

# Applying a Global Optimisation Algorithm to Fund of Hedge Funds Portfolio Optimisation

B. Minsky, M. Obradovic, Q. Tang and R. Thapar

**Abstract—** Portfolio optimisation for a Fund of Hedge Funds (“FoHF”) has to address the asymmetric, non-Gaussian nature of the underlying returns distributions. Furthermore, the objective functions and constraints are not necessarily convex or even smooth. Therefore traditional portfolio optimisation methods such as mean-variance optimisation are not appropriate for such problems and global search optimisation algorithms could serve better to address such problems. Also, in implementing such an approach the goal is to incorporate information as to the future expected outcomes to determine the optimised portfolio rather than optimise a portfolio on historic performance.

In this paper, we consider the suitability of global search optimisation algorithms applied to FoHF portfolios, and using one of these algorithms to construct an optimal portfolio of investable hedge fund indices given forecast views of the future and our confidence in such views.

**Index Terms—** direct search, fund of hedge funds, global search optimisation, portfolio optimisation

## I. INTRODUCTION

The motivation for this paper was to develop a more robust approach to constructing portfolios of hedge fund investments that takes account of the issues that confront portfolio managers:

- 1) The non-Gaussian, asymmetric nature of hedge fund returns;
- 2) The tendency of optimisation algorithms to find corner solutions;
- 3) The speed in computation and efficiency in finding the solution; and
- 4) The desire to incorporate forecast views into the problem specification.

We describe here how each of these issues was addressed and illustrate with reference to the optimising of a portfolio of investable hedge fund indices. This paper synthesises a review of the applicability of global search optimisation algorithms for financial portfolio optimisation with the development of a Monte Carlo simulation approach to forecasting hedge fund returns and implementing the methodology into an integrated forecasting and optimisation application.

In Section 2, we summarise the review of global search

optimisation algorithms and their applicability to the FoHF portfolio optimisation problem. In Section 3, we describe the Monte Carlo simulation technique adopted using resampled historical returns data of hedge fund managers and also how we incorporated forecast views and confidence levels, expressed as probability outcomes, into our returns distribution data. In Section 4, we report the results of applying the methodology to a FoHF portfolio optimisation problem and in Section 5; we draw our conclusions from the study.

## II. REVIEW OF GLOBAL SEARCH OPTIMISATION ALGORITHMS

The FoHF portfolio optimisation problem is an example of the typical minimisation problem in finance

$$\begin{array}{l} \min \quad f(x) \\ g(x) \leq g_0 \\ \vdots \\ h(x) \leq h_0 \end{array}$$

Where  $f$  is non-convex and maybe non-smooth, called the objective function. The  $g, \dots, h$  are constraint functions, with  $g_0, \dots, h_0$  as minimum thresholds. The variable  $x$  usually denotes the weights assigned to each asset and the constraints will usually include the buying and shorting limits on each asset.

It is well known that many of the objective functions and constraints specified in financial minimisation problems are not differentiable. Traditional asset management has relied on the Markowitz specification as a mean-variance optimisation problem which is soluble by classical optimisation methods. However, in FoHF portfolio optimisation the distribution of hedge fund returns are non-Gaussian and the typical objective functions and constraints are not limited to simple mean, variance and higher order moments of the distribution. We have previously. Reference [23] discussed the use of performance and risk statistics such as maximum drawdown, downside deviation, co-drawdown, and omega as potential objective functions and constraint functions which are not obviously differentiable.

With the ready availability of powerful computing abilities and less demand on smoothness, it is possible to look for global optimisation algorithms which do not require regularity of the objective (constraint) functions to solve the financial minimisation problem.

In our review of the literature [23], we found that there are three main ideas of global optimisation; Direct, Genetic Algorithm, and Simulated Annealing. In addition, there are a number of other methods which are derived from one or more of the ideas listed above.

A key characteristic of fund of fund portfolio optimisation,

Manuscript received March 03, 2010.

B.Minsky and R.Thapar, International Asset Management Ltd, 7 Clifford Street, London W1S 2FT, UK, rthapar@iam.uk.com, bminsky@iam.uk.com  
Q. Tang and M. Obradovic, School of Mathematical and Physical Sciences, Sussex University, Brighton BN1 9RF, UK

in common with other portfolio optimisation problems, is that the dimensionality of the problem space is large. Typically, a portfolio of hedge funds will have between 20 and 40 assets with some commingled funds having significantly more assets. This means that the search algorithm cannot conduct an exhaustive test of the whole space efficiently. For example, if we have a portfolio of 40 assets we have a 40-dimensional space, and an initial grid of 100 points on each axis produces 1040 initial test points to evaluate the region where the global minimum might be found. This would require considerable computing power and would not be readily feasible.

Each of the methods we considered in our review requires an initial search set. The choice of the initial search set is important as the quality of the set impacts the workload required to find the global minimum. The actual approach to moving from the initial set to finding better and better solutions differs across the methods and our search also revealed some approaches that combine the methods to produce a hybrid algorithm.

In [23] we evaluated seven algorithms across the methods to identify which method and specific algorithm was best suited to our FoHF portfolio optimisation problem. The algorithms considered are described here:

#### **1) PGSL: Probabilistic Global Search Lausanne**

PGSL is a hybrid algorithm, proposed by Benny Raphael [29], drawing on the Simulated Annealing method that adapts its search grid to concentrate on regions in the search space that are favourable and to intensify the density of sampling in these attractive regions.

The search space is sampled using a probability distribution function for each axis of the multi-dimensional search space. At the outset of the search process, the probability distribution function is a uniform distribution with intervals of constant width. During the process, a probability distribution function is updated by increasing probability and decreasing the width of intervals of the regions with good functional values. A focusing algorithm is used to progressively narrow the search space by changing the minimum and maximum of each dimension of the search space.

#### **2) MCS: Multi-level Co-ordinate Search**

MCS belongs to the family of branch and bound methods and it seeks to solve bound constrained optimisation problems by combining global search (by partitioning the search space into smaller boxes) and local search (by partitioning sub-boxes based on desired functional values). In this way, the search is focused in favour of sub-boxes where low functional values are expected. The balance between global and local parts of the search is obtained using a multi-level approach. The sub-boxes are assigned a level, which is a measure of how many times a sub-box has processed. The global search part of the optimisation process starts with the sub-boxes that have low level values. At each level, the box with lowest functional value determines the local search process. The optimisation method is described in the paper by Huyer and Neumaier [13].

Some of the finance papers that have examined MCS include [7] and [16]. In [7] Value-at-Risk of a portfolio is calculated using marginal distributions of the risk factors and

MCS is employed to search for the best-possible lower bound on the joint distribution of marginal distributions of the risk factors. The Reference [16] uses MCS to optimise for Omega ratio, a non-smooth performance measure, of a portfolio.

#### **3) MATLAB Direct:**

The Direct Search algorithm, available in MATLAB's Genetic Algorithm and Direct Search Toolbox, uses a pattern search methodology for solving bound linear or non-linear optimisation problems [22]. The algorithms used are Generalised Pattern Search (GPS) and Mesh Adaptive Search (MADS) algorithm.

The pattern search algorithm generates a set of search directions or search points to approach an optimal point. Around each search point, an area, called a mesh, is formed by adding the current point to a scalar multiple of a set of vectors called a pattern. If a point in the mesh is found that improves the objective function at the current point, the new point becomes the current point for the next step and so on. The GPS method uses fixed direction vectors and MADS uses random vectors to define a mesh.

#### **4) MATLAB Simulated Annealing:**

The Simulated Annealing method uses probabilistic search algorithm models that model the physical process of heating a material and then slowly lowering the temperature to decrease defects, thus minimising the system energy [22]. By analogy with this physical process, each step in the Simulated Annealing algorithm replaces the current point by another point that is chosen depending on the difference between the functional values at the two points and the temperature variable, which is systematically decreased during the process.

#### **5) MATLAB Genetic Algorithm:**

The MATLAB's Genetic Algorithm is based on the principles of natural selection and uses the idea of mutation to produce new points in the search for an optimised solution [22]. At each step, the Genetic Algorithm selects individuals at random from the current population to be parents and uses them to produce the children for the next generation. In this way, the population evolves toward an optimal solution.

#### **6) TOMLAB LGO:**

Tomlab's Global Optimiser, TOMLAB/LGO, combines global and local search methodologies [33]. The global search is implemented using the branch and bound method and adaptive random search. The local search is implemented using a generalised reduced gradient algorithm.

#### **7) NAG Global Optimiser:**

NAG's Global Optimiser, E05JBF, is based on MCS, as described above. E05JBF is described in [26] and [16].

The above algorithms were evaluated on the three constrained optimisation problems. The constraints consisted of both linear constraints on the allocation weights to the assets and constraints on the level of functions that characterise the portfolio's performance or risk. The algorithms were measured regarding time to run, percentage of corners in the optimal solution, and the deviation from the average optimal solution. A simple scoring rule combining these three factors as a weighted sum was constructed. There was considerable variation in relative performance between the algorithms across the different tests. Two

algorithms, MATLAB Annealing and MATLAB Genetics, were found to be unstable giving rise to different results when repeated runs of the same problem and environment were performed. They also produced widely different results, from very good to very bad, across the tests and were rejected from consideration easily. The other five algorithms all produced acceptable results with MATLAB Direct scoring best across the constrained optimisation examples. PGSL, the adaptive Simulated Annealing algorithm performs reasonably in most tests and has been used by International Asset Management Limited for the past four years. Therefore, we chose to compare MATLAB Direct with PGSL in our portfolio optimisation implementation.

### III. IMPLEMENTING THE GLOBAL SEARCH OPTIMISATION ALGORITHM

Traditional optimisation of portfolios has focused on determining the optimal portfolio given the history of asset returns and assuming that the distribution of returns is Gaussian and stationary over time. Our experience is that these assumptions do not hold and that any optimisation should use the best forecast we can make of the horizon for which the portfolio is being optimised. When investing in hedge funds, liquidity terms are quite onerous with lock ups and redemption terms from monthly to annual frequencies, and notice periods ranging from a few days to six months. This means that the investment horizon tends to be six to twelve months ahead to reflect the minimum time any investment will be in a portfolio.

The forecast performance of the assets within the portfolio is produced using Monte Carlo simulation and re-sampling. The objective is to produce a random sample of likely outcomes period by period for the forecast horizon based on the empirical distributions observed for the assets modified by our views as to the likely performance of the individual assets. This is clearly a non-trivial exercise, further complicated by our wish to maintain the relationship between the asset distributions and any embedded serial correlation within the individual asset distribution.

The approach implemented has three components:

- 1) Constructing a joint distribution of the asset returns from which to sample;
- 2) Simulating the returns of the assets over the forecast horizon; and
- 3) Calculating the relevant objective function and constraints for the optimisation.

#### A. Constructing the Joint Distribution of Asset Returns

We used bootstrapping in a Monte-Carlo simulation framework to produce the distribution of future portfolio returns. Bootstrapping is a means of using the available data by resampling with replacement. This generates a richer sample than would otherwise be available. To preserve the relationship between the assets we treat the set of returns for the assets in a time period as an observation of the joint distribution of the asset returns. An enhancement to this sampling scheme to capture any serial correlation is to block sample a group contiguously, say three periods together. Block sampling of three periods at a time offers around 10 million distinct samples of blocks of three time periods.

As we used bootstrapping to sample from the distribution and we wished to preserve the characteristics of the joint distribution, we needed to define a time range over which we have returns for all of the assets in the portfolio. Hedge funds report returns generally on a monthly basis, which means that we needed to go back a reasonable period of time to obtain a sufficiently large number of observations to enable the bootstrap sampling to be effective. For hedge funds this is complicated because many of the funds have not been in existence for very long, with the median life of a hedge fund being approximately three years. Although the longer the range that can be used for the joint distribution the greater the number of points available for sampling, the lack of stationarity within the distribution leads us to select a compromise period, typically five years, as the desired range. Where a hedge fund does not have a complete five year history, we employed a backfill methodology to provide the missing data.

There are a number of approaches to backfilling asset return time series such as selecting a proxy asset to fill the series; using a strategy index with a random noise component; constructing a factor model of the asset returns from the available history and using the factor return history and model to backfill; or to randomly select an asset from a set of candidate assets that could have been chosen for the portfolio for the periods that the actual asset did not exist. We adopted this last method, selecting an asset from a set of available candidates within a peer group for the missing asset. Where the range for which returns are missing was long, we repeated the exercise of selecting an asset at random from the available candidates within the peer group, say, every six periods. Our reasoning for applying this approach is that we assume as portfolio managers, given the strategy allocation of the portfolio, that we would have chosen an asset from the candidate peer group available at that time to complete the portfolio. Using this process we constructed a complete set of returns for each of the assets going back, say, five years. The quality of the backfill depends on how narrowly defined the candidate peer group is defined. At International Asset Management Limited (IAM), we have defined our internal set of strategy peer groups that reflect best our own interpretation of the strategies in which we invest. This is because hedge fund classifications adopted by most of the index providers tend to be broad, and can include funds that would not feature in IAM's classifications.

#### B. Simulating returns over the forecast horizon

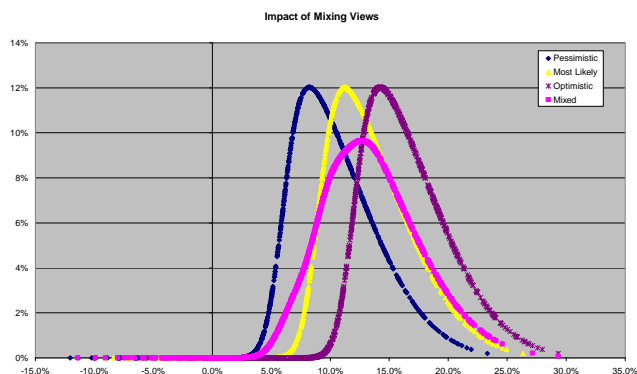
We simulated the returns of the assets using a block bootstrap of the empirical joint distributions, which are modified by probabilistically shifting the expected return of the sample according to our assessment of the likely return outcomes for the assets. First we describe the process of incorporating forecast views by expectation shifting and then we describe the block bootstrapping method.

The desire to include forecast views, expressed as expected annual returns, and confidence, expressed as probabilities, within a portfolio optimisation problem has been addressed in a number of ways. Black and Litterman developed an approach where the modeller expressed a view as to the expected mean of a returns series and attached a

confidence to each view. This approach is Bayesian and allows the traditional Mean-Variance approach to be adapted to allow for more stable and intuitive allocations which do not favour corner solutions. However, we have chosen an empirical approach, of mixing probabilistically mean shifted versions of the empirical distribution, to include views that allows a range of outcomes to be specified with a confidence associated with the views.

Figure 1 shows how applying a probabilistic shift to the mean of a distribution not only repositions the distribution but changes the higher order moments as the spread, skew and kurtosis all change.

1) Fig. 1: Probabilistic Shifting of Expected Mean



In Figure 2 the forecast views for a number of strategies are set out with associated confidence. The optimistic, pessimistic and most likely views are the best assessment of the potential expected return of the mean fund within the strategy. The confidence level represents the likelihood of that view prevailing. We note the sum of the three confidence levels is one. We use these likelihoods to determine for each asset, according to its strategy, which shift should be applied to the distribution for that simulation. This is implemented by simply sampling from the uniform distribution and dividing the distribution into three segments according to the confidence levels associated with the three views. Recognising that each asset does not track its strategy with certainty we calculate the beta for the asset with respect to the strategy and adjust the return by the randomly chosen shift (“k”) multiplied by the asset beta calculated. So the return in any period (“t”) for an asset (“a”) which follows strategy (“s”) for simulation trial (“n”) is:

$$m_{r_{a,t,n}} = raw_{r_{a,t,n}} + \beta \times shift_{s,k}$$

2) Fig. 2: Forecast views and confidence by hedge fund strategy

Strategy	Optimistic View	Conf./ Probability	Pessimistic View	Conf./ Probability	Most Likely View	Conf./ Probability
Convertible Bond Arbitrage	17.5%	25%	7.5%	25%	12.5%	50%
Credit	17.5%	25%	7.5%	25%	12.5%	50%
Event Driven	10.0%	25%	0.0%	25%	5.0%	50%
Fixed Income Relative Value	15.0%	25%	10.0%	25%	12.5%	50%

### C. Calculating the objective function and constraint functions

In implementing the bootstrapped Monte Carlo simulation we simulate 500 trials or scenarios for the assets in the portfolio. This produces a distribution of returns of each asset and the distributions of any statistics we may wish to compute. Our objective and constraint functions are statistics based on the distribution of portfolio returns. With a set of asset allocation weights, the distribution of portfolio returns

and statistics distributions may be calculated. It is worth discussing how we use this information within the optimisation algorithm. To do this we shall use as an example maximising expected return subject to a maximum level of maximum drawdown.

As we have chosen to optimise expected return, our objective function is simply the median of the distribution of portfolio returns. If we set our objective to ensure performance is at an acceptable level in most circumstances we might choose the bottom five percentile of return as the objective function so as to maximise the least likely (defined as fifth percentile) return. This reflects the flexibility we have with using a simulated distribution as the data input into the optimisation process.

In PGSL, as with almost all of the global search optimisation algorithms, both the linear and non-linear constraints are defined as penalty functions added to the objective function and hence are soft constraints rather than hard constraints that must be satisfied. The weight attached to each penalty function determines how acceptable a constraint violation is. In our example, we define the penalty function as the average of the maximum drawdown for the lowest five percentile of the maximum drawdown distribution less the constraint boundary assuming the conditional average exceeds the constraint level multiplied by an importance factor:

$$\begin{aligned} \text{Max\_dd\_penalty} = & \\ & - \text{Max}(\text{Constraint\_dd} - \text{average}(\text{Max\_ddn} | \text{Lower 5\%ile}), 0) \\ & / \text{No. of Trials} * \text{Importance} \end{aligned}$$

This measure is analogous to an expected tail loss or Conditional VaR (CVaR) in that it is an estimate of the conditional expectation of the maximum drawdown for the lower tail of the distribution of drawdowns.

## IV. RESULTS OF OPTIMISING A FOHF PORTFOLIO

The approach to optimising a FoHF portfolio has been implemented in MATLAB and applied to a portfolio of eight RBC Hedge 250 hedge fund strategy indices. The monthly returns for indices from July 2005 are available from the RBC website. As the simulation requires five years of monthly returns the series were backfilled from the IAM’s pre-determined group of candidate assets within the relevant investment strategy peer group, using random selection as previously described.

The portfolio was optimised with an objective function to maximise median returns subject to constraints on the maximum and minimum allocations to each asset, a constraint on the maximum and minimum allocation to Long/Short Equity strategies and a maximum allowable maximum drawdown of 5% over the forecast horizon.

Thus the optimisation problem is as set out in Figure 3:

1) Fig. 3: FoHF Portfolio Optimisation Problem

<b>Objective:</b>	
<b>Maximise median portfolio return</b>	
<b>Subject to:</b>	
<b>Maximum drawdown over forecast period</b>	<b>Less than 5%</b>
• Total allocations for full investment	100%
• Cash	10%
<b>Within the following constraints:</b>	
• RBC Hedge 250 Equity Market Neutral	between 10% to 16%
• RBC Hedge 250 Equity Long/Short Directional	between 14% to 20%
• All Long/Short Equity	between 24% to 36%
• RBC Hedge 250 Fixed Income Arbitrage	between 7% to 13%
• RBC Hedge 250 Macro	between 10% to 20%
• RBC Hedge 250 Managed Futures	between 10% to 20%
• RBC Hedge 250 Credit	between 5% to 15%
• RBC Hedge 250 Mergers & Special Situations	between 0% to 10%
• RBC Hedge 250 Multi-Strategy	between 0% to 10%


First we noted that the total allocations satisfying the equality constraint of all capital is deployed with both PGSL and MATLAB Direct, and that all the asset allocation constraints are satisfied including the constraint on all Long/Short Equity strategies by MATLAB Direct, but not by PGSL. Secondly we noted that with PGSL only one other allocation is near its lower or upper bounds whereas with MATLAB Direct five allocations are at or near either the lower or upper bounds. Thirdly we compared the results to a portfolio where the allocation of capital to the different assets was chosen to be the midpoint between the lower and upper bounds placed on each asset (the naïve allocation). We noted that both optimisers improved median returns (7.8% and 8.0% vs. 7.40%) and that MATLAB Direct reduced the breach of the maximum drawdown constraint (1.71% vs. 2.70%) whereas the PGSL optimisation failed to improve on this condition (3.22% vs. 2.70%). The overall performance of the three portfolios is shown in Figures 4, 5, 6 and 7. Figure 5 shows the median and lower five percentile of the return distribution for both the two optimised portfolios. In Figure 6, we compare the distributions of maximum drawdowns for the two optimised portfolios and the naïve portfolio. The graph shows that the MATLAB Direct portfolio had the better maximum drawdown distribution both in terms of worst case and general performance. Figure 7 shows the performance of the portfolios over the backtest period used in generating the data set. Again the MATLAB Direct optimised portfolio performs the best of the three portfolios. Finally we noted that PGSL optimisation terminated on maximum iterations and this might explain why it failed to meet all the allocation criteria.

2) Fig. 4: Optimal Allocations and Results


Asset	Lower Bound	Upper Bound	Naïve	PGSL	Direct
Cash	10%	10%	10.0%	10.0%	10.0%
RBC Hedge 250 Equity Market Neutral	10%	16%	13.0%	15.0%	16.0%
RBC Hedge 250 Equity Long/Short	14%	20%	17.0%	16.1%	20.0%
RBC Hedge 250 Fixed Income Arbitrage	7%	13%	10.0%	12.5%	13.0%
RBC Hedge 250 Macro	10%	20%	15.0%	12.7%	13.9%
RBC Hedge 250 Managed Futures	10%	20%	15.0%	12.7%	15.4%
RBC Hedge 250 Credit	5%	15%	10.0%	16.7%*	11.5%
RBC Hedge 250 Mergers & Special Situations	0%	10%	5.0%	0.6% (out of range)	10%
RBC Hedge 250 Multi-Strategy	0%	10%	5.0%	5.8%	0.1%
<b>Median Return</b>			7.40%	7.80%	7.96%
<b>Excess Tail Maximum Drawdown</b>			2.70%	3.22%	1.71%

\* In breach of upper allocation constraint


3) Fig. 5: MATLAB Direct, PGSL and Naïve Portfolios Returns



4) Fig. 6: MATLAB Direct, PGSL and Naïve Max Drawdown Distributions



5) Fig. 7: MATLAB Direct, PGSL and Naïve Backtested Returns



## V. CONCLUSION

The review of Global Search Optimisation algorithms showed that there is a range of methods available, but their relative performance is variable. The specifics of the problem and initial conditions can impact the results significantly. In applying MATLAB Direct and PGSL to the FoHF portfolio optimisation problem, we observed that we improved on the naïve solution in both cases, but each method presented solution characteristics that might be less desirable. PGSL was unable to find a solution that met its threshold stopping criterion whilst MATLAB Direct found a solution with many corner points. Further research studies are required to evaluate the stability of the optimiser outputs and sensitivity analysis of salient optimisation parameters.

REFERENCES

- [1] <http://www.mat.univie.ac.at/~neum/ms/StArt.pdf>
- [2] <http://www.andreassteiner.net/performanceanalysis/index.php>
- [3] F Black and R Litterman, Asset allocation: combining investor views with market equilibrium, Goldman, Sachs & Co., Fixed Income Research, September 1990
- [4] S Ciliberti, I Kondor and M Mezard, On the feasibility of portfolio optimisation under expected shortfall, *Quantitative Finance*, 7(4), 2007, 389-396.
- [5] A Chabaane, J P Laurent, Y Malevergne and F Turpin, Alternative risk measures for alternative investments, 21e Conference International en Finance, 2004.
- [6] A Chekhlov, S Uryasev and M Zabarankin, Drawdown measure in portfolio optimisation, *International Journal of Theoretical and Applied Finance* Vol. 8, No. 1 (2005) 13–58
- [7] Paul Embrechts, Giovanni Puccetti, Aggregating risk capital, with an application to operational risk, *The GENEVA Risk and Insurance*, Vol 31, Issue 2
- [8] I Ekeland and Roger Temam, *Convex analysis and variational problem*, Society for Industrial and Applied Mathematics, 1999.
- [9] D A Freedman, Bootstrapping regression models, *Ann Statist.* 9, 1981, 1218-1228
- [10] G Fisher and A B Sim, Some finite sample theory for bootstrap regression estimates, *Journal of Statistical Planning and Inference* 43, 1995, 289-300
- [11] C I Fabian and A Veszpremi, Algorithms for cvar optimisation in dynamic stochastic programming models with applications to finance, *Rutcor Research Report*, RRR 24-2006, 09/2006
- [12] M Gilli, E Kellezi and H Hysi, A data-driven optimisation heuristic for downside risk minimisation, *Journal of Risk*, 8(3), 2006, 1-19.
- [13] W. Huyer and A Neumaier, Global optimisation by multilevel coordinate search, *J. Global Optimisation* 14, 1999, 331-355.
- [14] W Huyer, A comparison of some algorithms for bound constrained global optimisation, working paper, 2004.
- [15] J. He, L. T. Watson, N. Ramakrishnan, C. A. Shaffer, A. Verstack, J Jiang, K. Bae and W. H. Tranter, Dynamic data structures for a direct search algorithm, *Computational Optimisation and Applications*, 23, 2002, 5–25
- [16] S.J. Kane and M.C. Bartholomew-Biggs, Optimising Omega, *Journal of Global Optimisation*, Vol 45, Number 1.
- [17] A I Kibzun and E A Kuznetsov, Comparison of VaR and CVaR criteria, *Automation and Remote Control*, 64(7), 2003, 1154-1164
- [18] P D Kaplan and J A Knowles, Kappa: A generalised downside risk-adjusted performance measure, *Journal of Performance Measurement*, Spring 2004.
- [19] P Krokmal and S Uryasev, Portfolio optimisation with conditional value-at-risk objective and constraints, *The Journal of Risk*, 4(2), Winter 2001/02.
- [20] C Keating and W F Shadwick, A universal performance measure, <http://www.spgshop.com/index.asp?PageAction=VIEWPROD&ProdID=413>.
- [21] H Kazemi, T Schneeweis and R Gupta, Omega as a performance measure, [http://www.edhec-risk.com/site\\_edhecrisk/public/research\\_news/choice/RISKReview1063631261806712604](http://www.edhec-risk.com/site_edhecrisk/public/research_news/choice/RISKReview1063631261806712604), 2003.
- [22] Genetic Algorithm and Direct Search Toolbox™ 2 User's Guide, Mathworks, [http://www.mathworks.com/access/helpdesk/help/pdf\\_doc/gads/gads\\_tb.pdf](http://www.mathworks.com/access/helpdesk/help/pdf_doc/gads/gads_tb.pdf)
- [23] B Minsky, M Obradovic, Q Tang and R Thapar, Global optimisation algorithms for financial portfolio optimisation, Working Paper University of Sussex, 2008
- [24] D Martinger and P Parpas, Global optimisation of higher order moments in portfolio selection, *J Glob Optim*, 2007.
- [25] L Mitra, X Sun, D Roman, G Mitra & K Yu, Mixture distribution scenarios for investment decisions with downside risk, Working Paper, SSRN, June 2009
- [26] NAG Library Routine Document E05JBF, NAG, [http://www.nag.co.uk/numeric/FL/nagdoc\\_fl22/pdf/E05/e05jbf.pdf](http://www.nag.co.uk/numeric/FL/nagdoc_fl22/pdf/E05/e05jbf.pdf)
- [27] G Ch Pflug, Some remarks on the value-at-risk and the conditional value-at-risk, in *Probabilistic Constrained Optimisation, Methodology and Applications*, ed. S Uryasev, Kluwer, 2000.
- [28] U Raber, A simplicial branch-and-bound method for solving nonconvex all-quadratic programs, *Journal of Global Optimisation*, 13, 1998, 417-432.
- [29] B Raphael and I F C Smith, A direct stochastic algorithm for global search, *Applied Mathematics and Computation*, 146, 2003, 729-758
- [30] R T Rockafellar and S Uryasev, Optimisation of conditional value-at-risk, *Journal of Risk*, 2(3), 2000, 21-41.
- [31] R T Rockafellar, S Uryasev and M Zabarankin, Deviation measure in risk analysis and optimisation (December 22, 2002). University of Florida, Department of Industrial & Systems Engineering Working Paper No. 2002-7. Available at SSRN: <http://ssrn.com/abstract=365640>.
- [32] R T Rockafellar, S Uryasev and M Zabarankin, Optimality conditions in portfolio analysis with general deviation measures (May 10, 2005). University of Florida Industrial and Systems Engineering Working Paper No. 2004-7. Available at SSRN: <http://ssrn.com/abstract=615581>
- [33] User's guide for TOMLAB/ LGO1TOMLAB [http://tomopt.com/docs/TOMLAB\\_LGO.pdf](http://tomopt.com/docs/TOMLAB_LGO.pdf)
- [34] P Winker and D Maringer, The threshold accepting optimisation algorithms in economics and statistics, in: Kontoghiorges, E.J., Gatu, Chr. (eds.), *Optimisation, Econometric and Financial Analysis*, Springer, 2007.
- [35] X Yang, Z Yang, Z Shen and G Lu, Gray-encoded hybrid accelerating genetic algorithm for global optimisation of water environmental model, Springer, 2006.