Expert Systems with Applications 39 (2012) 10722–10730

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

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

Time-stamped resampling for robust evolutionary portfolio optimization

Sandra García [1], David Quintana, Inés M. Galván *, Pedro Isasi

Computer Science Department, Carlos III University of Madrid, Avda. Universidad 30, 28911 Leganes, Spain

ARTICLE INFO

Keywords:
Financial portfolio optimization
Robust portfolio
Multiobjective evolutionary algorithms

ABSTRACT

Traditional mean–variance financial portfolio optimization is based on two sets of parameters, estimates for the asset returns and the variance–covariance matrix. The allocations resulting from both traditional methods and heuristics are very dependent on these values. Given the unreliability of these forecasts, the expected risk and return for the portfolios in the efficient frontier often differ from the expected ones. In this work we present a resampling method based on time-stamping to control the problem. The approach, which is compatible with different evolutionary multiobjective algorithms, is tested with four different alternatives. We also introduce new metrics to assess the reliability of forecast efficient frontiers.

1. Introduction

Asset allocation is one of the core topics in financial management. The search for the optimal choice of financial assets to be included in a portfolio has been the subject of research for a long time and it is one of the most active lines in finance.

The academic literature on this subject is very large and mostly based on the work of Markowitz (1952) and Markowitz (1959). Under this framework, the problem is introduced as a multiobjective optimization problem where the investor tries to figure out the weight that each of the investment alternatives should carry in the portfolio. The target of this investor would be both minimizing risk and maximizing return at the same time. The solution to the problem of optimizing for these two objectives in conflict defines a set of solutions called the efficient frontier. This Pareto front consists on portfolios that are neither better or worse than the rest. For each level of risk or return, there is no better alternative in terms of the other objective. This makes the election of one of them a choice to be made by investors according to the way they weight risk and rewards.

The basic version of the problem can be solved using Quadratic Programming (QP). Unfortunately, this approach is built on a set of assumptions that are unlikely to hold in the real world. The quest for alternatives has driven attention to metaheuristics that might not suffer this limitation. This is the reason why the framework of evolutionary computation is getting traction on this area.

The solutions based on metaheuristics mostly fall into one this two categories: solutions that transform the multiobjective problem into a single-objective form, and those who deal with it using a mutiobjective approach.

Among the first group, we would mention the work by Soleimani, Golmakani, and Salimi (2009). These authors use a genetic algorithm and extend the mentioned classic mean–variance optimization model and consider constraints on transaction costs, round lots or cardinality. They tackle the multiobjective nature of the problem minimizing the risk objective while setting a range of minimum acceptable returns in the constraints. Chiranjeevi and Sastry (2007) use another popular approach to transform the multiobjective problem into a version that can be handled by single-objective algorithms. Instead of keeping one of the objectives on the objective function and the other in the constraint, they consider the objectives in the fitness function weighting them. In this particular case, they manage five objectives that are obtained breaking down the basic two. Chang, Yang, and Chang (2009) suggest a similar solution and optimize a function that weights risk and return using a risk aversion parameter. Zhu, Wang, Wang, and Chen (2011) target a popular metric, the Sharpe Ratio, that combines both elements into a single expression.

The development of multiobjective evolutionary algorithms (MOEAs) has resulted into a number researchers exploring their performance in this area. Among these we could mention Skolpadungket, Dahal, and Harnpornchai (2007), who test a set of multi-objective algorithms (VEGA, SPEA2, NSGA-II etc.) on a constrained version of the two objective problem. More recently, Anagnostopoulos and Mamanis (2011) compare the performance of different multiobjective algorithms, and Deb, Steuer, Tewari, and Tewari (2011) introduce a customized hybrid version of NSGA-II to tackle the problem. Finally, we will mention the work

* Corresponding author.
E-mail addresses: sgrodrig@inf.uc3m.es (S. García), dquintan@inf.uc3m.es (D. Quintana), igalvan@inf.uc3m.es (I.M. Galván), isasi@ia.uc3m.es (P. Isasi).
URL: http://www.evannai.inf.uc3m.es (I.M. Galván).
[1] Principal corresponding author.

0957-4174/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2012.02.195

of Radziukyniene and Xilinskas (2008), where authors compare FastPGA, MOCELL, AbYSS, and NSGA-II on both of the basic problems, and an extended version that considers the dividend yield as a third objective.

Despite of the amount of research on portfolio optimization, there are still open issues. A key one is the robustness of results provided by algorithms. Among the most important factors that asset managers have to consider when evaluating the results provided by any of the above-mentioned methods, there is reliability. Very often, the expected efficient frontier lies far from the real one. This is due to the fact that the estimates for the expected risk and returns for the assets in the solution, and the portfolios derived from them, are very inaccurate. Understandably, this results on mistrust by some practitioners. The search for solutions for this problem has cleared the way for the field of robust portfolio optimization. It is in this area where we focus our contribution. We introduce a time-stamping method to control the population that enhances the reliability of the solutions provided by MOEAs.

The process of optimizing the risk and return of a portfolio relies on two parameters: the estimates for the expected asset returns and the variance–covariance matrix. The values of these parameters are usually based on past data and they might be inaccurate due to, for instance, the presence of outliers. In this context, there are several potential ways to tackle the problem. The two most prevalent alternatives found in the literature rely on either putting an emphasis on having robust estimates for the parameters, or managing the optimization process itself. The first one usually tries to filter the estimates to control, for instance, the influence of extreme past events on their computation (Perret-Gentil et al., 2005). The authors focusing on the second alternative design approaches handle uncertainty in the parameters during the optimization process (Pflug & Wozabal, 2007; Tütüncü & Koenig, 2004). The alternative suggested in this paper falls in the latter category. We will enhance the solutions of MOEAs testing the population for different values for the parameters, and selecting the portfolios that consistently offer a good performance.

Optimizing for a single scenario, a set of expected asset returns and the use of a single variance–covariance matrix, bears the risk of getting solutions that might be extremely sensitive to deviations. This could potentially be a problem as it is almost certain that the estimates will not be accurate. We have to bear in mind that having perfect estimates for the expected returns for instance, implies that we can make perfect predictions for future prices, which is highly unlikely. For this reason we consider that assessing the candidate solutions in different likely scenarios and favoring those that consistently offer a good performance, might be a promising approach. The requirement of consistency is key. In order to achieve it, we introduce a time-stamping mechanism that, together with performance in terms of risk and return, will drive the evolution process to find robust and stable solutions.

The approach introduced in this paper is related to alternatives based on resampling (Ruppert, 2006; Shiraishi, 2008). The most comparable among the traditional approaches is the one described by Idzorek (2006). This author suggests using combining traditional QP with Monte Carlo simulation to derive a set of fronts that are merged into a single solution at a later stage. This idea is very interesting but, unfortunately, the approach suffers the shortcomings of QP, namely, the ability to deal with real-world constraints. This is the reason why we feel that adapting the idea to the framework of MOEAs is very promising, as they do not suffer from this limitation. There is a previous effort based of evolutionary algorithms along the lines of this work, but is based on a very simple resampling approach that optimizes for a different scenario in each generation (Garcia, Quintana, Galvan, & Isasi, 2011). The time-stamping mechanism that we introduce in this work is based on the previous one extending the problem with a third implicit objective that favors solutions that are consistently reliable. As we will see in the experimental section, this results on significantly higher robustness.

The proposed approach is compatible with a wide array of MOEAs. Given their different nature and behavior, we will test the approach on a set of popular algorithms. Specifically, the experimental section will consider NSGA-II, SPEA2, SMPSO and GDE3. NSGA-II (Deb, Pratap, Agarwal, & Meyarivan, 2002) is one the most referenced algorithms in the field of multiobjective optimization. This one, together with SPEA2 (Zitzler, Laumanns, & Thiele, 2001) have been confirmed (Skolpadungket et al., 2007) to offer good performance in portfolio optimization. Apart from the mentioned two, we consider GDE3 (Kukkonen & Lampinen, 2005), a differential evolution strategy, and SMPSO, (Nebro et al., 2009) a multiobjective algorithm based on particle swarm optimization.

In the context of multiobjective problems, an important issue is the metric used to evaluate the solutions. It is generally admitted that there is no ideal single metric that should be used to evaluate different objectives simultaneously in every circumstance. The metrics that are most commonly used in this field, such as Hypervolume, Spread or SetCoverage are not appropriate indicators of stability in this context. For this reason, we define a set of metrics that capture different aspects of robustness in efficient frontiers. These, Estimation Error, Stability, Extreme Risk and Unrealized Returns draw on the basic principle that the expected risk and returns for the portfolios in the solution should be close to the observed ones ex-post.

The rest of the paper is organized as follows. First, we make a formal introduction to the financial portfolio optimization problem. After, we describe in detail the evolutionary approach proposed in this work. This section includes a brief description of the MOEAs, the solution encoding and the fitness mechanism in order to find robust and stable portfolios. Next, the different metrics used in this work to evaluate robustness of solutions are described. That will be followed by the experimental results and, finally, there will be a section devoted to summary and conclusions.

2. Financial portfolio optimization problem

Financial portfolios can be defined as a collection of investments or assets held by an institution or a private individual. The Modern Portfolio Theory was originated in the article published by Markowitz (1952). It explains how to use the diversification to optimize the portfolio. In general, the portfolio optimization problem is the choice of an optimum set of assets to include in the portfolio and the distribution of investor's wealth among them. Markowitz (1959) assumed that solving the problem requires the simultaneous satisfaction of maximizing the expected portfolio return $E(R_p)$ and minimizing the portfolio risk (variance) σ_p^2, that is, solving a multiobjective optimization problem with two output objective functions (Chiranjeevi & Sastry, 2007; Radziukyniene & Xilinskas, 2008; Skolpadungket et al., 2007; Soleimani et al., 2009). The portfolio optimization problem can be formally defined as:

- Minimize the **risk** (variance) of the portfolio:

$$\sigma_p^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j \sigma_{ij} \tag{1}$$

- Maximize the **return** of the portfolio:

$$E(R_p) = \sum_{i=1}^{n} w_i \mu_i \tag{2}$$

- Subject to:

$$\sum_{i=1}^{n} w_i = 1 \tag{3}$$

$$0 \leqslant w_i \leqslant 1; \quad i = 1 \ldots n \tag{4}$$