Applied Soft Computing 36 (2015) 125–142

Contents lists available at ScienceDirect

Applied Soft Computing

journal homepage: www.elsevier.com/locate/asoc

A memetic algorithm for cardinality-constrained portfolio optimization with transaction costs

Rubén Ruiz-Torrubiano [a,*], Alberto Suárez [b]

[a] Gruber&Petters, Belvederegasse 11, 2000 Stockerau, Austria
[b] Computer Science Department, Escuela Politécnica Superior, Universidad Autónoma de Madrid, 28049 Madrid, Spain

ARTICLE INFO

Article history:
Received 13 August 2013
Received in revised form 25 February 2015
Accepted 30 June 2015
Available online 17 July 2015

Keywords:
Genetic algorithms
Combinatorial optimization
Portfolio selection
Transaction costs

ABSTRACT

A memetic approach that combines a genetic algorithm (GA) and quadratic programming is used to address the problem of optimal portfolio selection with cardinality constraints and piecewise linear transaction costs. The framework used is an extension of the standard Markowitz mean–variance model that incorporates realistic constraints, such as upper and lower bounds for investment in individual assets and/or groups of assets, and minimum trading restrictions. The inclusion of constraints that limit the number of assets in the final portfolio and piecewise linear transaction costs transforms the selection of optimal portfolios into a mixed-integer quadratic problem, which cannot be solved by standard optimization techniques. We propose to use a genetic algorithm in which the candidate portfolios are encoded using a set representation to handle the combinatorial aspect of the optimization problem. Besides specifying which assets are included in the portfolio, this representation includes attributes that encode the trading operation (sell/hold/buy) performed when the portfolio is rebalanced. The results of this hybrid method are benchmarked against a range of investment strategies (passive management, the equally weighted portfolio, the minimum variance portfolio, optimal portfolios without cardinality constraints, ignoring transaction costs or obtained with L_1 regularization) using publicly available data. The transaction costs and the cardinality constraints provide regularization mechanisms that generally improve the out-of-sample performance of the selected portfolios.

1. Introduction

The classical framework for the selection of optimal portfolios was established by Markowitz in [25]. The problem consists in finding the allocation of a fixed budget in a universe of assets that maximizes the expected return from the investment in a given period while minimizing the corresponding risk. Since the future evolution of stock returns is uncertain, these returns are modeled as random variables. In the standard Markowitz formulation the risk is quantified in terms of the variance of the portfolio returns. Portfolio selection is therefore a multiobjective optimization task with two conflicting goals: the maximization of profit and the minimization of risk. It can be formulated as a quadratic programming (QP) problem that can be readily solved using standard quadratic optimization techniques. One of the shortcomings of the standard (unconstrained) Markowitz framework is that small variations in the inputs of the model (i.e. in the vector of expected values or in the covariance matrix of the asset returns) often lead to large changes in the composition of the resulting portfolios [7]. Another important drawback is that the portfolios selected generally have poor out-of-sample performance (see [29] and references therein). To address these issues additional constraints can be considered in the model. For instance, it is possible to include no short-selling constraints [16], which restrict all portfolio weights to be non-negative. Besides improving the robustness and performance of the portfolios, these additional constraints reflect actual restrictions in real-world applications. In particular, it is necessary to take into account the impact of the transaction costs incurred when the portfolio is rebalanced. Furthermore, limiting the number of assets in which the portfolio invests makes the management of the portfolio simpler. When cardinality constraints are included the problem becomes NP-Complete [31]. Standard QP solvers can no longer be used to address the portfolio selection problem. Therefore, one needs to resort to other types of methods to find near-optimal solutions at a reasonable computational cost.

In this paper, we propose to use a memetic approach that combines a genetic algorithm (GA) [15] with an extended set encoding

* Corresponding author. Tel.: +43 2266 62241.
E-mail addresses: ruben.rtorrubiano@grupet.at (R. Ruiz-Torrubiano), alberto.suarez@uam.es (A. Suárez).

http://dx.doi.org/10.1016/j.asoc.2015.06.053
1568-4946/

and quadratic programming (QP) to address the problem of portfolio optimization, taking into account transaction costs and other realistic constraints, such as cardinality constraints (restrictions on the maximum number of assets in the portfolio), minimum trading size constraints (restrictions on the minimum amount of assets that can be bought or sold), minimum and maximum bounds on variables or groups of variables (to limit the fraction of investment in a particular asset or group of assets) and no short-selling constraints (the portfolio weights are non-negative).

The proposed memetic approach is benchmarked against other portfolio selection algorithms in experiments that quantify both the in-sample and out-of-sample performance. In-sample performance measures are used to assess how effective the optimization algorithm is and to what extent do the constraints considered affect the value of the objective function that is being optimized. However, practitioners are primarily interested in the out-of-sample performance of the portfolio: using the information that is available at the time of the investment, how does one allocate a fixed budget among different assets so that the expected future return of the portfolio is maximized, while minimizing the corresponding risk? In this respect, the results of the current investigation are in agreement with the observation that in-sample performance is generally not a good predictor of out-of-sample gains [28,10,2]. The situation is analogous to the problem of *overfitting* in supervised learning [3]: high predictive accuracy in the training data does not guarantee a good generalization performance (i.e. high predictive accuracy in unseen instances).

A novel contribution of the memetic approach proposed in this work is the use of a set encoding for the candidate solutions that specifies not only which assets are included in the portfolio but also the type of trade (buy/hold/sell) that is carried out for each asset to rebalance the portfolio. The RAR crossover operator, which was introduced in [38], is adapted in this work to this extended set representation. The adapted RAR crossover operation produces individuals that satisfy all the constraints in the problem, so that no penalty functions or repair mechanisms are needed. An additional contribution of the current investigation is to illustrate how cardinality constraints and transaction costs act as regularization terms. Including these constraints in the optimization problem generally improves the robustness and out-of-sample performance of the portfolios.

The paper is organized as follows: previous approaches to the problem are reviewed in Section 2. In Section 3 the problem of mean–variance portfolio selection under transaction costs is described. Section 4 introduces the memetic approach proposed in this work to address the problem. The effectiveness of this approach is illustrated in experiments whose results are presented and discussed in Section 5. Finally, Section 6 summarizes the conclusions and perspectives of the current investigation.

2. Related work

The problem of optimal portfolio selection has been extensively studied in the literature [6]. It can be formulated in two different settings: single-period (static) and multi-period (dynamic) portfolio selection. Single-period portfolio selection consists in allocating a fixed budget among a collection of assets during a period in which the composition of the portfolio is held constant. The objective is to achieve an optimal tradeoff between the expected return and the risk of the investment [26]. In a multi-period setting the goal is to identify optimal investment policies that involve dynamically trading of the portfolio assets up to a specified time horizon [27,44,4]. In this case one typically maximizes an utility function that takes into account the expected future returns and the risk of the porftolio up to the investment horizon. It is also possible to design a dynamic investment strategy using a myopic approach, in which the multi-period problem is cast as a sequence of consecutive single-period problems. If the distribution of asset returns is predictable an optimized multi-period strategy should outperform a myopic one. These types of advantages are known in the financial literature as hedging demands [6]: by deviating from the single-period portfolio choice one tries to hedge against future changes in the investment conditions. However, hedging demands appear to be difficult to realize in practice: as a consequence of the uncertainty in the prediction of the time-varying distribution of asset returns [35] the differences in the out-of-sample performance of a myopic and a truly dynamic approach in actual multi-period portfolio selection problems are fairly small [10,2].

The goal of this work is to extend the Markowitz mean–variance framework [26] to consider realistic constraints and, specially, transaction costs. The effects of transaction costs in a single period setting are generally small. However, because costs are non-negative, their effects accrue with time. In consequence, they eventually become a significant factor in the design of long-term investment strategies. Therefore, to assess the out-of-sample performance of the portfolios identified by the different investment strategies we adopt the “rolling window” procedure used in [28]: a sequence of investment decisions is considered. At the beginning of each investment period the composition of the portfolio is determined on the basis of the information on the asset returns in a time-window of fixed size that immediately precedes that period. We then compute and store the portfolio returns (including transaction costs) in the period under consideration. After that we slide the fixed-size window forward and select the portfolio that will be held in the following investment period. These steps are repeated until the final investment horizon is reached. Although the objective of these rolling window experiments is not to address the multi-period portfolio selection problem, empirical studies have shown that the differences between the performance of sequential single-period investment strategies and the corresponding optimized dynamic trading strategies are generally small [10,2]. Therefore, myopic strategies, which are simpler to compute, are often preferred to truly multi-period ones in practice.

The extensions of the standard Markowitz framework to consider realistic constraints and transaction costs generally lead to complex (NP-complete) optimization problems that cannot be addressed using exact numerical techniques. In this work we propose metaheuristics to address this problem. The application of evolutionary and biologically inspired methods to financial problems has been the object of recent interest in the soft computing community [5]. One of the first investigations in which a genetic algorithm was used to address this problem in [8]. In that work the space of cardinality constrained portfolios was searched using a combination of a genetic algorithm [15], tabu search [12] and simulated annealing [19]. The candidate solutions were encoded using chromosomes composed of both discrete and real genes. No transaction costs were considered. As a consequence of the mixed encoding used, the crossover and neighborhood operators needed to handle both the discrete and the continuous constraints of the problem. This complicated the design of these operators. In the approach proposed in the current work the continuous part of the problem is handled separately by specialized techniques. In this manner the genetic algorithm can focus on the combinatorial search.

In [53] the cardinality constraints were handled using a clustering algorithm to reduce the size of the investment universe: after performing K-means clustering, the investor selects one asset from each cluster. The optimization is then carried out in the restricted space of the selected assets using a genetic algorithm with real-valued encoding, arithmetic variable-point crossover and real uniform mutation. Lower and upper bounds on the fraction of