Optimal robust and consistent active implementation of a pension fund's benchmark investment strategy

Received (in revised form): 12th June, 2007

Tim van Hest*
is a part-time PhD student at Tilburg University, The Netherlands, and the Head of credit risk modelling ABB department of Rabobank Netherlands. His current research interests include robust portfolio optimisation and credit risk modelling.

Anja De Waegenaere
is an associate professor of actuarial sciences and accountancy at Tilburg University. She obtained her PhD in mathematics in 1993. Her current research interests include actuarial modelling, robust portfolio optimisation, longevity risk and fair valuation.

*Department of Econometrics and OR, Tilburg University, Tilburg 5000 LE, The Netherlands.
Tel: +31 13 466 29 13; Fax: +31 13 466 32 80; E-mail: T.vanHest@uvt.nl

Abstract The benchmark investment strategy of a pension fund typically consists of a number of benchmark categories, each of which is assigned a weight in the overall investment budget. In this paper we assume that the benchmark strategy is given, and determine a model for its optimal active implementation. Active implementation involves a number of investment managers each of whom are assigned a specific benchmark category. We present a mean–variance approach to determine, for each investment manager, the optimal budget as well as the fraction of that budget that can be used for deviations from the benchmark. The emphasis is on robustness of the optimal allocation with respect to parameter misestimation, and on consistency in terms of risk-return preferences between active implementation and benchmark investment strategy.
Journal of Asset Management (2007) **8,** 176–187. doi:10.1057/palgrave.jam.2250072

Keywords: *asset liability management, pension funds, active investment decisions, robust optimisation*

Introduction

In the context of asset liability management (ALM), two types of investment decisions can be distinguished:

1. The choice of the ALM benchmark investment strategy: passive risk management.
2. Tactical and operational investment decisions: active risk management.

The ALM benchmark yields the allocation of the total investment budget over the different benchmark categories such as equity, bonds, real estate, etc. Usually, this allocation is given in (monetary) weights. Since for a pension fund the match between assets and liabilities is of utmost importance, benchmark weights are typically determined by means of large-scale simulation models, see for example Boender (1997). Practical implementation of the benchmark investment strategy usually involves several investment managers. Each of these investment managers is assigned a specific benchmark category, and is responsible for part of the total investment budget. In order to be able to benefit from short-term opportunities, investment managers would typically be allowed to spend a fraction of the budget assigned to them on active management, that is deviations from the benchmark are allowed. They will try to take advantage of market conditions to outperform the ALM benchmark and generate excess return. These deviations from the benchmark, however, also affect the risk-return profile of the overall investment portfolio. Roll (1992), Waring *et al.* (2000) and Baierl and Chen (2000) determine the efficient risk-return frontier for active investment decisions through Markowitz mean–variance optimisation.

Our goal is to determine the optimal active implementation of a given benchmark strategy. Our approach adds to the literature in the following sense. First, whereas the existing literature mainly focusses on the active weights in isolation, we explicitly distinguish weights for active and passive investment decisions, so that the optimal mean–variance allocation takes into account possible correlation between excess returns and benchmark returns. We incorporate costs for active and passive management, and introduce a restriction that guarantees consistency of the active implementation with the predefined benchmark investment strategy. Secondly, there is extant evidence in the literature (eg Best and Grauer, 1991; Chopra and Ziemba, 1993; Michaud, 1998; Ceria and Stubbs, 2006) that mean–variance optimisation is extremely sensitive to parameter misestimation. Small deviations in the estimated means and covariances can lead to large differences in the optimal allocation. Several approaches have been developed to successfully mitigate this problem (eg Goldfarb and Iyengar, 2003; Ceria and Stubbs, 2006; Schöttle and Werner, 2006). Kritzman (2006) argues that 'misallocations' due to parameter misestimation are only problematic if they give rise to non-negligible errors in the return distribution, and illustrates that estimation errors in the means not necessarily lead to such problematic misallocations. We allow for errors in the covariances as well as in the means of the excess returns, and show that parameter uncertainty is particularly relevant in determining optimal active weights. Misallocations due to parameter uncertainty potentially do lead to large differences in the returns distribution. We therefore draw on techniques for robust optimisation recently developed by Ben-Tal and Nemirovski (1998) and Goldfarb and Iyengar (2003) to obtain solutions with a pre-specified degree of robustness with respect to parameter uncertainty for both the means and the covariances. We find that because active management would typically increase expected return at the cost of higher variance, robust optimisation significantly decreases the weights assigned for active management as compared to nonrobust optimisation, and leads to non-negligible changes in the returns distribution of optimal allocation.

Thirdly, determining mean–variance optimal robust allocations requires specification of a risk aversion parameter. As argued above, in an ALM-context, benchmark weights are typically determined by means of large-scale simulation models. In order to guarantee consistency between the risk-return trade-offs of the benchmark and of the active investment decisions, we require the level of risk aversion used to determine optimal active budgets to be as close as possible to the level of risk aversion consistent with the benchmark weights.

The model

Let us start by introducing some notation and terminology. As discussed in the introduction, the benchmark investment strategy would typically consist of a number of benchmark categories. For the practical implementation of the strategy, there would be a number of investment managers and each of them would be assigned a particular benchmark category. We denote:

l = the number of categories in the ALM benchmark strategy

k = the number of available investment managers

and

$\bar{w}_j$: weight of category j in the ALM benchmark, $\bar{w}_j \in [0, 1]$
$I(j)$: set of investment managers with benchmark category j, $I(j) \subset \{1, \ldots, k\}$
$j(i)$: the benchmark category of manager i, $j(i) \in \{1, \ldots, 1\}$.

Now, active implementation requires determination of:

w_i: budgetary weight of manager i, $w_i \in [0, 1]$
γ_i: fraction of budgetary weight of manager i available for active management, $\gamma_i \in [0, 1]$.

Our goal in the paper is to determine values for w_i and γ_i, for $i = 1, \ldots, k$, such that an optimal trade-off between risk and return is obtained, with emphasis on robustness of the optimal allocation with respect to parameter misestimation. The specific objective function will be defined in the next section. In the remainder of this section, we first specify a number of constraints on the decision variables. First, we assume that short-selling of budgets is not allowed, that is we impose the following restrictions:

$$w_i \geq 0 \quad \text{for } i = 1, \ldots, k \tag{1}$$

$$0 \leq \gamma_i \leq 1 \quad \text{for } i = 1, \ldots, k \tag{2}$$

Moreover, in order to achieve consistency with the ALM benchmark strategy, the budgetary weights allocated to all managers with benchmark category j should sum up to the weight of benchmark category j, that is

$$\sum_{i \in I(j)} w_i = \bar{w}_j \quad \text{for } j = 1, \ldots, l \tag{3}$$

In order to determine the total return as a function of the decision variables, we distinguish the following random variables:

R_i^a: active return of manager i, $i = 1, \ldots, k$
$\bar{R}_j$: return of benchmark category j, $j = 1, \ldots, l$

so that

$\bar{R}_{bm} := \sum_{j=1}^{l} \bar{w}_j \bar{R}_j$: return of the benchmark investment strategy
$R_i^e := R_i^a - \bar{R}_{j(i)}$: excess return of manager i with respect to his benchmark category
$R_i := (1-\gamma_i)\bar{R}_{j(i)} + \gamma_i R_i^a$: return generated by manager i.

Now, the total return can be written as:

$$\begin{aligned} R_{tot}(w, \gamma) &= \sum_{i=1}^{k} w_i R_i, \\ &= \sum_{i=1}^{k} w_i((i = \gamma_i)\bar{R}_{j(i)} + \gamma_i R_i^a), \\ &= \sum_{i=1}^{k} w_i \bar{R}_{j(i)} + \sum_{i=1}^{k} \gamma_i w_i R_i^e \end{aligned} \tag{4}$$

$$= \bar{R}_{bm} + \sum_{i=1}^{k} \gamma_i w_i R_i^e \tag{5}$$

where the last equality follows from the fact that (3) implies that in the absence of active management, that is when $\gamma_i = 0$ for all i, the total return of the investment strategy is equal to the return of the ALM benchmark strategy.

Now only the cost function remains to be specified. The cost of active as well as passive management is usually assumed to be proportional to the budgets assigned for active and passive management, respectively (eg Scherer, 2002). Let us denote a_i (and b_i) for the cost per dollar of active (and passive) management for manager i, and $a = (a_1, \ldots, a_k)$ and $b = (b_1, \ldots, b_k)$, where $b_i = \bar{b}_j$ for all $i \in I(j)$. Then, the total cost of active and passive management equals:

$$F(w,\gamma)=\sum_{i=1}^{k}\left[\underbrace{a_i\gamma_i w_i}_{\text{active fee}}+\underbrace{b_i(1-\gamma_i)w_i}_{\text{passive fee}}\right]=\sum_{i=1}^{k}(a_i-b_i)\gamma_i w_i+\sum_{j=1}^{l}\bar{b}_j\bar{w}_j \quad (6)$$

where the second equality follows from (3).

Mean–variance optimisation

To determine the optimal weights for active and passive management, we consider a mean–variance approach where a trade-off is made between:

1. The expected total return, corrected for the cost of active and passive management.
2. The risk of the total return, as measured by the variance.

Our optimisation problem differs from the standard Markowitz mean–variance optimisation problem in the sense that we explicitly distinguish weights for active and passive investment decisions, incorporate costs for active and passive management, and introduce a restriction that guarantees consistency with the predefined benchmark investment strategy. Specifically, the optimal weights of active and passive management are found by solving the following optimisation problem:

$$(w^*,\gamma^*)\in \arg\max\{E[R_{tot}(w,\gamma)]-\lambda V[R_{tot}(w,\gamma)]-F(w,\gamma)\}\quad \text{s.t.}\quad (1),\ (2)\ \text{and}\ (3) \quad (7)$$

where $\lambda \geq 0$ represents the degree of risk aversion. Note that the total return $R_{tot}(w,\gamma)$ and the cost function $F(w,\gamma)$ depend only on the weights for active management, $\gamma_i w_i$, and not on the budgetary weights w_i and the active fractions γ_i separately. Let us therefore introduce

$$\tilde{w}_i := \gamma_i w_i$$

for the active weight of manager i, $i=1,\ldots,k$. Then,

$$E[R_{tot}(w,\gamma)]=(1,\tilde{w}')\cdot\mu \quad (8)$$

$$V[R_{tot}(w,\gamma)]=(1,\tilde{w}')\cdot\Omega\cdot\begin{pmatrix}1\\ \tilde{w}\end{pmatrix} \quad (9)$$

$$F(w,\gamma)=(1,\tilde{w}')\cdot cost \quad (10)$$

where $(1,\tilde{w}')$ denotes the row vector $(1,\tilde{w}_1,\ldots,\tilde{w}_k)$, μ denotes the column vector containing the expected return of the benchmark investment strategy μ^{bm} and the expected excess returns of the investment managers, $\alpha_i=E[R_i^e]$, that is $\mu=(\mu^{bm},\alpha_1,\ldots,\alpha_k)'$, Ω denotes the corresponding covariance matrix, that is

$$\Omega=\begin{bmatrix} V(\bar{R}_{bm}) & \sigma(\bar{R}_{bm},R_1^e) & \cdots & \sigma(\bar{R}_{bm},R_k^e)\\ \sigma(R_1^e\,\bar{R}_{bm}) & V(R_1^e) & \cdots & \sigma(R_1^e,R_k^e)\\ \vdots & \vdots & \ddots & \vdots\\ \sigma(R_k^e\,\bar{R}_{bm}) & \sigma(R_k^e,R_1^e) & \cdots & V(R_k^e)\end{bmatrix}$$

and $cost=(\sum_{j=1}^{l}\bar{b}_j\bar{w}_j,\ a_1-b_1,\ldots,a_k-b_k)'$ denotes the column vector containing the cost of passive management (first element) and the costs of active management in excess of passive management for the k investment managers.

Note that the first element in the weight vector $(1,\tilde{w}')$ corresponds to a weight of 100 per cent in the benchmark strategy. The existing literature mainly focuses on the allocation of active risk budgets in isolation, that is the benchmark return is left out of consideration (eg Waring *et al.*, 2000; Roll, 1992). It is clear however from (9) that, if the benchmark return and the returns of active management are correlated, the overall risk may be under- or overestimated if active risk budgets are determined in isolation from the passive investment strategy. By explicitly including the benchmark return in (9), the optimal allocation of active and passive weights takes into account possible correlation between excess returns and benchmark returns.

Now, given (8)–(10), optimisation problem (7) is equivalent to the following quadratic optimisation problem:

$$\tilde{w}^* \in \arg\max\left\{(1, \tilde{w}') \cdot (\mu - cost) - \lambda \cdot (1, \tilde{w}') \cdot \mathbf{\Omega} \cdot \begin{pmatrix} 1 \\ \tilde{w} \end{pmatrix}\right\} \quad (11)$$

$$\text{s.t.} \begin{cases} \sum_{i \in I(j)} \tilde{w}_i \leq \bar{w}_j, & j = 1, \ldots, l, \\ \tilde{w}_i \geq 0, & i = 1, \ldots, k \end{cases}$$

Optimisation problem (11) allows to determine the optimal values of the weights for active management, $\tilde{w}^*$. Then, budgetary weights w and active fractions γ can be determined such that (1), (2), and (3) are satisfied with $\tilde{w}_i^* = \gamma_i w_i$ for $i = 1, \ldots, k$.

Robust and consistent risk allocation

Determination of the optimal active weights as in optimisation problem (11) obviously requires the mean vector μ and the covariance matrix $\mathbf{\Omega}$ to be known. In practice, however, these parameters would typically be unknown. The optimal portfolio would then be determined on the basis of estimators such as for example the maximum likelihood estimator. There, however, exists extant evidence in the literature that mean–variance optimisation can be extremely sensitive to parameter misestimation (see eg Best and Grauer, 1991; Chopra and Ziemba, 1993; Michaud, 1998; Ceria and Stubbs, 2006). This implies that when the optimal weights would be determined on the basis of the maximum likelihood estimators, these weights could be highly suboptimal with respect to the true parameter values, and could potentially lead to significant differences in the returns distribution. We argue that this problem is particularly relevant in a context where optimal weights for active management need to be determined, because it is unlikely that, for a given investment manager, data would be available for a large number of years. We show how the approach developed by Goldfarb and Iyengar (2003) can be applied to obtain robust optimal active risk budgets. Then a risk aversion parameter λ remains to be specified. We determine the parameter such that overall risk management, active and passive, is consistent with the long-term strategic goals of the pension fund, as reflected by the benchmark weights.

The robust optimisation problem

To mitigate the sensitivity of optimisation outcomes to parameter misestimation, Ben-Tal and Nemirovski (1998) propose to introduce *uncertainty sets* for the unknown parameters, and to solve the robust counterpart of the optimisation problem. This robust counterpart optimises the worst-case value of the objective function over all parameter values in the uncertainty set. Formally, in the case of mean–variance optimisation, one would determine a set $M \subset \mathbb{R}^{k+1}$ of potential values of the mean vector μ, and a set $S \subset \mathbb{R}^{(k+1)\times(k+1)}$ of potential values for the covariance matrix $\mathbf{\Omega}$, and solve the following optimisation problem:

$$\tilde{w}^* \in \arg\max_{\tilde{w}} \min_{(\mu, \Omega) \in M \times S} \left\{(1, \tilde{w}') \cdot (\mu - cost) - \lambda(1, \tilde{w}') \cdot \mathbf{\Omega} \cdot \begin{pmatrix} 1 \\ \tilde{w} \end{pmatrix}\right\}$$

$$\text{s.t.} \begin{cases} \sum_{i \in I(j)} \tilde{w}_i \leq \bar{w}_j, & j = 1, \ldots, l, \\ \tilde{w}_i \geq 0, & i = 1, \ldots, k \end{cases} \quad (12)$$

The uncertainty sets M and S then remains to be determined. Typically, one would like these sets to satisfy the following two requirements:

1. provide an adequate description of uncertainty and
2. yield efficiently computable solutions.

In this section, we show how the approach developed by Goldfarb and Iyengar (2003) to obtain uncertainty sets that satisfy these criteria can be applied to obtain robust optimal active risk budgets.

Let $\hat{\mu}$ and $\hat{\Omega}$ denote the maximum likelihood estimators of the mean vector and the covariance matrix, respectively, and assume that $\hat{\Omega}$ is positive definite. Let us further assume that the returns of the underlying benchmark and excess returns of the k active investment managers follow a multivariate normal distribution with unknown parameters $\bar{\mu}$ and $\bar{\Omega}$, that is

$$(\bar{R}_{bm},\, R_1^e, \ldots, R_k^e) \sim N_{k+1}(\bar{\mu},\, \bar{\Omega})$$

Following Goldfarb and Iyengar (2003), we suggest the following choice for the uncertainty sets M and S[1]:

$$M = \{\mu : (\mu - \hat{\mu})' \cdot \hat{\Omega}^{-1} \cdot (\mu - \hat{\mu}) \leq \theta^2\} \quad (13)$$

$$S = \Big\{ \Omega : \Omega = \hat{\Omega} + \Delta \geq 0, \Delta = \Delta', \\ \| \hat{\Omega}^{-1/2} \cdot \Delta \cdot \hat{\Omega}^{-1/2} \| \frac{\beta}{1-\beta} \Big\} \quad (14)$$

This particular choice of uncertainty sets indeed satisfies the above-mentioned criteria:

1. *Adequate description of uncertainty*: The parameters θ and β can be chosen such that the probability that the true mean vector $\bar{\mu}$ and covariance matrix $\bar{\Omega}$ are elements of their corresponding uncertainty sets is sufficiently high. Indeed, let $\theta \in \mathbb{R}$ and $\beta \in [0, 1)$ be such that

$$F_{\chi^2}(\theta^2) = \alpha_M \quad (15)$$

$$F_\Gamma(1+\beta) - F_\Gamma(1-\beta) = \sqrt[k+1]{\alpha_S} \quad (16)$$

where F_{χ^2} denotes the CDF of a χ^2_{k+1} distributed random variable and F_Γ denotes the CDF of a $\Gamma(T+1/2,\, 2/T-1)$ distributed random variable. Then (see Goldfarb and Iyengar, 2003)

$$P(\bar{\mu} \in M) = \alpha_M \quad (17)$$

$$P(\bar{\Omega} \in S) \geq \alpha_S \quad (18)$$

Thus, the probability that the mean–variance objective with respect to the true means $\bar{\mu}$ and covariances $\bar{\Omega}$ is lower than the robust mean variance objective with respect to the estimated parameters $\hat{\mu}$ and $\hat{\Omega}$ is at most $2-\alpha_M-\alpha_S$.

2. *Efficiently computable solutions*: Optimisation problem (12) is equivalent to:

$$\tilde{w}^* \in \arg\max_{\tilde{w}} \left\{ (1, \tilde{w}') \cdot (\hat{\mu} - cost) - \theta \sqrt{(1, \tilde{w}') \cdot \hat{\Omega} \cdot \begin{pmatrix} 1 \\ \tilde{w} \end{pmatrix}} - \frac{\lambda}{1-\beta} (1, \tilde{w}') \cdot \hat{\Omega} \cdot \begin{pmatrix} 1 \\ \tilde{w} \end{pmatrix} \right\} \quad (19)$$
$$\text{s.t.} \begin{cases} \sum_{i \in I(j)} \tilde{w}_i \leq \bar{w}_j, & j = 1, \ldots, l, \\ \tilde{w}_i \geq 0, & i = 1, \ldots, k \end{cases}$$

Numerical solutions for the optimal robust allocation of active weights $\tilde{w}^*$ can therefore be computed efficiently using SOCP (Second Order Cone) optimisation software (see eg Sturm, 1999).

Note that it follows from (19) that robustness with respect to uncertainty in the mean vector (θ>0) is achieved by reducing the expected return with a factor that is proportional to the standard deviation of the return. Similarly, robustness with respect to uncertainty in the covariance matrix (β>0) is achieved by scaling the variance of the return by a factor $1/1-\beta>1$. This implies that increasing the level of robustness is equivalent to performing mean-variance optimisation with a higher degree of risk aversion.

Consistency in terms of risk-return trade-off

In order to determine the robust optimal active weights in optimisation problem (19), a risk aversion parameter λ needs to be specified. It is clearly important that the active and passive implementation of the benchmark strategy is consistent, in terms of risk-return trade-off, with the long-term strategic goals of the pension fund as reflected by the benchmark weights. We therefore require the risk aversion parameter λ used in (19) to represent as closely as possible the degree of risk aversion consistent with the benchmark weights. To achieve this, we determine the risk aversion parameter for which the *optimal mean–variance benchmark weights* are as close as possible (in Euclidian distance) to the *actual benchmark weights* $\bar{w} = (\bar{w}_1, \ldots, \bar{w}_l)$. Determining this ALM-consistent risk aversion parameter thus requires solving the following optimisation problem:

$$\min_{\lambda \in \mathbb{R}} U^c(\lambda) = \| \omega^* - \bar{w} \|^2$$
$$\text{s.t.} \begin{cases} \omega^* \in & \arg\max U^{bm}(\omega|\lambda) \\ & \text{s.t.} \begin{cases} \sum_{j=1}^{l} \omega_j = 1 \\ \omega_j \geq 0 \end{cases} \end{cases} \qquad (20)$$

where $\mu^{bm} \in R^l$ and $\Omega^{bm} \in R^{l \times l}$ denote the benchmark returns and covariances, and

$$U^{bm}(\omega|\lambda) = \omega' \cdot \mu^{bm} - \lambda \omega' \cdot \Omega^{bm} \cdot \omega$$

In the technical sense, problem (20) takes the form of a bi-level optimisation problem, since the constraint that the weights ω^* yield the optimal risk-return trade-off is an optimisation problem itself. In general, bi-level optimisation problems can be hard to solve due to the nonconvexity of the feasible area. Efficient algorithms for solving bi-level optimisation problems have, however, recently been developed. See for example Dempe (2002) for an overview.

A numerical illustration

Data

Consider a benchmark that consists of two categories: a US equity index and a US government bond index, with weights 25 and 75 per cent, respectively. The expected returns, standard deviations and covariances are given in Table 1.

This benchmark ALM investment policy yields an expected return of 5.75 per cent with a standard deviation of 8 per cent.

Now consider four active managers: two with respect to the US equity index and two with respect to the bond index. Table 2 presents the true expected excess return (third column), the maximum likelihood estimator of the expected excess return (fourth column) and the proportional cost for active and passive management for each of the four investment managers. Table 3 presents the true covariance matrix of the excess returns of the four managers and the return of the passive benchmark (BM) (left matrix), as well as the maximum likelihood estimator of that covariance matrix (right matrix). The maximum likelihood estimates $\hat{\mu}$ and $\hat{\Omega}$ are based on a sample of 15-years data, simulated from the true multivariate normal distribution with mean $\bar{\mu}$ and covariance matrix $\bar{\Omega}$.

ALM-consistent risk aversion parameter

We now first need to determine the ALM-consistent risk aversion parameter, that is the bi-level optimisation problem (20) needs to be solved for $\bar{w} = (0.25, 0.75)$, and μ^{bm} and Ω^{bm} as given in Table 1. Figure 1 displays $U^c(\lambda)$, that is the Euclidian distance between the actual benchmark weights $\bar{w}$ and the mean–variance optimal benchmark weights given the risk aversion parameter λ. Since this objective function is well behaved, the minimum can be found through an enumerative procedure.

Table 1 Benchmark strategy

Asset class	Weight (%)	Expected return (%)	Covariances	
Equity	25	8	279	25
Bonds	75	5	25	64

Table 2 Expected excess returns (true and estimated) and proportional costs

Manager	Category	$\bar{\nu}$ (%)	$\hat{\nu}$ (%)	*a* (%)	*b* (%)
1	Equity	1.5	1.01	1	0.3
2	Equity	1.0	0.77	0.6	0.3
3	Bonds	0.75	0.70	0.45	0.25
4	Bonds	0.50	0.30	0.5	0.25
BM		5.75	5.90		

Table 3 Covariance matrix (true and estimated)

$\bar{\Omega}_{i,j}$	1	2	3	4	BM	$\hat{\Omega}_{i,j}$	1	2	3	4	BM
1	25	8	3	1.3	−1	1	13	3	1	0	−1
2	8	16	1.2	3	16	2	3	19	−1	5	17
3	3	1.2	9	4.5	7.2	3	1	−1	9	4	2
4	1.3	3	4.5	6.3	0	4	0	5	4	8	−2
BM	−1	16	7.2	0	64	BM	−1	17	2	−2	68

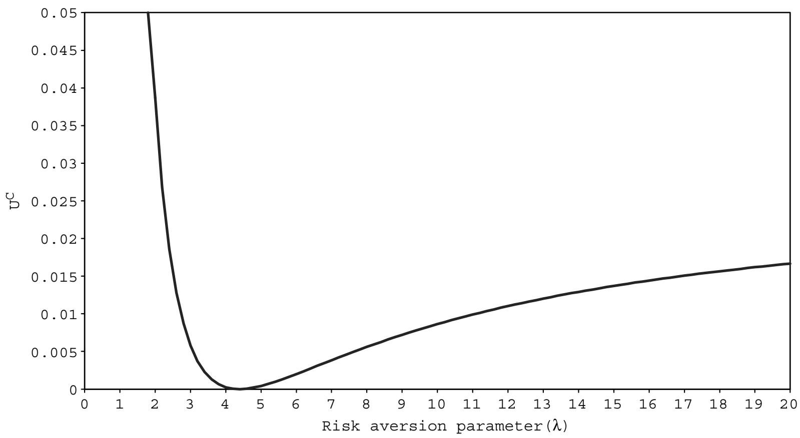

Figure 1 The function $U^c(\lambda)$

The optimum of the bi-level optimisation problem is reached at $\lambda = 4.4$, with corresponding optimal weights $\omega^* = (0.249, 0.751)$. These weights minimise the Euclidian distance to the benchmark weights $\bar{w}$ over all mean–variance optimal weights corresponding to a non-negative risk aversion parameter λ.

Robust vs nonrobust optimisation

Given the true parameters $\bar{\mu}$ and $\bar{\Omega}$ in Tables 2 and 3, the optimal consistent allocation of active weights $\tilde{w}$ can be determined by solving optimisation problem (11) for the ALM-consistent risk aversion parameter $\lambda = 4.4$. The true mean and covariance matrix is, however, typically not known, and the optimal allocations would be determined on the basis of the maximum likelihood estimators. As noted before, mean–variance optimisation is highly sensitive to parameter misestimation. We therefore also analyse the optimal allocation of active and passive weights with respect to the maximum likelihood estimators $\hat{\mu}$ and $\hat{\Omega}$ when a certain degree of robustness is required, and compare it to the optimal solution that is obtained when robustness is not taken into account. Let us first consider the case where the degrees of robustness as defined in (15) and (16) equal $\alpha_M = \alpha_S = 0.8$. Table 4

Table 4 Optimal allocation for $\bar{\mu}$ and $\bar{\Omega}$, nonrobust and robust allocation for $\hat{\mu}$ and $\hat{\Omega}$.

Manager	$\tilde{w}^*$ (%)	$\tilde{w}^*_{nr}$ (%)	$\tilde{w}^*_r$ (%)
1	5	25	9
2	0	0	0
3	0	31	0
4	18	17	25
Aggregate	23	73	34

Table 5 Comparison of optimal nonrobust and robust solutions

	True ($\tilde{w}$=$\tilde{w}^*$) (%)	Nonrobust ($\tilde{w}$=$\tilde{w}_{nr}^*$) (%)	Robust ($\tilde{w}$=$\tilde{w}_r^*$) (%)	BM ($\tilde{w}$=0) (%)
$E[R_{tot}]$	5.94	6.65	5.98	5.75
$\sigma[R_{tot}]$	8.00	8.47	8.01	8.00

presents the optimal allocation $\tilde{w}^*$ for the true parameters, as well as the optimal nonrobust allocation $\tilde{w}^*_{nr}$ from (11), and the optimal robust allocation $\tilde{w}^*_r$ from (19) for the estimated parameters $\hat{\mu}$ and $\hat{\Omega}$.

Comparison of the nonrobust and the robust allocation as given in Table 4 yields the following observations:

- The nonrobust allocation overallocates 50 per cent active weight (73 vs 23 per cent). The robust allocation overallocates only 11 per cent active weight (34 vs 23 per cent).
- The robust solution allocates relatively more active weight to the fourth manager and relatively less active weight to the first manager, because the former provides more diversification with the benchmark.

These observations are related to the fact that requiring a certain degree of robustness is equivalent to increasing the level of risk aversion, as can be seen from (19). As a consequence, diversification among active managers and with respect to the benchmark becomes more important.

Table 5 compares the expected total return $E[R_{tot}] = (1, \tilde{w}') \cdot (\bar{\mu} - cost)$ and the standard deviation of the total return $\sigma[R_{tot}] = \sqrt{(1, \tilde{w}') \cdot \bar{\Omega} \cdot \left(\begin{smallmatrix} 1 \\ \tilde{w} \end{smallmatrix}\right)}$ given the true mean and the true covariance matrix of the three optimal allocations given in Table 4, and of the benchmark allocation, that is $\tilde{w} = 0$ (last column).

Note first that the true optimal allocation $\tilde{w}^*$ outperforms the benchmark, because active management allows to increase the expected return without increasing the variance. The nonrobust allocation for the estimated parameters, $\tilde{w}^*_{nr}$, overallocates active weight, and leads to an allocation that is too risky. Since the robust allocation $\tilde{w}_r^*$ overallocates active weight to a lesser extent, the corresponding performance is much closer to the optimal one.

Sensitivity analysis

The above comparison between the nonrobust and the robust allocation already indicates that robust optimisation implies that there is more emphasis on diversification, and, as a consequence, the optimal allocation would typically be less risky. Clearly, this effect will depend to a large extent on the degrees of robustness α_M and α_S. Figure 2 displays the expected total return $(1, \tilde{w}') \cdot (\bar{\mu} - cost)$ (right axis, solid line) and the total active fraction $\sum_{i=1}^{k} \tilde{w}_i$ (left axis, dashed line) for the optimal robust allocation, as a function of α_M, the level of robustness with respect to errors in the expected excess returns. The level of robustness with respect to uncertainty in the covariance matrix is set equal to $\alpha_S = 0$.

Increasing the level of robustness with respect to errors in the mean vector leads to a lower expected total return and less active weight. Note the sharp decrease in active weight and in expected return when the robustness level α_M is slightly increased from 0 to 5 per cent.[2] Introducing a modest level of robustness (a value of α_M close to 0) implies non-negligible changes in the expected return. This illustrates that in our setting, contrary to the setting described in Kritzman (2006), the effect of parameter

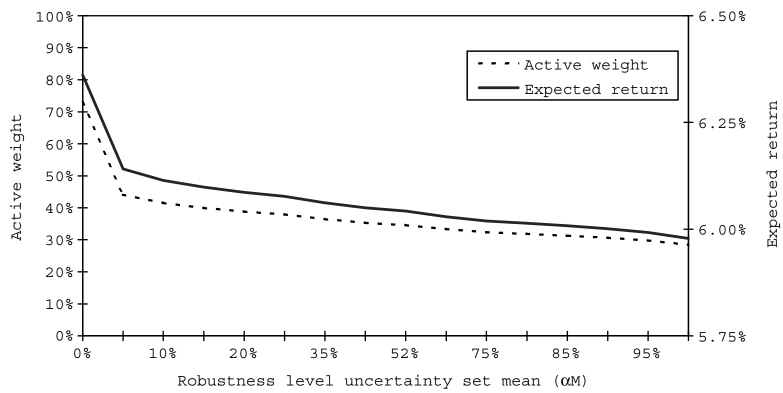

Figure 2 Sensitivity with respect to robustness-level mean vector

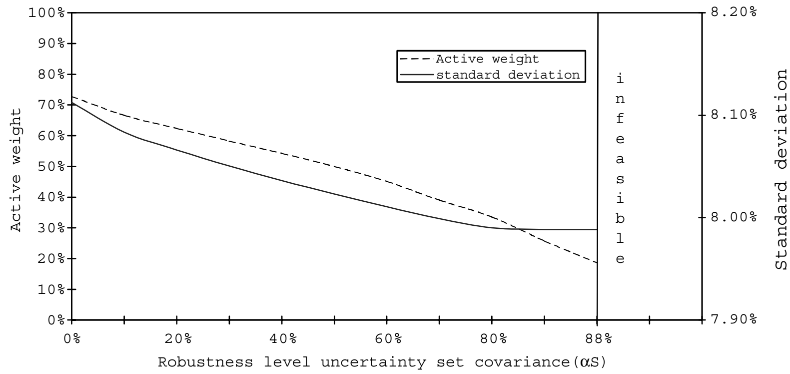

Figure 3 Sensitivity with respect to robustness-level covariance matrix

misestimation on the distribution of the returns can be substantial.

Figure 3 displays the standard deviation of the total return $\sqrt{(1,\ \tilde{w}')\cdot\bar{\Omega}\cdot\left(\begin{smallmatrix}1\\ \tilde{w}\end{smallmatrix}\right)}$ (right axis, solid line) and the total active fraction $\sum_{i=1}^{k}\tilde{w}_i$ (left axis, dashed line) for the optimal robust allocation, as a function of α_S, the level of robustness with respect to errors in the covariance matrix. The level of robustness with respect to errors in the mean vector is set equal to $\alpha_M = 0$. An increasing robustness level yields a more risk-averse attitude and therefore less active weight. The 'infeasible' area in Figure 3 indicates the levels of robustness that are not possible to achieve.[3] With four managers and a data sample of 15 years, the maximal achievable robustness level equals $\alpha_S = (F_\Gamma(2))^5 = 88$ per cent. The feasible area is smaller in case of small sample sizes and/or a large number of available managers. For example, a historical sample of 10 years (and four managers) implies a robustness confidence level of 71 per cent. If also a fifth manager is available (and a sample of 10 years), the maximum robustness level for uncertainty with respect to the covariances is approximately 85 per cent.

Performance of robust optimum

The example shows that robust allocation typically leads to lower weights for active management. Question remains however on whether the optimal allocation obtained through robust optimisation outperforms the one obtained through nonrobust optimisation. The answer to this question obviously depends on the performance criterion. We illustrate this in Figure 4,

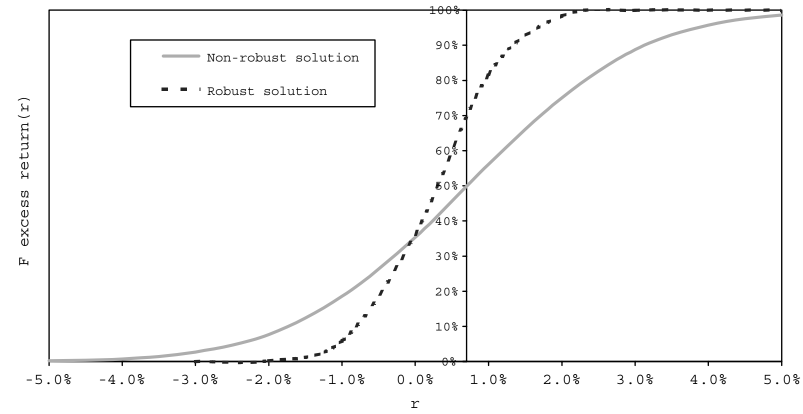

Figure 4 Cumulative excess return distribution for nonrobust optimum (solid) and robust optimum (dashed)

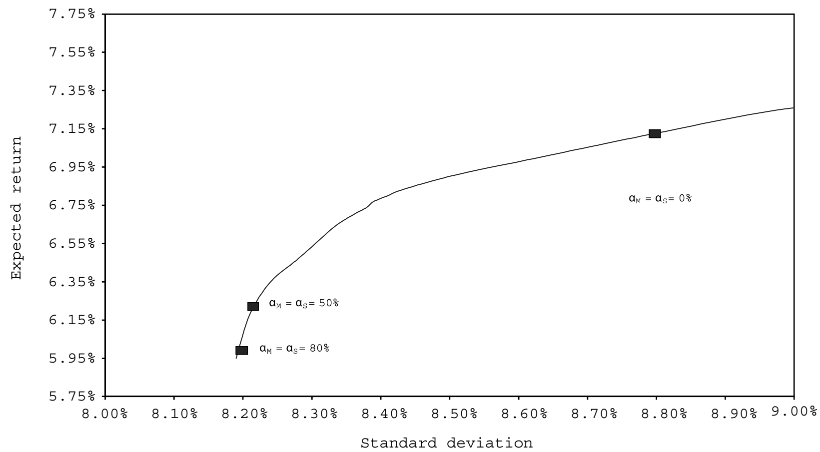

Figure 5 Efficient frontier for $\hat{\mu}$ and $\hat{\Omega}$

which yields the simulated cumulative probability distribution with respect to the true mean $\bar{\mu}$ and covariance matrix $\bar{\Omega}$ of the return generated by the nonrobust optimal weights (solid line) and by the robust optimal weights (dashed line). We see that robust optimisation implies that the probability of unfavourable returns is lower than with nonrobust optimisation, and the converse holds for very favourable returns.

Finally, we illustrate how robust optimisation affects the location of the optimal risk-return outcome on the efficient frontier.

Figure 5 displays the efficient frontier of active and passive management. First, note that the combination of the benchmark return (8 per cent) and standard deviation (5.75 per cent) is not an element of the efficient frontier, because allowing for active management yields mean–variance solutions that dominate the benchmark. The figure illustrates how the level of robustness affects the trade-off between return and risk. The boxes correspond to consistent and robust optimal risk budgeting solutions calculated with the ALM-consistent risk-aversion parameter $\lambda = 4.4$, for different robustness levels α_M and α_S. Higher levels of robustness lead to less-risky active allocations and a lower expected total return.

Conclusion

We developed a framework that allows optimal robust and consistent active implementation of a pension fund's given benchmark investment strategy. The mean–variance optimisation approach incorporates weights for active as well as passive management, and takes into account the different fees for active and passive management. In order to guarantee consistency between the risk-return trade-offs of the benchmark and of the active investment decisions, we choose the level of risk aversion to be as close as possible to the level of risk aversion consistent with the benchmark weights. Then, we draw on the recent literature on robust optimisation to determine optimal active implementations that are less sensitive to parameter misestimation. We find that, because requiring a certain degree of robustness is equivalent to increasing the level of risk aversion, robust optimal active implementation would typically lead to a significantly lower weight for active management.

Notes

1. The norm $\|A\|$ in (14) equals $\| A \| := \sqrt{\sum_i \lambda_i^2(A)}$, where $\lambda_i(A)$ denote the eigenvalues of A. Moreover, we use the notation $A \geq 0$ to indicate that A is positive semi-definite.
2. It follows from (19) that robustness implies that the expected return is reduced with a factor θ times the standard deviation of the return, where θ^2 satisfies $\theta^2 = (F_{\chi^2})^{-1}(\alpha_M)$. It can be verified that (F_{χ^2}) features a sharp increase for low values.
3. Since in (14) β is required to be strictly less than 1, it follows from (16) that the maximum level of robustness with respect to uncertainty in the covariance matrix equals $\alpha_S = (F_\Gamma(2))^{k+1}$.

References

Baierl, G. T. and Chen, P. (2000) 'Choosing Managers and Funds', *The Journal of Portfolio Management*, 26(2), 47–53, 264.

Ben-Tal, A. and Nemirovski, A. (1998) 'Robust Convex Optimization', *Mathematics of Operations Research*, 23, 769–805.

Best, M. J. and Grauer, R. R. (1991) 'On the Sensitivity of Mean–Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results', *Review of Financial Studies*, 4, 315–342.

Boender, G. (1997) 'A Hybrid Simulation/Optimisation Scenario Model for Asset/Liability Management', *European Journal of Operational Research*, 99, 126–135.

Ceria, S. and Stubbs, R. A. (2006) 'Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction', *Journal of Asset Management*, 7(2), 109–127.

Chopra, V. K. and Ziemba, W. T. (1993) 'The Effect of Errors in Means, Variances and Covariances on Optimal Portfolio Choice', *The Journal of Portfolio Management*, 19, 6–11.

Dempe, S. (2002) *Foundations of Bilevel Programming*, Kluwer Academic Publishers, Dordrecht.

Goldfarb, D. and Iyengar, G. (2003) 'Robust Portfolio Selection Problems', *Mathematics of Operations Research*, 28, 1–38.

Kritzman, M. (2006) 'Portfolio Optimization — Are Optimizers Error Maximizers?' *The Journal of Portfolio Management*, 32(4), 66–69.

Michaud, R. O. (1998) *Efficient Asset Management: A Practical Guide to Stock Portfolio Management and Asset Allocation*, Financial Management Association, Survey and Synthesis Series, Harvard Business School Press, Boston, MA.

Roll, R. (1992) 'A Mean-Variance Analysis of Tracking Error', *The Journal of Portfolio Management*, 18, 13–22.

Scherer, B. (2002) *Portfolio Construction and Risk Budgeting*, Risk Books, New York.

Schöttle, K. and Werner, R. (2006) 'Towards Reliable Efficient Frontiers', *Journal of Asset Management*, 7, 128–141.

Sturm, J. (1999) 'Using SeDuMi 1.02, a Matlab Toolbox for Optimization Over Symmetric Cones', *Optimization Methods and Software*, 11–12, 625–653.

Waring, B., Whitney, D., Pirone, J. and Castille, C. (2000) 'Optimizing Manager Structure and Budgeting Manager Risk', *The Journal of Portfolio Management*, 26(3), 90–104.

