Computers and Operations Research 90 (2018) 33–41

Contents lists available at ScienceDirect

Computers and Operations Research

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

Minimizing the tracking error of cardinality constrained portfolios

Purity Mutunge, Dag Haugland*

Department of Informatics, University of Bergen, Norway

ARTICLE INFO

Article history:
Received 1 July 2016
Revised 19 June 2017
Accepted 4 September 2017
Available online 6 September 2017

Keywords:
Portfolio management
Index tracking
Integer quadratic programming
Heuristics

ABSTRACT

We study the problem of selecting a restricted number of shares included in a stock market index, such that the portfolio resembles the index as closely as possible. To measure the difference between the portfolio and the index, referred to as the tracking error, we use a quadratic function with the covariance matrix of the index returns as coefficient matrix. The problem is proved to be strongly NP-hard, and we give theoretical evidence that continuous relaxations of mixed integer quadratic programming (MIQP) formulations are likely to produce poor lower bounds on the tracking error. For fast computation of near-optimal portfolios, we demonstrate how the best-extension-by-one construction heuristic can be designed to run in time bounded by a fourth order polynomial. We also show that the running time of one iteration of the best-exchange-by one improvement heuristic is of the same order. Computational experiments applied to real-life stock market indices show that in instances where an index of less than 500 assets is to be tracked by a portfolio of 10 assets, a commercially available MIQP solver fails to reduce the integrality gap below 94% in 30 CPU-minutes. In contrast, the construction heuristic under study needs less than 30 CPU-seconds to produce a portfolio of 100 assets tracking an index of nearly 2000 assets.

1. Introduction

In fund management, where wealth from numerous investors is managed, different strategies are employed. Making an investment is in itself a risky affair as one is not guaranteed of a gain on wealth. However, managers have the task of making decisions on how to achieve this goal while avoiding losses.

The objective of active fund management is to identify stocks (assets) that do considerably better than the market in general. While the expected return from the fund is higher than from the market at large, also the probability of substantial losses is relatively high. In contrast, a passive fund manager aims to compose a portfolio that brings approximately the average market return. The expected return from the portfolio is therefore lower than the one from active fund management, whereas passive fund management involves lower risk. Active and passive fund managers can therefore be said to have different attitudes to risk, with the passive manager being the most risk averse.

Passive fund management is commonly implemented by composing funds that simulate a chosen benchmark, typically a stock market index. This strategy is referred to as *index tracking*. A fund, or portfolio, that tracks a stock market index is composed by a selection of, typically few, shares that are present in the index. In theory, the portfolio may contain all the stocks in the tracked benchmark index, in their respective proportions, as the index. Such a full replication approach is however neither practical nor cost effective. Besides lower risks, index tracking has several cost-saving advantages when compared to active fund management. Management costs involved in stock picking and market timing can be considerable for active funds. Especially when the number of stocks is restricted, administration costs of index tracking portfolios can be kept low. Cutting costs is therefore an important motivation for designing index trackers containing only a small subset of the assets in a tracked stock market index.

Portfolio size being determined, choosing which stocks to include becomes a problem of vital importance to the passive fund manager. Some measure of difference between a portfolio and the benchmark which it mimics is referred to as the *tracking error*, and the problem becomes to find a portfolio minimizing this measure. Different formulations of the index tracking error are proposed, all of which are functions of the weights of each asset in the index and the tracking portfolio, respectively. Statistical parameters related to the stock returns are typical parameters of the various definitions of the tracking error, which are briefly reviewed in Section 2.1.

When the portfolio minimizing the tracking error is to be found, the restriction in portfolio size is known to represent a serious computational challenge. Regardless of the tracking error def-

* Corresponding author.
E-mail addresses: purity.mutunge@student.uib.no (P. Mutunge), dag@ii.uib.no (D. Haugland).

http://dx.doi.org/10.1016/j.cor.2017.09.002
0305-0548/

inition and what additional constraints that apply, the cardinality constraint is notorious for being the source of high computational costs. In Section 2.2, we review some models from the scientific literature for minimizing the tracking error, and discuss briefly their computational complexity and suggested solution methods.

The objective of the current work is to present new theoretical results and computational procedures for cardinality constrained index tracking. To this end, we first give a straightforward mathematical formulation of the problem. Our model has a convex quadratic objective function representing the tracking error to be minimized, binary variables representing the selection of stocks, and continuous variables representing the stock weights in the portfolio. Besides the cardinality constraint, we only introduce logical constraints expressing that the weight of unselected stocks must be zero, and that the sum of all weights must be one. We verify theoretically that even a plain index tracking model of this kind has unfavorable computational properties, and show that exact algorithms have to evaluate a large proportion of all feasible portfolios. Further, we formulate fast construction and improvement heuristics for the problem, and give experimental evidence for their applicability.

The remainder of the text is organized as follows: the literature review in the next section, we present in Section 3 a theoretical analysis of the computational properties of the model. Inexact solution methods are developed in Section 4 and evaluated experimentally in Section 5, while conclusions are drawn in Section 6. Throughout the text, we write vectors and matrices in boldface letters. If N is a finite set, notation $\mathbf{x} \in \mathbb{R}^N$ ($\mathbf{x} \in \mathbb{R}^N_+$) says that $\mathbf{x}$ is a vector of (nonnegative) real numbers corresponding to the elements of N. For $i \in N$ and $M \subseteq N$, we let x_i denote the component of $\mathbf{x}$ corresponding to i, and we let $\mathbf{x}_M$ denote the vector with components x_j, where $j \in M$. If p is a real number, $\mathbf{p}_N \in \mathbb{R}^N$ denotes the vector with one component with value p for each element in N, and we write $\mathbf{p}$ in place of $\mathbf{p}_N$ whenever it is clear from the context to what set the components correspond. A submatrix of $\mathbf{A} \in \mathbb{R}^{N_1 \times N_2}$, where N_1 and N_2 are finite sets, consisting of the rows corresponding to $M_1 \subseteq N_1$ and the columns corresponding to $M_2 \subseteq N_2$, is denoted $\mathbf{A}_{M_1 M_2}$. If $M_1 = \{k\}$ ($M_2 = \{k\}$), then notation $\mathbf{A}_{kM_2}$ ($\mathbf{A}_{M_1 k}$) is used, while the element in row $i \in N_1$ and column $j \in N_2$ is denoted a_{ij}. For integers m and n, we also make use of notations $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{A} \in \mathbb{R}^{m \times n}$, meaning that $\mathbf{x}$ and $\mathbf{A}$ are, respectively, a real vector of length n, and a real matrix with m rows and n columns. Vector $\mathbf{e}_i \in \mathbb{R}^n$ is the unit vector with a 1-entry in position i. Matrices and vectors superscripted by $\top$ give the corresponding transpose. Notations analogous to all the above are adopted for vectors and matrices over other number fields.

2. Literature review

2.1. Tracking error definitions

Roßbach and Karlow (2011) consider a multi-period model, and formulate the tracking error in time period t as $|R_t - v_t|$, where R_t and v_t denote the returns of the tracking portfolio and the index, respectively, in period t, given as the sum of the total weights of the assets multiplied by their respective returns in period t. Averaged over the time periods $t = 1, \dots, T$, this results (Roßbach and Karlow, 2011) in the error: $\frac{1}{T}\sum_{t=1}^{T} |R_t - v_t|$.

Rudolf et al. (1999) also consider the tracking error defined as a function of the absolute values of historical differences between the returns from the tracked index and the tracking portfolio. They also analyze a version including only time periods where the portfolio return is below the index return, and suggest to replace the unit norm deviation by the maximum norm. These approaches contrast the traditional (Roll, 1992) definition of the tracking error as (the square of) the Euclidean norm of the difference over time.

As Beasley et al. (2003) remark, other p-norms can also be applied to the return differences. While the unit norm ($p = 1$) implies that all deviations have equal weight, the Euclidean norm ($p = 2$) says that more emphasis is on large deviations. Taking this to the extreme, the maximum norm ($p = \infty$) corresponds to neglecting all but the largest deviation.

Rather than including all return values observed in a given time interval, Jansen and van Dijk (2002) suggest to aggregate the observations by estimating the covariance matrix, $\mathbf{Q}$, of the stock returns. Their formulation of the tracking error thus becomes $(\mathbf{x} - \mathbf{w})^\top \mathbf{Q} (\mathbf{x} - \mathbf{w})$, where $\mathbf{x}$ and $\mathbf{w}$ are the vectors consisting of the weights of all stocks in, respectively, the tracking portfolio and the index. Such a tracking error function estimates the variance of the return differences in future time periods. As Coleman et al. (2006) remark, the tracking error definition in Jansen and van Dijk (2002) is mathematically more appealing as it is convex, and can be used in financial interpretations on the assumption that the covariance matrix is accurate for future returns.

As observed by Rudolf et al. (1999) and Konno and Yamazaki (1991), there exist similarities in performance when using the different approaches. Roßbach and Karlow (2011) present a comparative study of these different approaches.

In this work, we adopt the tracking error definition

$$f(\mathbf{x}) = (\mathbf{x} - \mathbf{w})^\top \mathbf{Q} (\mathbf{x} - \mathbf{w}),$$

giving us a convex quadratic objective function to be minimized, provided that the selection of stocks is determined.

2.2. Models and solution methods

When a cardinality constraint restricting the number of stocks is introduced, the problem of optimizing the composition of a portfolio tends to become NP-hard (Moral-Escudero et al., 2006; Ruiz-Torrubiano and Suárez, 2009; Shaw et al., 2008). This means that exact solutions to instances of realistic sizes are computationally intractable, and thus inexact solution methods are the only practical ones.

Let N denote the set of stocks in the benchmark index, and let m denote the maximum number of stocks in the portfolio. Jansen and van Dijk (2002) and Coleman et al. (2006) present two different approaches to dealing with the discontinuity introduced by the cardinality constraint $|\{i \in N\colon x_i > 0\}| \le m$.

Jansen and van Dijk (2002) focus on minimizing the tracking error with a relatively small number of stocks. When $n = |N|$ is relatively small, it is argued that because $\lim_{p \to 0^+} x_i^p = 1$, $|\{i \in N\colon x_i > 0\}|$ can be approximated by the continuous function $\sum_{i=1}^{n} x_i^p$ for some parameter $p > 0$.

Building on Jansen and van Dijk (2002), Coleman et al. (2006) approximate the portfolio cardinality $|\{i \in N\colon x_i > 0\}|$ by a continuously differentiable non-convex function $\Sigma_{i \in N} h_\lambda(x_i)$, where $h_\lambda(x_i) = \lambda x_i^2$ if $x_i \le \sqrt{\frac{1}{\lambda}}$ and $h_\lambda(x_i) = 1$, otherwise. When the parameter λ is assigned a large value, it is argued that the approximation becomes close. Further improvements in terms of an objective function that also is differentiable is achieved by refining this idea.

Shaw et al. (2008) propose a model where the objective function combines the tracking error with a linear term representing expected revenues. In addition, the model contains a set of arbitrary linear side constraints, as well as lower and upper bounds on each index weight. The resulting model is solved by a procedure based on Lagrangian relaxation.

Ruiz-Torrubiano and Suárez (2009) define the tracking error in terms of a quadratic function of the asset weights. In addition to