An EA for Portfolio Selection over Multiple Investment Periods with Exponential Transaction Costs

Matthew J. Craven
University of Abertay Dundee
Bell Street, Dundee, DD1 1HG, UK
m.craven@abertay.ac.uk

Henri C. Jimbo
American University of the Middle East
P. O. Box 220, Dasman 15453, Kuwait
jimbo.claver@aum.edu.kw

ABSTRACT

We study a matrix representation for an EA attack on the CCPOP with transaction costs. The representation is based on portfolio sequences which change over the investment lifetime in response to asset price changes. We show the approach is effective and that EA performance is directly related to asset price correlation. We compare the EA with a matrix hillclimber and show some common results of vector representations do not hold for a matrix one, potentially providing a step forward in performance of such algorithms.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; G.1.6 [**Optimization**]: Constrained optimization; I.6.3 [**Simulation and Modeling**]: Applications

Keywords

Portfolio optimization; evolutionary algorithm; matrix representation; transaction cost

1. INTRODUCTION

The Portfolio Optimization Problem (POP) seeks optimal numbers of different assets that may be owned at a given time, with certain criteria met. Often the criteria are risk and return, investors seeking maximum return for minimum risk. Optimization subject to such criteria has been studied since the 1950s, when the POP was posed as a QOP. Indeed, simple POP flavors were tackled by quadratic programs [3]. The cardinality constrained POP (CCPOP) [3, 4] is a variant where, to spread risk, only k out of a total n assets are bought per period. Attacks range from time series prediction to evolutionary methods [1, 3, 4]. The work [4], via Sharpe's Index, compares three EAs that optimize weight vectors, and [1] attacks the CCPOP with rounded lots. A GA and hillclimber, both with vector representations, are compared, showing negligible difference in results.

This work extends our work in [5], using a matrix EA representation to optimize a sequence of vectors (portfolios), one for each period in the investment lifetime. This representation makes the CCPOP of higher complexity in number of assets. We compare our EA with a hillclimber to determine if the conclusions of [1] hold for our representation. Matrices have been used in multiperiod asset allocation, but, to our knowledge, not as an EA representation. For example, [2] used matrices to calculate fitness values but the EA representation was a string. Trading *transaction costs* tend to be ignored in simpler POP treatments, but are an issue. They are hard to model as, in practice, there are many costs involved and social costs are not considered. There have been various cost models, from constant to piecewise linear [8]. We take transaction cost as a function of asset risk, echoing practice where investors pay more for higher risk in expectation of larger returns. Reflecting this and future social levies, we use an exponential model with small coefficients.

GECCO'13 Companion, July 6–10, 2013, Amsterdam, The Netherlands.
ACM 978-1-4503-1964-5/13/07.

2. THE PROBLEM AND EA

We extend standard notation from [3]. The proportion of an asset to buy at a given period is its *weight*. The $n \times n$ weight matrix is $W = (w_{ij})$, with w_{ij} the asset j weight at period i. Row i of W is a vector denoted $W(i,:)$. This vector is the standard representation in POP models, but our representation is matrix-based. Let $k \leq n$ be the constant number of assets bought; the choice of which k assets are bought may change at each period. The initial data, X, is the normalized prices of n assets over n periods. The investment value matrix over all periods is $P = X.*W$, with total *value* at period i being $v_i(W) = \sum_{j=1}^{n} P_{ij}$. The vectors of expected return and risk of W over all periods are respectively written $\overline{R}(W)$ and $r(W)$ [5]. Extending the COP formulation of [1, 3] to include matrices, the constraints are:

$\sum_{j=1}^{n} w_{ij} = 1$ for all i;

$w_{ij} \in [\gamma_i, \delta_i]$ with $0 \leq \gamma_i \leq \delta_i \leq 1$ for all i;

$-\Delta v_i(W) \leq 0.01 v_i(W)$ for all i;

$Z_i = n - k$ for all i.

Our CCPOP objective function, to be minimized, is

$$C(W) = \left\| \begin{array}{c} \lambda r(W) - (1-\lambda)\overline{R}(W) \\ + (1-\lambda)\left(\alpha e^{\beta \overline{R}(W)} - 1\right) \end{array} \right\|_2 - V(W).$$

An optimal weight matrix is denoted W^*. Constant λ represents investor risk preference, and transaction cost is taken from return $\overline{R}$ [6]. The number of zero entries in $W(i,:)$ is Z_i, and $V(W) = \sum_{i=1}^{n} v_i(W)$ is taken from the norm to favor acceptance of larger-valued portfolios. The EA returns the best found weight matrix by the above constraints. Inputs are the normalized price matrix X (the *instance*) and initial (*naive*) weight matrix W_0 (produced at random [7] with $n-k$ zeros in each row at random positions and each row sum as 1). The main operator on W is *perturbation*:

for a given value m, add a random vector ε, with $\sum \varepsilon = 0$ and $\sum(W(m,:) + \varepsilon) = 1$, to the vector of non-zero entries of row $W(m,:)$. There is also a *feasibility* operator. Selection is by a partially-elitist 3-tournament scheme. Mutation perturbs a random row, and crossover perturbs every row, of a weight matrix W chosen u.a.r. from the top 10% of the population. We chose population size 100 with $n_c = 30$ crossovers, $n_s = 10$ selections and $n_m = 60$ mutations.

Instance RW-8 covers eight real-world indices from 2003-10, sampled yearly [5], and RW-20 covers twenty, sampled monthly from 03/2011-10/2012. Simulated instances are matrices of random numbers in $[0, 1]$. Intuitively, real-world instances likely have high correlation between assets over subsequent periods, but simulated instances may not. To test this, for each instance we calculated the mean absolute correlation between all assets over all periods (Table 1). Observe that correlation in real-world instances is notably larger than in simulated instances. We chose medium risk aversion $\lambda = 0.5$, cardinality constraint $k = \frac{n}{2}$ and constants $\alpha = 1, \beta = 0.5$. All trials ran for 5000 generations on a Core Duo 1.8GHz computer with 1GB RAM running MATLAB.

Table 1: Mean absolute correlation by instance.

Instance	RW-8		RW-20	
Mean abs. corr.	0.5316		0.4794	
Instance	Sim-8	Sim-20	Sim-30	Sim-40
Mean abs. corr.	0.3221	0.1809	0.1524	0.1283

3. RESULTS

Each instance was trialled on both algorithms thirty times, every trial using a distinct random naive weight matrix W_0. Let $I(W^*, W_0)$ be the percent gain in value from naive to best found weight matrix, with μ_I, σ_I and σ_C respectively the mean of I and standard deviations of I and C over the given instance. Table 2 summarizes EA results.

Table 2: Summary statistics for the EA.

Instance	RW-8		RW-20	
$\mu_{I(W^*,W_0)}$	120.21		331.32	
$\sigma_{I(W^*,W_0)}$	56.37		59.65	
$\sigma_{C(W^*)}$	0.66		0.96	
Instance	Sim-8	Sim-20	Sim-30	Sim-40
$\mu_{I(W^*,W_0)}$	33.52	40.87	41.23	40.40
$\sigma_{I(W^*,W_0)}$	10.77	7.98	4.63	3.68
$\sigma_{C(W^*)}$	0.36	0.31	0.41	0.36

The results show that, in all simulated instances, we obtain large mean gain from naive to best found weight matrices, and these gains increase with the number of assets. Strikingly, in real-world instances, mean gains are much larger. We conjecture this is due to increased mean absolute correlation compared to the simulated instances (Table 1), and, moreover, EA performance follows the level of asset correlation. Also, as the number of simulated assets increases, EA performance standard deviation decreases. We may interpret this result as evidence of EA robustness.

Emulating a hillclimber, we take population size two and parameters $(n_s, n_m, n_c) = (1, 1, 0)$. Table 3 summarizes hillclimber results. Gains for real-world over simulated instances are smaller than those exhibited by the EA. Table 3 also implies that gains in value are directly related to the absolute correlation between assets in X.

Table 3: Summary statistics for the hillclimber.

Instance	RW-8		RW-20	
$\mu_{I(W^*,W_0)}$	79.78		61.53	
$\sigma_{I(W^*,W_0)}$	33.51		15.09	
$\sigma_{C(W^*)}$	0.57		0.61	
Instance	Sim-8	Sim-20	Sim-30	Sim-40
$\mu_{I(W^*,W_0)}$	27.25	21.48	11.59	7.23
$\sigma_{I(W^*,W_0)}$	7.90	4.50	1.53	0.89
$\sigma_{C(W^*)}$	0.42	0.45	0.42	0.47

4. COMPARISON AND CONCLUSION

Comparing both algorithms, Tables 2–3 show that the difference in μ_I from hillclimber to EA increases with n. Thus the advantage of EA over hillclimber increases with instance difficulty. In all instances, the difference in σ_I between EA and hillclimber also generally increases with n, implying the EA produces a wider range of solutions than the hillclimber. Hence the importance of the crossover operator (and a larger population) increases with n. Work [1] compares a vector EA and hillclimber, showing the EA offers negligible performance gain over the hillclimber. We conclude the opposite for the matrix representation, especially for larger numbers of simulated assets. For real-world assets, the performance difference is very clear. We also showed a direct relationship between asset correlation and algorithm performance. Our analysis shows the time-dependency of individual portfolios in the sequence; vector representation approaches do not. Overall, we have a promising technology for improving the accuracy of multiperiod portfolio optimization.

5. REFERENCES

[1] J. Arriaga and M. Valenzuela-Rendón. Steepest ascent hill climbing for portfolio selection. *LNCS*, 7248:145–154, 2012.

[2] S. Baglioni, C. de Costa Pereira, D. Sorbello, and A. Tettamanzi. An evolutionary approach to multiperiod asset allocation. *LNCS*, 1802:225–236, 2000.

[3] T. J. Chang, N. Meade, J. E. Beasley, and Y. M. Sharaiha. Heuristics for cardinality constrained portfolio optimisation. *Comp. Op. Res.*, 27:1271–1302, 2000.

[4] F. E. Colomine Duran, C. Cotta, and A. J. Fernández-Leiva. A comparative study of multi-objective evolutionary algorithms to optimize the selection of investment portfolios with cardinality constraints. *LNCS*, 7248:165–173, 2012.

[5] H. C. Jimbo, M. J. Craven, S. I. Ngongo, and T. Suzuki. Optimization of investment stock portfolios with stochastic constraints. In *7th Intl. Conf. on Nonl. Analysis and Conv. Analysis - I*, pages 127–141, 2013.

[6] J. M. Mulvey. Incorporating transaction costs in models for asset allocations. In *Financial Optimization*, pages 241–259. CUP, 1993.

[7] R. Stafford. Randfixedsum, 2006. From `http://www.mathworks.com/matlabcentral/fileexchange/9700`.

[8] A. Yoshimoto. The mean-variance approach to portfolio optimization subject to transaction costs. *J. Op. Res. Soc. Japan*, 39(1):99–117, 1996.