

FRONTIERS OF STOCHASTICALLY NONDOMINATED PORTFOLIOS

ANDRZEJ RUSZCZYŃSKI AND ROBERT J. VANDERBEI

ABSTRACT. We consider the problem of constructing a portfolio of finitely many assets whose returns are described by a discrete joint distribution. We propose mean–risk models which are solvable by linear programming and generate portfolios whose returns are nondominated in the sense of second-order stochastic dominance. Next, we develop a specialized parametric method for recovering the entire mean–risk efficient frontiers of these models and we illustrate its operation on a large data set involving thousands of assets and realizations.

KEYWORDS: Portfolio optimization, stochastic dominance, mean–risk analysis, least absolute deviations, linear programming, parametric simplex method, robust statistics.

1. INTRODUCTION

The problem of optimizing a portfolio of finitely many assets is a classical problem in theoretical and computational finance. Since the seminal work of Markowitz (1952) it is generally agreed that portfolio performance should be measured in two distinct dimensions: the *mean* describing the expected return, and the *risk* which measures the uncertainty of the return. In the mean–risk approach, we select from the universe of all possible portfolios those that are *efficient*: for a given value of the mean they minimize the risk or, equivalently, for a given value of risk they maximize the mean. Such an approach has many advantages: it allows one to formulate the problem as a parametric optimization problem, and it facilitates the trade-off analysis between mean and risk.

Markowitz used the variance of the return as the measure of the risk. It is easy to compute, and it reduces the portfolio selection problem to a parametric quadratic programming problem. One can, however, construct simple counterexamples that show the imperfection of the

variance as the risk measure: it treats over-performance equally as under-performance, and more importantly its use may suggest a portfolio which is always outperformed by another portfolio. The use of the semivariance rather than the variance was already recommended by Markowitz (1959). But even in this case significant deficiencies remain, as we shall explain.

Another theoretical approach to the portfolio selection problem is that of *stochastic dominance* (see Whitmore and Findlay (1978), Levy (1992)). The usual (first order) definition of stochastic dominance gives a partial order in the space of real random variables. More important from the portfolio point of view is the notion of second-order dominance which is also defined as a partial order but which is equivalent to this statement: a random variable Y dominates the random variable Z if $\mathbb{E}[U(Y)] \geq \mathbb{E}[U(Z)]$ for all nondecreasing concave functions $U(\cdot)$ for which these expected values are finite. Thus, no risk-averse decision maker will prefer a portfolio with return Z over a portfolio with return Y . While theoretically attractive, stochastic dominance order is computationally very difficult, as a multi-objective model with a continuum of objectives.

We shall, therefore, concentrate on mean–risk portfolio models, but we shall look for such models whose efficient frontiers consist of stochastically nondominated solutions, at least above a certain modest level of mean return.

The general question of constructing mean–risk models which are in harmony with the stochastic dominance relations has been the subject of the analysis by Ogryczak and Ruszczyński (1999, 2001, 2002). We shall apply and specialize some of the results obtained there to the portfolio optimization problem. We shall show that the resulting mean–risk models can be formulated as linear programming problems. In this sense, our work has a different motivation than the classical paper by Sharpe (1971), who develops a linear programming approximation to the mean–variance model. We do not want to approximate the mean–variance model, but rather to construct a linear programming model that has better theoretical properties than the mean–variance model and its approximations. Moreover, we develop a highly effective algorithm for recovering the entire efficient frontiers of our models. Our numerical results show that our approach is capable of solving portfolio problems of large sizes in a reasonable time. This, combined with the theoretical property of stochastic efficiency of

the solutions obtained constitutes a strong argument for the use of our models in practical portfolio optimization.

2. THE MEAN–RISK PORTFOLIO PROBLEM

Let R_1, R_2, \dots, R_n be random returns of assets $1, 2, \dots, n$. We assume that the returns have a discrete joint distribution with realizations r_{jt} , $t = 1, \dots, T$, $j = 1, \dots, n$, attained with probabilities p_t , $t = 1, 2, \dots, T$. Our aim is to invest our capital in these assets in order to obtain some desirable characteristics of the total return on the investment. Denoting by x_1, x_2, \dots, x_n the fractions of the initial capital invested in assets $1, 2, \dots, n$ we can easily derive the formula for the total return:

$$(1) \quad R(x) = R_1x_1 + R_2x_2 + \dots + R_nx_n.$$

Clearly, the set of possible asset allocations can be defined as follows:

$$X = \{x \in \mathbb{R}^n : x_1 + x_2 + \dots + x_n = 1, x_j \geq 0, j = 1, 2, \dots, n\}.$$

Our analysis will not depend on the detailed way this set is defined; we shall only use the fact that it is a convex polyhedron. So, in some applications one may introduce the possibility of *short positions*, i.e., allow some x_j 's to become negative. One can also limit the exposure to particular assets or their groups, by imposing upper bounds on the x_j 's or on their partial sums. One can also limit the absolute differences between the x_j 's and some reference investments \bar{x}_j , which may represent the existing portfolio, etc. All these modifications define some convex polyhedral feasible sets, and are, therefore, covered by our approach.

With each portfolio allocation x we can associate the mean return

$$\mu(x) = \mathbb{E}[R(x)] = \sum_{j=1}^n \sum_{t=1}^T r_{jt}x_jp_t$$

and some risk measure $\rho(x)$ representing the variability of the return $R(x)$. At this moment we may think of $\rho(x)$ being the variance of the return, although later we shall work with other risk measures which, as we shall argue, are superior to the variance.

The mean–risk portfolio optimization problem is formulated as follows:

$$(2) \quad \max_{x \in X} \mu(x) - \lambda\rho(x)$$

Here, λ is a nonnegative parameter representing our desirable exchange rate of mean for risk. If $\lambda = 0$, the risk has no value and the problem reduces to the problem of maximizing the mean. If $\lambda > 0$ we look for a compromise between the mean and the risk. The results of the mean–risk analysis are usually depicted on a mean–risk graph, as illustrated in Figure 1.

Artzner et al. (1999) introduced the concept of *coherent risk measures* by means of several axioms. In our terminology their measures correspond to composite objectives of the form $-\mu(x) + \lambda\rho(x)$. The risk measures that we discuss in this paper are coherent.

3. CONSISTENCY WITH STOCHASTIC DOMINANCE

The concept of *stochastic dominance* is related to an axiomatic model of risk-averse preferences (Fishburn (1964)). It originated from the theory of majorization (Hardy, Littlewood and Polya (1934), Marshall and Olkin (1979)) for the discrete case, was later extended to general distributions (Quirk and Saposnik (1962), Hadar and Russell (1969), Hanoch and Levy (1969), Rotschild and Stiglitz (1969)), and is now widely used in economics and finance (Bawa (1982), Levy (1992)).

In the stochastic dominance approach, random returns are compared by a point-wise comparison of some performance functions constructed from their distribution functions. For a real random variable V , its first performance function is defined as the right-continuous cumulative distribution function of V : $F_V(\eta) = \mathbb{P}\{V \leq \eta\}$ for $\eta \in \mathbb{R}$. A random return V is said (Lehmann (1955), Quirk and Saposnik (1962)) to *stochastically dominate* another random return S to the first order, denoted $V \succeq_{FSD} S$, if $F_V(\eta) \leq F_S(\eta)$ for all $\eta \in \mathbb{R}$. The second performance function $F^{(2)}$ is given by areas below the distribution function F ,

$$F_V^{(2)}(\eta) = \int_{-\infty}^{\eta} F_V(\xi) d\xi \quad \text{for } \eta \in \mathbb{R},$$

and defines the weak relation of the *second-order stochastic dominance* (SSD). That is, random return V stochastically dominates S to the second order, denoted $V \succeq_{SSD} S$, if $F_V^{(2)}(\eta) \leq F_S^{(2)}(\eta)$ for all $\eta \in \mathbb{R}$ (see Hadar and Russell (1969), Hanoch and Levy (1969)). The corresponding strict dominance relations \succ_{FSD} and \succ_{SSD} are defined in the usual way: $V \succ S$ if and only if $V \succeq S$ and $S \not\preceq V$. For portfolios, the random variables in question are the returns defined by (1). To avoid placing the decision vector, x , in a subscript expression, we shall simply write $F(\eta; x) = F_{R(x)}(\eta)$ and $F^{(2)}(\eta; x) = F_{R(x)}^{(2)}(\eta)$. It will not lead to any

confusion, we believe. Thus, we say that portfolio x *dominates* portfolio y *under the FSD rules*, if $F(\eta; x) \leq F(\eta; y)$ for all $\eta \in \mathbb{R}$, where at least one strict inequality holds. Similarly, we say that x *dominates* y *under the SSD rules* ($R(x) \succ_{SSD} R(y)$), if $F^{(2)}(\eta; x) \leq F^{(2)}(\eta; y)$ for all $\eta \in \mathbb{R}$, with at least one inequality strict.

Stochastic dominance relations are of crucial importance for decision theory. It is known that $R(x) \succeq_{FSD} R(y)$ if and only if $\mathbb{E}[U(R(x))] \geq \mathbb{E}[U(R(y))]$ for any nondecreasing function $U(\cdot)$ for which these expected values are finite. Also, $R(x) \succeq_{SSD} R(y)$ if and only if $\mathbb{E}[U(R(x))] \geq \mathbb{E}[U(R(y))]$ for every nondecreasing and concave $U(\cdot)$ for which these expected values are finite (see, e.g., Levy (1992)).

For a set X of portfolios, a portfolio $x \in X$ is called *SSD-efficient* (or *FSD-efficient*) in X if there is no $y \in X$ such that $R(y) \succ_{SSD} R(x)$ (or $R(y) \succ_{FSD} R(x)$).

We shall focus our attention on the SSD relation, because of its consistency with risk-averse preferences: if $R(x) \succ_{SSD} R(y)$, then portfolio x is preferred to y by all risk-averse decision makers. By changing the order of integration we can express the function $F^{(2)}(\cdot; x)$ as the expected shortfall: for each target value η we have $F^{(2)}(\eta; x) = \mathbb{E}[\max(\eta - R(x), 0)]$. The function $F^{(2)}(\cdot; x)$ is continuous, convex, nonnegative and nondecreasing.

Our main concern is the following: could the mean-risk efficient frontier, as illustrated in Figure 1, contain portfolios which are dominated in the SSD sense? It is unfortunately true for the mean-risk model using the variance as the risk measure (see, e.g., Yitzhaki (1982)).

Following Ogryczak and Ruszczyński (1999), we introduce the following definition.

DEFINITION 1. The mean-risk model (μ, ρ) is *consistent* with SSD with coefficient $\alpha > 0$, if the following relation is true

$$R(x) \succeq_{SSD} R(y) \quad \Rightarrow \quad \mu(x) - \lambda\rho(x) \geq \mu(y) - \lambda\rho(y)$$

for all $0 \leq \lambda \leq \alpha$.

In fact, it is sufficient to have the above inequality satisfied for α ; its validity for all $0 \leq \lambda \leq \alpha$ follows from that.

The concept of consistency turns out to be fruitful. Ogryczak and Ruszczyński (1999) proved the following result.

THEOREM 1. *The mean–risk model in which the risk is defined as the absolute semideviation, $\bar{\delta}(x) = \mathbb{E} \{ \max(\mu(x) - R(x), 0) \}$, is consistent with SSD with coefficient 1.*

An identical result (under the condition of finite second moments) has been obtained by Ogryczak and Ruszczyński (1999) for the *standard semideviation* and later extended in Ogryczak and Ruszczyński (2001) to central semideviations of higher orders and stochastic dominance relations of higher orders (see also Gotoh and Konno (2000)).

Elementary calculations show that for any distribution $\bar{\delta}(x) = \frac{1}{2}\delta(x)$, where $\delta(x)$ is the mean absolute deviation from the mean:

$$(3) \quad \delta(x) = \mathbb{E} |R(x) - \mu(x)|.$$

Thus, $\delta(x)$ is a consistent risk measure with the coefficient $\alpha = \frac{1}{2}$, which provides a theoretical support for the mean–absolute deviation model of Konno and Yamazaki (1991).

Another useful class of risk measures can be obtained by using quantiles of the distribution of the return $R(x)$. Let $q_p(x)$ denote the p -th quantile¹ of the distribution of the return $R(x)$, i.e., $\mathbb{P}[R(x) < q_p(x)] \leq p \leq \mathbb{P}[R(x) \leq q_p(x)]$. We may define the risk measure

$$(4) \quad \rho_p(x) = \mathbb{E} \left[\max \left(\frac{1-p}{p} (q_p(x) - R(x)), R(x) - q_p(x) \right) \right].$$

In the special case of $p = \frac{1}{2}$ the measure above represents the mean absolute deviation from the median. For small p , deviations to the left of the p -th quantile are penalized in a much more severe way than deviations to the right.

Although the p -th quantile $q_p(x)$ might not be uniquely defined, the risk measure $\rho_p(x)$ is a well defined quantity. Indeed, it is the optimal value of a certain optimization problem:

$$(5) \quad \rho_p(x) = \min_z \mathbb{E} \left[\max \left(\frac{1-p}{p} (z - R(x)), R(x) - z \right) \right].$$

It is well known that the optimizing z will be one of the p -th quantiles of $R(x)$ (Bloomfield and Steiger (1983)). Ogryczak and Ruszczyński (2002) proved the following result.

THEOREM 2. *The mean–risk model with the risk defined as $\rho_p(x)$ is consistent with SSD with coefficient 1, for all $p \in (0, 1)$.*

¹In the financial literature, the quantity $-q_p(x)W$, where W is the initial investment, is sometimes called the *Value at Risk*.

Our composite objective, $G(p; x) = \mu(x) - \rho_p(x)$, as a function of p , is the classical *absolute Lorenz curve* (Lorenz (1905)) divided by p . It is the Fenchel dual of the shortfall function $F^{(2)}$. The function $G(p; x)$ can also be interpreted as the Conditional Value at Risk of Uryasev and Rockafellar (2001).

4. LINEAR PROGRAMMING FORMULATIONS

The second major advantage of the risk measures discussed in the previous section, in addition to being consistent with second-order stochastic dominance, is the possibility of formulating the models (2) as linear programming problems, if the underlying distributions are discrete.

Let us start from the risk measure defined as the expected absolute deviation from the mean, as defined in (3). The resulting linear programming model takes on the form:

$$\begin{aligned}
 (6) \quad & \text{maximize} && \sum_{j=1}^n \sum_{t=1}^T r_{jt} p_t x_j - \lambda \sum_{t=1}^T p_t (u_t + v_t) \\
 & \text{subject to} && \sum_{j=1}^n \left(r_{jt} - \sum_{t'=1}^T r_{jt'} p_{t'} \right) x_j = u_t - v_t, \quad t = 1, 2, \dots, T, \\
 & && \sum_{j=1}^n x_j = 1, \quad x_j \geq 0, \quad j = 1, 2, \dots, n, \\
 & && u_t \geq 0, \quad v_t \geq 0, \quad t = 1, 2, \dots, T.
 \end{aligned}$$

in which the decision variables are x_j , $j = 1, \dots, n$, and u_t and v_t , $t = 1, \dots, T$. By elementary argument, at the optimal solution we shall have $u_t v_t = 0$ for all t , and $u_t + v_t$ will be the absolute deviation from the mean in realization t .

It follows from Theorem 1 that for every $\lambda \in (0, \frac{1}{2})$ the set of optimal solutions of this problem contains a portfolio which is nondominated in the SSD sense. So, if the solution is unique, it is nondominated. If it is not unique, there may be another solution y of (6) that dominates it, but it will have exactly the same values of the mean and the absolute deviation: $\mu(y) = \mu(x)$ and $\delta(y) = \delta(x)$.

To prove it, suppose that the optimal portfolio x in the above model is dominated (in the SSD sense) by another portfolio y . It follows from Theorem 1 that $\mu(y) - \beta\delta(y) \geq \mu(x) - \beta\delta(x)$, for all $\beta \in [0, \frac{1}{2}]$. For $\beta = \lambda$ both sides must be equal, because x is the solution

of (6). Since $\lambda \in (0, \frac{1}{2})$, the above inequality must in fact be an equation for all $\beta \in [0, \frac{1}{2}]$, and the result follows.

The model with weighted absolute deviations from the quantile can be formulated in a similar way

$$(7) \quad \begin{aligned} & \text{maximize} && \sum_{j=1}^n \sum_{t=1}^T r_{jt} p_t x_j - \lambda \sum_{t=1}^T p_t \left(u_t + \frac{1-p}{p} v_t \right) \\ & \text{subject to} && \sum_{j=1}^n r_{jt} x_j - z = u_t - v_t, \quad t = 1, 2, \dots, T, \\ & && \sum_{j=1}^n x_j = 1, \quad x_j \geq 0, \quad j = 1, 2, \dots, n, \\ & && u_t \geq 0, \quad v_t \geq 0, \quad t = 1, 2, \dots, T. \end{aligned}$$

in which the decision variables are x_j , $j = 1, \dots, n$, u_t and v_t , $t = 1, \dots, T$, and z . Indeed, when x and z are fixed, the best values of u_t and v_t are the excess and the shortfall of the return with respect to z in realization t . Thus, for a fixed x , the best value of z is the p -th quantile of $R(x)$, as follows from (5). Consequently, the expression $\sum_{t=1}^T p_t (u_t + ((1-p)/p)v_t)$ represents the risk measure $\rho_p(x)$.

Again, Theorem 2 implies that for every $\lambda \in (0, 1)$ the set of optimal solutions of this problem contains a portfolio which is nondominated in the SSD sense. If the solution is unique, it is nondominated. If it is not unique, there may be another solution y of (6) that dominates it, but it will have exactly the same values of the mean, and the average deviation from the p -th quantile: $\mu(y) = \mu(x)$ and $\rho_p(y) = \rho_p(x)$. The argument that supports this is the same as in the previous case.

5. THE PARAMETRIC METHOD FOR CONSTRUCTING THE EFFICIENT FRONTIER

Before defining the parametric simplex method, we start by reminding the reader of some terminology from linear programming. First, a specific choice of values for the variables in a linear programming problem is called a *solution*. It is a *feasible solution* if it satisfies the constraints. Whenever one identifies a set of variables, one for each equality constraint, called *basic* variables, that can be expressed as a linear combination of the remaining, *nonbasic*, variables, the solution obtained by setting the nonbasic variables to zero and reading off the values of the basic variables is called a *basic solution*. If the basic variables satisfy the

nonnegativity constraints of the problem, then the solution is called a *basic feasible solution*. The best basic feasible solution, if it exists, is optimal for the problem.

If $\lambda = 0$, both (6) and (7) have obvious optimal solutions: make a portfolio consisting of just one asset with the highest expected return. In problem (6), the other basic variables are chosen as exactly one from each pair (u_t, v_t) . In problem (7), z is basic and one pair (u_t, v_t) has both variables nonbasic.

Let us focus on problem (6) and denote by $(x^{(0)}, u^{(0)}, v^{(0)})$ the optimal basic solution for $\lambda^{(0)} = 0$. We shall show how to construct a sequence of feasible basic solutions $(x^{(k)}, u^{(k)}, v^{(k)})$, $k = 0, 1, 2, \dots$, such that each of them is optimal for (6) for λ in some interval $[\lambda^{(k)}, \lambda^{(k+1)}]$.

Introducing Lagrange multipliers μ_0 and μ_t , $t = 1, \dots, T$, we can write the Lagrangian:

$$\begin{aligned} L(x, u, v, \mu; \lambda) = & \sum_{j=1}^n \sum_{t=1}^T r_{jt} p_t x_j - \lambda \sum_{t=1}^T p_t (u_t + v_t) + \mu_0 \left(1 - \sum_{j=1}^n x_j \right) \\ & + \sum_{t=1}^T \mu_t \left(\sum_{j=1}^n \left(r_{jt} - \sum_{t'=1}^T r_{jt'} p_{t'} \right) x_j - u_t + v_t \right). \end{aligned}$$

At an optimal basic feasible solution $(x^{(k)}, u^{(k)}, v^{(k)})$ for $\lambda = \lambda^{(k)}$ the values of the Lagrange multipliers $\mu^{(k)}$ are such that the gradient of the Lagrangian with respect to the basic variables is zero. Since λ appears in the objective only, these equations define an affine function $\mu^{(k)}(\lambda)$. Thus the gradient of the Lagrangian, $\nabla_{(x,u,v)} L(x^{(k)}, u^{(k)}, v^{(k)}, \mu^{(k)}(\lambda); \lambda)$, is an affine function of λ , too. The condition that the gradient with respect to the nonbasic variables is non-positive, which is necessary for optimality, can be now used to find the critical value $\lambda^{(k+1)} \geq \lambda^{(k)}$, above which the current basic solution $(x^{(k)}, u^{(k)}, v^{(k)})$ ceases to be optimal. Also, a nonbasic variable will be identified, whose introduction to the set of basic variables will become profitable in terms of the mean-risk objective. A simplex pivot will determine the next basic solution $(x^{(k+1)}, u^{(k+1)}, v^{(k+1)})$. The new basic variable may correspond to an asset, in which case the set of securities in the portfolio will change, or it may be one of the deviations, in which case the proportions within the current portfolio will change. One of the old assets may be removed. The new portfolio will be optimal for some interval $[\lambda^{(k+1)}, \lambda^{(k+2)}]$ of λ . Continuing in this way we can identify the entire efficient frontier by moving from one efficient solution to the next one. The procedure in problem (7) is identical.

The parametric simplex method can be defined for any linear programming problem. Sometimes, as in our portfolio optimization problems, the parameter appears intrinsically in the formulation. At other times, a second linear objective, scaled by a parameter λ , is added artificially. In fact, this is an efficient variant of the simplex method (see, Chapters 7 and 12 in Vanderbei (2001)). Since, as a method for solving for the solution for a fixed large value of λ , it compares favorably with other variants of the simplex method, it can construct the entire efficient frontier with about the same amount of computational effort as is normally required just to find one point on the efficient frontier.

6. NUMERICAL ILLUSTRATION

Using a data file consisting of daily return data for 719 securities from January 1, 1990 to March 18, 2002 ($T = 3080$), we computed the entire efficient frontier using both the deviation-from-mean (Figure 2) and the deviation-from-quantile (Figure 3) risk measures. For the deviation from quantile case, we used the $p = 0.05$ quantile. The clustered marks in each graph represent mean–risk characteristics of individual securities.

Although we don't show it here, we also computed the efficient frontier for the deviation from median risk measure (i.e, $p = 0.5$). Normally, one would expect the deviation-from-median measure to be superior to the deviation-from-mean measure for two reasons: (i) medians are more robust estimators as evidenced by their fundamental role in nonparametric statistics and (ii) the deviation from median provides stochastically nondominated portfolios for $0 \leq \lambda \leq 1$ whereas for deviation from the mean the interval is only $0 \leq \lambda \leq 1/2$. However, the data set at hand, and perhaps it is true for this type of data in general, does not contain significant outliers that would skew the results obtained using the mean measure. Furthermore, most portfolios fall in the range $0 \leq \lambda \leq 1/2$ so not much is lost by stopping here. Hence, these two efficient frontiers, while containing different portfolios, look very similar in the usual risk-reward plots such as shown in Figure 2.

6.1. Deviation from the mean. For the deviation from the mean, the efficient frontier consists of 23446 distinct portfolios. At the risky extreme, the first portfolio consists of putting 100% into the highest return security. Other high-risk portfolios consist of mixtures of only a few high return securities. As λ increases, the portfolios get more complicated as

they incorporate more and more hedging. The most risk averse portfolio contains a mixture of 80 securities—this represents a typical size for a portfolio at the highly-hedged risk-averse end of the efficient frontier. All portfolios corresponding to $\lambda \in [0, \frac{1}{2})$ are nondominated in the second-order stochastic dominance sense, because each of them is the unique solution for some values of λ . There are 18982 such portfolios and they cover almost the entire efficient frontier displayed in Figure 2.

As we can see from the figure, dramatic improvements in the mean and in the risk are possible, in comparison to the individual securities.

The entire frontier was computed in 23446 pivots taking 1 hour and 46 minutes of cpu time on a Windows 2000 laptop computer having a 1.2 GHz clock. To put this into context, we note that it took an analogous code (i.e., a code built from the same linear algebra subroutines) that implements the usual two-phased simplex method about one and a half hours to compute a single portfolio on the efficient frontier.

6.2. Deviation from the quantile. For the deviation from the 0.05-quantile, the efficient frontier consists of 5023 distinct portfolios and was computed in 5052 pivots taking 18 minutes of cpu time. As with the deviation-from-mean case, risky portfolios contain only a few securities whereas risk-averse portfolios are richer in content. For example, the most risk-averse portfolio consists of 52 securities. All portfolios corresponding to $\lambda \in [0, 1)$ are nondominated in the second-order stochastic dominance sense. There are 4816 of them and they cover almost the entire efficient frontier displayed in Figure 3.

When the risk is measured by the weighted deviation from the 0.05-quantile, the deviations to the left of it are penalized about 20 times more strongly than the deviations to the right. Improving the shape of the left tail of the distribution has therefore an even more dramatic effect on the risk measure than in the deviation-from-mean case.

Acknowledgement. We would like to thank Humbert Suarez at Goldman-Sachs for providing us with a real-world data set on which to base our experiments. The second author was supported by grants from the ONR (N00014-98-1-0036) and NSF (CCR-0098040) which he gratefully acknowledges.

7. REFERENCES

- ARTZNER, P., F. DELBAEN, J.-M. EBER AND D. HEATH (1999): “Coherent Measures of Risk,” *Mathematical Finance*, 9, 203–228.
- BAWA, V. S. (1982): “Stochastic Dominance: a Research Bibliography,” *Management Science*, 28, 698–712.
- BLOOMFIELD, P. AND W. L. STEIGER (1983): *Least Absolute Deviations*, Boston: Birkhäuser.
- FISHBURN, P. C. (1964): *Decision and Value Theory*, New York: John Wiley & Sons.
- GOTOH, J. AND H. KONNO (2000): “Third Degree Stochastic Dominance and Mean–Risk Analysis,” *Management Science*, 46, 289–301.
- HADAR, J. AND W. RUSSELL (1969): “Rules for Ordering Uncertain Prospects,” *The American Economic Review*, 59, 25–34.
- HANOCH, G. AND H. LEVY (1969): “The Efficiency Analysis of Choices Involving Risk,” *Review of Economic Studies*, 36, 335–346.
- HARDY, G. H., J. E. LITTLEWOOD AND G. PÓLYA (1934): *Inequalities*, Cambridge, U.K.: Cambridge University Press.
- KONNO, H. AND H. YAMAZAKI (1991): “Mean–Absolute Deviation Portfolio Optimization Model and its Application to Tokyo Stock Market,” *Management Science*, 37, 519–531.
- LEHMANN, E. (1955): “Ordered Families of Distributions,” *Annals of Mathematical Statistics*, 26, 399–419.
- LEVY, H. (1992): “Stochastic Dominance and Expected Utility: Survey and Analysis,” *Management Science*, 38, 555–593.
- LORENZ, M. O. (1905), “Methods of Measuring Concentration of Wealth,” *Journal of the American Statistical Association*, 9, 209–219.
- MARKOWITZ, H. M. (1952): “Portfolio Selection,” *Journal of Finance*, 7, 77–91.
- (1959): *Portfolio Selection*, New York: John Wiley & Sons.
- MARSHALL, A. W. AND I. OLKIN (1979): *Inequalities: Theory of Majorization and Its Applications*, San Diego: Academic Press.

- OGRYCZAK, W. AND A. RUSZCZYŃSKI (1999): “From Stochastic Dominance to Mean–Risk Models: Semideviations as Risk Measures,” *European Journal of Operational Research*, 116, 33–50.
- (2001): “On Consistency of Stochastic Dominance and Mean–Semideviation Models,” *Mathematical Programming*, 89, 217–232.
- (2002): “Dual Stochastic Dominance and Related Mean–Risk Models,” *SIAM Journal on Optimization*, 13, 60–78.
- QUIRK, J.P AND R. SAPOSNIK (1962): “Admissibility and Measurable Utility Functions,” *Review of Economic Studies*, 29, 140–146.
- ROTHSCHILD, M. AND J. E. STIGLITZ (1969): “Increasing Risk: I. A Definition,” *Journal of Economic Theory*, 2, 225–243.
- SHARPE, W. F. (1971): “A Linear Programming Approximation for the General Portfolio Analysis Problem,” *Journal of Financial and Quantitative Analysis*, 6, 1263–1275.
- URYASEV, S. AND R. T. ROCKAFELLAR (2001), “Conditional Value-at-Risk: Optimization Approach,” in *Stochastic Optimization: Algorithms and Applications (Gainesville, FL, 2000)*, Dordrecht: Kluwer Academic Publishers, 411–435.
- VANDERBEI, R.J. (2001): *Linear Programming: Foundations and Extensions, 2nd Ed.*, Dordrecht: Kluwer Academic Publishers.
- WHITMORE, G. A. AND M. C. FINDLAY, EDS. (1978): *Stochastic Dominance: An Approach to Decision–Making Under Risk*, Lexington, MA: D. C. Heath.
- YITZHAKI, S. (1982): “Stochastic Dominance, Mean Variance, and Gini’s Mean Difference,” *American Economic Review*, 72, 178–185.

Rutgers University, Department of Management Science and Information Systems and RUTCOR, Piscataway, NJ 08854, USA; rusz@rutcor.rutgers.edu; <http://www.rusz.rutgers.edu/>.

Princeton University, Department of Operations Research and Financial Engineering, and Bendheim Center for Finance, Princeton, NJ 08544, USA; rvd@princeton.edu; <http://www.princeton.edu/~rvdb>.

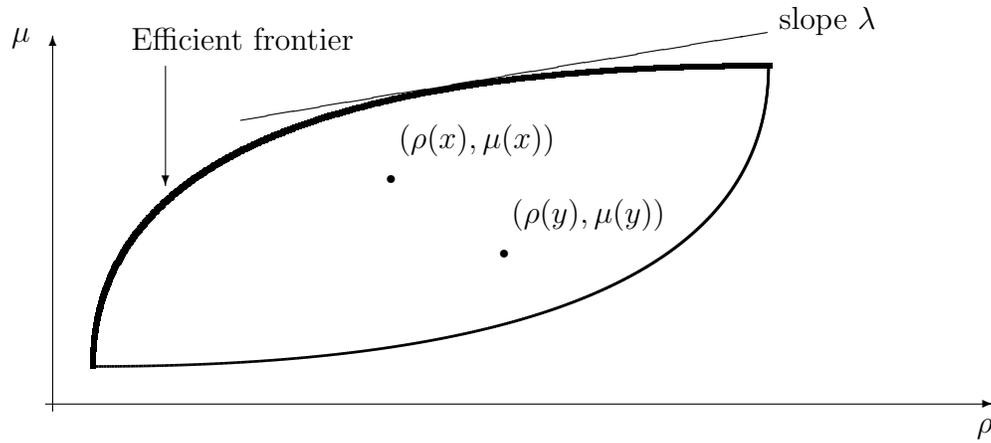


FIGURE 1. Mean–risk analysis. Portfolio x is better than portfolio y in the mean–risk sense, but neither is efficient.

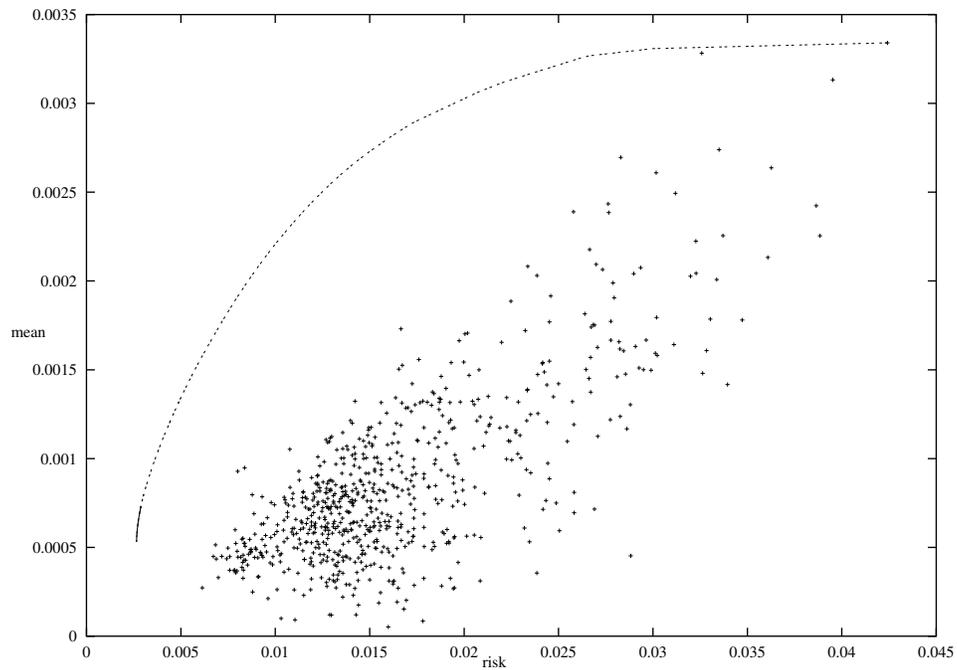


FIGURE 2. The efficient frontier for the deviation-from-mean risk measure. The dashed portion of the frontier represents stochastically nondominated portfolios ($\lambda \leq 1/2$) whereas the small solid portion in the lower left corner corresponds to $\lambda > 1/2$.

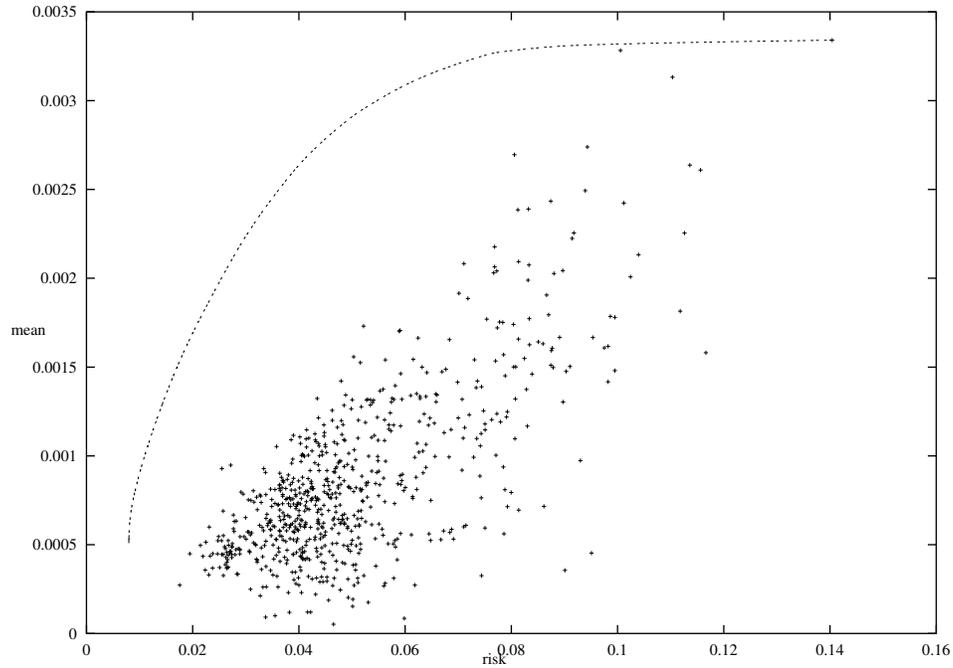


FIGURE 3. The efficient frontier for the deviation-from-quantile risk measure using the $p = 0.05$ quantile. The dashed portion of the frontier represents stochastically nondominated portfolios ($\lambda \leq 1$) whereas the small solid portion in the lower left corner corresponds to $\lambda > 1$.