



# Kent Academic Repository

**Alentorn, Amadeo and Moraglio, Alberto (2010) *Heuristic Portfolio Optimisation for a Hedge Fund Strategy using the Geometric Nelder-Mead Algorithm*. In: 2010 UK Workshop on Computational Intelligence (UKCI 2010). . ISBN 978-1-4244-8774-5.**

## Downloaded from

<https://kar.kent.ac.uk/71013/> The University of Kent's Academic Repository KAR

## The version of record is available from

## This document version

Author's Accepted Manuscript

## DOI for this version

## Licence for this version

UNSPECIFIED

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

## Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

# Heuristic Portfolio Optimisation for a Hedge Fund Strategy using the Geometric Nelder-Mead Algorithm

Amadeo Alentorn, Alberto Moraglio, and Colin Johnson

**Abstract**—This paper presents a framework for heuristic portfolio optimisation applied to a hedge fund investment strategy. The first contribution of the paper is to present a framework for implementing portfolio optimisation of a market neutral hedge fund strategy. The paper also illustrates the application of the recently developed Geometric Nelder-Mead Algorithm (GNMA) in solving this real world optimization problem, compared with a Genetic Algorithm (GA) approach.

## I. INTRODUCTION

Given a set of stock returns and risk forecasts, and any number of constraints, the portfolio optimization problem typically consists in maximizing return or minimizing risk by finding the optimal set of stock weights (i.e., percentages of invested capital) that satisfies a set of constraints. There are a number of optimization models for portfolio construction. The classic one is the mean-variance optimisation introduced by Markowitz [7][8]. His seminal work made a number of simplifying assumptions in order to solve the problem using classical mathematical optimization techniques, such that returns are normally distributed, and that investors have quadratic preferences. But it is now well understood that stock returns exhibit non-normal characteristics, including skewness and fat tails, and that most investors are loss averse, thus, invalidating those original simplifying assumptions. Heuristic optimisation methods provide a more flexible toolset where no simplifying assumptions are needed. These methods have been applied to tackle portfolio optimisation for some time now [12]. Some of the applications have successfully achieved the use of non-quadratic risk measures such as Value at Risk [4], or have allowed the incorporation of integer constraints, such as cardinality constraints [2]. Several different heuristic algorithms have been used, including Genetic Algorithms [12], and Threshold Accepting algorithms [3],[6].

A hedge fund is a special type of investment fund that typically uses more complex strategies than a traditional investment fund (i.e., portfolio). One of the most common characteristics of hedge funds is the use of *short-selling* to reduce part of the risk of their strategies. In practice, the

mechanics of short-selling works as follows. A hedge fund manager with a negative view about a stock would like to include it in his portfolio with a *negative* weight (i.e., hold a certain negative amount of the stock). This can be implemented as follows. Brokers/dealers on behalf of clients try to identify stock owners willing to lend the stock for a fee. The hedge fund manager borrows the stock, and sells it on the market. At a future date, usually when the stock price has fallen, the hedge fund manager buys back the stock on the market at the lower price, making a profit, and then returns the stock to the original owner. It is important to note that the use of shorting (i.e., short-selling) in an investment process introduces a number of new risks for the portfolio manager. The main risk is the potential for unlimited losses. The maximum loss from a long position (i.e., without short-selling) in a stock is limited to the initial investment value. That is, if you buy a share in a company worth 100, and that company goes bankrupt, you would lose the entire 100 investment. However, the maximum loss from a short position (i.e., with short-selling) can be unlimited. If you “short” that same company worth 100, and its stock price for example triples in value, you would lose 200. This risk of losing more than the initial investment, together with the fact that stock returns are not normally distributed, motivates the need to attempt to control for the risk of extreme events, i.e. tail risk. Value-at-Risk (VaR) is one of the most widely used risk measures to model and control for tail risk [1].

There is a large number of different hedge fund strategies, with *market neutral equity* being one of the most popular strategies. The strategy consists of selecting a set of stocks to buy and a set of stocks to short-sell, from the constituents of an equity index (i.e., a fixed large set of stocks), such that the amount invested in the buys and sells is the same, thus, neutralising the risk of changes in value of the equity index. Usually for a hedge fund strategy, the objective is to maximise expected returns for a given a level of expected risk.

The contribution of this paper is two-fold. Firstly, it presents a framework for implementing a market neutral hedge fund strategy, where short-selling is allowed, and using VaR as the risk measure, instead of variance. To the author’s best knowledge, to date all literature on heuristic portfolio optimisation enforces the no-shorting constraint. Secondly, it presents the first application of the recently developed Geometric Nelder-Mead Algorithm (GNMA) [10], described in section III, to a portfolio optimisation problem.

Amadeo Alentorn is with Old Mutual Asset Managers (UK) Ltd, 2 Lambeth Hill, London, EC4P 4WR, UK, email: Amadeo.Alentorn@omam.co.uk.

Alberto Moraglio is with the School of Computing and Centre for Reasoning, University of Kent, Canterbury, UK, email: A.Moraglio@kent.ac.uk.

Colin Johnson is with the School of Computing, University of Kent, Canterbury, UK, email: C.G.Johnson@kent.ac.uk.

Disclaimer: the views and results given in this paper are those of the authors and do not necessarily reflect the views of any company or institution affiliated with the authors.

## II. MARKET NEUTRAL PORTFOLIO OPTIMIZATION

For a traditional long-only investment fund (i.e., without short-selling), the portfolio optimisation problem consists of selecting a subset of stocks from a given universe of stocks, and the amount to invest in each of them, subject to a number of constraints and a number of objectives. On a traditional long-only portfolio problem, the sum of weights (i.e., amounts to be invested in the stocks hold in the portfolio) must equal to one, and no negative weights are allowed (i.e., no-shorting constraint). In formulas:

$$\sum w_j = 1 \quad (1)$$

$$w_j \geq 0 \quad (2)$$

for  $j = 1, \dots, n_A$  where  $w_j$  is the weight of stock  $j$  in the portfolio and  $n_A$  is the number of stocks available in that given universe.

To implement a market neutral strategy, we relax the no-shorting constraint, and aim at constructing a portfolio with two books: a “long book” with positive weights adding to 1, and a “short book” with negative weights adding to -1. Such weights are characterized by the following constraints:

$$\sum |w_j| = 2 \quad (3)$$

$$\sum w_j = 0 \quad (4)$$

for  $j = 1, \dots, n_A$ . The resulting portfolio is referred to as a market neutral portfolio because it has a zero exposure to the equity market.

### A. The mean-variance model

Mean-variance optimization is the most widely used framework, both in the academia and in the industry, to construct portfolios. The utility function for defining optimality is based on the original framework proposed by Markowitz:

$$u(\mathbf{w}) = \mathbf{w}\mathbf{r} - \lambda \mathbf{w}'\mathbf{Q}\mathbf{w} \quad (5)$$

where  $\mathbf{w}$  is the vector of portfolio weights,  $\mathbf{r}$  is the vector of stock return forecasts,  $\mathbf{Q}$  is the forecasted covariance matrix, and  $\lambda$  is the risk aversion parameter. This risk aversion parameter defines the preferences of the investor, in terms of the trade-off between taking additional expected risk for additional expected return. The objective of the optimiser is to maximise utility, by finding portfolios that give this optimal trade-off between risk and return according to the investor preferences.

### B. The bounded-VaR model

The variance measure above is a symmetric risk measure, but given the non-normal properties of stock returns, and the potential for unlimited losses when using short-selling, a risk measure that focuses on modeling tail risk such as VaR may be more appropriate. VaR is an estimate, with a given degree of confidence, of how much can be lost from a portfolio over a given time horizon. We could easily formulate a mean-VaR optimization with a risk aversion parameter, similar to the above mean-variance approach. However, in practice, VaR is

usually a constraint set by either the financial regulator or a risk management policy, aimed at limiting the amount of risk that an investment manager takes. Therefore, a practically more relevant utility function is the expected return,  $\mathbf{r}\mathbf{w}$ , to be maximized with an extra constraint limiting VaR:

$$u(\mathbf{w}) = \mathbf{r}\mathbf{w} \quad (6)$$

$$VaR(\mathbf{w}, c) < VaR_{max} \quad (7)$$

where  $VaR(\mathbf{w}, c)$  is the forecasted VaR at a given confidence level  $c$  of a portfolio defined by the vector of weights  $\mathbf{w}$ , and  $VaR_{max}$  is the limit to VaR. This type of problem including as a hard constraint the limit on VaR, which is a non-linear function of  $\mathbf{w}$ , cannot be easily solved with a traditional optimiser, and thus, it exemplifies the benefits in terms of flexibility provided by heuristic optimisation methods.

### C. Solution encoding and fitness function

In the following, we cast the optimization models above in a form that can be solved by a traditional Genetic Algorithm with binary representation by choosing appropriate solution encoding and fitness function that can handle the various types of constraints. The same encoding and fitness function will be used with the Geometric Nelder-Mead Algorithm, so making the two algorithms directly comparable.

Every stock in the considered universe is encoded using two bits, the first bit determines whether or not the stock has to be part of the portfolio, and the second bit determines whether to buy or to short-sell the stock when the stock is part of the portfolio (i.e., it indicates the book the stock belongs to). So a complete portfolio is encoded with string of  $2 \cdot n_A$  bits. The weights of the stocks are assigned in a way that all selected stocks in a book are given the same weight, and the sum of the weights for each book equals one. This encoding guarantees that constraints (3) and (4) hold by construction. This encoding is not expressive enough to represent all possible portfolios. It was chosen to simplify the optimization problem and to be able to apply the GNMA heuristic method. However, given that market neutral strategies tend to hold highly diversified portfolios with large number of stocks, restricting the optimisation process from its natural continuous space to a discrete representation should still result in realistic portfolio solutions.

The fitness function for the mean-variance model corresponds to the utility function (equation 5). To enforce the constraint (7), the fitness function for the bounded-VaR model corresponds to the utility function (equation 6) when the constraint holds and the utility is larger than zero, and zero otherwise. In other words, fitness zero indicates both very poor and infeasible solutions.

## III. BINARY NELDER-MEAD ALGORITHM

The Nelder-Mead Algorithm (NMA) [11] is an almost half-century old method for numerical optimization, and it is a close relative of Particle Swarm Optimization (PSO) and Differential Evolution (DE). In recent work, PSO, DE and

NMA have been generalized using a formal geometric framework [9] that treats solution representations in a uniform way. These formal algorithms can be used as templates to derive rigorously specific PSO, DE and NMA for both continuous and combinatorial spaces retaining the same geometric interpretation of the search dynamics of the original algorithms across representations. In previous work, a geometric NMA was formally derived for the binary string representation. To the authors’s best knowledge, apart from very recent work of the authors [10], there are no generalizations of the NMA to combinatorial spaces. The geometric NMA was then preliminary tested on NK-landscapes [5], which are a well-known benchmark of artificially constructed problems, on which it performs well in the comparison with a Genetic Algorithm. In this paper, we extend the experimental analysis of the geometric NMA to the case of Portfolio Optimisation Problem, which is an interesting and challenging real-world problem.

### A. Classic Nelder-Mead Algorithm

In this section, we describe the traditional NMA [11]. The NMA uses  $n + 1$  points in  $\mathbf{R}^n$ . These points form a type of  $n$ -dimensional polygon, a simplex, which has  $n + 1$  points as vertices in  $\mathbf{R}^n$ . For example, the simplex is a triangle in  $\mathbf{R}^2$  and a tetrahedron in  $\mathbf{R}^3$ . The initial simplex has to be non-degenerate, i.e., the points must not lie in the same hyperplane. This allows the NMA to search in all  $n$  dimensions. The method then performs a sequence of transformations of the simplex, which preserve non-degeneracy, aimed at decreasing the function values at its vertices. At each step, the transformation is determined by computing one or more test points and comparing their function values. In Figure 1, we illustrate the NMA transformations for the two-dimensional case, where the simplex  $S$  consists of three points.

The optimization process starts with creating a sample of  $n + 1$  random points in the search space. Notice that apart from the creation of the initial simplex, all further steps are deterministic and do not involve random choices. In each loop iteration, the points in the simplex  $S$  are arranged in ascending order according to their corresponding objective values. Hence, the best solution candidate is  $S[0]$  and the worst is  $S[n]$ . We then compute the center  $m$  of the  $n$  best points and then reflect the worst candidate solution  $S[n]$  through this point, obtaining the new point  $r$  as also illustrated in Fig. 1(a). The reflection parameter  $\alpha$  is usually set to 1. In the case that  $r$  is neither better than  $S[0]$  nor as worse as  $S[n]$ , we directly replace  $S[n]$  with it. If  $r$  is better than the best solution candidate  $S[0]$ , we expand the simplex further into this promising direction. As sketched in Fig. 1(b), we obtain the point  $e$  with the expansion parameter  $\gamma$  set to 1. We now take the best of these two points to replace  $S[n]$ . If  $r$  is no better than  $S[n]$ , the simplex is contracted by creating a point  $c$  somewhere in between  $r$  and  $m$ . In Fig. 1(c), the contraction parameter  $\rho$  was set to  $1/2$ . We substitute  $S[n]$  with  $c$  only if  $c$  is better than  $r$ . When everything else fails, we shrink the whole simplex by moving all points (except  $S[0]$ ) into the direction of the current optimum  $S[0]$ . The

shrinking parameter  $\sigma$  normally has the value  $1/2$ , as is the case in the example outlined in Fig. 1(d).




Fig. 1. One step of the NMA in  $\mathbf{R}^2$  (figure modified from [13])

### B. Geometric Nelder-Mead Algorithm

The generalization of the classic Nelder-Mead Algorithm to the binary representation was done by following a formal generalization methodology that allows us to generalize search algorithms defined on continuous spaces to combinatorial spaces in a systematic way, without arbitrary choices in the transition from continuous to combinatorial spaces (e.g., to deal with the “discreteness” of combinatorial spaces). The interested reader is referred to [10] where the generalization methodology and the formal derivation of the Nelder-Mead for the Hamming space are described in details. Briefly, the generalization methodology is as follows. Firstly, the search operations described in the previous section (reflection, expansion, contraction and shrinking) have to be rewritten expressing them as functions of the Euclidean distance. Then, in these definitions, the Euclidean distance is substituted with the Hamming distance associated with binary strings. Finally, the search operators for the binary space are formally derived by rewriting the (declarative) definitions of the search operators in terms of Hamming distance, in an equivalent but operational form, in terms of manipulation of binary strings. In principle, the same technique can be used to formally derive the Nelder-Mead Algorithm for any representation associated with a well-defined notion of distance between solutions.

Center of mass, convex combinations and extension rays in the Euclidean space can be expressed solely in terms of distance relations between points in space, hence, these geometric concepts can be naturally generalized to the Hamming space associated with binary strings, by replacing the Euclidean distance with the Hamming distance. The graphical description of the search operations of NMA (Fig. 1) leads directly to their geometric interpretation in terms of appropriate compositions of center of mass, convex combination and extension ray. Hence all the search operators of the NMA can be formally generalized and formally specialized to binary strings using only these three operators once they are formally instantiated for binary strings in the Hamming space. In this space, the convex combination corresponds to a form of biased uniform crossover for binary strings. The center of mass for binary strings corresponds to a multi-parent recombination where the offspring is determined by position-wise majority voting of their parents. The extended

ray recombination is the “inverse” recombination operator of the convex combination: given that we know a parent string and the offspring string obtained by the convex combination of the known parent and a second unknown parent, the extension ray recombination reconstructs the unknown parent string. Detailed descriptions of these operators, their formal derivations, and of the complete binary GNMA can be found in [10].

#### IV. EXPERIMENTS

In this section we present a description of the data used and the experiments carried out.

We constructed a data-set of daily stock returns for the 100 names that constitute the FTSE 100 index between January 2007 and December 2009, over a total of 500 trading days.

The forecasted covariance matrix, used for the mean-variance problem was calculated as a full covariance matrix with a rolling window of the last 250 days. This approach of using historical returns in order to forecast future volatility is widely used.<sup>1</sup> The Value at Risk (VaR) forecasts are based on the historical method. Given a set of stock weights, defining a portfolio at time  $t$ , we take a historical observation period of 1 year, 250 trading days, from  $t-1$  to  $t-250$ , and compute for each day in the period the daily return that such portfolio would have delivered based on historic stock returns. This gives an empirical distribution of historical daily returns, and can be used to estimate with a degree of confidence, how much can be lost. With a 99 percent confidence level, and a sample of 250 days, the historical 1-day VaR estimate is obtained by taking the third highest lost in that window. Then, a given portfolio satisfies the VaR constraint if its third largest loss over that historic period is smaller than the VaR limit  $VaR_{max}$ .

The forecasts for stock returns were constructed in such a way as to reflect some level of forecasting ability, but at the same time, to also reflect the high level of noise and uncertainty in industry models for stock return forecasting. We achieved this by calculating forecasted returns as a weighted sum between a normally distributed random variable (95 percent weight) and one day ahead known stock returns (5 percent weight).<sup>2</sup> Portfolio simulations are performed by re-balancing the portfolio in each from time  $t=251$  to  $t=500$ . We need to start at  $t=251$  as we need 250 historical days to calculate the covariance matrix and to calculate the VaR estimate. The ex-post performance of the portfolio is calculated as the portfolio return over one holding period. Note we ignore transaction costs, as the aim

<sup>1</sup>In practice, a factor based approach may lead to better risk forecasts. However, for simplicity, and given the number of dates is larger than the number of stocks, a full covariance matrix approach was used.

<sup>2</sup>Obviously, this would not be possible to implement in practice, but it is a simple way of introducing some forecasting power to our ex-ante stock returns. A 5 percent information content reflects the explanatory power of a typical industry cross-sectional multi factor model, and despite of these low implied R-squared, successful investment strategies can be build which such levels of explanatory power. Obviously, in practice it would not be possible to implement this strategy, as we don't know one day ahead stock returns. However, as our aim is to illustrate the effectiveness of the heuristic methods in portfolio optimisation, this is not a problem.

of the paper is not to propose a profitable trading strategy, but rather, to illustrate the portfolio construction framework, and implications of the choice of risk measure.

#### V. ANALYSIS AND DISCUSSION

We present two types of analysis. The first one aims at comparing the the GNMA with a GA in terms of quality of solutions found and efficiency. The second one aims at validating the new features of the portfolio optimization models proposed (i.e., allowing for short-selling and using VaR as a measure of risk) by checking if the solutions found make sense from a financial point of view.

##### A. Comparison of GNMA and GA

Two algorithms have been used. The GNMA above (referred to in the diagrams as NM) and a genetic algorithm (referred to in the diagrams as GA). For both algorithms, we have used standard parameter values from the literature, as they produced good results in preliminary trials. The GA uses uniform crossover with probability 0.8, bitflip mutation with a probability of  $1/n$ , where  $n = 2 \cdot n_A$  is the size of the binary string, population size of  $n + 1$  individuals, elitism and roulette-wheel selection. The GNMA uses canonical parameters for the geometric transformations ( $\alpha = 1.0, \gamma = 2.0, \rho = 0.5, \sigma = 0.5$ ) and population size of  $n + 1$  individuals. Both algorithms return the best solution found after 100,000 fitness evaluations, which are sufficient to reach high quality solutions. This allows for a fair comparison of the algorithms in terms of quality of the solution found. To be able to compare the algorithms also in terms of efficiency, the number of fitness evaluations needed to find the solution are also considered.

We have tested the GNMA and GA on the two types of problem models, mean-variance with risk aversion  $\lambda = 10000$  and bounded-VaR with threshold  $VaR_{max} = 0.020$  (referred in the diagrams as MV and BV, respectively). For each combination of algorithm and problem instance, the table I reports summary statistics on the fitness of the best solution found (fitness) and how many evaluations were needed to find it (time). The statistics refer to averages and variances on 250 days, and the values of each day are the means of 10 independent runs (as the algorithms considered are stochastic). Overall, the performance of GA and GNMA is similar, both in terms of quality of solution produced and time taken.

##### B. Financial Analysis of Optimal Portfolios

We can analyse whether the assumptions made in the financial models used were appropriate. The performance chart in Figure 2 below displays the average cumulative realised return of each of the four strategies, with the average calculated over the 10 runs. The annualised realised returns are shown in the legend. From that, we can conclude the model for stock return forecasts is appropriate, given that all four strategies deliver positive returns over the sample period, despite of forecasts only containing 5 percent of information, with the remaining 95 percent being just noise. This proves

TABLE I

SUMMARY STATISTICS OVER 250 DAYS, EACH DAY IS AVERAGED OVER TEN RUNS. EXPECTED AND REALISED RISK NUMBERS ARE IN ANNUALISED PERCENTAGES, VaR NUMBERS ARE IN DAILY PERCENTAGES.

Alg-Prob	Fitness	Time	Expected Risk	Realised Risk	Expected VaR	Realised VaR
GA-MV	$2.63 \pm 0.24$	$89410 \pm 9955$	$14.22 \pm 0.17$	$13.06 \pm 0.52$	n/a	n/a
NM-MV	$2.63 \pm 0.24$	$89179 \pm 10431$	$14.23 \pm 0.24$	$13.15 \pm 0.65$	n/a	n/a
GA-BV	$2.24 \pm 0.28$	$91710 \pm 7764$	$16.16 \pm 0.24$	$14.08 \pm 1.09$	$2.00 \pm 0.17$	$2.61 \pm 0.48$
NM-BV	$2.25 \pm 0.28$	$91603 \pm 8039$	$16.14 \pm 0.24$	$14.48 \pm 1.06$	$2.00 \pm 0.14$	$2.54 \pm 0.22$

the point that having a model with relatively low explanatory power of next day's stock return can be efficiently used to implement profitable investment strategies, when applied over a large enough set of stocks.




Fig. 2. Portfolio cumulative performance

We can also compare expected risk with realised risk of the portfolios, to assess whether the model used for forecasting risk was appropriate. Table I shows the average expected risk and the average realised risk in annualised percentages, over the 10 runs, for each of the 4 problems. Here, realised risk is calculated as the standard deviation of the realised daily portfolio returns. Expected risk is only needed in the utility function for the MV problem, but for comparison purposes we also report it here for the BV problem. We can see that the risk model used based on a historic 250 day rolling window appears to be appropriate, as the forecasted risk levels are generally in line with the realised risk exhibited by the portfolios. The same conclusion can be drawn when looking at expected VaR versus realised VaR numbers in Table I. Despite of realised VaR being slightly higher than the target 2 percent, and given that this is a small sample of only 10 runs, we can conclude the VaR constraint is achieving its objective of limiting the downside risk of the portfolios appropriately.

We used different risk measures in each of the two problems, variance for the MV problem, and VaR for the BV problem. The risk aversion parameter  $\lambda$  in the MV problem and the VaR threshold  $VaR_{max}$  in the BV problem

were calibrated such that the realised risk of both sets of portfolio returns were similar. This can be seen in Table I, which shows realised risk levels for the two problems being around 13.5 percent. By having two sets of portfolios with comparable risk levels, we can draw some conclusions about the impact of using the different risk measures. An interesting observation from an investment point of view is how the realised return for the mean-VaR portfolios is substantially higher than the mean-variance portfolios. Variance, being a symmetric measure, appears to limit the risk of the portfolio both on the downside as well as on the upside. On the other hand, VaR only limits the risk of the portfolio on the downside (i.e., losses). This has implications for investment in practice, assuming expected returns contain directional information, and also assuming that historic based VaR forecasts are a good predictor of extreme events in realised performance. If that is the case, the use of VaR as a risk measure should deliver higher returns as compared to the use of variance, everything else being equal.

## VI. CONCLUSIONS

In this paper we have proposed an interesting application of heuristic methods for portfolio optimisation in implementing a market neutral hedge fund strategy. We have shown how the recently developed GNMA method is suitable for tackling this type of problem, delivering solutions that are in line with a GA approach. We have shown for the first time that the GNMA performs well when applied to a real world problem.

The analysis of using VaR instead of variance as a risk measure has shown that VaR, by limiting downside risk, is able to deliver higher returns, while variance, by limiting both downside and upside symmetrically, delivers lower returns.

Further work will consist of testing the sensitivity of the results presented here to the choice of risk parameters, as well as testing this portfolio optimisation framework with larger datasets. We also plan to replace the return forecasting model with a more realistic one, such as one based on typically used stock selection factors in the investment industry, such as valuation, momentum, fundamentals and technical indicators.

## REFERENCES

- [1] F. R. R. Airgeadais. *Guidance Note 3/03, Undertakings for Collective Investment in Transferable Securities (UCITS) Financial Derivative Instruments*. www.financialregulator.ie, 2008.
- [2] T.-J. Chang, N. Meade, J. B. Beasley, and Y. Sharaiha. Heuristics for cardinality constrained portfolio optimization. *Computers and Operations Research*, 27:1271–1302, 2000.

- [3] G. Dueck and P. Winker. New concepts and algorithms for portfolio choice. *Applied Stochastic Models and Data Analysis*, 8:159–178, 1992.
- [4] M. Gilli and E. Killezi. A global optimization heuristic for portfolio choice with var and expected shortfall. Technical report, University of Geneva, 2001.
- [5] S. A. Kauffman and S. Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128:11–45, 1987.
- [6] D. Maringer. Portfolio management with heuristic optimization. *Advances in Computational Management Science*, 8, 2005.
- [7] H. M. Markowitz. Portfolio selection. *Journal of Finance*, 1 (7):77–91, 1952.
- [8] H. M. Markowitz. *Portfolio Selection: Efficient diversification of investments*. John Wiley and Sons, 1959.
- [9] A. Moraglio. *Towards a geometric unification of evolutionary algorithms*. PhD thesis, University of Essex, 2007.
- [10] A. Moraglio and C. Johnson. Geometric generalization of the nelder-mead algorithm. In *Proceedings of the 10th European Conference on Evolutionary Computation in Combinatorial Optimization*, 2010.
- [11] J. A. Nelder and R. A. Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.
- [12] M. Speranza. A heuristic algorithm for a portfolio optimization model applied to the milan stock market. *Computers and Operations Research*, 23:433–441, 1996.
- [13] T. Weise. *Global Optimization Algorithms - Theory and Application*. on-line ebook, 2009.