of several Ξ matrices.

we design an experiment showing that good matrices Ξ exist, even among diagonal matrices with only two different diagonal entries. We construct Ξ based on partial information on the true expected return $\boldsymbol{\mu}$. In particular, in Figures 2a – 2b, we let some entries of the estimation-error matrix equal to $\epsilon = 0.001$ ($\kappa = 1$). In Figure 2a, the entry corresponding to the asset with the highest true return is set to ϵ , and in Figure 2b, the entries corresponding to the four assets with highest true return are set to ϵ .

The overall picture that we see in Figure 2 is that there is room for improvement if one is able to choose a good matrix Ξ . It seems promising that the robust optimization can close the gap between

the actual Markowitz frontier and the true frontier.

4. The class of diagonal estimation-error matrices In this section, we focus on the class of diagonal estimation-error matrices. This choice is partly motivated by the study of [Stubbs and Vance \(2005\)](#). The authors argue that it is difficult to generate estimation-error matrices accurately, and they suggest that the use of simple diagonal matrices is beneficial. A natural question to ask is: Do we loose in performance by restricting our attention to diagonal matrices Ξ ?

The issue boils down to the following: Given an estimate $\hat{\boldsymbol{\mu}}$ of $\boldsymbol{\mu}$, can we always find a diagonal matrix Ξ such that the resulting robust portfolio $\hat{\mathbf{x}}^R$ has an actual expected return very close to the true expected return? To this end, we start examining the optimality conditions of (4).

4.1 Analysis of the robust portfolio optimization problem We first write (4) as a convex optimization problem of the following form.

$$\underset{\mathbf{x}}{\text{minimize}} \quad -\hat{\boldsymbol{\mu}}^\top \mathbf{x} + \sqrt{\mathbf{x}^\top \Xi \mathbf{x}} \quad (5a)$$

$$\text{subject to} \quad \mathbf{x}^\top \Sigma \mathbf{x} \leq v, \quad (5b)$$

$$\mathbf{1}^\top \mathbf{x} = 1, \quad (5c)$$

$$-\mathbf{x} \leq \mathbf{0}. \quad (5d)$$

We next make some observations about (5). The proof of these observations are deferred to the Appendix.

Proposition 1. *Suppose that Ξ and Σ are positive definite. Then, the following statements are true about (5).*

1. *It is a convex optimization problem with a strictly convex objective function. Therefore, the optimal solution $\hat{\mathbf{x}}^R$ is unique.*
2. *It satisfies Slater's condition.*
3. *It is sufficient to consider the optimality conditions for differentiable functions.*

We next derive the optimality conditions of (5). Let $\lambda^1 \in \mathbb{R}$, $\lambda^2 \in \mathbb{R}$, and $\boldsymbol{\lambda}^3 \in \mathbb{R}^n$ respectively denote the Lagrangian multipliers of the constraints of (5). The Lagrangian function associated with this problem is given below.

$$L(\mathbf{x}, \lambda^1, \lambda^2, \boldsymbol{\lambda}^3) = -\hat{\boldsymbol{\mu}}^\top \mathbf{x} + \sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}} + \lambda^1 (\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} - v) + \lambda^2 (\mathbf{1}^\top \mathbf{x} - 1) - (\boldsymbol{\lambda}^3)^\top \mathbf{x}$$

Using the Lagrangian function, we can write the optimality conditions of (5).

$$-\hat{\boldsymbol{\mu}} + \frac{\boldsymbol{\Xi} \mathbf{x}}{\sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}}} + 2\lambda^1 \boldsymbol{\Sigma} \mathbf{x} + \lambda^2 \mathbf{1} - \boldsymbol{\lambda}^3 = \mathbf{0} \quad (6a)$$

$$\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} \leq v \quad (6b)$$

$$\mathbf{1}^\top \mathbf{x} = 1 \quad (6c)$$

$$-\mathbf{x} \leq \mathbf{0} \quad (6d)$$

$$\lambda^1 \geq 0 \quad (6e)$$

$$\boldsymbol{\lambda}^3 \geq \mathbf{0} \quad (6f)$$

$$\lambda^1 (\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} - v) + (\boldsymbol{\lambda}^3)^\top \mathbf{x} = 0 \quad (6g)$$

When $\boldsymbol{\Xi}$ is a diagonal matrix, (6a) can be written as

$$-\hat{\mu}_i + \frac{\xi_i x_i}{\sqrt{\sum_{j=1}^n \xi_j x_j^2}} + 2\lambda^1 \sum_{j=1}^n \sigma_{ij} x_j + \lambda^2 - \lambda_i^3 = 0, \quad i \in [1, n] \quad (7)$$

Here, $\boldsymbol{\xi}$ denotes the vector of diagonal entries of $\boldsymbol{\Xi}$ and $[1, n] := \{1, \dots, n\}$. Now, we introduce additional variables and constraints to rewrite the square-root term in (7). Let

$$z_i = \frac{\xi_i x_i}{\sqrt{\sum_{j=1}^n \xi_j x_j^2}}, \quad i \in [1, n]. \quad (8)$$

Using this substitution, we rewrite (7) as

$$-\hat{\mu}_i + z_i + 2\lambda^1 \sum_{j=1}^n \sigma_{ij} x_j + \lambda^2 - \lambda_i^3 = 0, \quad i \in [1, n], \quad (9a)$$

$$\sum_{i=1}^n x_i z_i = \alpha, \quad (9b)$$

$$\xi_i x_i = z_i \alpha, \quad i \in [1, n], \quad (9c)$$

$$z_i \geq 0, \quad i \in [1, n], \quad (9d)$$

$$\xi_i > 0, \quad i \in [1, n], \quad (9e)$$

$$\alpha > 0 \quad (9f)$$

We have the following result, its proof is deferred to Appendix B.2.

Proposition 2. *The system $\{(6b) - (6g), (9)\}$ is a correct reformulation of the system (6).*

4.2 Analysis of the loss due to estimation error Given a solution estimate $\hat{\mathbf{x}}$, we define the *loss* in objective value in the following way.

$$\text{loss}(\hat{\mathbf{x}}) = \boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \hat{\mathbf{x}}$$

A solution $\hat{\mathbf{x}}^R$ to the robust problem (4) depends on the expected return estimate $\hat{\boldsymbol{\mu}}$ and the estimation-error matrix Ξ . In this case we write

$$\text{loss}(\hat{\boldsymbol{\mu}}, \Xi) = \boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \hat{\mathbf{x}}^R$$

The first question we want to answer is the following: Given any $\hat{\boldsymbol{\mu}}$, does there exist a diagonal Ξ matrix such that $\text{loss}(\hat{\boldsymbol{\mu}}, \Xi)$ is equal to zero. The following example shows that it is not always possible to find such a diagonal matrix Ξ .

Example 3. *Consider a two asset portfolio optimization problem. Let $\sigma_{11} < v$, $\sigma_{22} > v$. Let $\mu_1 > \mu_2$. Clearly $(x_1 = 1, x_2 = 0)$ is the true optimal solution. Suppose that $\hat{\mu}_2 > \hat{\mu}_1$.*

It follows that we have $\lambda^1 = 0$ because the risk constraint is not tight. Furthermore, $\lambda_1^3 = 0$ because the non-negativity constraint of x_1 is not active. We see that the optimality conditions are satisfied

except (6a). Using our reformulation idea, we must find a feasible solution to the following system.

$$\begin{aligned}
 -\hat{\mu}_1 + z_1 + \lambda^2 &= 0 \\
 -\hat{\mu}_2 + z_2 + \lambda^2 - \lambda_2^3 &= 0 \\
 x_1 z_1 + x_2 z_2 &= \alpha \\
 \xi_1 x_1 &= z_1 \alpha \\
 \xi_2 x_2 &= z_2 \alpha \\
 \xi_1, \xi_2 &> 0 \\
 z_1, z_2 &\geq 0 \\
 \alpha &> 0
 \end{aligned}$$

We see that because $x_2 = 0$ and $\alpha > 0$, it must be that $z_2 = 0$. Furthermore, we can rewrite the first two equations in the following way by treating z_1 and λ_2^3 as slack variables.

$$\begin{aligned}
 \lambda^2 &\leq \hat{\mu}_1 \\
 \lambda^2 &\geq \hat{\mu}_2
 \end{aligned}$$

This is impossible since $\hat{\mu}_2 > \hat{\mu}_1$. Because the true optimal solution is unique, we conclude there does not exist ξ such that the robust portfolio optimization problem yields zero loss.

Our next result shows that even though it is not possible to reach zero loss, we can get arbitrarily close to it.

Theorem 4. *Given $\epsilon > 0$, for every $\hat{\mu}$ there exists a diagonal Ξ matrix such that $\text{loss}(\hat{\mu}, \Xi) \leq \epsilon$.*

Proof. Let $\mathbf{x}^{\text{mv}} := \arg \min_{\mathbf{x}} \{\mathbf{x}^T \Sigma \mathbf{x} : \mathbf{x} \in \Delta\}$ denote the minimum variance portfolio. Let \mathbf{x}^{eq} denote the equal portfolio, i.e. $x_i = 1/n$, $i \in [1, n]$. Pick $\delta \in (0, 1)$ such that $\tilde{\mathbf{x}} = \delta \mathbf{x}^{\text{mv}} + (1 - \delta) \mathbf{x}^{\text{eq}}$ has a lower variance than the true optimal solution \mathbf{x}^* . Clearly, all components of $\tilde{\mathbf{x}}$ are strictly greater than zero. Let $\tilde{\mathbf{x}}^* = (1 - \bar{\epsilon}) \mathbf{x}^* + \bar{\epsilon} \tilde{\mathbf{x}}$ for some $\bar{\epsilon} \in (0, 1)$. Note that $\tilde{\mathbf{x}}^* \in \mathcal{X}$ and has a lower variance than \mathbf{x}^* , therefore it is a feasible solution to the portfolio problem.

Now, observe that the loss corresponding to $\tilde{\mathbf{x}}^*$ is given by

$$\begin{aligned} \text{loss}(\tilde{\mathbf{x}}^*) &= \sum_{i=1}^n \mu_i (x_i^* - \tilde{x}_i^*) \\ &= \sum_{i=1}^n \mu_i (x_i^* - [(1 - \bar{\epsilon})x_i^* + \bar{\epsilon}\tilde{x}_i]) \\ &= \bar{\epsilon} \sum_{i=1}^n \mu_i (x_i^* - \tilde{x}_i) \\ &= \bar{\epsilon} (\boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \tilde{\mathbf{x}}) > 0. \end{aligned}$$

Therefore, given ϵ , we can pick $\bar{\epsilon} \in (0, 1)$ such that

$$\bar{\epsilon} \leq \frac{\epsilon}{(\boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \tilde{\mathbf{x}})}$$

This pick ensures that $\text{loss}(\tilde{\mathbf{x}}^*) \leq \epsilon$.

It suffices to find $\boldsymbol{\xi}$, λ^1 , λ^2 and λ^3 such that together with $\tilde{\mathbf{x}}^*$ we have a solution that satisfies the optimality conditions.

Let $\lambda^1 = 0$, and $\lambda_i^3 = 0$ for all $i \in [1, n]$. We see that all optimality conditions are satisfied except (6a). Using the reformulation idea, we see that finding such $\boldsymbol{\xi}$ is equivalent to finding a solution to the below system.

$$\begin{aligned} -\hat{\mu}_i + z_i + \lambda^2 &= 0, \quad i \in [1, n] \\ \sum_{i=1}^n z_i \tilde{x}_i^* &= \alpha \\ \xi_i \tilde{x}_i^* &= z_i \alpha, \quad i \in [1, n] \\ \xi_i &\geq 0, \quad i \in [1, n] \\ z_i &\geq 0, \quad i \in [1, n] \\ \alpha &\geq 0 \end{aligned}$$

Let $\lambda^2 < \min_i \{\hat{\mu}_i\}$. Let $z_i = \hat{\mu}_i - \lambda^2$ for all $i \in [1, n]$. Then $\alpha = \sum_{i=1}^n z_i \tilde{x}_i^*$. We let $\xi_i = \frac{z_i \alpha}{\tilde{x}_i^*}$ for all $i \in [1, n]$. Clearly, all ξ_i , z_i , α are greater than zero. Therefore, $\boldsymbol{\xi}$ satisfies the optimality conditions

and this concludes the proof. □

Finally we conclude this section by stating a sufficient condition for obtaining zero loss.

Theorem 5. *If all the assets are active in the true optimal solution, then for every $\hat{\boldsymbol{\mu}}$ there exists a diagonal $\boldsymbol{\Xi}$ matrix such that $\text{loss}(\hat{\boldsymbol{\mu}}, \boldsymbol{\Xi}) = 0$.*

Proof. Without loss of generality, we assume that all elements of the given $\hat{\boldsymbol{\mu}}$ are greater than 0. It suffices to find $\boldsymbol{\xi}$, λ^1 , λ^2 and λ^3 such that together with \mathbf{x}^* we have a solution that satisfies the optimality condition.

Let $\lambda^1 = 0$, $\lambda^2 = 0$, and $\lambda_i^3 = 0$ for all $i \in [1, n]$. Using the reformulation idea, we see that finding such $\boldsymbol{\xi}$ is equivalent to finding a solution to the below system.

$$\begin{aligned} -\hat{\mu}_i + z_i &= 0, & i \in [1, n] \\ \sum_{i=1}^n z_i x_i^* &= \alpha \\ \xi_i x_i^* &= z_i \alpha, & i \in [1, n] \\ \xi_i &> 0, & i \in [1, n] \\ z_i &\geq 0, & i \in [1, n] \\ \alpha &> 0 \end{aligned}$$

Let $z_i = \hat{\mu}_i$ for all $i \in [1, n]$. Then we have that $\alpha = \sum_{i=1}^n \hat{\mu}_i x_i^*$. We let $\xi_i = \frac{z_i \alpha}{x_i^*}$ for all $i \in [1, n]$. Clearly, all ξ_i , z_i , α are greater than zero. Therefore, $\boldsymbol{\xi}$ satisfies the optimality conditions and this concludes the proof. □

5. Finding the best estimation-error matrix In this section, we generalize the results of the previous section to the situation where we have several estimates $\hat{\boldsymbol{\mu}}^1, \dots, \hat{\boldsymbol{\mu}}^T$ of $\boldsymbol{\mu}$. We show that we can always find a diagonal matrix $\boldsymbol{\Xi}$ such that the resulting robust portfolios $\hat{\mathbf{x}}^1, \dots, \hat{\mathbf{x}}^T$ all have an actual expected return very close to the optimal expected return $\boldsymbol{\mu}^\top x^*$ where x^* is the optimal

portfolio computed using the true $\boldsymbol{\mu}$. To this end, we write a mathematical program that outputs a matrix $\boldsymbol{\Xi}$ that achieves the minimum loss. We define the loss given a collection of return estimates $\{\hat{\boldsymbol{\mu}}^t\}_{t=1}^T$ and estimation-error matrix $\boldsymbol{\Xi}$ as

$$\text{loss}(\{\hat{\boldsymbol{\mu}}^t\}_{t=1}^T, \boldsymbol{\Xi}) = \sum_{t \in [1, T]} (\boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \hat{\mathbf{x}}^{R,t})$$

where $\hat{\mathbf{x}}^{R,t}$ denotes the optimal robust portfolio for $t \in [1, T]$.

5.1 A bilevel programming formulation In this section, we describe the bilevel programming formulation that computes the best estimation-error matrix $\boldsymbol{\Xi}$ for a given collection of return estimates $\{\hat{\boldsymbol{\mu}}^t\}_{t=1}^T$.

The following parameters are used in the description of the bilevel model.

- $\boldsymbol{\mu}$: True expected return vector.
- $\boldsymbol{\Sigma}$: True covariance matrix of the asset returns.
- $\hat{\boldsymbol{\mu}}^t$: Estimated return vector under trial $t \in [1, T]$.
- v : Risk threshold.
- $\mathbf{x}^*(v) = \arg \max_{\mathbf{x} \in \Delta^n} \left\{ \boldsymbol{\mu}^\top \mathbf{x} : \mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} \leq v \right\}$, The portfolio with maximum expected return for a given risk threshold value v .

Let $\hat{\mathbf{x}}^t(v, \boldsymbol{\Xi}, \hat{\boldsymbol{\mu}}^t)$ denote the optimal portfolio given that the estimated return vector is $\hat{\boldsymbol{\mu}}^t$, estimation-error matrix is $\boldsymbol{\Xi}$, and the risk threshold is v under trial $t \in [1, T]$. Note that this value is obtained by solving a convex optimization problem of the form (4). Therefore, we have the following bilevel program to compute the best $\boldsymbol{\Xi}$ that minimizes the loss.

$$\underset{\boldsymbol{\Xi}}{\text{minimize}} \quad \sum_{t=1}^T \left(\boldsymbol{\mu}^\top \mathbf{x}^*(v) - \boldsymbol{\mu}^\top \hat{\mathbf{x}}^t(v, \boldsymbol{\Xi}, \hat{\boldsymbol{\mu}}^t) \right) \quad (10a)$$

$$\text{subject to} \quad \hat{\mathbf{x}}^t(v, \boldsymbol{\Xi}, \hat{\boldsymbol{\mu}}^t) = \arg \max_{\mathbf{x} \in \Delta^n} \left\{ (\hat{\boldsymbol{\mu}}^t)^\top \mathbf{x} - \sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}} : \mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} \leq v \right\}, \quad t \in [1, T], \quad (10b)$$

$$\boldsymbol{\Xi} \in \mathcal{S}_{++}^n. \quad (10c)$$

In this formulation, Ξ is the upper-level decision variable, and each $\hat{\mathbf{x}}^t(v, \Xi, \hat{\boldsymbol{\mu}}^t)$ is a lower-level decision variable. Note that we have as many lower-level problems as the number of trials we have.

The above bilevel program minimizes the sum of the difference between the value of the true frontier, and the value of the actual frontier under each trial. Alternatively, one can use other performance measures, for example minimizing the maximum of the differences. In this formulation, we assume the risk threshold v is large enough that the lower-level problem is feasible. Our next goal is to reformulate problem (10) as a single level optimization program.

5.2 Reformulating (10) as a single level program We know that robust portfolio optimization problem is a convex optimization program, it satisfies Slater's condition, and it has a unique solution, see Proposition 1. Therefore, we can use the optimality conditions (6) to reformulate (10) as a single level program (Bard, 1998).

$$\begin{aligned} & \underset{\Xi, \hat{\mathbf{x}}^t, \lambda^{1,t}, \lambda^{2,t}, \lambda^{3,t}}{\text{minimize}} && \sum_{t=1}^T \left(\boldsymbol{\mu}^\top \mathbf{x}^* - \boldsymbol{\mu}^\top \mathbf{x}^t \right) \end{aligned} \quad (11a)$$

$$\text{subject to} \quad -\hat{\boldsymbol{\mu}}^t + \frac{\Xi \mathbf{x}^t}{\sqrt{(\mathbf{x}^t)^\top \Xi \mathbf{x}^t}} + 2\lambda^{1,t} \Sigma \mathbf{x}^t + \lambda^{2,t} \mathbf{1} - \boldsymbol{\lambda}^{3,t} = \mathbf{0}, \quad t \in [1, T], \quad (11b)$$

$$(\mathbf{x}^t)^\top \Sigma \mathbf{x}^t \leq v, \quad t \in [1, T], \quad (11c)$$

$$\mathbf{1}^\top \mathbf{x}^t = 1, \quad t \in [1, T], \quad (11d)$$

$$-\mathbf{x}^t \leq \mathbf{0}, \quad t \in [1, T], \quad (11e)$$

$$\lambda^{1,t} \geq 0, \quad t \in [1, T], \quad (11f)$$

$$\boldsymbol{\lambda}^{3,t} \geq \mathbf{0}, \quad t \in [1, T], \quad (11g)$$

$$\lambda^{1,t} \left((\mathbf{x}^t)^\top \Sigma \mathbf{x}^t - v \right) + (\boldsymbol{\lambda}^{3,t})^\top \mathbf{x}^t = 0, \quad t \in [1, T], \quad (11h)$$

$$\Xi \in \mathcal{S}_{++}. \quad (11i)$$

Note that in the above formulation, the optimality conditions are appended for each trial $t \in [1, T]$.

The bilevel model (11) is a non-convex optimization problem due to the stationary point equations (11b) and the complementary slackness conditions (11h). We next present a reformulation for (11) when Ξ is a positive definite diagonal matrix.

Using the earlier reformulation idea in Section 4, we can reformulate (11) so as to avoid the square-root terms.

$$\begin{aligned} & \underset{\boldsymbol{\xi}, \mathbf{x}^t, \lambda^{1,t}, \lambda^{2,t}, \lambda^{3,t}}{\text{minimize}} & (11a) & (12a) \end{aligned}$$

$$\text{subject to} \quad (11c) - (11h) \quad (12b)$$

$$-\hat{\mu}_i^t + z_i^t + 2\lambda^{1,t} \sum_{j=1}^n \sigma_{ij} x_j^t + \lambda^{2,t} - \lambda_i^{3,t} = 0, \quad i \in [1, n], t \in [T], \quad (12c)$$

$$\sum_{i=1}^n x_i^t z_i^t = \alpha^t, \quad t \in [T], \quad (12d)$$

$$\xi_i x_i^t = z_i^t \alpha^t, \quad i \in [n], t \in [T], \quad (12e)$$

$$z_i^t \geq 0, \quad i \in [1, n], t \in [T], \quad (12f)$$

$$\alpha^t > 0, \quad t \in [T], \quad (12g)$$

$$\xi_i > 0. \quad (12h)$$

5.3 Analysis of the bilevel model In this section, we show that the bilevel model (12) has always an optimal solution that yields arbitrarily small loss.

Theorem 6. *For every $\epsilon > 0$ and $\{\hat{\boldsymbol{\mu}}^t\}_{t=1}^T$, there exists a diagonal matrix Ξ such that $\text{loss}(\{\hat{\boldsymbol{\mu}}^t\}_{t=1}^T, \Xi) \leq \epsilon$.*

Proof. It suffices to find $\boldsymbol{\xi}$, $\lambda^{1,t}$, $\lambda^{2,t}$, $\lambda^{3,t}$, and \mathbf{x}^t such that they are feasible for (12), and all the vectors \mathbf{x}^t are within a small neighborhood of \mathbf{x}^* .

Let $\lambda^{1,t} = 0$, and $\lambda_i^{3,t} = 0$ for all $i \in [1, n]$ and $t \in [1, T]$. Then (12c) can be written as

$$-\hat{\mu}_i^t + z_i^t + \lambda^{2,t} = 0, \quad i \in [1, n], t \in [T], \quad (13)$$

Using (12h), we substitute $x_i^t = z_i^t \alpha^t / \xi_i$ in (12g). We see that it suffices to solve the following system

to complete the proof.

$$\sum_{i \in [1, n]} \frac{(\hat{\mu}_i^t - \lambda^{2,t})^2}{\xi_i} = 1, \quad t \in [1, T]. \quad (14a)$$

$$z_i^t \geq 0, \quad i \in [1, n], \quad t \in [T], \quad (14b)$$

$$\alpha^t > 0, \quad t \in [T], \quad (14c)$$

$$\xi_i > 0. \quad (14d)$$

Let $\tilde{\mathbf{x}}^*$ be as in the proof of Theorem 4.

That is, $\tilde{\mathbf{x}}^*$ is a feasible solution of (1) that satisfies the constraints $\mathbf{x}^\top \Sigma \mathbf{x} \leq v$ and $\mathbf{x} \geq \mathbf{0}$ strictly and is within a small ball of the true optimum \mathbf{x}^* .

Pick $\xi_i = M/\tilde{x}_i^*$ for some large M . With this choice, the equations in (14a) are written as

$$\sum_{i \in [1, n]} \frac{(\hat{\mu}_i^t - \lambda^{2,t})^2}{\tilde{x}_i^*} = M, \quad t \in [1, T]. \quad (15)$$

For each $t \in [1, T]$, start by letting $\lambda^{2,t} = \min_{i \in [1, n]} \{\hat{\mu}_i^t\}$ and decrease $\lambda^{2,t}$ until the equation in (15) is satisfied. This will happen, because the left hand side in (15) is continuous and monotonically increasing.

Now, let $z_i^t = \hat{\mu}_i^t - \lambda^{2,t}$. Let $\eta > 0$ be any positive real. Note that by choosing M sufficiently large, we can guarantee that, for each t , all z_i^t are within $1 + \eta$ of the smallest. Namely, for all $i \in [1, n]$, $k_t \leq z_i^t \leq (1 + \eta)k_t$, where $k_t = \min_j z_j^t$. This is because the $\hat{\mu}_i^t$ are fixed whereas $-\lambda^{2,t}$ increases as M increases.

For each $t \in [1, T]$, set $\alpha^t = \frac{1}{\sum_i \frac{z_i^t}{\xi_i}}$. Now let $x_i^t = \alpha^t \times \frac{z_i^t}{\xi_i}$. Note that we have $\sum_{i \in [1, n]} x_i^t = 1$ and $x_i^t > 0$. To show that \mathbf{x}^t is feasible, we only need to show that it satisfies the variance constraint. This will follow from showing that \mathbf{x}^t is in a small neighborhood of $\tilde{\mathbf{x}}^*$, which satisfies it strictly. We have $x_i^t = \alpha^t \times \frac{z_i^t}{\xi_i} = \alpha^t \times \frac{z_i^t \tilde{x}_i^*}{M}$. Using our bounds on z_i^t , we get

$$\frac{\alpha^t k_t}{M} \tilde{x}_i^* \leq x_i^t \leq (1 + \eta) \frac{\alpha^t k_t}{M} \tilde{x}_i^*.$$

Adding over i we get $\frac{\alpha^t k_t}{M} \leq 1 \leq (1 + \eta) \frac{\alpha^t k_t}{M}$. This implies

$$(1 - \eta) \tilde{x}_i^* \leq x_i^t \leq (1 + \eta) \tilde{x}_i^*.$$

This translates into a vanishing loss in the objective value as η goes to 0. And this concludes the proof.

□

6. Using the identity matrix as estimation-error matrix The results of the previous sections demonstrate that one can focus on diagonal estimation-error matrices and consequently deal with calibrating only n parameters in constructing the uncertainty sets. In this section, we restrict our attention to the simplest diagonal estimation-error matrix, the identity matrix. In this case, we only need to calibrate one parameter (κ), which is advantageous for decision makers.

Our goal is to investigate whether using an identity matrix as the estimation-error matrix can improve the performance of the classical Markowitz model. To this end, we perform a simulation where we compare the out-of-sample performances of the two approaches. We generate samples from the distribution $\text{Normal}(\boldsymbol{\mu}, \boldsymbol{\Sigma}/n)$. The value of n determines the accuracy of the samples we generate (increasing n results in better estimates of the expected return). Then, we solve the classical Markowitz model and the robust portfolio model for varying κ values and compare their actual performances.

		$\kappa \times n$									
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
n	1	13.0	18.3	19.3	19.4	19.3	19.2	19.1	19.0	18.9	18.9
	3	7.8	13.7	16.3	17.2	17.4	17.4	17.2	17.0	16.8	16.7
	6	5.3	10.1	13.1	14.6	15.2	15.3	15.2	15.0	14.8	14.6
	12	3.6	7.0	9.6	11.3	12.2	12.5	12.6	12.4	12.2	11.9
	24	2.4	4.6	6.4	7.8	8.7	9.1	9.2	9.0	8.7	8.3
	120	1.2	2.0	2.7	3.1	3.3	3.2	2.8	2.3	1.6	0.8

Table 1: Percentage gap closed by the robust model when $\boldsymbol{\Xi} = \mathbf{I}$ ($v = 0.002$).

Table 1 gives the result of our simulation. Each cell in Table 1 denotes the percentage gap closed by the robust solution estimates compared to the Markowitz solution estimates. In particular, in each

cell we report $(\bar{R} - \bar{M}) / (T - \bar{M})$ where \bar{R} denotes the actual performance of the robust solution estimates, \bar{M} denotes the actual performance of the Markowitz solution estimates, and T denotes the true return. Note that for each cell, \bar{R} and \bar{M} values are averages over 10000 trials and we report the standard errors for percentage gap closed values in Appendix C.

We see from Table 1 that the robust model can significantly outperform the Markowitz model when κ is chosen correctly. This is interesting, because the identity matrix contains no information about the problem or the relationship between the assets. We observe that the percentage gap closed is significantly larger when n is smaller. This shows that robustification helps more when the return estimates are poor.

Furthermore, the results in Table 1 illustrate that the $\kappa \times n$ value that performs the best varies between 0.4 and 0.7. This implies that keeping $\kappa \times n$ constant when n varies works well for improving the actual performance. We interpret this result as the need for robustification diminishes when we have better estimates as n increases.

In order to understand the results in Table 1 better, we focus on two particular distributions $n = 3$ and $n = 24$ when $\kappa \times n = 0.5$. We present the averages \bar{M} and \bar{R} as a histogram. Figure 3 provides




Figure 3: Histograms of the actual performances of Robust and Markowitz models.

a clearer picture for the superior performance of the robust model. We see that the empirical distributions of the actual returns under the robust model have a smaller variance. On the other

hand, the actual returns under the Markowitz model features outliers that perform very poorly. As a result, the robust model outperforms the Markowitz model, and the difference is significant.

7. Conclusion In this paper, we take a theoretical look at the robust portfolio optimization problem. Our main contribution is to show that the family of diagonal estimation-error matrices is sufficiently rich to achieve an arbitrarily small loss compared to the optimum expected return. We perform illustrative computational experiments to show that even using an identity matrix can improve on the actual performance of the Markowitz model when the size of the uncertainty set is correctly calibrated.

Our work is a step towards constructing better estimation-error matrices for robust portfolio optimization. Our bilevel model can be used to construct good estimation-error matrices using the available data. This will require developing a more efficient method to solve the bilevel model. Another research direction would be a theoretical analysis on the use of identity matrices as estimation-error matrices.

A. Data set We use the data set provided in [Kocuk and Cornuéjols \(2020\)](#). The data set includes 360 monthly returns of 11 sectors based on the Global Industrial Classification Standard. We calculate the true return vector $\boldsymbol{\mu}$ and the true covariance matrix $\boldsymbol{\Sigma}$ as the sample average and the sample covariance matrix of these returns. We refer the reader to [Kocuk and Cornuéjols \(2020\)](#) for more details about the dataset.

B. Additional proofs In this section, we present the additional proofs.

B.1 Proof of Proposition 1

1. We know that there exists a matrix \mathbf{L} such that $\boldsymbol{\Xi} = \mathbf{L}\mathbf{L}^\top$, because $\boldsymbol{\Xi}$ is a positive definite matrix. Therefore, the objective function f of (5) can be written as

$$\begin{aligned} f(\mathbf{x}) &= -(\hat{\boldsymbol{\mu}})^\top \mathbf{x} + \sqrt{\mathbf{x}^\top \boldsymbol{\Xi} \mathbf{x}} \\ &= -(\hat{\boldsymbol{\mu}})^\top \mathbf{x} + \|\mathbf{x}\mathbf{L}\|_2 \end{aligned} \tag{16}$$

To conclude the proof, it suffices to show that the norm function $\|\cdot\|_2$ is strictly convex on the feasible region of (5), because the composition with an affine mapping of a strictly convex function is strictly convex.

To this end, we first note that any two distinct point in the feasible region of (5) are linearly independent because of $\mathbf{1}^\top \mathbf{x} = 1$. Let $\lambda \in (0, 1)$. We want to show that

$$\|\lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2\|_2 < \lambda \|\mathbf{x}_1\|_2 + (1 - \lambda) \|\mathbf{x}_2\|_2$$

We take the square of both sides and observe that

$$\begin{aligned} \|\lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2\|_2^2 &= (\lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2, \lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2) \\ &= \lambda^2 (\mathbf{x}_1, \mathbf{x}_1) + (1 - \lambda)^2 (\mathbf{x}_2, \mathbf{x}_2) + 2\lambda(1 - \lambda) (\mathbf{x}_1, \mathbf{x}_2) \\ &= \lambda^2 \|\mathbf{x}_1\|_2^2 + (1 - \lambda)^2 \|\mathbf{x}_2\|_2^2 + 2\lambda(1 - \lambda) (\mathbf{x}_1, \mathbf{x}_2) \\ &< \lambda^2 \|\mathbf{x}_1\|_2^2 + (1 - \lambda)^2 \|\mathbf{x}_2\|_2^2 + 2\lambda(1 - \lambda) \|\mathbf{x}_1\|_2 \|\mathbf{x}_2\|_2 \\ &= (\lambda \|\mathbf{x}_1\|_2 + (1 - \lambda) \|\mathbf{x}_2\|_2)^2 \end{aligned}$$

The strict inequality above is true due to the fact that Cauchy-Schwarz inequality is strict when \mathbf{x}_1 and \mathbf{x}_2 are linearly independent.

2. It is clear that the robust portfolio optimization problem (5) is a convex program. It suffices to exhibit an interior point of the feasible region. We assume that v is strictly greater than the minimum variance portfolio. Let \mathbf{x}^{mv} and \mathbf{x}^{eq} denote the minimum variance and the equal portfolios, respectively. There exists $\delta \in (0, 1)$ such that the variance of $\tilde{\mathbf{x}} = \delta\mathbf{x}^{\text{mv}} + (1 - \delta)\mathbf{x}^{\text{eq}}$ is less than v . Clearly, all components of $\tilde{\mathbf{x}}$ is greater than zero. Thus, $\tilde{\mathbf{x}}$ is an exterior point of the feasible region of (5).
3. The square-root function is non-differentiable only at the origin. We have that the objective function (16) is always differentiable in the feasible region of (5), because Ξ is positive definite, and $\mathbf{1}^\top \mathbf{x} = 1$.

B.2 Proof of Proposition 2 We want to show the equivalence of the system (6) and the system $\{(6b) - (6g), (9)\}$. We only need to show given $(\Xi, \mathbf{x}, \lambda^1, \lambda^2, \lambda^3)$ that satisfies (6) with Ξ is positive definite and diagonal, there exists $(\mathbf{z}, \alpha, \boldsymbol{\xi}, \mathbf{x}, \lambda^1, \lambda^2, \lambda^3)$ that satisfies $\{(6b) - (6g), (9)\}$. The other direction is trivial since the $\boldsymbol{\xi}$ vector in any feasible solution to $\{(6b) - (6g), (9)\}$ is a diagonal estimation-error matrix that satisfies the optimality conditions.

Let $\boldsymbol{\xi}$ denote the diagonal entries of Ξ . Because Ξ is positive definite, we know that all elements of $\boldsymbol{\xi}$ are strictly greater than zero. It suffices to find (\mathbf{z}, α) that satisfies (9), since $(\mathbf{x}, \lambda^1, \lambda^2, \lambda^3)$ automatically satisfies (6b) – (6g).

Consider any given $\mathbf{x} \geq 0$. Let $P := \{i \in [1, n] : x_i > 0\}$. Let $Z = [1, n] \setminus P$. Since $\mathbf{1}^\top \mathbf{x} = 1$, we know that P is non-empty. This implies that the right hand side of (9c) is positive for all $i \in P$. Therefore, α must be positive.

Now note that we can write (9b) as

$$\begin{aligned} \alpha &= \sum_{i \in P} x_i z_i + \sum_{i \in Z} x_i z_i \\ &= \sum_{i \in P} x_i z_i \\ &= \sum_{i \in P} x_i \frac{\xi_i x_i}{\alpha} \end{aligned}$$

Thus we have that $\alpha = \sqrt{\sum_{i \in P} x_i \xi_i x_i} > 0$. Accordingly, each z_i assumes the value $\frac{\xi_i x_i}{\alpha}$ for all $i \in P$. Letting each $z_i = 0$ for all $i \in Z$ concludes the proof.

C. Standard errors of Table 1 We use bootstrapping to calculate the standard errors for the statistic $(\bar{R} - \bar{M})/(T - \bar{M})$ described in Section 6.

		$\kappa \times n$									
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
n	1	0.29	0.33	0.35	0.38	0.39	0.39	0.41	0.40	0.41	0.42
	3	0.26	0.32	0.34	0.36	0.38	0.39	0.40	0.41	0.41	0.43
	6	0.23	0.30	0.34	0.36	0.38	0.39	0.41	0.42	0.42	0.43
	12	0.21	0.28	0.32	0.35	0.37	0.39	0.41	0.42	0.44	0.45
	24	0.19	0.26	0.31	0.34	0.37	0.39	0.41	0.43	0.44	0.46
	120	0.16	0.23	0.28	0.32	0.35	0.37	0.41	0.43	0.45	0.47

Table 2: Standard errors for percentage gap closed statistic in Table 1.

References

- Ban, G., El Karoui, N., and Lim, A. (2018). Machine learning and portfolio optimization. *Management Science*, 64(3):1136–1154.
- Bard, J. F. (1998). *Practical Bilevel Optimization*, volume 30 of *Nonconvex Optimization and Its Applications*. Springer US, Boston, MA.
- Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13.
- Best, M. J. and Grauer, R. R. (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.
- Broadie, M. (1993). Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, 45(1):21–58.
- Ceria, S. and Stubbs, R. A. (2006). Incorporating estimation errors into portfolio selection: Robust portfolio construction. *Journal of Asset Management*, 7(2):109–127.
- Chopra, V. and Ziemba, W. (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19(2):6–11.
- Chopra, V. K. (1993). Mean-variance revisited: Near-optimal portfolios and sensitivity to input variations. *Journal of Investing*.
- DeMiguel, V., Garlappi, L., Nogales, F., and Uppal, R. (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5):798–812.
- Fabozzi, F. J., Kolm, P. N., Pachamanova, D. A., and Focardi, S. M. (2007). *Robust portfolio optimization and management*. John Wiley & Sons.
- Goldfarb, D. and Iyengar, G. (2003). Robust Portfolio Selection Problems. *Mathematics of Operations Research*, 28(1):1–38.

- Gotoh, J.-y. and Takeda, A. (2011). On the role of norm constraints in portfolio selection. *Computational Management Science*, 8(4):323–353.
- Horst, J. R. t., Roon, F. A. d., and Werker, B. J. M. (2001). Incorporating estimation risk in portfolio choice. *Technical report, Department of Finance, Tilburg University, Netherlands*.
- Jagannathan, R. and Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance*, 58:1651–1684.
- Kocuk, B. and Cornuéjols, G. (2020). Incorporating black-litterman views in portfolio construction when stock returns are a mixture of normals. *Omega*, 91:102008.
- Lim, A., Shanthikumar, J., and G.-Y., V. (2012). Robust portfolio choice with learning in the framework of regret: Single-period case. *Management Science*, 58(9):1732–1746.
- Markowitz, H. (1952). Portfolio analysis. *Journal of Finance*, 8:77–91.
- Michaud, R. (2008). *Efficient asset management*. Oxford University Press.
- Olivares-Nadal, A. V. and DeMiguel, V. (2018). Technical Note—A Robust Perspective on Transaction Costs in Portfolio Optimization. *Operations Research*, 66(3):733–739.
- Scherer, B. (2007). Can robust portfolio optimisation help to build better portfolios? *Journal of Asset Management*, 7(6):374–387.
- Stubbs, R. A. and Vance, P. (2005). Computing return estimation error matrices for robust optimization. *Axioma Research Paper*.
- Tutuncu, R. and Koenig, M. (2004). Robust asset allocation. *Annals of Operations Research*, 132:157–187.