Cardinality constrained portfolio optimisation

Jonathan E. Fieldsend[1] and John Matatko[2] and Ming Peng[2]

[1] Department of Computer Science, University of Exeter, UK
[2] Department of Accounting and Finance, University of Exeter, UK

Abstract The traditional quadratic programming approach to portfolio optimisation is difficult to implement when there are cardinality constraints. Recent approaches to resolving this have used heuristic algorithms to search for points on the cardinality constrained frontier. However, these can be computationally expensive when the practitioner does not know *a priori* exactly how many assets they may desire in a portfolio, or what level of return/risk they wish to be exposed to without recourse to analysing the actual trade-off frontier.
This study introduces a parallel solution to this problem. By extending techniques developed in the multi-objective evolutionary optimisation domain, a set of portfolios representing estimates of all possible cardinality constrained frontiers can be found in a single search process, for a range of portfolio sizes and constraints. Empirical results are provided on emerging markets and US asset data, and compared to unconstrained frontiers found by quadratic programming.

1 Introduction

Given constraints on the number of stocks to hold in a portfolio, due to the costs of monitoring and portfolio re-weighting, the question arises as to how to choose the 'best' portfolio given particular risk/return preferences and these cardinality constraints. The number of stocks needed to achieve a particular diversification gain depends on the correlation among stock returns; the lower the correlation, the more stocks are needed. Campbell *et al.* [1] show that although overall market volatility has not increased in recent years, individual stock returns have become less correlated with each other. Consequently, more stocks are thought to be needed in a portfolio than in the period studied by [2] (1926-1965).

Markowitz [3] defined the set of optimal portfolios that are on the efficient frontier, based on estimated moments of the sampled distribution. Ignoring the uncertainty inherent in the underlying probability model, the portfolio that maximizes expected return, through to the one which minimises risk, is in this set. This approach reduces the task of choosing an optimal portfolio to choosing a portfolio from this efficient frontier.

When there are no cardinality constraints, quadratic programming (QP) can be effectively used to determine an optimal portfolio's weights, given 'a' different assets. With μ_i being the expected return of the ith asset, and $V_{i,j}$ the covariance between assets i and j, this takes the general form: $\min\{\sigma_P\}$, where

$\sigma_P = \sqrt{\Sigma_{i=1}^{a}\Sigma_{i=1}^{a} w_i w_j V_{i,j}}$, subject to $\Sigma_j^a w_i \mu_i = r_P^*$ and $\Sigma_i^a w_i = 1$, where w_i is the portfolio weight of the ith asset, and r_P^* is the desired expected return of the portfolio. In more formal notation, expected return of a portfolio (r_P) equals $\boldsymbol{\mu}^T \mathbf{w}$ and risk (σ_P) equals $\mathbf{w}^T V \mathbf{w}$. Commonly there is also a no short selling constraint $0 \leq w_i \leq 1$. However, when realistic cardinality constraints are imposed QP cannot be applied to find optimal subsets.

We introduce a heuristic model to search for all the cardinality constrained (CC) efficient portfolio frontiers available for any set of assets. We illustrate the differences that arise in the shape of this efficient frontier when such constraints are present, and test for significant difference between these frontiers and the unconstrained efficient (UC) frontier found using quadratic programming.

2 Heuristic search methods

Heuristic search approaches addressing the portfolio search problem when cardinality constraints are imposed have recently been developed [4,5] . These have taken the form of optimising a composite weighting of σ_P and r_P, typically $\max\{\lambda r_P - (1-\lambda)\sigma_P\}$, where $0 \leq \lambda \leq 1$. Problematic however is that the cardinality to be searched needs to be defined *a priori*, and a separate run is needed for each portfolio to be optimised. If a range of cardinalities need to be compared then obviously this increases the computational cost further.

Multi-objective evolutionary algorithms (MOEAs) represent a popular approach to confronting these types of problem by using evolutionary search techniques [6]. Here we investigate the use of a modified MOEA to optimise CC portfolio frontiers in parallel (optimising all plausible values of k in a single process). Novel search processes are incorporated in this algorithm to enable it to maintain these disparate frontier sets and efficiently compare new portfolios during the search process. Prior to this however the concepts of Pareto optimality (central to modern MOEAs) and non-dominance will be briefly described.

2.1 Pareto optimality

The multi-objective optimisation problem seeks to simultaneously extremise D objectives: $y_i = f_i(\mathbf{w})$, where $i = 1, \ldots, D$ and where each objective depends upon a vector $\mathbf{w}$ of n parameters or decision variables. The parameters may also be subject to the m constraints: $e_j(\mathbf{w}) \geq 0$ where $j = 1, \ldots, m$. In the context of portfolio optimisation, these may be constraints on the maximum/minimum proportion of a portfolio that can be derived from a particular market or sector, or a minimum/maximum weight an asset can have in a portfolio.

Without loss of generality it is assumed that these objectives are to be minimised (minimising $-1 \times r_P$ is analogous to maximising r_P), as such the problem can be stated as: minimise $\mathbf{y} = \mathbf{f}(\mathbf{w}) = (f_1(\mathbf{w}), f_2(\mathbf{w}), \ldots, f_D(\mathbf{w}))$, subject to $\mathbf{e}(\mathbf{w}) = (e_1(\mathbf{w}), e_2(\mathbf{w}), \ldots, e_m(\mathbf{w}))$. A decision vector $\mathbf{u}$ is said to *strictly dominate* another $\mathbf{v}$ (denoted $\mathbf{u} \prec \mathbf{v}$) if $f_i(\mathbf{u}) \leq f_i(\mathbf{v})\ \forall i = 1, \ldots, D$ and $f_i(\mathbf{u}) < f_i(\mathbf{v})$ for some i. A set of M decision vectors is said to be a *non-dominated set*

Algorithm 1 Algorithmic description of the multi-objective optimiser.

g, *maximum number of algorithm iterations*
$\mathcal{H}$, *Set of sets of portfolios defining the c different estimated frontiers*
$t := 0$, $\mathcal{H}^t_k = \emptyset \ \forall k = 1, \ldots, a$
$\mathcal{H}^t_{k,1} := random\_portfolio(k) \ \forall k = 1, \ldots, a$
while $(t < g)$
    $k = U(1, a)$
    $\mathbf{w} := select(\mathcal{H}^t, k)$
    $\mathbf{w} := adjust(\mathbf{w})$
    $\mathbf{y} := evaluate(\mathbf{w})$
    $\mathcal{H}^{t+1} = check\_insert\_remove(\mathcal{H}^t, \mathbf{w}, \mathbf{y})$
    $t := t + 1$
end

(an estimated Pareto front) if no member of the set is dominated by any other member. The *true* Pareto front is the non-dominated set of solutions which are not dominated by any feasible solution. In the context of portfolio optimisation, the efficient frontier can be seen as an example of a Pareto optimal set.

2.2 The proposed model

If we are not concerned with constraints other than cardinality, then cardinality can be incorporated as a third objective to be minimised, therefore aiming to find the 3-dimensional surface defining the trade-off between risk, return and cardinality minimisation. We can then extract the 2-dimensional cardinality constrained frontier for any particular k. According to finance theory for higher cardinality levels (more assets) the CC front extracted will be short, as identical risk/return levels may be available at a lower cardinality for high r_P and σ_P. With no other constraints this is not a problem as the lower cardinality portfolios are equivalent to higher cardinality portfolios with some weights equal to zero. If however there are other constraints (as mentioned in Section 2.1), this transformation may no longer be possible. As such CC portfolio optimisation with MOEAs is interesting, as we need to maintain a separate estimated Pareto set for each cardinality, k. (NB, in the empirical results shown here there are no additional constraints, so the cardinality constraint is effectively a maximum one).

One solution would be to run separate 2-objective MOEAs for each k, (extending the approach used previously in [7] for UC optimisation). However, although more computationally efficient than the existing heuristic methods used in the application domain, this may still be computationally expensive for a large number of k. Instead here we search for each k constrained front in parallel, and constructively use information from each of these fronts to improve the search process of the others. The decision vector $\mathbf{w}$ here consists of the weights of the a different available assets. A description of the MOEA used, based on a simple (1+1)-evolution strategy [8] is provided in Algorithm 1.

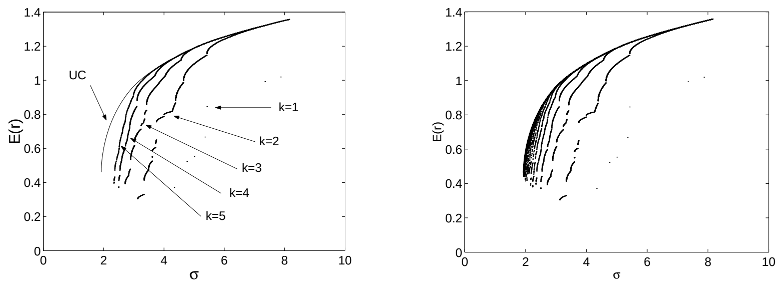

Figure 1. *Left*: EM UC frontier found exactly using QP and the first 5 optimised CC fronts. *Right*: All 40 optimised EM CC fronts and the UC frontier.

The algorithm maintains a set of sets $\mathcal{H}$ of the a different CC efficient frontiers. Each of these $\mathcal{H}_k$ cardinality sets is initialised with a random portfolio (line 2), with random non-negative w_is, where $\Sigma_{i=1}^{a} w_i = 1$ and the number of non-zero w_i equals k. The algorithm proceeds at each iteration by first randomly selecting an archive with cardinality k at that generation t, $\mathcal{H}_k^t$, and copying a portfolio from it (using partitioned quasi-random selection [8]). This copied portfolio $\mathbf{w}$ is then adjusted (line 6); this takes the form of weight adjustment 50% of the time, otherwise a weight adjustment plus dimensionality change. When only weights are adjusted $\mathbf{w} :\sim Dir(b\mathbf{w})$, where Dir denotes the Dirichlet distribution. By sampling a Dirichlet with parameters equal to the current weights, and a large enough multiplying constant b (set at 400 here), the new weights will be close to the original values [9]. In addition, sampling from a Dirichlet ensures $\Sigma_{i=1}^{a} w_i = 1$, $w_i \geq 0$, and any w_i that were previously zero will remain zero. When dimensionality change is also implemented either a non-zero value in $\mathbf{w}$ set to zero (asset removal) or a zero valued w_i is assigned a value drawn from $U(0, 1/k)$ (where U denotes the Uniform distribution and k is the number of active assets in the *new* portfolio). In both cases new weights are then drawn from $Dir(b\mathbf{w})$. The new portfolio $\mathbf{w}$ is evaluated on line 7, r_P and σ_P are then assigned to $\mathbf{y}$. Using $\mathbf{y}$ the new portfolio can be compared to the relevant $\mathcal{H}_k^t$ (line 8), to see if it is non-dominated, and if so, any dominated portfolios are removed when it is inserted into $\mathcal{H}_k^{t+1}$.

3 Empirical results

Empirical evaluation of this new method is provided here on weekly stock data from the US S&P 100 index and emerging markets (EM) stock. Both sets were obtained from Datastream and were those stocks within the index that persisted from January 1992 to December 2003. Using the first 500 points of the returns data, Algorithm 1 was run for 10^7 portfolio evaluations.

Figure 1 shows the CC frontiers found by the MOEA on the EM asset set, the CC frontiers rapidly approaching the UC frontier (also shown) as k is increased. Using the methodology described in [10], coupled with methods developed in the

Table 1. Cardinality for which the CC is not significantly better than the UC segment.

Partition (low risk to high risk)	1	2	3	4	5	6	7	8	9	10
EM opt. CC level when no sig difference to UC	-	8	5	4	4	3	3	2	2	2
S&P opt. CC level when no sig difference to UC	-	9	6	5	4	3	3	3	2	2

Table 2. *Ex ante* performance of minimum variance CC portfolios ($k = 10$).

	EM					S&P				
t (weeks)	4	12	26	52	104	4	12	26	52	104
σ_P of opt. CC portfolio	0.93	3.26	2.58	2.42	2.20	0.55	2.86	2.35	2.98	3.34
Mean of bootstrap CCs σ_P	2.13	3.39	3.05	2.77	2.55	1.44	2.98	2.63	3.08	3.33
Rank of opt CC (/1001)	279	414	120	134	96	45	481	307	579	623

MOEA community, we can test if the optimised CC frontiers are significantly different from the UC frontier. We do this by sampling means and covariances from the posterior distribution of the S&P and EM data, and evaluating the optimised portfolios with respect to these. The UC frontier is sliced into 10 evenly spaced segments (with respect to σ_P). The portfolios lying in each segment, and those on the CC frontier defined by the σ_P bounds of the relevant UC segment are then re-evaluated with new covariances and means sampled from the posterior distribution of the stocks. The resultant sets of points are then compared by calculating the proportion of points $\mathbf{y}$ in each set which are strictly dominated by those in the other set [8]. This is performed for 1000 different samples from the posterior, and the difference in the respective proportions assessed as to whether they are statistically different from 0. Table 1 shows that for both asset sets, apart from the lowest risk levels, the optimised UC portfolios are not significantly different from the CC frontier except for a relatively small k.

The previous analysis has shown that it is possible to replicate closely the mean and variance of an efficient portfolio with a portfolio composed of a relatively small number of stocks. The analysis is however *ex post* and might be of little help to portfolio managers concerned with choosing investment positions *ex ante*. Thus while we have shown that optimising a (CC) portfolio could take us close to the UC frontier, it might be that the benefits that we know exist *ex post* are impossible to realise *ex ante*. In the preliminary proof of concept study here, we solely concentrate on the *global minimum variance portfolio*. This frees us from the difficulty of having to estimate expected returns on the portfolio [10]. To actually evaluate the performance of our historically optimal CC portfolio we use a bootstrapping technique. We draw without replacement k random integers, all between 1 and a. The global MV portfolio formed from stocks in our dataset corresponding to the k integers is then calculated using QP. The standard deviation of the weekly returns for the portfolio is derived, beginning with week 501 for a period of t weeks, t = 4, 12, 26, 52, 104. This is repeated 1000 times. Descriptive statistics are reported in Table 2 for $k = 10$.

In both asset sets the largest reduction in volatility, compared to the mean bootstrapped CC portfolio, is in the shortest period (t=4) after portfolio formation. After this, performance is mixed with the US data showing only a small or no improvement for longer horizons, while the EM *ex ante* optimal portfolio does offer 14.7% risk reduction at the 104 week horizon. The short term result does appear consistent with ARCH type modelling of volatility movements.

4 Discussion

A new approach to discovering cardinality constrained portfolios has been described here, using MOEAs to search for and maintain a number of different CC frontiers. Empirical results show that *ex post* with a relatively small cardinality level one can attain performance that is not significantly different than the unconstrained frontier and *ex ante* there is some level of persistence in the portfolio weights found. The authors are currently applying the methodology to both within and across markets, with additional constraints recommended by fund managers. In addition the authors are investigating non-Markovitz formulations of risk, and the use of higher moments within the general framework.

Acknowledgements

We thank Richard Everson for his helpful comments. JEF is supported by the EPSRC, grant GR/R24357/01. MP is supported by a university scholarship.

References

1. Campbell, J., Lettau, M., Malkiel, B., Xu, Y.: Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. J. Finance **56** (2001) 1–43
2. Fisher, L., Lorie, J.: Some studies of variability of returns on investments in common stocks. J. Bus. **43** (1970) 99–134
3. Markowitz, H.: Portfolio Selection: Efficient Diversification of Investments. John Wiley & Sons (1959)
4. Chang, T.J., Meade, N., Beasley, J., Sharaiha, Y.: Heuristics for cardinality constrained portfolio optimisation. Comp. & Operations Res. **27** (2000) 1271–1302
5. Doerner, K., Gutjahr, W., Hartl, R., Strauss, C., Stummer, C.: Ant colony optimization in multiobjective portfolio selection. In: Proceedings of the 4th Metaheuristics Intl. Conf., MIC'2001, July 16–20, Porto, Portugal. (2001) 243–248
6. Deb, K.: Multi-objective optimization using evolutionary algorithms. Wiley (2001)
7. Lin, D., Wang, S., Yan, H.: A multiobjective genetic algorithm for portfolio selection. In: Proceedings of ICOTA 2001, December 15-17, Hong Kong,. (2001)
8. Fieldsend, J., Everson, R., Singh, S.: Using Unconstrained Elite Archives for Multi-Objective Optimisation. IEEE Trans. on Evol. Comp. **7** (2003) 305–323
9. Larget, B., Simon, D.: Markov chain Monte Carlo algorithms for the Bayesian analysis of phylogenetic trees. Mol. Bio. & Evol. **16** (1999) 750–759
10. Li, K., Sarkar, A., Wang, Z.: Diversification benefits of emerging markets subject to portfolio constraints. J. Empirical Finance **10** (2003) 57–80