

Robust Asset Allocation

R. H. Tütüncü* and M. Koenig†

August 20, 2002
Revised July 3, 2003

Abstract

This article addresses the problem of finding an optimal allocation of funds among different asset classes in a robust manner when the estimates of the structure of returns are unreliable. Instead of point estimates used in classical mean-variance optimization, moments of returns are described using *uncertainty sets* that contain all, or most, of their possible realizations. The approach presented here takes a conservative viewpoint and identifies asset mixes that have the best worst-case behavior. Techniques for generating uncertainty sets from historical data are discussed and numerical results that illustrate the stability of robust optimal asset mixes are reported.

Key words: Robust optimization, mean-variance optimization, saddle-point problems.

1 Introduction

Portfolio optimization is one of the best known and most widely used methods in financial portfolio selection. Developed by Harry Markowitz five decades ago [4], this approach quantifies the trade-off between the expected return and the risk of portfolios of financial securities using mathematical techniques and offers a method for determining a frontier of optimal (Pareto-efficient) portfolios. Since risk is measured by the variance of the random portfolio return in this approach, it is also called mean-variance optimization (MVO). Here, we address asset allocation problems, i.e., the problem of finding an optimal allocation of funds among different asset classes which can be formulated in an identical manner to MVO.

Often, the set of optimal or efficient portfolios is described using a two-dimensional graph called the *efficient frontier* that plots their expected returns and standard deviations. Each efficient portfolio can be identified by solving an associated convex

*Associate Professor, Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania.

†Director of Quantitative Analysis, National City Investment Management Company, Cleveland, Ohio.

quadratic programming (QP) problem, a well-studied problem in the optimization literature. To determine the entire efficient frontier ranging from the portfolio with the smallest overall variance to the portfolio with the highest expected return, one has to solve a *parametric QP* problem using, for example, Markowitz' method of critical lines.

Despite the elegance of the model developed by Markowitz, the powerful optimization theory supporting this model, and the availability of efficient software to solve the resulting problem, MVO continues to encounter skepticism among investment practitioners. One reason usually cited for this skepticism is the counter-intuitive nature of the optimal portfolios generated by the MVO approach. Optimal portfolios tend to concentrate on a small subset of the available securities, and appear not to be well diversified. Furthermore, optimal portfolios are often sensitive to changes in the input parameters of the problem (expected returns and the covariance matrix) and lead to large turnover ratios with periodic readjustments of the input estimates; see for example [8, 9].

This last observation indicates that the inputs to the MVO model need to be very accurately estimated. However, this is a very difficult task, especially in the case of expected return estimates. Different techniques used in moment estimation can and do generate significantly different point estimates of MVO inputs which, in turn, lead to large variations in the composition of efficient portfolios. Using estimates from a particular source in the MVO model introduces an *estimation risk* in portfolio choice, and methods for optimal selection of portfolios must take this risk into account [6].

Robust optimization, an emerging branch of the field of optimization, offers vehicles to incorporate estimation risk into the decision making process in portfolio choice/asset allocation. Generally speaking, robust optimization refers to finding solutions to given optimization problems with uncertain input parameters that will achieve good objective values for all, or most, realizations of the uncertain input parameters. It should be noted, however, that there are different interpretations of robustness that lead to different mathematical formulations—see [7] for at least 17 different definitions of robustness in different contexts. Here, we take the pessimistic view of robustness and look for a solution that has the best performance under its worst case.

In our approach, uncertainty is described using an *uncertainty set* which includes all, or most, possible realizations of the uncertain input parameters. Given a problem with uncertain inputs and an uncertainty set for these inputs, our robust optimization approach addresses the following problem: What choice of the variables of the problem will optimize the worst case objective value? That is, for each choice of the decision variables, we consider the worst case realization of the data and evaluate the corresponding objective value, and then pick the set of values for the variables with the best worst-case objective. We apply this approach to the portfolio selection problem using a judicious choice of the uncertainty set. We demonstrate that the resulting robust optimization problem is *simple* in some cases meaning that it can be solved as a standard quadratic programming problem. In most cases, however, this simplification is not possible. For such cases we formulate the robust optimization problem as a saddle-point problem and apply an interior-point algorithm to this saddle-point problem.

While the approach in this article is related to the methods in [3, 10], the formulation

and the algorithm used here are based on those developed by Halldórsson and Tütüncü [2]. In addition to presenting a variation of their formulation and discussing a new formulation for identifying robust portfolios with the largest Sharpe ratio, we discuss an implementation of these algorithms.

We also address the issue of generating uncertainty sets and describe two approaches; one based on bootstrapping and the other on moving averages. Since our worst-case based approach can be detrimentally influenced by outliers in the data, uncertainty sets need to be carefully chosen. This is why we may prefer to include most rather than all possible realizations of the uncertain parameters in uncertainty sets. To minimize outlier effects, we eliminate some of the lowest and highest quantiles of the processed data in both the bootstrapping and moving averages strategies and use the remaining data to define uncertainty sets.

Our numerical experiments indicate that robust asset allocation is indeed a valuable alternative for conservative investors. Robustness is achieved at relatively little cost—robust efficient portfolios are only marginally inefficient when faced with nominal inputs. In contrast, efficient portfolios derived from nominal inputs can be severely inefficient under worst-case realizations of the uncertain parameters. We further demonstrate that robust optimal allocations are stable in the sense that re-solving the robust asset allocation problem periodically as new data is collected results in essentially unchanged portfolios. This type of low turnover is often attractive for long-term investors.

The remainder of this article is organized as follows: Section 2 presents formulations of problems to find robust optimal allocation of assets and robust portfolios with the maximum Sharpe ratio. In Section 3, we present a rigorous description of the method we implemented to determine the robust efficient frontier. A detailed description of a key subroutine is given in the Appendix. Numerical experiments and their results are discussed in Section 4. We present our conclusions in Section 5.

2 Robust Optimization Problems

2.1 The Robust MVO Problem

Optimal portfolio selection/asset allocation problems can be formulated mathematically as quadratic programming (QP) problems. Convex QP refers to minimizing a convex quadratic function (or, equivalently, maximizing a concave quadratic function) subject to linear equality and inequality constraints. Solution of a convex QP associated with an asset allocation problem generates an efficient portfolio on the efficient frontier. To generate the entire efficient frontier, the QP has to be parametrized and this can be done in three essentially equivalent ways: (i) maximize expected return subject to an upper limit on the variance, (ii) minimize the variance subject to a lower limit on the expected return, (iii) maximize the risk-adjusted expected return. These three problems are parametrized by the variance limit, expected return limit, and the risk-aversion parameter, respectively. Since the variance constraint is nonlinear, the

first formulation is not a QP. Hence, we focus on the last two of these formulations:

$$\begin{aligned} \min_{x \in \mathfrak{R}^n} \quad & x^T Q x \\ \text{s.t.} \quad & \mu^T x \geq R, \\ & x \in \mathcal{X} \end{aligned} \tag{1}$$

$$\begin{aligned} \max_{x \in \mathfrak{R}^n} \quad & \mu^T x - \lambda x^T Q x \\ \text{s.t.} \quad & x \in \mathcal{X} \end{aligned} \tag{2}$$

Above, μ_i —the i^{th} component of the vector μ —denotes the estimated expected return of security i . The matrix Q is the covariance matrix of these returns. Diagonal elements q_{ii} of the Q matrix denote the variance of the return on security i while off-diagonal elements q_{ij} denote the covariance between the returns of securities i and j . The components x_i of the variable vector x denote the proportion of the portfolio to be invested in security i . R in the right-hand-side of (1) is the lower limit on the expected return one would like to achieve. \mathcal{X} represents the polyhedral set of *feasible* portfolios, i.e., portfolios that satisfy certain linear constraints imposed by the investor. The scalar λ in (2) represents the risk-aversion parameter mentioned above. The objective function of problem (2) represents a risk-adjusted expected return function. Since the covariance matrix Q is always positive semidefinite, problems (1) and (2) are convex quadratic programming problems solvable in polynomial time.

By solving (1) and (2) for different values of R and λ , one can generate a sequence of optimal portfolios on the efficient frontier ranging from the portfolio with the smallest overall variance to the portfolio with the highest expected return. Problems (1) and (2) are equivalent in the following sense: If $x^*(\lambda)$ solves (2) for a fixed value of λ , it also solves (1) for $R = \mu^T x^*(\lambda)$.

We will make the reasonable assumption that the set \mathcal{X} of feasible portfolios is non-empty and bounded. For example,

$$\mathcal{X} = \{x \in \mathfrak{R}^n \mid \sum_{i=1}^n x_i = 1, x \geq 0\} \tag{3}$$

corresponds to the case where short-sales are not allowed and satisfies this assumption. In a more general setting, \mathcal{X} may contain additional constraints corresponding to sector distribution requirements, minimal investment restrictions, bounds on the total number of assets in the portfolio, etc.

We argued above that one of the main criticisms against the MVO approach centers on the observation that the optimal portfolios generated by this approach are often quite sensitive to the input parameters— μ and Q in our notation. To make matters worse, these parameters can never be observed and one has to settle for estimates found using some particular technique. In the presence of, say, equally reliable multiple estimates for μ and Q , the best way to integrate all available information is no longer clear. Our strategy is to represent all the available information on the unknown input parameters in the form of an *uncertainty set*, i.e., a set that contains most of the possible values for these parameters.

In the case of the MVO problems (1) and (2), an uncertainty set for the expected return vector μ and the covariance matrix Q may take the form of intervals:

$$\mathcal{U}_\mu = \{\mu : \mu^L \leq \mu \leq \mu^U\}, \quad (4)$$

$$\mathcal{U}_Q = \{Q : Q^L \leq Q \leq Q^U, \quad Q \succeq 0\}, \quad (5)$$

$$\mathcal{U} = \{(\mu, Q) : \mu \in \mathcal{U}_\mu, Q \in \mathcal{U}_Q\}. \quad (6)$$

Above, μ^L, μ^U, Q^L, Q^U are the extreme values of the intervals we just mentioned. The restriction $Q \succeq 0$ indicates that Q is a symmetric positive semidefinite matrix which is a necessary property for the uncertain Q to be a covariance matrix. The uncertainty set \mathcal{U} need not be separable as above, but we focus on separable sets in this paper as they appear more natural than compound uncertainty sets \mathcal{U} that are not separable.

There are, of course, other ways to describe uncertainty sets. For example, they could be discrete sets representing a collection of estimates for the unknown input parameters. Goldfarb and Iyengar consider ellipsoidal uncertainty sets [10]. We prefer intervals and there are several reasons for this preference. For example, the end-points of the interval may correspond to the extreme values of the corresponding statistic in historical data, in analyst estimates, in simulated scenarios, etc. Alternatively, a modeler may choose a confidence level and then generate estimates of return and covariance parameters in the form of prediction intervals. For the examples we present in Section 4, we generated the bounds μ^L, μ^U, Q^L, Q^U using percentiles of bootstrapped samples of historical data as well as the percentiles of moving averages.

Given the uncertainty set \mathcal{U} , the robust versions of the mean-variance optimization problems (1) and (2) can be expressed as follows:

$$\begin{aligned} \min \quad & \{\max_{Q \in \mathcal{U}_Q} x^T Q x\} \\ \text{s.t.} \quad & x \in \mathcal{X} \\ & \min_{\mu \in \mathcal{U}_\mu} \mu^T x \geq R, \end{aligned} \quad (7)$$

$$\max_{x \in \mathcal{X}} \left\{ \min_{\mu \in \mathcal{U}_\mu, Q \in \mathcal{U}_Q} \mu^T x - \lambda x^T Q x \right\}. \quad (8)$$

The minimax problem in (7) was formulated by Goldfarb and Iyengar [10]. The max-min problem in (8) was formulated in Halldórsson and Tütüncü [2] where a saddle-point representation and a solution algorithm were also provided. We discuss their algorithm in the Appendix.

We have argued above that the problems (1) and (2) are equivalent. There is a similar equivalence between the robust optimization problems (7) and (8) as well:

Proposition 1 *Let $x^*(\lambda)$ denote an optimal solution of (8) for a given positive value of λ . Then, $x^*(\lambda)$ is also an optimal solution of (7) for*

$$R = \min_{\mu \in \mathcal{U}_\mu} \mu^T x^*(\lambda). \quad (9)$$

Proof:

Since objective function of the inner minimization problem in (8) is separable, this problem can be solved by solving the two smaller problems $\min_{\mu \in \mathcal{U}_\mu} \mu^T x$ and $\max_{Q \in \mathcal{U}_Q} x^T Q x$. Let $\mu^*(\lambda)$ and $Q^*(\lambda)$ be the optimal solutions to these problems corresponding to $x^*(\lambda)$.

Now, if $x^*(\lambda)$ is not optimal for (7) with R as in (9), there must exist $\hat{x} \in \mathcal{X}$ and $\hat{Q} \in \mathcal{U}_Q$ satisfying the following:

$$\begin{aligned} \min_{\mu \in \mathcal{U}_\mu} \mu^T \hat{x} &\geq R = \min_{\mu \in \mathcal{U}_\mu} \mu^T x^*(\lambda) = \mu^*(\lambda)^T x^*(\lambda), \\ \max_{Q \in \mathcal{U}_Q} \hat{x}^T Q \hat{x} = \hat{x}^T \hat{Q} \hat{x} &< x^*(\lambda)^T Q^*(\lambda) x^*(\lambda). \end{aligned}$$

But this would imply that

$$\min_{(\mu, Q) \in \mathcal{U}} \left(\mu^T \hat{x} - \lambda \hat{x}^T Q \hat{x} \right) = \min_{\mu \in \mathcal{U}_\mu} \mu^T \hat{x} - \lambda \hat{x}^T \hat{Q} \hat{x} > \mu^*(\lambda)^T x^*(\lambda) - \lambda x^*(\lambda)^T Q^*(\lambda) x^*(\lambda)$$

contradicting the optimality of $x^*(\lambda)$ for (8). Thus, $x^*(\lambda)$ must be optimal for (7) as well when $R = \min_{\mu \in \mathcal{U}_\mu} \mu^T x^*(\lambda)$. \square

2.2 Robust Portfolios with the Maximum Sharpe Ratio

There are assets such as US Treasury Bills that can be considered essentially riskless for investment or borrowing purposes. If such an asset is included in the asset pool for optimal portfolio selection, all efficient portfolios turn out to be linear combinations of this riskless asset and a unique *optimal risky portfolio*. This optimal risky portfolio is the efficient portfolio with the highest Sharpe ratio, which is defined as follows for a portfolio x :

$$h(x) = \frac{\mu^T x - r_f}{\sqrt{x^T Q x}}. \quad (10)$$

Here, r_f represents the known return on the riskless asset. This quantity represents the reward to variability (measured by the standard deviation) ratio of a zero-investment portfolio funded by riskless borrowing. We assume that there are feasible portfolios with positive Sharpe ratios—otherwise, the riskless asset is the only efficient portfolio.

In this section, we assume that Q is positive definite so that $x^T Q x > 0$ and, therefore, $h(x)$ is defined for all non-zero portfolios. Since the covariance matrix Q is always positive semidefinite, this assumption is equivalent to assuming that Q is nonsingular and is essentially equivalent to assuming that there are no redundant assets in our collection.

The portfolio with the highest Sharpe ratio is obtained by solving the optimization problem

$$\max h(x) \text{ s.t. } x \in \mathcal{X}. \quad (11)$$

This maximization problem, however, has a nonlinear and non-concave objective function and may be difficult to solve directly. In what follows, we generalize an approach introduced by Goldfarb and Iyengar [10], to reduce this problem into a convex minimization problem, which is easier to solve. Then, we will present the robust formulation corresponding to this problem.

The elegant argument in [10] is based on the simple observation that $e^T x = 1$ whenever $x \in \mathcal{X}$ (e represents an n -dimensional vector of 1's) since proportions in

all securities must sum to 1, and therefore, $h(x)$ can be rewritten as a homogeneous function of x —we call this function $g(x)$:

$$h(x) = \frac{\mu^T x - r_f}{\sqrt{x^T Q x}} = \frac{(\mu - r_f e)^T x}{\sqrt{x^T Q x}} =: g(x) = g\left(\frac{x}{\kappa}\right), \quad \forall \kappa > 0.$$

The vector $\mu - r_f e$ is the vector of returns in excess of the risk-free lending rate. Goldfarb and Iyengar demonstrate that when \mathcal{X} has the form in (3), using the argument above, one can replace the normalization constraint $e^T x = 1$ with the alternative normalization constraint $(\mu - r_f e)^T x = 1$ without affecting the optimal solution. But then, the objective function is equivalent to minimizing $x^T Q x$, a strictly convex quadratic function of x (recall our assumption that Q is a positive definite matrix).

We show that, a similar reduction can be achieved even when \mathcal{X} is not in the form in (3), as long as $x \in \mathcal{X}$ implies that $e^T x = 1$. To achieve the desired reduction, we first homogenize \mathcal{X} applying the *lifting* technique to it, i.e., we consider a set \mathcal{X}^+ that lives in a one higher dimensional space than \mathcal{X} and is defined as follows:

$$\mathcal{X}^+ := \{x \in \mathbb{R}^n, \kappa \in \mathbb{R} | \kappa > 0, \frac{x}{\kappa} \in \mathcal{X}\} \cup (0, 0). \quad (12)$$

We add the vector $(0, 0)$ to the set to achieve a closed set. Note that \mathcal{X}^+ is a cone. For example, when \mathcal{X} is a circle, \mathcal{X}^+ resembles an ice-cream cone. When \mathcal{X} is polyhedral, e.g., $\mathcal{X} = \{x | Ax \geq b, Cx = d\}$, we have $\mathcal{X}^+ = \{(x, \kappa) | Ax - b\kappa \geq 0, Cx - d\kappa = 0, \kappa \geq 0\}$. Now, using the observation that $h(x) = g(x), \forall x \in \mathcal{X}$ and that $g(x)$ is homogeneous, we conclude that (11) is equivalent to

$$\max g(x) \text{ s.t. } (x, \kappa) \in \mathcal{X}^+. \quad (13)$$

Again, using the observation that $g(x)$ is homogeneous in x , we see that adding the normalizing constraint $(\mu - r_f e)^T x = 1$ to (13) does not affect the optimal solution—among a ray of optimal solutions, we will find the one on the normalizing hyperplane. Note that for any $x \in \mathcal{X}$ with $(\mu - r_f e)^T x > 0$, the normalizing hyperplane will intersect with an $(x^+, \kappa^+) \in \mathcal{X}^+$ such that $x = x^+ / \kappa^+$. In fact, $x^+ = \frac{x}{(\mu - r_f e)^T x}$ and $\kappa^+ = \frac{1}{(\mu - r_f e)^T x}$. The normalizing hyperplane will miss the rays corresponding to points in \mathcal{X} with $(\mu - r_f e)^T x \leq 0$, but since they can not be optimal, this will not affect the optimal solution. Therefore, substituting $(\mu - r_f e)^T x = 1$ into $g(x)$ we obtain the following equivalent problem:

$$\max \frac{1}{\sqrt{x^T Q x}} \text{ s.t. } (x, \kappa) \in \mathcal{X}^+, (\mu - r_f e)^T x = 1. \quad (14)$$

Thus, we proved the following result:

Proposition 2 *Given a set \mathcal{X} of feasible portfolios with the property that $e^T x = 1, \forall x \in \mathcal{X}$, the portfolio x^* with the maximum Sharpe ratio in this set can be found by solving the following problem with a convex quadratic objective function*

$$\min x^T Q x \text{ s.t. } (x, \kappa) \in \mathcal{X}^+, (\mu - r_f e)^T x = 1, \quad (15)$$

with \mathcal{X}^+ as in (12). If $(\hat{x}, \hat{\kappa})$ is the solution to (15), then $x^* = \frac{\hat{x}}{\hat{\kappa}}$. \square

As in [10], we observe that the normalizing constraint in (15) can be relaxed to $(\mu - r_{fe})^T x \geq 1$ by recognizing that this constraint will always be tight at an optimal solution. The relaxed problem will be in the form (1) and therefore, its robust version can be formulated as follows:

$$\begin{aligned} \min \quad & \{\max_{Q \in \mathcal{U}_Q} x^T Q x\} \\ \text{s.t.} \quad & (x, \kappa) \in \mathcal{X}^+ \\ & \min_{\mu \in \mathcal{U}_\mu} (\mu - r_{fe})^T x \geq 1. \end{aligned} \tag{16}$$

3 Finding Robust Portfolios

In this section, we discuss methods for solving the robust formulation (8) of the MVO problem presented in the introduction. First, we discuss a special case of the robust optimization formulation that can be solved as a standard QP problem and therefore, does not require the development of any new solution techniques:

3.1 The Simple Case

In most asset allocation problems, short sales are not allowed. Money managers often look for a nonnegative portfolio of mutual funds representing different asset classes. If there are no additional considerations for the asset allocation problem, the feasible set of portfolios has precisely the description given in (3):

$$\mathcal{X} = \{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1, x \geq 0\}.$$

Above, $x \geq 0$ represent the “no-short-sales” constraint and the restriction $\sum_{i=1}^n x_i = 1$ is necessary to ensure that all the money available for investment is allocated.

Now, we consider an uncertainty set \mathcal{U} of the form (6) with the property that the matrix Q^U is positive semidefinite. In this case, we have the following result that simplifies the search for robust portfolios:

Proposition 3 *Let $x \in \mathbb{R}^n$ be a nonnegative vector and let \mathcal{U} be as in (4)–(6) with a positive semidefinite matrix Q^U . Then, an optimal solution of the problem*

$$\min_{(\mu, Q) \in \mathcal{U}} \mu^T x - \lambda x^T Q x \tag{17}$$

is $\mu^ = \mu^L$ and $Q^* = Q^U$ regardless of the values of the nonnegative scalar λ and the vector x .*

Proof:

First note that both the objective function $\mu^T x - \lambda x^T Q x$ and the constraint set \mathcal{U} are separable in μ and Q . Therefore, as in the proof of Proposition 1, the problem (17) can be solved by solving the following two smaller problems:

$$\begin{aligned} \min \quad & \mu^T x \\ \text{s.t.} \quad & \mu^L \leq \mu \leq \mu^U, \end{aligned} \quad \begin{aligned} \max \quad & x^T Q x \\ \text{s.t.} \quad & Q^L \leq Q \leq Q^U \\ & Q \succeq 0 \end{aligned} . \tag{18}$$

Since $x \geq 0$, the objective value of the first problem in (18) is minimized when each element of the vector μ is at its lower bound, i.e., when $\mu = \mu^L$. Consider the relaxation of the second problem obtained by ignoring the positive semi-definiteness constraint $Q \succeq 0$. For this relaxation, since $x_i x_j \geq 0$ for all i and j , $x^T Q x = \sum_{i,j} q_{ij} x_i x_j$ will be maximized when all q_{ij} attain their largest feasible values, i.e., when $Q = Q^U$. Since Q^U is assumed to be a positive semidefinite matrix, it must be optimal for the unrelaxed problem as well. \square

The proposition above indicates that when short sales are not allowed and when the upper bounds on the covariance matrix form an acceptable covariance matrix, then the worst-case realization of the data is the same regardless of what portfolio is chosen—expected returns are realized at their lowest possible values and the covariances are realized at their highest possible values. In this scenario, the maximin problem given in (8) reduces to the following maximization problem:

$$\max_{x \in \mathcal{X}} (\mu^L)^T x - \lambda x^T Q^U x. \quad (19)$$

This is a standard MVO problem and the associated efficient frontier can be determined using methods such as Markowitz' method of critical lines [4]. A similar argument shows that the minimax problem (7) reduces to the following minimization problem when Q^U is positive semidefinite:

$$\begin{aligned} \min \quad & x^T Q^U x \\ \text{s.t.} \quad & x \in \mathcal{X} \\ & (\mu^L)^T x \geq R. \end{aligned} \quad (20)$$

3.2 The General Case

In the previous subsection we saw that the robust asset allocation problem can be reduced to a simple MVO problem under certain assumptions. When these assumptions do not hold, the worst-case realization of the uncertain inputs is no longer the same for all possible portfolios. This means that we cannot expect to solve the robust asset allocation problem in a sequential manner, i.e., by first finding the worst-case input data and then finding the best allocation with this data, as we did in the previous subsection. Fortunately, the robust problem can still be solved using a saddle-point formulation as we describe below. The arguments below follow the construction in Halldórsson and Tütüncü [2]:

First we need to introduce some notation to represent the objective function of the robust optimization problem (8). Let

$$\psi_\lambda(x, \mu, Q) := \mu^T x - \lambda x^T Q x, \quad x \in \mathcal{X}, (\mu, Q) \in \mathcal{U}. \quad (21)$$

For fixed $(\mu, Q) \in \mathcal{U}$ and a given $\lambda \geq 0$, the function ψ_λ is a concave quadratic function of x . Similarly, for fixed x and a given λ , the function ψ_λ is a linear function of μ and Q . This last fact follows from the observation that $x^T Q x = \sum_{i,j} (x_i x_j) q_{ij}$.

Combining these observations with the assumptions that the sets \mathcal{X} and \mathcal{U} are nonempty and bounded, and Lemma 2.3 from [2], we have the following conclusion: Optimal values of the following pair of primal and dual problems,

$$\max_{x \in \mathcal{X}} \left\{ \min_{(\mu, Q) \in \mathcal{U}} \psi_\lambda(x, \mu, Q) \right\}, \quad \text{and} \quad \min_{(\mu, Q) \in \mathcal{U}} \left\{ \max_{x \in \mathcal{X}} \psi_\lambda(x, \mu, Q) \right\} \quad (22)$$

are equal and are obtained at a saddle-point of the function $\psi_\lambda(x, \mu, Q)$. In other words, there exists a vector $\bar{x} \in \mathcal{X}$ and a vector-matrix pair $(\bar{\mu}, \bar{Q}) \in \mathcal{U}$ such that

$$\psi_\lambda(x, \bar{\mu}, \bar{Q}) \leq \psi_\lambda(\bar{x}, \bar{\mu}, \bar{Q}) \leq \psi_\lambda(\bar{x}, \mu, Q), \quad \forall x \in \mathcal{X}, (\mu, Q) \in \mathcal{U}, \quad (23)$$

and $\bar{x} \in \mathcal{X}$, $(\bar{\mu}, \bar{Q}) \in \mathcal{U}$ collectively solve both problems in (22).

Therefore, the maximin problem (8) we formulated for the robust solution of the asset allocation problem (2) is equivalent to finding a saddle-point of the function $\psi_\lambda(x, \mu, Q)$. This equivalence is significant because we can now use the rich literature on saddle-point problems and in particular, the work of Halldórsson and Tütüncü [2] where an interior-point algorithm for their solution is developed.

In a similar manner, we can develop a saddle-point formulation for the minimax problem (7). First, we note that the constraint $\min_{\mu \in \mathcal{U}_\mu} \mu^T x \geq R$ can be simplified into $(\mu^L)^T x \geq R$ when \mathcal{U}_μ is given by (4) and $x \in \mathcal{X}$ implies that $x \geq 0$, which is natural for asset allocation problems. Now, defining $\phi(x, Q) = x^T Q x$ and $\mathcal{X}_R := \{x \in \mathcal{X} : (\mu^L)^T x \geq R\}$, we obtain the following saddle-point formulation for (7): Find $\bar{x} \in \mathcal{X}_R$ and $\bar{Q} \in \mathcal{U}_Q$ such that

$$\phi(\bar{x}, Q) \leq \phi(\bar{x}, \bar{Q}) \leq \phi(x, \bar{Q}), \quad \forall x \in \mathcal{X}_R, Q \in \mathcal{U}_Q. \quad (24)$$

One can deduce conditions characterizing saddle-points of the function $\phi(x, Q)$ using the optimality of \bar{x} for $\min_{x \in \mathcal{X}_R} \phi(x, \bar{Q})$ and of \bar{Q} for $\max_{Q \in \mathcal{U}_Q} \phi(\bar{x}, Q)$, which forms the basis of the algorithm we present next.

3.3 The Algorithm

Given μ , Q , and \mathcal{X} , the efficient frontier is the collection of portfolios that are optimal solutions to the problem (1) for all possible values of R (or to problem (2) for all possible values of λ). By representing each optimal portfolio as a point in two-dimensional space with coordinates equal to the standard deviation and expected return of the portfolio, one obtains the familiar depiction of this efficient frontier.

Since there are potentially infinitely many points on the efficient frontier, obtaining a precise description of this set may appear impossible. Markowitz' *method of critical lines* [4], which is essentially a parametric quadratic programming algorithm, recognizes that there are a finite number of critical values of R in (1) with corresponding critical optimal portfolios. As R varies between two of its consecutive critical values, say R_k and R_{k+1} , the optimal solutions to corresponding problems (1) can be obtained as a convex combination of the two critical portfolios corresponding to R_k and R_{k+1} . Therefore, the infinite number of points on the efficient frontier can be generated by a finite algorithm, such as the method of critical lines.

In the robust case, however, such methods are either not available or not implemented. While there are algorithms to solve the robust portfolio selection/asset allocation problems [2, 10], these algorithms find a single point on the robust efficient frontier. The literature on parametric min-max problems or parametric saddle-point problems appears to be surprisingly sparse—we are aware of only one reference, and there the emphasis is on continuity properties of the solutions of such problems rather than algorithms [5]. Therefore, to generate the efficient frontier, we will first determine the robust efficient portfolios with the lowest and highest expected returns, discretize the range between these two extremes to obtain a finite number of set levels of the expected return and solve the problem (7) for each set level of the expected return. As we discussed at the end of Section 3.2, assuming that $x \in \mathcal{X}$, the min-max problem (7) is equivalent to the saddle-point problem (24).

To obtain the robust efficient portfolios with highest and lowest expected returns as well as to solve the problem (7) (or, equivalently, the problem (24)) for each intermediate value we use the saddle-point algorithm (SP Algorithm, for short) developed by Halldórsson and Tütüncü [2]. This is an interior-point path-following method with computationally attractive polynomial-time convergence guarantees. We include a description of this method in the Appendix of this article.

We are now ready to formally present the algorithm that generates a discrete approximation to the robust efficient frontier:

Robust Efficient Frontier Algorithm

1. Solve problem (7) without the expected return constraint using the SP Algorithm. Let x_{\min} denote its optimal solution. Set $R_{\min} = (\mu^L)^T x_{\min}$.
2. Solve problem (8) with $\lambda = 0$. Let x_{\max} denote its optimal solution. Set $R_{\max} = (\mu^L)^T x_{\max}$ and $\Delta = R_{\max} - R_{\min}$.
3. Choose K , the number of desired points on the efficient frontier. For $R \in \{R_{\min} + \frac{\Delta}{K-1}, R_{\min} + 2\frac{\Delta}{K-1}, \dots, R_{\min} + (K-2)\frac{\Delta}{K-1}\}$ solve problem (7) with the expected return constraint using the SP algorithm.

As we mentioned above, in Steps 1 and 3 of the Robust Efficient Frontier Algorithm, we actually solve the saddle-point problem (24) that is equivalent to the corresponding min-max problem (7). Step 1 generates the robust efficient portfolio that has the overall minimum worst-case variance, without any regard for the expected returns. The worst-case expected return of this portfolio is the the minimum expected return for a robust efficient portfolio since no portfolio with a smaller return (and necessarily higher variance) would be efficient.

In contrast, Step 2 finds the robust efficient portfolio with the highest (worst-case) expected return without giving any consideration to the variance of the portfolio. This turns out to be a trivial problem. We want to solve the following maximin problem:

$$\max_{x \in \mathcal{X}} \{ \min_{\mu \in \mathcal{U}_\mu} \mu^T x \}. \quad (25)$$

When \mathcal{U}_μ is as in (4) and $x \geq 0$ when $x \in \mathcal{X}$, this problem is equivalent to the following LP and can be solved easily:

$$\max_{x \in \mathcal{X}} (\mu^L)^T x. \quad (26)$$

Once we determined the robust efficient portfolios with the lowest and highest worst-case expected returns, in Step 3, we generate a sequence of robust efficient portfolios whose expected returns lie in between these two extremes with equal increments.

4 Computational Results

In this section, we apply the robust asset allocation methods discussed in the previous sections to market data and compare the behavior of the solutions obtained by the robust optimization technique and the solutions obtained by standard approaches using nominal data.

For our first experiment, we use a universe of 5 asset classes: large cap growth stocks, large cap value stocks, small cap growth stocks, small cap value stocks, and fixed income securities. To represent each asset class, we use a monthly log-return time series of corresponding market indices: Russell 1000 growth and value indices for large cap stocks, Russell 2000 growth and value indices for small cap stocks, and Lehman Brothers US Intermediate Government/Credit Bond index for fixed income securities. Lehman Brothers U.S. Intermediate Government/Credit Bond Index is an unmanaged index generally representative of government and investment-grade corporate securities with maturities of 1-10 years. Our time series data spans the period January 1979 to July 2002, a total of $n = 283$ months.

Using this data, we computed the historical means and covariances of the five indices mentioned above. Further, we computed lower and upper bound vectors and matrices on means and covariances using a bootstrapping strategy. Namely, a time series of length n was chosen for each index by bootstrapping from the available observations and means and covariances were computed for these series. This process was repeated 3000 times and the quantiles of the statistics were computed. Table 1 lists the 2.5, 50, and 97.5 percentiles for means of monthly returns and covariances of these returns.

The ranges for the return and covariance estimates presented in Table 1 can be considered wide. However, most of these ranges are not wide enough to include the negative or very high returns and/or negative correlations often observed in shorter observation periods. By bootstrapping over the entire history of the time series, we obtained bounds that can be considered reliable for average behavior of returns over long investment periods. We will discuss an alternative method to generate lower and upper bounds on returns and covariances below which may be more suitable for robust asset allocation over shorter investment horizons.

Using the data presented above, we generated the classical and robust efficient frontiers. Figure 1 depicts the classical efficient frontier obtained by using the 50 percentile values for expected returns and covariances as inputs as well as the robust efficient frontiers obtained by using the 2.5 and 97.5 percentile values for expected returns and covariances as inputs. Note that the robust efficient frontier lies below the

$10^{-2} \times$	Ru 1000 Gr	Ru 1000 Va	Ru 2000 Gr	Ru 2000 Va	LB IT Gov/Cre
2.5 percentile	0.3398	0.6330	-0.1358	0.5866	0.5868
50 percentile	0.9644	1.1135	0.7221	1.1726	0.7449
97.5 percentile	1.5602	1.5825	1.5497	1.7145	0.9029

$10^{-3} \times$	Ru 1000 Gr	Ru 1000 Va	Ru 2000 Gr	Ru 2000 Va	LB IT Gov/Cre
Ru 1000 Gr	2.2147 2.8891 3.6629				
Ru 1000 Va	1.3493 1.8417 2.4820	1.3060 1.7427 2.3011			
Ru 2000 Gr	2.3928 3.2870 4.3749	1.4138 2.1361 3.0965	3.8449 5.1551 6.7911		
Ru 2000 Va	1.2949 1.9204 2.7833	1.1212 1.6879 2.4465	2.1245 3.0847 4.4034	1.6247 2.4182 3.5308	
LB IT Gov/Cre	0.0477 0.1346 0.2162	0.0628 0.1441 0.2224	-0.0332 0.0859 0.1950	0.0152 0.1158 0.2116	0.1337 0.1848 0.2500

Table 1: 2.5, 50, and 97.5 percentiles of mean monthly log-returns as well as the entries of the covariance matrix obtained from bootstrapped samples. Only the lower diagonal entries in the covariance matrix are listed for brevity.

classical efficient frontier since it depicts the *worst-case* values of the expected returns and standard deviations. In contrast, the classical efficient frontier shows nominal values of these quantities. Therefore, a direct comparison of the classical and robust efficient portfolios should not be made based on this figure. They are presented together to show how different the efficient frontier may look if optimization inputs are inexact/incorrect.




Figure 1: The classical efficient frontier found using nominal data and the robust efficient frontier found using worst-case data.

We mentioned that we used the 2.5 and 97.5 percentiles listed in Table 1 as the

lower and upper bounds μ^L , μ^U , Q^L , and Q^U when computing the robust efficient portfolios. We observe that Q^U obtained in this manner is a positive definite matrix. Therefore, using the result of Proposition 2, the robust efficient portfolios were found using the classical mean-variance optimization approach with inputs μ^L and Q^U .

Next, we compare the compositions of the classical and robust efficient portfolios. They are presented in Figure 2. On the classical efficient frontier, the lowest risk efficient portfolios are obtained, as expected, using the fixed income securities. As one moves along the efficient frontier toward the efficient portfolio with the highest expected return, fixed income securities are gradually replaced by a mixture of large-cap and small-cap value stocks. Close to the high-return end of the frontier, large-cap stocks are also phased out and one gets a portfolio consisting entirely of small cap value stocks. In contrast to the classical efficient portfolios, robust efficient portfolios focus almost all equity holdings in large cap value stocks. A mixture of small cap growth and value stocks are used only in very small amounts and only at the low-risk end of the frontier. One might be surprised that the robust efficient portfolios are even less inclusive (i.e., contain fewer asset classes) than the classical efficient portfolios—we will comment on this observation in our conclusion.

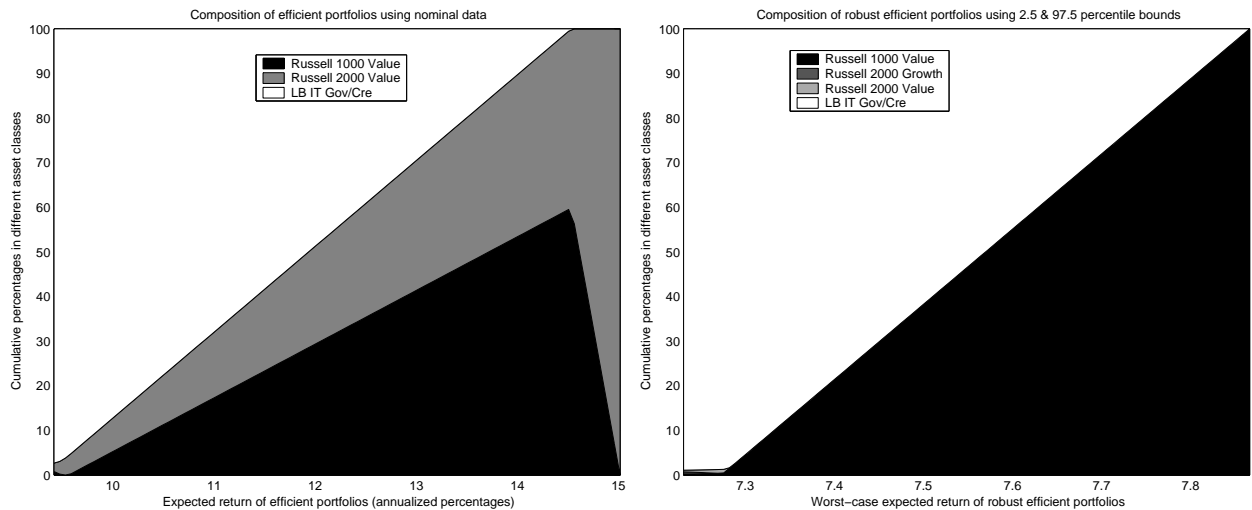


Figure 2: The composition of the classical and robust efficient portfolios. 2.5 and 97.5 percentiles of means and covariances of bootstrapped samples were used to describe the uncertainty intervals for robust portfolios.

We then repeated the robust asset allocation algorithm on the same set of asset classes using a different method to generate the lower and upper bounds for the input parameters. We considered moving windows of four years and computed mean returns and covariances in each such window. Then, we computed different quantiles of the moving window time series. Table 2 lists the 5 and 95 percentiles for means of monthly log-returns and covariances of these returns. Interestingly, Q^U obtained with this method is also positive definite and is the unique worst-case covariance matrix within the interval for all possible feasible portfolios. As can be seen from these tables, this

process generates wider bounds on the returns and covariances than the bootstrapping strategy discussed above. We have the following explanation for this observation. The bootstrapping strategy we used considers (and averages out) 23 years of return data and therefore is not able to identify significantly volatile nature of returns over short and mid-term horizons. Moving averages with two-to-four year windows detect such behavior and result in wider interval estimates for the returns and covariances.

$10^{-2} \times$	Ru 1000 Gr	Ru 1000 Va	Ru 2000 Gr	Ru 2000 Va	LB IT Gov/Cre
5 percentile	0.3093	0.5333	-0.0666	0.2560	0.4570
95 percentile	2.1511	1.9658	1.7425	2.0545	1.2664

$10^{-3} \times$	Ru 1000 Gr	Ru 1000 Va	Ru 2000 Gr	Ru 2000 Va	LB IT Gov/Cre
Ru 1000 Gr	0.8309/5.2834				
Ru 1000 Va	0.5621/3.1729	0.6309/2.8234			
Ru 2000 Va	0.9239/5.4941	0.5957/3.8920	1.7901/8.9472		
Ru 2000 Gr	0.5265/3.7516	0.4963/3.2523	0.9815/5.2682	0.7193/4.4271	
LB IT Gov/Cre	-0.0555/0.3909	-0.0313/0.3582	-0.0737/0.3976	-0.0307/0.3819	0.0582/0.5433

Table 2: Percentiles of 4-year moving averages of monthly log-returns and covariances of monthly log-returns.

Figure 3 illustrates the robust efficient frontier and the composition of robust efficient portfolios when we use the percentiles given in Table 2 as the uncertainty bounds. Efficient portfolios are very similar to those we found with the previous set of bounds, except that some of the low-risk efficient portfolios of the previous set are inefficient for the problem with the current bounds.




Figure 3: The efficient frontier and the composition of the efficient portfolios found using the robust asset allocation approach. 5 and 95 percentiles of 4-year moving averages means and covariances were used to describe the uncertainty intervals for these inputs.

We conclude this experiment by comparing the (σ, μ) frontiers for classical and robust efficient portfolios under two scenarios. Since robust efficient portfolios generated

with two different sets of bounds were very similar, we just focus on the first set and ignore the other in this comparison. First, we plot the standard deviation-expected return profiles of the generated portfolios assuming that the expected returns and covariances were actually equal to the point estimates used for the classical mean-variance optimization approach. Under this scenario, portfolios coming from the classical MVO approach only slightly outperform the portfolios generated with the worst-case in mind. Next, on the same graph we plot the standard deviation-expected return profiles of the portfolios generated assuming that actual expected returns and covariances were the worst-case values within the lower and upper bounds used for robust optimization. Figure 4 shows these plots.

Compared to the nominal case, the difference in the worst-case performances of the two sets of efficient portfolios is greater. We observe that the performance of classical efficient portfolios deteriorate significantly at the high-return end with worst-case inputs. Furthermore, Figure 4 suggests that one of the most significant benefits of the robust approach is in risk reduction for worst-case scenarios. This is a consequence of the (σ, μ) frontiers being relatively flat with worst-case inputs. For example, the robust efficient portfolio achieving a 7.5% worst-case annualized expected return has an 8% standard deviation, while the classical efficient portfolio with a 7.5% worst-case annualized expected return has approximately 12% standard deviation indicating that it is significantly riskier.




Figure 4: (σ, μ) -profiles of classical and robust efficient portfolios when actual moments are (i) equal to their point estimates, (ii) equal to their worst possible values within given bounds.

In our second experiment, we used a wider set of asset classes: growth and value stocks in large-cap, mid-cap, and small-cap categories, intermediate term fixed-income securities, international stocks, real estate securities, and high-yield corporate bonds. To represent the domestic equity classes, we used the log-return time series for Wilshire

Target indices corresponding to each category. We use the Lehman Brothers Intermediate Government/Credit index, MSCI EAFE (Europe, Australasia, Far East) index, Wilshire Real Estate Securities index, and Lehman Brothers High-Yield Bond index for the remaining categories. Our time series data spans the period July 1983 to July 2002, a total of $n = 229$ months.

Using this data, we computed the historical means and covariances of the log-returns of the ten indices mentioned above. Further, we computed lower and upper bound vectors and matrices on means and covariances using the 4-year moving averages as was done in the first experiment. Table 3 lists the 5, 50, and 95 quantiles for means of monthly log-returns and covariances of these returns.

$10^{-2} \times$	Ws LC Gr	Ws LC Va	Ws MC Gr	Ws MC Va	Ws SC Gr	Ws SC Va	LB IT G/C	MS EAFE	Ws RE Sec	LB Hi-Yld
5 %	0.5711	0.2586	0.0314	0.3622	0.0352	0.2076	0.4534	-0.2109	-0.3670	0.0644
50 %	1.2486	1.0909	1.0480	1.1863	0.8681	1.1460	0.6251	0.7180	0.6537	0.8603
95 %	2.3590	1.8205	1.5684	1.6851	1.5460	1.7965	0.9702	2.9500	1.2550	1.2857

$10^{-3} \times$	Ws LC Gr	Ws LC Va	Ws MC Gr	Ws MC Va	Ws SC Gr	Ws SC Va	LB IT G/C	MS EAFE	Ws RE Sec	LB Hi-Yld
Ws LC Gr	0.7633 2.2616 4.2952									
Ws LC Va	0.4118 1.5653 2.7668	0.6743 1.5586 2.5193								
Ws MC Gr	0.8451 2.4285 4.3482	0.5159 1.8221 3.0201	1.5223 3.1778 5.4658							
Ws MC Va	0.3776 1.3810 2.7408	0.5221 1.4195 2.4136	0.5781 1.9812 3.3073	0.5698 1.5851 3.1725						
Ws SC Gr	0.8666 2.6279 4.7417	0.5180 1.7503 3.2677	1.6709 3.5112 5.9247	0.5966 2.0355 3.6341	1.9342 4.0516 6.6948					
Ws SC Va	0.3554 1.2184 2.5611	0.4364 1.2428 1.9703	0.5842 1.8745 3.1768	0.4969 1.4211 2.3507	0.6284 2.0104 3.5407	0.4784 1.3795 2.2155				
LB IT G/C	-0.0680 0.1180 0.2240	-0.0616 0.1525 0.2006	-0.0902 0.0623 0.1925	-0.0250 0.1250 0.1801	-0.0584 0.0336 0.1761	-0.0070 0.0977 0.1516	0.0577 0.0924 0.1618			
MS EAFE	0.1637 1.3035 2.5973	0.3085 1.0202 1.4259	0.4374 1.5154 2.7885	0.2542 0.9073 1.3379	0.5027 1.5756 2.6630	0.2275 0.8434 1.3036	-0.0856 0.0585 0.2000	1.2500 2.3682 4.0717		
Ws RE Sec	0.1861 0.7061 2.3900	0.3166 1.0485 1.8004	0.4188 1.3376 3.0029	0.3598 1.2646 2.0014	0.5484 1.3905 3.4410	0.3975 1.1763 1.9695	-0.0213 0.0672 0.0979	0.2433 0.5687 1.3112	0.9526 1.6691 2.3312	
LB Hi-Yield	0.1131 0.5910 0.9148	0.1520 0.4289 0.7006	0.1566 0.6901 1.1368	0.1376 0.4658 0.7611	0.1641 0.7478 1.2796	0.1336 0.4635 0.7498	0.0062 0.0633 0.1519	0.1315 0.3022 0.5425	0.1185 0.3264 0.8800	0.1345 0.3088 0.8082

Table 3: 5, 50, and 95 percentiles of 4-year moving averages of monthly log-returns and covariances of monthly log-returns.

In Figures 5 and 6 we illustrate the efficient frontiers and the composition of the efficient portfolios found using the classical and robust asset allocation approaches¹. The composition graphs illustrate the similarities as well as the stark differences between the choices made by the two approaches: Both approaches invest a substantial proportion of their portfolios in fixed income securities at the low-risk end of their spectrum and gradually phase them out of the portfolios as one moves toward the high

¹50 percentile values of the covariances results in a matrix with a negative eigenvalue. Since this eigenvalue is very small in absolute value, we simply modified the 50 percentile matrix with a small multiple of the identity to obtain a legitimate (positive semidefinite) covariance matrix. This modified matrix was used to determine the classical efficient portfolios.

return end of the frontier. High yield bonds that constitute more than 60% of the mid-range portfolios in the nominal case are too risky for the robust approach and are replaced by government bonds. For the equity portion of the portfolios, the classical approach chooses a mixture of large cap growth stocks as well as mid and small-cap value stocks. In contrast, the robust approach focuses on the large-cap growth stocks and includes small amounts of real estate securities at the low-risk end of the frontier. Once again, we observe that the robust portfolios are less inclusive.




Figure 5: The classical efficient frontier found using nominal data and the robust efficient frontier found using worst-case data. 5 and 95 percentiles of 4-year moving averages for means and covariances were used to describe the “worst-case inputs”.

Next, we examine the effect of the length of the moving window used for the computation of lower and upper bounds describing uncertainty intervals. Instead of a 4-year window, we use 2-year windows and compute the corresponding percentiles. Using 5 and 95 percentiles as the lower and upper bounds leads to uncertainty sets that are too wide to be of practical interest. In fact, with such bounds, the robust efficient frontier degenerates into a very small set centered at the pure fixed-income portfolio. When 10 and 90 percentiles of 2-year moving averages are used, we obtain bounds that are similar to those found with 5 and 95 percentiles of 4-year moving averages and virtually identical robust efficient portfolios. We also used 25 and 75 percentiles of the 2-year moving averages to describe our uncertainty sets. This choice leads to intervals that are often, but not always, tighter than those we found with 5 and 95 percentiles of 4-year moving averages. We list the lower and upper bounds obtained in this manner in Table 4 and illustrate the corresponding robust efficient frontier as well as the composition of robust portfolios in Figure 7. As opposed to the previous experiments, the upper bound matrix Q^U obtained in this manner is not a positive semidefinite matrix, and therefore the robust efficient portfolios were generated using




Figure 6: The composition of the classical and robust efficient portfolios. 5 and 95 percentiles of 4-year moving averages for means and covariances were used to describe the uncertainty intervals for robust portfolios.

the algorithm presented in Section 3.3. The resulting efficient portfolios are different from the previous two sets, and mostly focus on large-cap value stocks for their equity allocations.

$10^{-2} \times$	Ws LC Gr	Ws LC Va	Ws MC Gr	Ws MC Va	Ws SC Gr	Ws SC Va	LB IT G/C	MS EAFE	Ws RE Sec	LB Hi-Yld
25 %	0.5952	0.6220	0.3428	0.5809	0.2408	0.4752	0.5063	0.3186	0.2384	0.3880
75 %	2.0878	1.6486	1.5027	1.7666	1.4305	1.8257	0.8516	1.4490	1.3859	1.1550

$10^{-3} \times$	Ws LC Gr	Ws LC Va	Ws MC Gr	Ws MC Va	Ws SC Gr	Ws SC Va	LB IT G/C	MS EAFE	Ws RE Sec	LB Hi-Yld
Ws LC Gr	1.4421 3.5834									
Ws LC Va	0.6879 2.1060	0.8653 2.5030								
Ws MC Gr	1.4091 3.5458	0.7722 2.6544	1.8239 4.7440							
Ws MC Va	0.6233 2.0468	0.6748 2.3619	0.8672 3.0221	0.6672 2.6682						
Ws SC Gr	1.3957 3.7728	0.6547 2.8021	1.9984 5.0969	0.6913 3.0908	2.3277 5.8410					
Ws SC Va	0.5588 1.8277	0.5527 1.7764	0.8355 2.7839	0.5982 2.1986	0.7996 2.9850	0.5763 2.0594				
LB IT G/C	0.0238 0.1935	0.0465 0.1876	-0.0300 0.1595	0.0514 0.1700	-0.0414 0.1540	0.0389 0.1286	0.0642 0.1203			
MS EAFE	0.3602 2.1599	0.3490 1.5159	0.5775 2.5945	0.3265 1.4203	0.6313 2.5283	0.2883 1.2517	-0.0285 0.1702	1.8514 2.8549		
Ws RE Sec	0.2420 1.3167	0.4558 1.5972	0.5471 1.9951	0.5378 1.8447	0.6385 2.2476	0.4930 1.5504	-0.0025 0.0932	0.2662 1.0387	0.8211 2.0566	
LB Hi-Yield	0.2444 0.7214	0.1786 0.5904	0.2498 0.8910	0.1597 0.5887	0.2525 0.9822	0.1570 0.5634	0.0224 0.0933	0.1147 0.4623	0.1467 0.4919	0.1793 0.4389

Table 4: 25 and 75 percentiles of 2-year moving averages of monthly log-returns and covariances of monthly log-returns.

We now compare the performance of the portfolios generated by the classical and robust asset allocation approaches under three scenarios. In Figures 8 and 9 we plot the

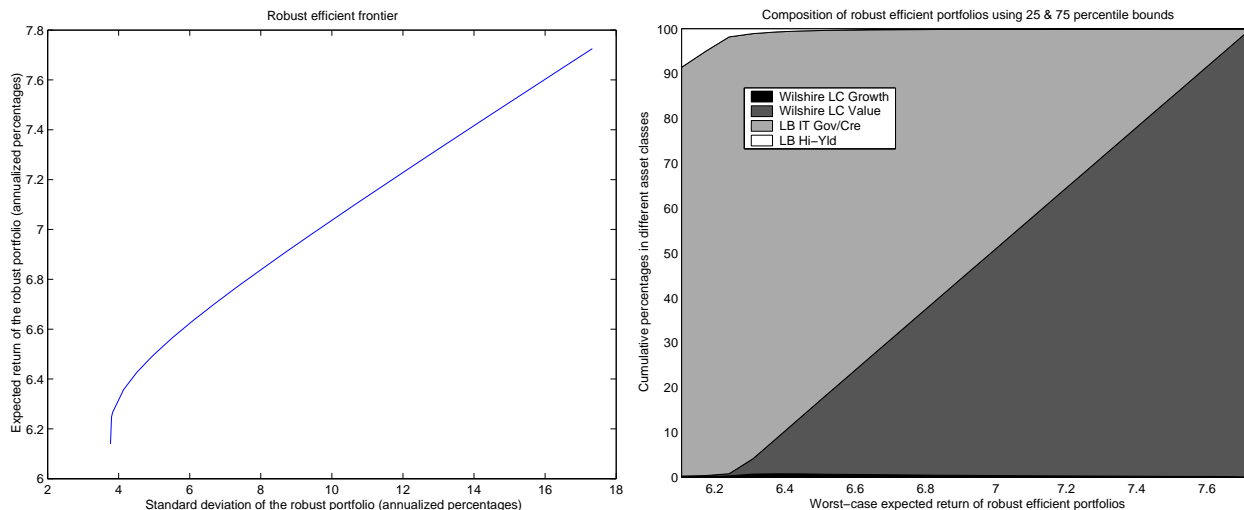


Figure 7: The efficient frontier and the composition of the efficient portfolios found using the robust asset allocation approach. 25 and 75 percentiles of 2-year moving averages for means and covariances were used to describe the uncertainty intervals for these inputs.

standard deviation-expected return profiles of the generated portfolios assuming that the realized returns and covariances were actually equal to the point estimates used for the classical mean-variance optimization approach. Under this scenario, portfolios coming from the classical MVO approach outperform the portfolios generated with the worst-case in mind.

In Figures 8 and 9, we also plot the standard deviation-expected return profiles of the generated portfolios assuming that the realized returns and covariances were among the worst possible from the uncertainty sets described in Tables 3 and 4, respectively. Note that the worst-case realization of the data depends on the particular portfolio one has (except when Q^U is positive semidefinite, see Proposition 2) and therefore, there is no unique realization of the uncertain parameters that is uniformly bad for all portfolios that we could use for these graphs. For the graphs in Figures 8 and 9, we used the worst case realization of the data for the robust minimum variance portfolio. As is clear from the graph, portfolios generated to perform optimally in the worst-case scenario, significantly outperform the portfolios generated using the classical MVO approach without considering data uncertainty. Portfolios generated using bounds derived from 4-year moving windows perform reasonably well under the two scenarios favoring the other portfolio sets and they perform much better than other portfolio sets under the scenario it was chosen to optimize over. As such, these portfolios provide a sensible alternative to classical efficient portfolios that suffer greatly when the input estimates are inaccurate.

As a final exercise, we test the stability of the efficient portfolios generated by the classical and robust approaches through time. For this purpose, we first consider the portion of our data that was available at the end of 1997. With this data, using the approaches outlined above, we generate mean and covariance estimates for classical




Figure 8: (σ, μ) -profiles of classical and robust efficient portfolios when actual moments are (i) equal to their point estimates, (ii) equal to their worst possible values within bounds found using 4-year moving windows.




Figure 9: (σ, μ) -profiles of classical and robust efficient portfolios when actual moments are (i) equal to their point estimates, (ii) equal to their worst possible values within bounds found using 2-year moving windows.

mean-variance optimization and use 4-year moving averages to generate lower and upper bounds for robust optimization. In addition to identifying the classical and

robust efficient frontiers, we also compute the portfolios that maximize the Sharpe ratio in the classical and robust senses. Then, we repeat this process with data available through the end of each year between 1998 and 2001. Each additional year of data changes the mean and covariance estimates, and we track the effects of these changes on efficient portfolio compositions. While the compositions of classical efficient portfolios vary significantly from year to year, robust efficient portfolios have remarkably small turnover ratios indicating that investors following such strategies will incur minimal trading costs at rebalancing dates. Figures 10 and 11 illustrate our arguments. The first pair of figures show how the composition of the portfolios that maximize the Sharpe ratio using classical and robust approaches change over the course of 5 years. The second pair of figures depict the composition of the classical and robust efficient portfolios that are at the “middle” of the efficient frontier (representing, approximately 13.5% annual return in the nominal case and 7.7% annual return in the worst-case).

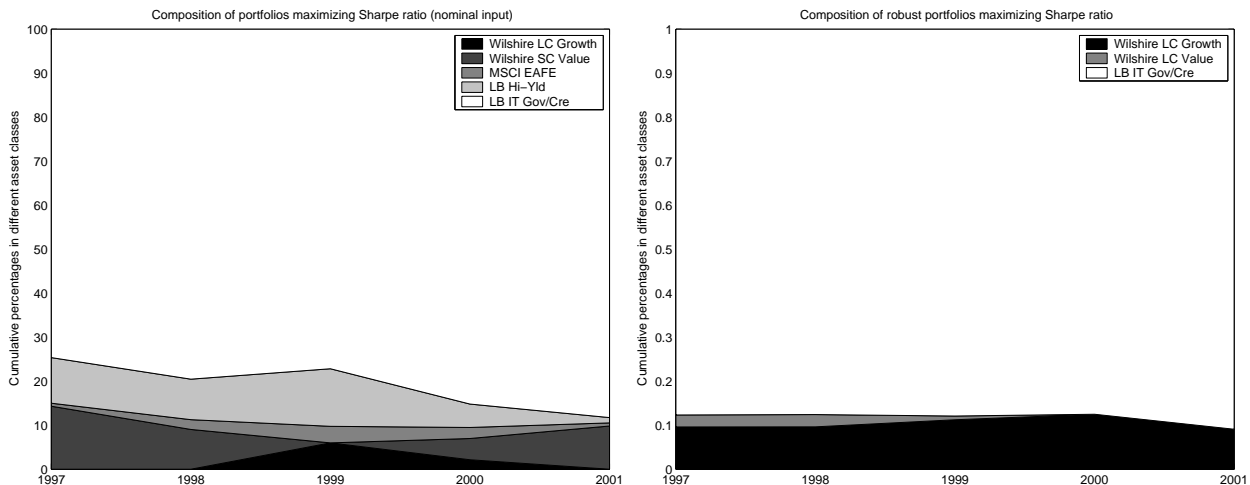


Figure 10: The composition of robust portfolios maximizing the Sharpe ratio remains relatively constant through a 5 year period.

5 Comments and Conclusion

Building on recent research in robust optimization, this article presents a novel approach to asset allocation problems under data uncertainty. As opposed to the classical approach, where one estimates the inputs (returns and covariances) to the asset allocation problem and then treats them as certain and accurate, the approach presented here advocates the description of input estimates in the form of *uncertainty sets*. Robust asset allocation refers to finding an asset allocation strategy whose behavior under the worst possible realizations of the uncertain inputs is optimized. The article presents an algorithm for generating robust efficient portfolios using a variant of an interior-point method for saddle-point problems, and discusses an implementation of this algorithm.




Figure 11: The composition of mid-range robust efficient portfolios remain relatively constant while corresponding classical efficient portfolios experience large turnovers through a 5 year period.

Our analysis demonstrates some key properties of asset mixes formed using robust optimization techniques: (i) significantly better worst-case behavior, (ii) stability over time, and (iii) concentration on a small set of asset classes. While the first two properties are expected, the third one appears surprising. We now comment on all three properties.

The numerical experiments with the algorithm presented in Section 4 suggest that the worst-case behavior of portfolios of different asset classes can be improved significantly using the robust asset allocation approach, often with small performance losses on more likely scenarios. As the size of uncertainty sets increases, both the benefits of robust portfolios under worst-case scenarios and their under-performance under most likely-scenarios appear to increase as well. This trade-off suggests that a cost-benefit analysis of the size of uncertainty sets needs to be performed and the right size will depend on the risk-attitudes of individual investors.

Another property of robust efficient portfolios demonstrated by our results is that they remain relatively unchanged over long periods of time. As illustrated in Figures 10 and 11, recomputing robust efficient portfolios as new data becomes available generates remarkably small turnover ratios and, therefore, such portfolios would experience only modest trading costs if they are rebalanced regularly. Moreover, since robust optimal portfolios computed at the beginning of a long investment horizon appear to remain optimal or near optimal throughout this horizon, they represent attractive alternatives for buy-and-hold investors.

An observation (and sometimes, complaint) expressed by users of the classical asset allocation methods/software is that the portfolios generated often concentrate on a small number (say, 2 or 3) of the asset classes available for investment. This appears contradictory to the well-known benefits of diversification. We now give a heuristic explanation of this phenomenon that also helps us interpret the behavior of robust asset allocation methods. Trying to negotiate the two competing objectives of maximizing

returns and minimizing risk, the MVO algorithm, in a way, generates two rankings of the assets (or asset classes)—one for high return potentials and one for low risk potentials. Highest return portfolios are obtained by picking one or a few of the assets topping the high-return list. As the algorithm moves from the high return end of the efficient frontier toward the low risk end, these assets are gradually replaced by those that top the low-risk list. Assets dominated by another asset in both lists almost never find a place in an efficient portfolio.

While one may expect the robust efficient portfolios to be “more diversified” or more inclusive than their classical counterparts—this comparison may be hard to make without an appropriate metric—we observed a similar phenomenon to the one described in the previous paragraph in our computational experiments with the robust asset allocation approach. This time, it appears, the method ranks the worst-case return potentials of the assets as well as their worst-case riskiness, and then proceeds in a similar manner as before, from the assets with highest worst-case returns to those with lowest worst-case variance. This behavior is apparent from Table 3 and Figures 5, 6. Instead of the small-cap value stocks (whose 50 percentile returns are the highest) that we see at the high-return end of the classical efficient portfolios, large-cap growth stocks, whose worst-case returns are highest, are seen at the high-return end of the robust efficient portfolio.

Overall, we conclude that by directly addressing some of the weaknesses of classical MVO, robust optimization provides a valuable asset allocation vehicle to conservative investors.

ACKNOWLEDGMENT

The first author would like to thank L. N. Vicente for discussions related to the material in Section 2.2. We also thank three anonymous referees for useful comments and suggestions. The first author was supported in part by NSF under Grant No. CCR-9875559 and Grant No. DMS-0139911.

References

- [1] BEN-TAL, A., and NEMIROVSKI, A., *Robust Solutions of Uncertain Linear Programs*, Operations Research Letters, Vol. 25, pp. 1–13, 1999.
- [2] HALLDORSSON, B. V., and TÜTÜNCÜ, R. H., *An interior-point method for a class of saddle-point problems*, Journal of Optimization Theory and Applications, Vol. 116, No. 3, pp. 559–590, 2003.
- [3] LOBO, M. S., and BOYD, S., *The Worst-Case Risk of a Portfolio*, Informations Systems Laboratory, Stanford University, 1999.
- [4] MARKOWITZ, H., *Portfolio Selection*, Journal of Finance, Vol. 7, pp. 77–91, 1952.
- [5] MORGAN, J., and RAUCCI, R., *Continuity properties of ϵ -solutions for generalized parametric saddle point problems and application to hierarchical games*, J. Math. Anal. Appl. Vol. 211, pp. 30–48, 1997.

- [6] BAWA, V. S., BROWN, S. J., and KLEIN, R. W., *Estimation Risk and Optimal Portfolio Choice*, North-Holland, Amsterdam, Netherlands, 1979.
- [7] JEN, E., *Working Definitions of Robustness*," SFI Robustness Site <http://discuss.santafe.edu/robustness>, RS-2001-009, 2001.
- [8] MICHAUD, R. O., *The Markowitz Optimization Enigma: Is Optimized Optimal?*, Financial Analysts Journal, Vol. 45, pp. 31–42, 1989.
- [9] MICHAUD, R. O., *Efficient Asset Management*, Harvard Business School Press, Boston, Massachusetts, 1998.
- [10] GOLDFARB, D. and IYENGAR, G., *Robust Portfolio Selection Problems*, Technical Report, IEOR Department, Columbia University, 2001.
- [11] NESTEROV, Yu., and NEMIROVSKI, A., *Interior-Point Polynomial Algorithms in Convex Programming*, SIAM, Philadelphia, Pennsylvania, 1994.

6 Appendix

Robust Efficient Frontier Algorithm we presented in Section 3.3 uses a saddle-point algorithm in each intermediate step. This saddle-point algorithm, originally developed in [2], is from the class of interior-point path-following algorithms. Such methods are based on the concept of a *central path* in the relative interior of the feasible set of points that converges to a saddle-point of the given function within this set (or, to an optimal point in the case of optimization problems). Iterates are generated in close proximity to this path (hence the name *path-following*) so that they can be guaranteed to converge to a saddle-point, just like the central path.

Let us describe the central path in more detail. First, we define the following *saddle-barrier function*:

$$\phi_t(x, Q) := t\phi(x, Q) + F(x) - G(Q),$$

where $F(x)$ and $G(Q)$ are self-concordant barrier functions [11] for the relative interiors \mathcal{X}_R^0 and \mathcal{U}_Q^0 of the sets \mathcal{X}_R and \mathcal{U}_Q :

$$F(x) := - \sum_{j=1}^n \log(x_j) - \log(\mu_L^T x - R), \quad x \in \mathcal{X}_R^0$$

$$G(Q) := - \sum_{1 \leq i \leq j \leq n} \log(Q_{ij}^U - Q_{ij}) - \sum_{1 \leq i \leq j \leq n} \log(Q_{ij} - Q_{ij}^L) - \log \det(Q), \quad Q \in \mathcal{U}_Q^0.$$

As long as \mathcal{X}_R^0 and \mathcal{U}_Q^0 are nonempty, there is a unique saddle-point of the function ϕ_t for each $t \geq 0$. These saddle-points form a continuous trajectory in the interior of the set of feasible points called the *central path* of the saddle-point problem (24). As t tends to ∞ , saddle-points of the functions ϕ_t , and therefore, the central path converge to a saddle-point of the problem (24). Our algorithm generates iterates that are in close proximity to the saddle-points of functions ϕ_t for increasing values of t and converge to a solution using this strategy.

The algorithm we present below starts from a trial point (x_0, Q_0) close to a point on the central path corresponding $t = t_0$ and using updates generated by Newton's method, determines iterates that are close to points on the central path corresponding to exponentially increasing values of t . Proximity of points to the central path is measured by the following function, called the *Newton decrement*:

$$\begin{aligned}\eta(\phi_t, x, Q) &:= \sqrt{\eta^2(\phi_Q^t, x) + \eta^2(-\phi_x^t, Q)}, \text{ with,} \\ \eta(\psi, z) &:= \sqrt{\nabla\psi^T(z)[\nabla^2\psi(z)]^{-1}\nabla\psi(z)}, \\ \phi_Q^t(x) &:= t\phi(x, Q) + F(x), \text{ and} \\ \phi_x^t(Q) &:= t\phi(x, Q) - G(Q).\end{aligned}$$

The measure $\eta(\phi_t, x, Q)$ is always nonnegative and is zero only on central points. Furthermore, a small value of the Newton decrement function implies close proximity to the central path. By a judicious choice of the parameters, the algorithm makes sure that the Newton decrement is small for all iterates generated, guaranteeing eventual convergence. We are now ready to present the saddle-point algorithm formally:

Saddle-Point Algorithm (SP Algorithm)[2]

1. *Initialization:*

Choose $\alpha > 0$ and $\beta > 0$. Find a $t_0 > 0$ and $(x_0, Q_0) \in \mathcal{X}_R^0 \times \mathcal{U}_Q^0$ that satisfies $\eta(\phi_{t_0}, x_0, Q_0) \leq \beta$. Set $k = 0$.

2. *Iteration:*

while $t_k < M$

Set

$$t_{k+1} = (1 + \alpha)t_k. \quad (27)$$

Take a full Newton step:

$$(x_{k+1}, Q_{k+1}) = (x_k, Q_k) - \left[\nabla^2 \phi_{t_{k+1}}(x_k, Q_k) \right]^{-1} \nabla \phi_{t_{k+1}}(x_k, Q_k). \quad (28)$$

Set $k = k + 1$.

end

The parameters α and β need to satisfy certain conditions as prescribed in [2] to ensure that

$$\eta(\phi_{t_k}, x_k, Q_k) \leq \beta$$

for all k . This implies that all iterates are close to the central path. The equation (27) indicates that t_k is growing exponentially. Therefore, it eventually exceeds M , which is a large positive number chosen in a way to ensure that the final iterate is close enough to the actual saddle-point. We leave out the remainder of the implementation details for the sake of brevity. The analysis presented in [2] can be adapted to problem (24) to show the polynomiality of this saddle-point algorithm.