Asymmetric risk measures and tracking models for portfolio optimization under uncertainty

Alan J. King
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598, USA

Traditional asset allocation of the Markowitz type defines risk to be the variance of the return, contradicting the common-sense intuition that higher returns should be preferred to lower. An argument of Levy and Markowitz justifies the mean/variance selection criteria by deriving it from a local quadratic approximation to utility functions. We extend the Levy–Markowitz argument to account for asymmetric risk by basing the local approximation on *piecewise linear–quadratic risk measures*, which can be tuned to express a wide range of preferences and adjusted to reject outliers in the data. The implications of this argument lead us to reject the commonly proposed asymmetric alternatives, the mean/lower partial moment efficient frontiers, in favor of the "risk tolerance" frontier. An alternative model that allows for asymmetry is the tracking model, where a portfolio is sought to reproduce a (possibly) asymmetric distribution at lowest cost.

Keywords: Portfolio optimization, asymmetric risk, lower partial moments, tracking models, quadratic programming, stochastic programming.

1. Introduction

In his seminal work [13], Markowitz introduced a mean/variance portfolio optimization model that identifies an "efficient frontier" of portfolios possessing minimum variance of return for given levels of expected return and showed how these can be computed using the technique of parametric quadratic programming. A major criticism of this model is that it does not account for asymmetries in return distributions and investor preferences. The present paper presents two extensions of the Markowitz model that address these concerns.

In the 1970s, many papers (e.g. Bawa [1], Fishburn [4]) recommended a *mean/lower partial mean* or *mean/lower partial variance* framework for portfolio optimization, whose justification proceeds from the observation that an investor's true risk is the downside risk. Levy and Markowitz [12] responded by arguing that mean/variance is to be viewed as a local second-order approximation to utility maximization provided the region of good approximation was large compared to the standard deviation of portfolio return. In the first part of this paper, we adapt the Levy–Markowitz argument to cover asymmetric risk. The key principles we follow in the derivation are (1) that the local approximations should themselves be utility functions, and (2) that the approximation not involve derivatives higher than the second order. The argument begins by introducing a class of linear–quadratic risk measures to replace the second-order terms in the utility function approximation, and then applies conditions to ensure that they be globally non-decreasing. The major aim of the discussion is to clarify the relationship between asymmetric risk measures and expected utility. This perspective leads to some surprising insights into the use of lower partial variance as a measure of risk – namely, that a mean/lower partial variance efficient frontier can *over-estimate* the investor's preference for upside versus downside performance – arguing for some caution in the use of asymmetric risk measures until these principles (and their value in practice) are better understood.

Asymmetry can also be treated in a slightly different extension of the Markowitz model. Here the idea is to replicate a desired return distribution through the market, by purchasing (at a cheaper cost, one hopes) a combination of assets that "track" the target distribution [2, 3]. We also discuss, in the final section, how quadratic programs can be constructed to generate efficient frontiers of risk versus reward for this class of asymmetric risk measures – the construction here recalls Konno's [11] use of piecewise linear functions to represent investor risk.

The empirical justification for the exploitation of asymmetry in investment programs awaits a comprehensive research program. There are some indications [6, 14] that asymmetry can be found and exploited to yield excess expected returns in certain cases. Certainly, the explosion of derivatives markets in response to investors' need to hedge risky positions would seem to offer a great opportunity for the use of asymmetric risk as a guide to portfolio diversification once the underlying statistical models are understood. A convincing linkage between a reliable statistical model for the estimation of asymmetric risk and a computationally tractable optimization model to diversify that risk has yet to be discovered.

2. Linear–quadratic risk approximations to utility functions

When all returns are known, the mathematical statement of the portfolio optimization problem is to maximize the total portfolio return

$$r^{\mathrm{T}}x \tag{2.1}$$

over all portfolios $x = (x_1, \ldots, x_n)$ with single period net returns $r = (r_1, \ldots, r_n)$ in a portfolio universe X defined by finitely many linear constraints

$$X = \{Ax = b,\ x \geq 0\}. \tag{2.2}$$

(The reader should note that in specifying the constraint $x \geq 0$, we do not rule out short-selling. Short sales of an asset – selling an asset one does not own – can be modeled by introducing short and long positions in each asset as separate variables and adjusting the corresponding "returns".)

The difficulty arises, of course, when the returns r are unknown. In such a case one wishes to design an investment program that results in the best possible *distribution* of portfolio returns. There are two steps in such a design: first, one must estimate the probability distribution of r, and second, one must design a mathematical procedure to find the portfolio giving the best distribution of return. We are concerned in this paper with the second step, the design of a mathematical procedure to select portfolio return distributions with favorable qualities.

The best known mathematical procedure for portfolio selection under uncertainty is due to Markowitz [13], who proposed modeling the selection problem as one of finding a desirable trade-off between mean return and variance of return and designed a procedure to display an "efficient frontier" of portfolios with minimum variance for a given level of mean return. This procedure has been criticized because it penalizes upside variations equally as severely as the downside variations, whereas it is the latter which are presumably of greater concern to the investor. These criticism are well-known in the literature; see, for example [1, 4–6, 14]. What seems less well-known is the connection between mean/variance and the theory of utility maximization.

This was explained in a paper of Levy and Markowitz [12], where it is argued that any utility function describing investor's risk is well approximated by a combination of mean and variance. The argument may be paraphrased roughly like this. Letting u be the investor's true utility function for wealth at the end of the investment period, the portfolio selection problem may be characterized as one of maximizing expected utility:

$$\text{maximize} \quad Eu(r^{\mathrm{T}}x) = \int u(r^{\mathrm{T}}x)\,\mathrm{d}P(r)$$

$$\text{subject to } x \in X.$$

In this problem, r is the random vector of portfolio returns with probability distribution P, and E represents the operation of taking the mathematical expectation over this distribution. Suppose, for the moment, that the utility-maximizing portfolio $\bar{x}$ has been found by some means or other. We can compute the expected return of this portfolio, say $\bar{R} = Er^{\mathrm{T}}\bar{x}$. Next, define the function u_2 to be the second-order Taylor expansion of u about the given point $\bar{R}$:

$$u_2(t) = u(\bar{R}) + u'(\bar{R})(t - \bar{R}) + \tfrac{1}{2}u''(\bar{R})(t - \bar{R})^2. \tag{2.3}$$

If the variation of total portfolio return about $\bar{R}$ is so low that u_2 is a good approximation to u over this range of variation, then we may find a portfolio that is nearly as good as $\bar{x}$ by substituting u_2 for u in the utility maximization and restricting the maximization to all portfolios with expected return $\bar{R}$. This leads to the following problem (after eliminating the constant term and observing that the linear term integrates to zero)

$$\begin{aligned} &\text{maximize } \tfrac{1}{2}u''(\bar{R})E(r^{\mathrm{T}}x - \bar{R})^2 \\ &\text{subject to } x \in X, \\ &\qquad Er^{\mathrm{T}}x = \bar{R}. \end{aligned}$$

Making the reasonable assumption that the utility function u is concave and twice differentiable, the factor u'' in the objective term will be negative. Clearing this term converts the maximization to a minimization, and thus, the optimization problem appears in its familiar form as none other than the traditional mean/variance quadratic program for one particular value $\bar{R}$ of the expected return. By running through all possible levels of return R, it follows that we will identify the entire class of portfolios that optimize the second-order approximation to concave twice-differentiable utility functions. This would be, of course, the mean/variance efficient frontier

$$\begin{aligned} &\text{minimize } E(r^{\mathrm{T}}x - \bar{r}^{\mathrm{T}}x)^2 \\ &\text{subject to } \bar{r}^{\mathrm{T}}x \geq R, \\ &\qquad x \in X. \end{aligned} \tag{2.4}$$

Levy and Markowitz end their argument by presenting evidence that this class of portfolios seems sufficiently rich that any risk-averse investor would not feel unduly restricted in choosing from among them.

This argument leaves one assumption in place: that the distribution of return r about its expected return be such that the approximation of u by u_2 remains valid. Of course, any asymmetry in the return distribution would render the approximation invalid. The simplest condition of validity we could impose, given that it is difficult to know one's utility function precisely enough to estimate the derivatives and given that we do not wish to make distributional assumptions concerning r, is that *the approximations themselves have the attributes of a utility function.*

We define now a general class of linear–quadratic risk measures that will allow us to retain the local quadratic feature of the Taylor approximation but give us the flexibility to adjust the slope so that the approximation is concave and nondecreasing.

DEFINITION

Functions of the type

$$\rho_{q^-,q^+}(t) = \begin{cases} q^- t - \frac{1}{2}(q^-)^2 & \text{if } t \le q^-, \\ \frac{1}{2}t^2 & \text{if } q^- < t < q^+, \\ q^+ t - \frac{1}{2}(q^+)^2 & \text{if } q^+ \le t \end{cases}$$

are called *linear–quadratic risk measures*. The parameters are the left slope q^-, and the right slope q^+. (We require $q^- \le q^+$ to preserve convexity.) A wide choice of linear–quadratic risk measures are available. One useful variant is the "robust" version of the lower partial variance

$$\rho_{-k,0}(t) = \begin{cases} -kt - \frac{1}{2}k^2 & \text{if } t \le -k, \\ \frac{1}{2}t^2 & \text{if } -k \le t \le 0, \\ 0 & \text{if } 0 \le t. \end{cases}$$

This risk measure is less influenced by "outliers" than the lower partial variance, and may perhaps be useful when the distribution of random returns is formed from sample data. We will not pursue this issue here; interested readers can consult Huber [7].

Returning now to the definition of the second order approximation u_2 in (2.3), let us modify it so that the local approximation is concave and nondecreasing:

$$u_2(t) = u(R) + u'(R)(t-R) + u''(R)\rho_{q^-,q^+}(t-R). \tag{2.5}$$

That is, we replace the quadratic term in (2.3) by the linear–quadratic term

$$u''(R)\rho_{q^-,q^+}(t-R).$$

In more detail, the replacement term is:

$$u''(R)\rho_{q^-,q^+}(t-R) = \begin{cases} u''(R)q^-(t-R) - \frac{1}{2}u''(R)(q^-)^2, & \text{if } (t-R) \le q^-, \\ \frac{1}{2}u''(R)(t-R)^2, & \text{if } q^- \le (t-R) \le q^+, \\ u''(R)q^+(t-R) - \frac{1}{2}u''(R)(q^+)^2, & \text{if } q^+ \le (t-R). \end{cases}$$

We now discuss the setting of the two parameters q^- and q^+ which represent the left and right slopes, respectively, of the linear–quadratic term. To fulfill u_2's role of providing a local approximation to the utility function u, it is clear we must have

$\rho_{q^-,q^+}(0) = 0$. This condition is satisfied if and only if $q^- \leq 0 \leq q^+$. Next, we want u_2 to be a nondecreasing, concave utility. (Note that $u''(R) \leq 0$, always, and thus, $q^- \leq q^+$ implies u_2 is concave.) To see what is required to have u_2 be nondecreasing, consider the terminal right slope of u_2

$$\lim_{t\to\infty} u_2'(t) = u'(R) + u''(R)q^+.$$

It follows from this expression that u_2 will be nondecreasing if and only if

$$q^+ \leq \frac{-u'(R)}{u''(R)}.$$

Combining these observations, it follows that the appropriate range of q^- and q^+ is

$$q^- \leq 0 \leq q^+ \leq \frac{-u'(R)}{u''(R)}.$$

The value of q^- is of less importance in the remainder of the discussion (although may be important in reference to statistical robustness, as noted above) and we shall henceforth assume that $q^- = -\infty$. Varying q^+ over the range $[0, -u'(R)/u''(R)]$ preserves the quality of u_2's approximation to u at the point R while keeping u_2 in the class of concave, nondecreasing utility functions.

The two ends of this range have special significance. If we place q^+ at the lower side of the range, the Levy–Markowitz argument leads directly to the *mean/lower partial variance* efficient frontier generated by

$$\begin{aligned} &\text{minimize } E\rho_{-\infty,0}(r^{\mathrm{T}}x - \bar{r}^{\mathrm{T}}x) \\ &\text{subject to } \bar{r}^{\mathrm{T}}x \geq R, \\ &x \in X. \end{aligned} \tag{2.6}$$

The term $E\rho_{-\infty,0}(r^{\mathrm{T}}x - \bar{r}^{\mathrm{T}}x)$ is the *lower partial variance* of the random total return $r^{\mathrm{T}}x$. Mean/lower partial moment efficient frontiers appear in the theory of stochastic dominance; see, for example, Fishburn [4] or Bawa [1] for details. However, in the present context where we are deriving the risk measure from principles of expected utility, there is a surprise.

At the lower side of the range where $q^+ = 0$, the approximating function u_2 has a slope equal to $u'(R)$ for *all* values of t greater than R. Such functions lie *on or above* the true utility function; in other words, they *over-estimate* the investor's utility of obtaining an increase in return above R. We may conclude that selecting assets solely on the basis of mean return and lower partial variance (or indeed, any other lower partial moment) leads to portfolios that are *riskier* than the investor may wish.

This surprising conclusion has been completely unanticipated by the literature on downside risk, much of which seems to have been inspired by Markowitz's observation [13] that in the presence of asymmetric return distributions the mean/semivariance frontier would yield a superior set of portfolios to those identified by the mean/variance frontier. In the light of Levy and Markowitz's (later) argument relating mean/variance to utility maximization, one sees that the issue is more complex than it at first seemed. The discussion of the previous paragraph shows that while one should try to modify the mean/variance model to reduce the penalization of upside variation implied by the minimization of variance, the Levy–Markowitz argument demonstrates that some upside penalty should be applied in order to reduce the impact of the linear preference implied by the maximization of expected return.

Indeed, in the design of an asymmetric risk penalty for a risk averse investor, it would seem prudent to ensure that the risk penalty be consistent with a utility approximation that understates as far as possible the advantage of upside return while still remaining a valid utility. This leads us to propose the upper limit of the range for q^+ as the proper choice of design parameter for asymmetric risk penalties, namely,

$$q^+ = \lambda(R) := \frac{-u'(R)}{u''(R)},$$

which is the value of q^+ that causes the approximate utility u_2 eventually to have slope 0 as $t \to \infty$. The inverse of this ratio is known to students of utility theory as the investor's "local risk aversion" parameter – the risk premium that an investor would wish to receive in order to be compensated for accepting an additional unit of risk measured by variance [15]. We may characterize $\lambda(R)$ as the investor's "risk tolerance" parameter – the amount of risk (variance) that an investor is willing to accept for an infinitesimal increase in expected return above R. The greater the value of $\lambda(R)$, the more risk tolerant the investor. For example, for an investor whose utility is the logarithm function, the risk tolerance parameter is

$$\lambda(R) = \frac{1}{R}.$$

At a return value of $R = 1.05$, the risk tolerance is $\lambda(R) = 0.95$ and at a return value of $R = 1.20$, the risk tolerance is $\lambda(R) = 0.83$. (The explanation is that an investor with this utility will accept some risk to raise a low level of expected return.) We now make a very important observation with respect to this utility: namely that over the range $0 < t \leq 2$ (approximately) *the asymmetric (risk-tolerance) risk measure is identical to the symmetric (variance) risk measure.* In other words, if this is your utility and you think that there is no reasonable prospect that any asset in your portfolio universe will return more than 100%, then going to an asymmetric model will not buy you much! On the other hand, if you do wish to consider assets with such a high potential upside return, then going to an asymmetric model is the only way to cause such assets to appear in an efficient frontier.

The risk-tolerance frontier may be drawn by generating the efficient frontier for the following optimization

$$\begin{aligned} &\text{minimize} \quad E\rho_{-\infty,\lambda(R)}(r^{\mathrm{T}}x - \bar{r}^{\mathrm{T}}x) \\ &\text{subject to } \bar{r}^{\mathrm{T}}x \geq R, \\ &\qquad\qquad x \in X. \end{aligned} \tag{2.7}$$

This imposes the additional burden of knowing the investor's risk tolerance ratios over the range of R. The examination of the logarithmic utility above shows that there is much to be gained, however, from even a simple exploration of an investors risk-tolerance ratios relative to the gains anticipated from the investor's portfolio universe.

Some general comments about the relevance of asymmetric risk measures are in order here. When the asset returns are symmetrically distributed, the mean/variance, mean/semivariance, and the risk tolerance efficient frontiers are all identical, of course. It makes sense to go to asymmetric risk measures only when the investor anticipates some asymmetry in the asset return distributions. There is some evidence that historical asset returns are asymmetric; see [5, 6, 14]. We comment briefly on two other possible sources of asymmetry.

First, and perhaps most importantly, many financial assets such as *options* are designed to have an asymmetric return when held for any reasonable length of time. Portfolios containing options with expiration dates on the order of the portfolio revision cycle will show asymmetric return distributions even when all underlying primary assets are symmetrically distributed. Alternatively, some financial assets contain embedded options, such as mortgage-backed securities. Programs that now price such assets can also be used to model a distribution of returns at a future date. Such distributions can be used for portfolio optimization, and they are not likely to be symmetric.

Second, many investor models of equity return distributions are designed around a linear factor model of the type $r = \beta v + \epsilon$. The factors v may relate to various macro-economic and/or policy indices. The investor may wish to design a simple asymmetric distribution to model the future trend of the factors; alternatively, careful statistical study may have lead the investor to the conclusion that some of the factors have historically displayed an asymmetry. Either would lead to asymmetric distributions for the equity returns through the factor model.

3. Tracking models and minimum regret

Asymmetry can be treated in a model that is closer in spirit to the sort of decisions that arise in practical financial situations. (The discussion in this section summarizes developments in the three refs. [2, 3, 10].) To motivate the ideas, we view the Markowitz mean/variance model in the form of a "tracking model"

$$\begin{aligned} &\text{minimize} \quad E(r^{\mathrm{T}}x - \tau)^2 \\ &\text{subject to } x \in X. \end{aligned} \tag{3.1}$$

(See [10], where it is shown that the mean/variance efficient frontier can be generated via a model of this type, for τ varying between suitable limits.) In [2], two simple modifications were proposed. First, the target τ in the objective may be generalized, for example, by letting it be a random variable. Second, the efficient frontier may be generalized by parametrizing a different criterion, for example, the cost of assembling the portfolio.

These generalizations lead to the following formulation of the tracking model:

$$\begin{aligned} &\text{minimize} \quad E(r^{\mathrm{T}}x - \tau)^2 \\ &\text{subject to } c^{\mathrm{T}}x \le d, \\ &\qquad\qquad\quad x \in X. \end{aligned} \tag{3.2}$$

An efficient frontier is drawn by letting d increase from zero up to the cost of purchasing the distribution τ in the marketplace. The interpretation of this tracking model is as follows. The investor desires a portfolio $x \in X$ with two competing characteristics:

$$\begin{aligned} r^{\mathrm{T}}x \ge \tau, &\quad \text{almost surely, and} \\ c^{\mathrm{T}}x &\quad \text{as low as possible.} \end{aligned}$$

Such may arise when a portfolio with given performance distribution τ is sought at lowest cost d. The efficient frontier reveals portfolios that optimally track this required performance for different levels of cost; portfolios with a suitable tradeoff of tracking quality and cost may be selected from this frontier. This formulation gives a natural way to introduce asymmetric distributions into the portfolio universe. If the distribution of the target return τ is asymmetric, then the mix of assets chosen to replicate the target will match the asymmetry (in the least squares sense) for the given cost d. Obviously, the more one is willing to spend, the closer the match.

But what performance would an investor require of a portfolio? One natural answer to this question was explored in Dembo and King [3]. Let us suppose that perfect foresight allows us to predict the actual value of returns r. Let $\tau(r)$ be the optimal return with perfect foresight; that is,

$$\tau(r) = \max\{r^{\mathrm{T}}x : x \in X\}.$$

This optimal return is a random variable, since it depends on the random vector of returns r. Tracking this optimal return leads to the following problem:

$$\begin{aligned} &\text{minimize} \quad E(r^{\mathrm{T}}x - \tau(r))^2 \\ &\text{subject to } x \in X. \end{aligned} \tag{3.3}$$

Because $\tau(r)$ is the best that can be achieved, the risk measure in effect penalizes only the downside risk. The solution to this problem is called the "minimum regret" portfolio. Regret is the performance sacrificed because of the impossibility of perfect foresight, namely $R(x) = \tau(r) - r^{\mathrm{T}}x$. One possible measure of the performance of a portfolio x is how closely its regret $R(x)$ matches perfect foresight. In [3] it is shown that the solution to (3.2) minimizes this regret measure.

Tracking as a measure of portfolio performance is virtually unlimited in its versatility. For example, one can specify performance goals over a number of time periods and design a tracking model to find a portfolio that tracks these goals, period by period.

It is interesting to compare the two philosophies, that of utility maximization as described in section 2 and that of tracking as discussed in the present section. In the decision problem that motivated the present discussion, the decision maker is seeking to hedge an exposure τ with a portfolio x returning $r^{\mathrm{T}}x$ and costing $c^{\mathrm{T}}x$. Thus the investor is in effect paying $c^{\mathrm{T}}x$ now to receive $r^{\mathrm{T}}x - \tau$ later. Framing the decision as a utility maximization problem, one would seek to maximize the expected utility of $(r - c)^{\mathrm{T}}x - \tau$. The procedure of utility maximization would determine the appropriate cost d of purchasing the hedging portfolio indirectly from the shape of the utility function and the distributions of r and τ, since this cost has an impact on the expected return of the portfolio. Framing the decision as a tracking problem, one is seeking to find a portfolio that matches (in some sense) the distribution of the exposure τ at least cost d. This has more the flavor of risk arbitrage, where one is selling τ, buying x at a cost d, and accepting a risk $r^{\mathrm{T}}x - \tau$. The question then is: what is the price one should set for τ in order to accept the risk? Presumably, the answer to this would be based on the net effect on the arbitrageur's total portfolio.

4. Computation of linear–quadratic efficient frontiers

We now presume that we have selected a generic linear–quadratic tracking function

$$\rho(t):=\rho_{q^-,q^+}(t)$$

and proceed to discuss how to compute solutions to the parametric linear–quadratic stochastic program

$$\begin{aligned} &\text{minimize} \;\; E\rho(r^{\mathrm{T}}x - \tau) \\ &\text{subject to} \;\; c^{\mathrm{T}}x \le d, \\ &\qquad\qquad x \in X. \end{aligned} \tag{4.1}$$

It is not possible in general to compute a closed-form expression for the objective function $x \mapsto E\rho(r^{\mathrm{T}}x - \tau)$. Indeed, it would not be clear how we should proceed even if we did find one, since this would place the problem in the realm of large-scale nonlinear programming for which reliable techniques are not yet available. The whole point of restricting attention to piecewise linear–quadratic tracking functions is to retain access to the technology of quadratic programming.

The procedure to form a quadratic program from (4.1) was explained in King and Jensen [10]. The first step in the conversion of (4.1) to a quadratic program is to discretize the probability measure if it is not already finitely supported. We suppose that the investor has in mind a finite set $\{r_s : s \in S\}$ of *scenarios* that describes (or approximates) the investor's view of the range of uncertain returns, and for each scenario a probability p_s that indicates the weight of belief to be ascribed to each scenario. We may suppose the probabilities are normalized so that $\sum_s p_s = 1$.

With this finite distribution, the objective function (4.1) is a finite sum of piecewise linear–quadratic functions and can be solved using large scale quadratic programming techniques. The quadratic program emerges from the usual trick of introducing variables to take the value of the various pieces of the objective. For each scenario s we introduce up to three variables v_s, v_s^-, and v_s^+, and the equations and inequalities

$$\begin{aligned} &v_s - v_s^- + v_s^+ = r_s^{\mathrm{T}}x - \tau_s, \\ &v_s \text{ free}, \quad v_s^- \ge q^-, \quad v_s^+ \ge q^+. \end{aligned} \tag{4.2}$$

The first variable takes the quadratic part of ρ and the latter two take the lower linear and upper linear parts, respectively. The contribution to the objective function from this scenario is

$$\tfrac{1}{2}v_s^2 - q^- v_s^- + q^+ v_s^+. \tag{4.3}$$

It is easy to verify that minimizing this objective in the new variables subject to (4.2) produces the value $\rho_{q^-,q^+}(r_s^{\mathrm{T}}x - R)$. Collecting all the rows (4.2) and objective terms (4.3) produces the quadratic program

$$\begin{aligned} \text{minimize } & \sum_{s\in S} p_s[\tfrac{1}{2}v_s^2 - q^- v_s^- + q^+ v_s^+]' \\ \text{subject to } & c^{\mathrm{T}}x \le d, \\ & v_s - v_s^- + v_s^+ = r_s^{\mathrm{T}}x - \tau_s, \qquad \forall s \in S, \\ & v_s \text{ free}, \quad v_s^- \ge q^-, \quad v_s^+ \ge q^+, \quad \forall s \in S, \\ & x \in X. \end{aligned} \tag{4.4}$$

The quadratic program (4.4) is entirely equivalent to the linear–quadratic tracking model (4.1) under the finite distribution for r, and it is in a form that can be solved by quadratic programming computer codes. Note that the quadratic contribution is a sum of squares – this is a rather simple quadratic form. The difficulty lies in the matrix whose rows are the samples r_s^{T}. When $|S|$ is large, this forms a large, dense matrix that can cause all sorts of numerical crises in the course of solution. Nevertheless, RISC hardware and specialized optimization software gives surprisingly good performance on problems of this type. To generate a complete efficient frontier for (4.4), with the right-hand side d varying over its entire feasible range, on a practical problem with 100 scenarios and 1015 assets takes about ninety seconds using currently available commercial optimization codes (IBM's Optimization Subroutine Library [8]) on a mid-level scientific workstation (IBM's RISC System/6000 Model 530). For further performance details, see [10]; and for a description of a software package "FRONTIER" designed to solve problems like (4.4) see Jensen and King [9].

References

[1] V.S. Bawa, Optimal rules for ordering uncertain prospects, J. Fin. Econ. 10 (1975) 849–857.
[2] R.S. Dembo, Scenario optimization, Ann. Oper. Res. 30 (1991) 63–80.
[3] R.S. Dembo and A.J. King, Tracking models and the optimal regret distribution in asset allocation, IBM Research Report RC 17156 (1991).
[4] P.C. Fishburn, Mean-risk analysis with risk associated with below-target returns, Amer. Econ. Rev. 67 (1977) 116–125.
[5] W.V. Harlow and R.K.S. Rao, Asset pricing in a generalized mean-lower partial moment framework: Theory and evidence, J. Fin. Quantit. Anal. 24 (1989) 285–311.
[6] W.V. Harlow, Asset allocation in a downside risk framework, Equity Portfolio Analysis Report, Salomon Brothers Inc. (1991).
[7] P.J. Huber, *Robust Statistics* (Wiley, New York, 1981).
[8] International Business Machines Corporation, *Optimization Subroutine Library Guide and Reference* (IBM Corp, Document SC23-0519-1, 1990).

[9] D.L. Jensen and A.J. King, Frontier: a graphical interface for portfolio optimization in a piecewise linear–quadratic risk framework, IBM Syst. J. 31 (1992) 62–70.
[10] A.J. King and D.L. Jensen, Linear–quadratic efficient frontiers for portfolio optimization, Appl. Stoch. Models Data Anal. 8 (1992) 195–207.
[11] H. Konno, Piecewise linear risk function and portfolio optimization, J. Oper. Res. Soc. Japan 33 (1990) 139–156.
[12] H. Levy and H.M. Markowitz, Approximating expected utility by a function of mean and variance, Amer. Econ. Rev. (June 1979) 308–317.
[13] H.M. Markowitz, *Portfolio Selection: Efficient Diversification of Investments* (Wiley, New York, 1959).
[14] T.J. Nantell, K. Price and B. Price, Mean-lower partial moment asset pricing model: Some empirical evidence, J. Fin. Quantit. Anal. 17 (1982) 763–782.
[15] J.W. Pratt, Risk aversion in the small and in the large, Econometrica 32 (1964) 122–137.