

Optimizing Omega

S. J. Kane · M. C. Bartholomew-Biggs · M. Cross ·
M. Dewar

Received: 11 December 2008 / Accepted: 17 December 2008 / Published online: 7 January 2009
© Springer Science+Business Media, LLC. 2009

Abstract This paper considers the Omega function, proposed by Cascon, Keating & Shadwick as a performance measure for comparing financial assets. We discuss the use of Omega as a basis for portfolio selection. We show that the problem of choosing portfolio weights in order to maximize Omega typically has many local solutions and we describe some preliminary computational experience of finding the global optimum using a NAG library implementation of the Huyer & Neumaier MCS method.

Keywords Portfolio selection · Omega function · MCS method

1 Introduction

The Omega function was introduced in [5, 11] as a measure for comparing the performance of financial assets. A brief description is as follows. Suppose we have an m -day history of returns, r , for an asset and that the observed returns lie in a range $r_{\min} \leq r \leq r_{\max}$. The cumulative distribution of returns will have a form like the one shown in Fig. 1 in which the horizontal axis represents the observed returns and the vertical scale shows $f(r) = P[\text{asset return} < r]$.

In Fig. 1, the point **A** represents some threshold value for return r_t . The Omega function associated with r_t is defined as

$$\Omega(r_t) = \frac{\int_{r_t}^{r_{\max}} (1 - f(r)) dr}{\int_{r_{\min}}^{r_t} f(r) dr} = \frac{\text{Area BCU}}{\text{Area LAB}}. \quad (1)$$

If r_t is close to r_{\min} then area **BCU** is much larger than area **LAB** and so Ω is large. Conversely, $\Omega \rightarrow 0$ as $r_t \rightarrow r_{\max}$. (Cascon et al. [5] show that $\Omega = 1$ when r_t is the

S. J. Kane · M. C. Bartholomew-Biggs (✉)
School of Physics Astronomy and Mathematics, University of Hertfordshire, Hatfield AL10 9AB, UK
e-mail: M.Bartholomew-Biggs@herts.ac.uk

M. Cross · M. Dewar
Numerical Algorithms Group, Oxford, UK

Fig. 1 Typical cumulative distribution of returns on an asset





Fig. 2 $\Omega(r_t)$ for $0.0\% \leq r_t \leq 0.65\%$




mean return for the distribution.) Hence the size of Ω is a measure of the extent to which the historical performance of an asset has exceeded the threshold return r_t . In particular, Keating and Shadwick in [11] describe the Omega function as a probability adjusted ratio of gains to losses and say that, for a given threshold r_t , the simple rule of preferring more to less implies that an asset with a high value of Omega is a better investment than one with a lower value.

Figure 2 shows how Ω typically varies with r_t .

Ω curves like the one in Fig. 2 can be used to compare several assets. Figure 3 plots Ω for three assets (referred to as A–C) when r_t is in the range 0.4–0.8%. Using the argument of Keating and Shadwick [11], we can say that Fig. 3 shows that asset A is always to be preferred to asset B because its Ω value is higher for all values of threshold return. Moreover, if the threshold return is less than about 0.6% then asset A is also better than asset C. However this situation changes as the threshold is increased and asset C becomes the best investment when target return is between 0.6% and 0.8%.

Fig. 3 $\Omega(r_t)$ plots for assets A–C



2 Using Ω for portfolio optimization

In this paper we regard Ω as more than a means of comparing individual assets and use it as a basis for portfolio selection. Consider a portfolio involving n assets in which y_1, y_2, \dots, y_n denote the invested fractions (or weights) which specify what proportion of the investment is assigned to each asset. Suppose also that we have an m -day history of asset returns $r_{ji}, j = 1, \dots, n, i = 1, \dots, m$. Then, for any values of the y_j we can use this historical data to determine how the corresponding portfolio would have performed over the previous m days. Specifically, the portfolio return on day i would have been

$$R_i(y) = \sum_{j=1}^n r_{ji} y_j \quad \text{for } i = 1, \dots, m. \tag{2}$$

From these values of R_i we can construct a cumulative return distribution similar to that in Fig. 1; and then for any threshold return r_t we can evaluate $\Omega(r_t)$. This Ω value obviously depends on y_1, \dots, y_n and so a possible way of optimizing the portfolio would be to choose the invested fractions to *maximize* $\Omega(r_t)$. For practical reasons the invested fractions must of course satisfy the constraint

$$\sum_{j=1}^n y_j = 1. \tag{3}$$

(We might also want to impose other conditions such as bounds on the y_j).

Maximizing Ω represents a significantly different portfolio selection strategy from the more familiar Markowitz approach [12] which involves choosing the invested fractions to *minimize* a measure of portfolio risk. Specifically, if \bar{r}_j denotes the mean of the values $r_{ji}, i = 1, \dots, m$ and if Q is the variance-covariance matrix for the historical data $r_{ji}, j = 1, \dots, n, i = 1, \dots, m$ then the expected portfolio return is

$$\bar{R}(y) = E(R_i(y)) = y^T \bar{r} \tag{4}$$

and portfolio risk can be defined as

$$V(y) = y^T Q y. \tag{5}$$

Hence an extremely risk-averse investor could choose a portfolio simply by minimizing $V(y)$ subject to the constraint (3). This strategy ignores return altogether and, more realistically, a portfolio with low risk and high return could be obtained by minimizing a composite function such as

$$-\lambda y^T \bar{r} + (1 - \lambda) y^T Q y$$

where λ is a parameter between 0 and 1 which controls the balance between return and risk. Alternatively, a portfolio could be chosen to maximize the *Sharpe ratio*

$$S = \frac{y^T \bar{r} - r_t}{y^T Q y}. \quad (6)$$

Clearly the maximum of S can be expected to yield a portfolio which gives low risk together with an expected return which significantly exceeds the threshold r_t . Finally, if we want to set a target value for portfolio return, we could minimize $V(y)$ subject to (3) and an additional constraint such as

$$\sum_{j=1}^n \bar{r}_j y_j = R_p = \text{a specified expected portfolio return}. \quad (7)$$

The use of Omega in portfolio selection has been proposed by a number of authors—see for instance [1, 6, 7, 14]—and the theoretical properties of Omega have been quite widely discussed. It is argued in [5] that Omega differs in an important way from other statistical estimators in that it is calculated directly from the historical data and hence it can be seen as equivalent to the actual returns distribution. This is in contrast to risk measures based only on the mean and variance which may not capture all the features of the data. Avouyi-Dovi et al. [1] agree with this observation and, in particular, they compare Omega with the Sharpe ratio, remarking that the use of expression (5) for risk in (6) is based on the prior assumption that the asset data is normally distributed. No such assumption is made in the definition of Omega and hence it should give a better representation of data with a non-normal distribution—e.g., asymmetric data—which can frequently be encountered. When dealing with non-normal distributions, Keating and Shadwick [11] say that the simplicity of the calculation of Ω gives it an advantage over more sophisticated statistical measures involving estimation of higher order moments.

Another advantage of Ω pointed out by Avouyi-Dovi et al. [1] is that the way r_t contributes to the definition (1) means that the choice of threshold value can be a useful way of taking into account an investor's preferences about loss and gain. Mausser et al. [14] agree with this general remark about the role of r_t ; but they also observe that there may be situations in which it is not obvious how to select a specific threshold value. They suggest that it may sometimes be advantageous to maximize Ω for a *range* of thresholds in order to combine the results in a more sophisticated higher-level portfolio selection process.

In later sections we shall consider the optimization methods used and results obtained by other authors when maximizing Omega. First however we describe our own approach to this problem.

3 Max- Ω portfolio selection—a simple example

Figure 4 shows some artificial 50-day return histories for three assets and Fig. 5 shows their cumulative distributions.




Fig. 4 Return histories for assets 1–3 (left to right)




Fig. 5 Cumulative distribution of returns for assets 1–3 (left to right)

In order to construct a maximum- Ω portfolio from these assets we can use the fact that $y_1 + y_2 + y_3 = 1$ and calculate returns (2) as functions of y_1 and y_2 only, i.e.,

$$R_i(y) = y_1 r_{1i} + y_2 r_{2i} + (1 - y_1 - y_2) r_{3i}, \quad i = 1, 2, \dots, 50.$$

We can obtain $\Omega(r_t)$ from a cumulative distribution of R_1, \dots, R_{50} . We wish to exclude solutions which involve short-selling and so we want to maximize $\Omega(r_t)$ subject to the constraints

$$y_1 \geq 0, \quad y_2 \geq 0, \quad 1 - y_1 - y_2 \geq 0. \tag{8}$$

In order to solve this inequality constrained problem we can apply an *unconstrained* minimization algorithm to a penalty function

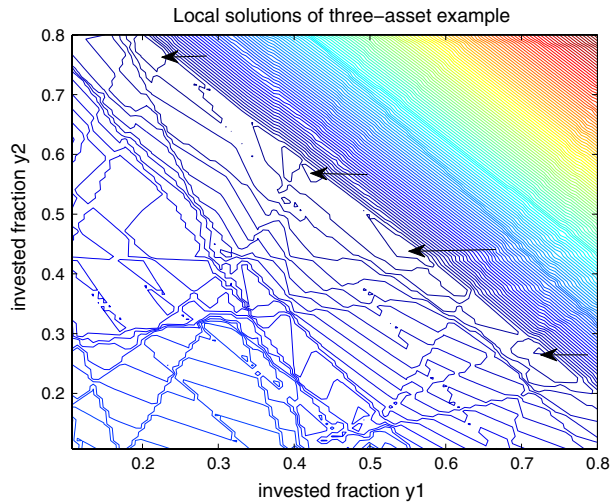
$$F(y_1, y_2) = \frac{1}{\Omega(r_t)} + \rho\{|\min(0, y_1)| + |\min(0, y_2)| + |\min(0, 1 - y_1 - y_2)|\} \tag{9}$$

where ρ is a positive weighting parameter. It is well-known (see for instance [3]) that the minimum of the exact penalty function (9) coincides with the solution of the original inequality constrained problem provided ρ is chosen sufficiently large.

For this demonstration example we take the threshold $r_t = 0$ and set $\rho = 1$. The contours of (9) are then as shown in Fig. 6. The jagged nature of these contours is due to the fact that the cumulative distribution functions in Fig. 5 are non-smooth and hence the numerical integral involved in computing Ω does not change smoothly as the invested fractions change. In the upper part of the figure the third constraint penalty term in (9) becomes active and there is a further discontinuity of slope along the line $y_1 + y_2 = 1$.

The contour lines in Fig. 6 suggest that there are a number of possible minima some of which are in the regions indicated by arrows. We can seek to locate these minima precisely by minimizing (9) using the Nelder and Mead Simplex algorithm [15] as implemented

Fig. 6 Contours of (9) when $r_t = 0.0$



in the MATLAB procedure `fminsearch` [13]. This direct search technique is a suitable method since the objective function $F(y)$ is non-differentiable. Results from the simplex method are consistent with Fig. 6 in that different local minima are found when the search is started from different initial guesses. The best solution is found with $y_1 \approx 0.57$, $y_2 \approx 0.43$ giving $\Omega \approx 5.49$. However the other apparent minima are not artefacts of the contour-plotting process and alternative solutions include $y_1 \approx 0.41$, $y_2 \approx 0.59$ (with $\Omega \approx 5.37$) and $y_1 \approx 0.44$, $y_2 \approx 0.56$ (with $\Omega \approx 5.27$).

It is interesting to note that all these solutions are feasible points with respect to the constraints (8). The fact that y_3 is always set to zero is easily explained when we consider the value of $\Omega(0.0)$ for the three assets individually. For asset one $\Omega(0.0) \approx 3.6$; for asset two $\Omega(0.0) \approx 4.2$ and for asset three $\Omega(0.0) \approx 2.2$. Hence it is not surprising that asset three does not contribute to a portfolio designed to maximize Ω .

The observation that (9) has many local minima immediately makes portfolio selection by maximizing Ω a more challenging problem. The non-convexity and non-smoothness of Ω has also been noted in [1, 7, 14]. Hence the calculation of a Max- Ω portfolio is more difficult than the Markowitz approaches. Minimizing (5) subject to constraints (3) and (7) is a convex quadratic programming problem while minimizing the Sharpe ratio (6) subject to (3) is a nonconvex problem but the function and constraint are both differentiable.

If we wish to find a maximum- Ω portfolio by minimizing a penalty function like (9) then the simplex method is not really suitable since it is only a local optimization procedure. We might seek a global optimum of F by using a *multistart* technique which would essentially involve running `fminsearch` from a large number of randomly distributed starting guesses. This seems unlikely to be an efficient process for a portfolio involving a large number of assets. Other approaches to the global optimization problem that have been proposed can be divided into two types: stochastic and deterministic. *Stochastic* approaches involve some element of random sampling of the space of the variables together with rules for accepting or rejecting new trial points. Examples are Simulated Annealing, Genetic Algorithms and Tabu Search. *Deterministic* algorithms make a more systematic exploration based on information accumulated about the objective function on previous iterations and among methods of this type we mention DIRECT [10] and the MCS algorithm [9] which is described in the next section.

Avouyi-Dovi et al. [1,2] use a non-smooth penalty function like (9) to solve the maximum- Ω problem and they employ a technique called Threshold Acceptance which they characterize as a deterministic version of Simulated Annealing. This method is described by Gilli et al. [8] and is also used by them for Omega-based portfolio optimization in [7]. Mausser et al. [14] prefer to consider the maximum- Ω problem in its constrained form, rather than reducing it to a penalty function minimization. They propose avoiding the difficulties of dealing with multiple solutions of this nonconvex problem by transforming it into a much simpler linear programming calculation. This only yields a solution to the original problem under certain circumstances which are fully discussed in [14].

We have a preference for deterministic rather than stochastic methods for global optimization. This preference is largely based on a study [4] in which DIRECT [10] was found to perform much better than two tabu search methods on a route-finding problem. DIRECT works by repeated splitting of the region of search into smaller hyperboxes, using the accumulated record of function values at points sampled so far to decide which of the current hyperboxes seem most promising candidates for further subdivision. The Multi-level Coordinate Splitting (MCS) method [9] works in a broadly similar manner and was our choice for seeking global solutions to the Max- Ω problem because of the ready availability of a beta test implementation [16] from the NAG MATLAB toolbox [17]. An outline of the MCS method is given in the next section.

4 Using the MCS algorithm for maximum- Ω problems

We consider problems involving a n -asset portfolio based on an m -day performance history. As before, we use the relation $\sum_j y_j = 1$ and express daily portfolio returns in terms of invested fractions as

$$R_i = \sum_{j=1}^{n-1} r_{ji} y_j + r_{jn} \left(1 - \sum_{k=1}^{n-1} y_k \right), \quad \text{for } i = 1, \dots, m.$$

From the cumulative distribution function for these R_i we can evaluate $\Omega(r_t)$ for a specified threshold return r_t . We calculate Max- Ω portfolios (in which short-selling is not permitted) by finding y_1, \dots, y_{n-1} to minimize the exact penalty function

$$F(y) = \frac{1}{\Omega(r_t)} + \rho \left\{ \sum_{i=1}^{n-1} |\min(0, y_i)| \right\} + \rho \left| \min \left(0, 1 - \sum_{i=1}^{n-1} y_i \right) \right|. \quad (10)$$

In order to find a global optimum of (10) we use the MCS algorithm [9] implemented as the NAG procedure e05jbb [16].

4.1 The MCS algorithm and its e05jbb implementation

The MCS algorithm seeks a global minimum of an n -variable function $F(x)$ in a hyperbox defined by $l \leq x \leq u$. (In what follows we shall simply use the term *box* when strictly we mean an n -dimensional hyperbox.)

MCS searches for a global minimizer using branching recursively in order to divide the search space in a nonuniform manner. It divides, or *splits*, the root box $[l, u]$ into smaller sub-boxes. Each sub-box contains a basepoint at which the objective function is sampled. The splitting procedure biases the search in favour of sub-boxes where low function values are

expected. The global part of the algorithm explores sub-boxes that enclose large unexplored territory, while the local part splits sub-boxes that have good function values.

A balance between the global and local parts of the method is achieved using a multilevel approach, where every sub-box is assigned a level $s \in \{0, 1, \dots, s_{\max}\}$. The value of s_{\max} can be specified by a user. A sub-box with level 0 has already been split; a sub-box with level s_{\max} will be split no further. Whenever a sub-box of intermediate level $0 < s < s_{\max}$ is split a descendant will be given the level $s + 1$ or $\min(s + 2, s_{\max})$. The child with the better function value is given the larger fraction of the splitting interval because then it is likely to be split again more quickly.

An initialization procedure generates a preliminary set of sub-boxes, using points input by the user or derived using a default generation procedure. The method ranks each coordinate based on an estimated variability of the objective function, computed by generating quadratic interpolants through the points used in the initialization. Then the algorithm begins sweeping through levels.

Each sweep starts with the sub-boxes at the lowest level, this process being the global part of the algorithm. At each level the sub-box with the best function value is selected for splitting; this forms the local part of the algorithm. A box is split either by *rank* (when it reaches a sufficiently high level; in particular, as $s_{\max} \rightarrow \infty$ this ensures each coordinate is split arbitrarily often) or by *expected gain* (along a coordinate where a maximal gain in function value is expected, again computed by fitting quadratics).

The splitting procedure as a whole is a variant of the standard coordinate search method: MCS splits along a single coordinate at a time, at adaptively chosen points. In most cases one new function evaluation is needed to split a sub-box into two or even three children. Each child is given a basepoint chosen to differ from the basepoint of the parent in at most one coordinate, and safeguards are present to ensure a degree of symmetry in the splits. In the NAG implementation `e05jb` an optional parameter `Local Searches` can be set to 'OFF' and this causes MCS to put the basepoints and function values of sub-boxes of maximum level s_{\max} into a 'shopping basket' of candidate minima. Turning `Local Searches` 'ON' will enable local searches to be started from these basepoints before they go into the shopping basket.

Local searches go ahead providing the basepoint is not likely to be in the basin of attraction of a previously-found local minimum. The search itself uses line searches along directions that are determined by minimizing quadratic models, subject to bound constraints. In particular, triples of vectors are computed using coordinate searches based on line searches. These triples are used in *triple search* procedures to build local quadratic models for the objective, which are then minimized using a trust-region-type approach. The quadratic model need not be positive definite, so it is minimized using a general nonlinear optimizer.

4.2 Calculating maximum Ω portfolios with `e05jb`

We now consider a 10-asset portfolio based on a 100-day history. We can summarize the properties of the 10 assets by showing their mean returns and displaying plots of their Ω functions. The mean returns are

$$\bar{r}_1 = 0.078\%, \bar{r}_2 = -0.066\%, \bar{r}_3 = 0.01\%, \bar{r}_4 = 0.048\%, \bar{r}_5 = -0.07\%$$

$$\bar{r}_6 = 0.044\%, \bar{r}_7 = 0.003\%, \bar{r}_8 = 0.02\%, \bar{r}_9 = 0.040\%, \bar{r}_{10} = -0.045\%.$$

For clarity we give Ω -plots in groups of four assets at a time. Figure 7 shows the four highest Ω curves. Asset 1 has the best Ω value for threshold returns in the range $0 \leq r_t \leq 0.1$.

Fig. 7 Ω plots for example assets 1, 9, 6 and 4





Fig. 8 Ω plots for example assets 4, 7, 3 and 8



Assets 9, 6 and 4 are the next best performers in terms of Ω and they remain consistently ranked in this order although the curves for assets 6 and 4 touch near $r_t = 0.04\%$ and the curves for assets 6 and 9 touch near $r_t = 0.03\%$.

Figure 8 shows Ω plots for Assets 4, 7, 3 and 8. There are some changes of ranking between these assets as r_t increases: asset 8 is the worst performer when r_t is near-zero but almost matches asset 7 in being the best performer of this group as $r_t \rightarrow 0.1\%$.

Finally, Figure 9 shows that assets 2, 10 and 5 have Ω -values which are well below asset 8. Assets 2 and 10 change places once or twice over the range of r_t but asset 5 is uniformly the worst choice.

We now consider the global minimization of the function (10) formed using data from our sample set of 10-assets and with various values for threshold return r_t . Table 1 shows Max- Ω portfolios obtained by applying $\epsilon = 0.5\%$ in the search box defined by $0 \leq y_i \leq 1$ for $i = 1, \dots, 9$. In the column for invested fractions we show only those which are non-zero. For the smaller values of r_t the solutions generally favour assets 1, 4, 6 and 9 which is con-

Fig. 9 Ω plots for example assets 8, 2, 10 and 5




Table 1 Max- Ω portfolios for varying r_t

r_t (%)	Invested fractions	Ω_{\max}
0	$y_1 = 0.147, y_2 = 0.040, y_3 = 0.213, y_4 = 0.208$ $y_6 = 0.286, y_7 = 0.007, y_9 = 0.099$	7.08
0.025	$y_1 = 0.118, y_2 = 0.036, y_4 = 0.225, y_6 = 0.361$ $y_7 = 0.003, y_9 = 0.219$	5.16
0.05	$y_1 = 0.235, y_4 = 0.099, y_6 = 0.357, y_7 = 0.051$ $y_9 = 0.21, y_{10} = 0.046$	4.11
0.075	$y_1 = 0.884, y_5 = 0.035, y_7 = 0.039, y_9 = 0.043$	3.48
0.1	$y_1 = 0.977, y_5 = 0.014, y_8 = 0.009$	2.97

Table 2 Expected return and risk for Max- Ω portfolios

r_t (%)	R_p (%)	V
0	0.038	0.058
0.025	0.040	0.064
0.05	0.045	0.061
0.075	0.068	0.211
0.1	0.075	0.261

sistent with the Ω curves in the graphs in Figs. 7–9. For $r_t = 0.075\%$ and $r_t = 0.1\%$ the portfolios are strongly dominated by asset 1.

Table 2 gives more information about the Max- Ω solutions in terms of their portfolio expected return R_p and risk V . Note that the optimized portfolio expected returns R_p are not the same as the threshold values r_t used in defining Ω .

We now consider how the portfolios in Table 1 compare with those obtained by minimizing the risk associated with the target returns R_p in Table 2.

Table 3 Min- V portfolios giving same expected return as Max- Ω portfolios

R_p (%)	Invested fractions	V_{\min}
0.038	$y_1 = 0.201, y_3 = 0.069, y_4 = 0.115, y_6 = 0.171$ $y_7 = 0.122, y_8 = 0.09, y_9 = 0.191, y_{10} = 0.039$	0.035
0.040	$y_1 = 0.214, y_3 = 0.065, y_4 = 0.118, y_6 = 0.173$ $y_7 = 0.116, y_8 = 0.095, y_9 = 0.192, y_{10} = 0.025$	0.037
0.045	$y_1 = 0.260, y_3 = 0.049, y_4 = 0.133, y_6 = 0.174$ $y_7 = 0.101, y_8 = 0.091, y_9 = 0.193$	0.043
0.068	$y_1 = 0.740, y_3 = 0.015, y_4 = 0.08, y_6 = 0.085$ $y_8 = 0.035, y_9 = 0.046$	0.156
0.075	$y_1 = 0.817, y_4 = 0.027, y_6 = 0.085$ $y_7 = 0.029, y_8 = 0.044, y_9 = 0.083$	0.184

4.3 Comparing maximum Ω and minimum-risk portfolios

We can compare Max- Ω portfolios with more conventional minimum-risk ones (denoted by Min- V). We calculate the Min- V portfolio by minimizing (5) subject to constraints (3) and (7) using the R_p achieved by the Max- Ω solutions in Table 2. We also include positivity constraints on the invested fractions to prevent short-selling. Results are shown in Table 3. Clearly the invested fractions y_i are substantially different from those in Table 1. The Min- V solutions tend to use non-zero contributions from more of the assets; and almost the only point of similarity is the growing dominance of asset 1 as R_p increases. Comparing the values V_{\min} with the values of V in Table 2 we can see that, for a given level of portfolio return, the Max- Ω portfolios are appreciably more risky than Min- V ones with the ratio

$$\frac{\text{risk of a Max-}\Omega \text{ portfolio}}{\text{risk of a Min-}V \text{ portfolio}}$$

lying between 1.4 and 1.7.

It is interesting to compare the cumulative distributions of returns for Max- Ω and Min- V portfolios. Figure 10 shows the two curves for the case when $R_p = 0.038\%$. (Figures for other values of R_p are generally similar to this one.)

Both curves are quite similar in the central part but the dotted curve for the Min- V portfolio has a shorter tail towards the upper end of the range of observed returns. This simply reflects the fact that the Max- Ω solution seeks to adjust the areas above and/or below the curve defined by (1) in order to maximize the ratio Ω .

A Min- V portfolio based on the risk definition (5) is calculated to minimize deviations both above and below expected portfolio return R_p . A Max- Ω solution, on the other hand, is concerned with seeking a portfolio to obtain *more than* the threshold return r_t . Hence a Max- Ω portfolio might be expected to have more in common with one which is designed to minimize *downside* risk—that is the risk of obtaining less than target return. If $R_i(y)$ is the portfolio return given by (2) then downside risk can be calculated as

$$DV = \frac{1}{m} \sum_{i=1}^m [\min(0, R_i(y) - R_p)] \tag{11}$$

Fig. 10 Comparing Max- Ω and Min- V for $R_p = 0.038\%$




Table 4 Min- DV portfolios giving same expected return as Max- Ω portfolios

R_p (%)	Invested fractions	DV_{\min}	DV_{Ω}
0.038	$y_1 = 0.193, y_2 = 0.17, y_3 = 0.062, y_4 = 0.150, y_6 = 0.172$ $y_7 = 0.119, y_8 = 0.083, y_9 = 0.18, y_{10} = 0.023$	0.017	0.027
0.040	$y_1 = 0.204, y_2 = 0.006, y_3 = 0.065, y_4 = 0.155, y_6 = 0.172$ $y_7 = 0.118, y_8 = 0.083, y_9 = 0.18, y_{10} = 0.017$	0.018	0.029
0.045	$y_1 = 0.253, y_3 = 0.064, y_4 = 0.165, y_6 = 0.17$ $y_7 = 0.096, y_8 = 0.073, y_9 = 0.178$	0.021	0.031
0.068	$y_1 = 0.782, y_4 = 0.148, y_6 = 0.008$ $y_8 = 0.015, y_9 = 0.021, y_{10} = 0.026$	0.09	0.11
0.075	$y_1 = 0.932, y_2 = 0.001, y_4 = 0.017, y_6 = 0.048$	0.13	0.14

(see [3] for instance). Table 4 shows the portfolios obtained by minimizing downside risk DV for the values of R_p in Tables 2 and 3. It also quotes the value DV_{Ω} of downside risk at the corresponding Max- Ω solution.

In most of the rows of Table 4, the invested fraction distribution in the Min- DV portfolios bears more resemblance to that of the Min- V ones than to that of the Max- Ω solutions. The only evidence of a relationship between Min- DV and Max- Ω portfolios is confined to the larger values of R_p , for which the Max- Ω portfolios have a downside risk that is not very much inferior to DV_{\min} .

4.4 Comments on the performance of e05jb

The global optimization algorithm implemented in e05jb has performed quite successfully on the problems considered above. It has certainly been more efficient than a rather crude multistart approach in which the simplex method in `fminsearch` was applied from 50 to 100 random starting points. This procedure was reasonably useful for the demonstration three-variable example in the previous section; but for the ten-variable case it was very time-consuming and seldom yielded as good an estimate of the global solution as e05jb.

Notwithstanding these positive comments, however, it must be noted that in order to obtain satisfactory results we have had to do some trial-and-error tuning of user-specified parameters of $\epsilon_{0.5j}$. These are listed below.

- Selection of initial points. These are used by the algorithm in deciding whereabouts along the coordinate axes to split the original box. These points may be user-defined; but there are a number of default options for selecting them automatically. We had most success with the default approach in which initial splitting points are selected on the basis of local searches along each coordinate axis.
- Balance between global and local searching. The user can specify how many iterations of a local minimization method are to be used to refine each point identified by the splitting procedure as a candidate global solution. These local searches use a trust-region approach based on quadratic interpolation. While this is often likely to be a good approach, it may not be very suitable for our application since the objective function (10) is nonsmooth. Consequently we have chosen to use rather few local search iterations compared with the suggested default. We have sought to compensate for this by increasing the rigour of the global search and making the number of splitting points per variable about twice the suggested default value.
- Setting termination conditions. The main tests for successful termination of $\epsilon_{0.5j}$ are based either on the search reaching a pre-specified target function value or on there being no decrease in the best function value for a pre-specified number of sweeps (the *static limit*). Furthermore $\epsilon_{0.5j}$ may terminate unsuccessfully—i.e. with a non-zero error flag—if a specified number of function evaluations or box splits is exceeded. We found that the suggested value for the static limit ($3n$ sweeps) was often too small and sometimes led to the method stopping well short of the global solution. We had to increase this to at least $5n$ to obtain acceptable solutions consistently. Basing successful termination on reaching a target function value is a more reliable way of ensuring that a satisfactory stopping point is reached: but often it is not possible to know in advance what the global minimum function is likely to be. We were largely able to avoid unsuccessful terminations by taking the maximum number of splits as being 120 and the maximum number of function calls as 15,000.

In the absence of a computable test for a *global* optimum (as opposed to a local one), algorithms like MCS must terminate on the basis of heuristic rules (like those mentioned in the previous paragraph) which suggest that there is little or no further reduction in function value to be obtained. Depending on how high the threshold is set for the number of sweeps without a function decrease, we can expect MCS to terminate within a fairly small box around—but not precisely *at*—the global optimum. MCS is designed to use quadratic local searches to improve on the sample points in each box; but these may not be very effective on the nonsmooth problem we are considering. As a consequence of these two factors, it is unlikely that the best point returned by $\epsilon_{0.5j}$ will give the global optimum to high precision. We have confirmed this by running the MATLAB implementation of the Nelder and Mead simplex method (`fminsearch`) from the best point found by MCS and observing that it is usually able to obtain a small further reduction in function value. It is these refined estimates of global solutions that are quoted in Table 1.

5 Conclusions

We have given a preliminary account of an investigation of portfolio selection methods based on seeking invested fractions which maximize Ω as defined in (1) [5, 11]. Because maximiz-

ing Ω leads to a non-convex problem, we have considered some Max- Ω portfolios produced using the MCS global optimization algorithm [9] as implemented in the NAG procedure `e05jb` [16]. We have compared these portfolios with ones produced using the well-known Markowitz approach [12] based on minimizing risk. Our small-scale sample calculations indicate that, for a given set of assets, a Max- Ω portfolio can be quite different from portfolios based on minimizing risk or downside risk. The occurrence of such differences is consistent with the original motivation [11] for considering the Omega ratio. Keating and Shadwick argue that $\Omega(r_t)$ is a better measure of performance than the Markowitz mean and variance measures because it does not rest on any prior assumptions (such as normality) about the distribution of asset returns. These arguments are quite persuasive: but we have not investigated whether Max- Ω portfolios are better than minimum-risk ones in any practical sense. To do this we would need to use back-testing based on real-life asset data which would allow us to compare subsequent performance of rival portfolios. Some Max- Ω results with real-life data are reported and discussed in [1, 7, 14]. In particular, Avouyi-Dovi et al. [1] compare a Max- Ω portfolio with one obtained by maximizing the Sharpe ratio (6). They comment that the Sharpe portfolio over-invests in an asset displaying negative asymmetry whereas the Omega portfolio avoids this undesirable feature.

Maximizing Ω has been found to lead to a non-convex and nonsmooth optimization problem. Applying the MCS algorithm [9] (using the NAG implementation `e05jb`) to an exact penalty function has proved quite successful. Since our main aim in this paper has been to show that the problem of portfolio selection by maximizing Ω is one that can be solved with off-the-shelf software, it is encouraging that we have been able to use the automatic procedures in `e05jb` for generating initial points. However this means that we have not obtained any systematic information about how the algorithm's behaviour can be affected by poor choices of starting guess. This question could be part of a more exhaustive numerical investigation of MCS/`e05jb` which has not yet been carried out. Such an investigation would, of course, also need to explore the practicality of solving the Max- Ω problem for much larger numbers of assets.

In view of the non-smoothness of Ω , one might also consider alternatives to MCS (which uses quadratic interpolation and hence assumes differentiability). Within the scope of the present work we have not attempted any comparison between MCS and other global optimization techniques. However it is worth mentioning that the DIRECT method [10] is a box-splitting approach which does not use quadratic models and hence might be more suitable for our nonsmooth problem. This remains an interesting topic for further work.

As a final remark, we mention that the issue of nonsmoothness of Ω may be alleviated if we deal with assets for which a long performance history is available. This may make the cumulative density functions appear rather less jagged.

References

1. Avouyi-Dovi, S., Morin, S., Neto, D.: Optimal Asset Allocation with Omega Function. Technical Report, Banque de France (2004)
2. Avouyi-Dovi, S., Morin, S., Neto, D.: Optimal Asset Allocation with Omega Function. EBC Presentation, Zurich, October 2004. <http://www.effas.ebc.org/EBC>
3. Bartholomew-Biggs, M.: Nonlinear Optimization with Financial Applications. Kluwer, Dordrecht (2005)
4. Bartholomew-Biggs, M., Parkhurst, S., Wilson, S.: Global optimization approaches to an aircraft routing problem. *Eur. J. Oper. Res.* **146**(2), 417–431 (2003)

5. Cascon, A., Keating, C., Shadwick, W.: *The Mathematics of the Omega Measure*. The Finance Development Centre, London (2002)
6. Favre-Bulle, A., Pache, S.: *The Omega Measure: Hedge Fund Portfolio Optimization*. MBF Master's Thesis, University of Lausanne—Ecole des HEC (2003)
7. Gilli, M., Schumann, E.: *Distributed Optimization of a Portfolio's Omega*, Swiss Finance Institute Research Paper 08-17 (2008). <http://ssrn.com/abstract=1156355>
8. Gilli, M., Kellezi, E., Hysi, H.: A data-driven optimization heuristic for downside risk minimization. *J. Risk* **8**(3), 1–19 (2006). <http://ssrn.com/abstract=900365>
9. Huyer, W., Neumaier, A.: MCS global optimization by multilevel coordinate search. *J. Glob. Optim.* **14**, 331–355 (1999)
10. Jones, D.R., Perttunen, C.D., Stuckman, B.E.: Lipschitzian Optimization without the Lipschitz Constant. *J. Opt. Theory Appl.* **79**, 157–181 (1993)
11. Keating, C., Shadwick, W.: *An Introduction to Omega*. AIMA Newsletter, April 2002
12. Markowitz, H.M.: *Portfolio Selection: Efficient Diversification of Investments*, 2nd edn. Wiley, New York (1959), Blackwell, Oxford (1991)
13. The Mathworks Inc. www.mathworks.com
14. Mausser, H., Saunders, D., Seco, L.: *Optimizing Omega*. Risk, November 2006, pp. 88–92. www.algorithmics.com/EN/media/pdfs/OptimizingOmega1106.pdf
15. Nelder, J.A., Mead, R.: A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965)
16. The Numerical Algorithms Group—Private Communication
17. The Numerical Algorithms Group. www.nag.co.uk