Expert Systems with Applications 41 (2014) 3901–3914

Contents lists available at ScienceDirect

Expert Systems with Applications

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

Multi-period portfolio selection using kernel-based control policy with dimensionality reduction

Yuichi Takano [a,*], Jun-ya Gotoh [b]

[a] Department of Industrial Engineering and Management, Graduate School of Decision Science and Technology, Tokyo Institute of Technology, 2-12-1-W9-77 Ookayama, Meguro-ku, Tokyo 152-8552, Japan

[b] Department of Industrial and Systems Engineering, Chuo University, 1-13-27 Kasuga, Bunkyo-ku, Tokyo 112-8551, Japan

ARTICLE INFO

Keywords:
Multi-period portfolio selection
Kernel method
Control policy
Dimensionality reduction

ABSTRACT

This paper studies a nonlinear control policy for multi-period investment. The nonlinear strategy we implement is categorized as a kernel method, but solving large-scale instances of the resulting optimization problem in a direct manner is computationally intractable in the literature. In order to overcome this difficulty, we employ a dimensionality reduction technique which is often used in principal component analysis. Numerical experiments show that our strategy works not only to reduce the computation time, but also to improve out-of-sample investment performance.

1. Introduction

Risk management based on diversified investment makes it possible to mitigate the risk of suffering a large loss while securing a certain level of profitability, and portfolio selection accordingly plays an important role in financial decision making (see, e.g., Cornuejols & Tütüncü, 2007). Portfolio selection is usually conducted in a single-period framework, as initially formulated by Markowitz (1952). It is, however, advantageous for investors to periodically adjust their portfolios by following an effective rebalancing strategy. In this respect, the traditional single-period model is not sufficient. Indeed, Mulvey, Pauling, and Madey (2003) state that multi-period models can enhance risk-adjusted performance and help investors evaluate the probability of reaching a certain target by linking asset and liability policies.

Among the various rebalancing strategies, constant rebalancing reverts the investment proportion to the original proportion at the beginning of every period. It is known that a constant rebalancing strategy achieves the optimal growth rate of wealth on the assumption that asset returns in each period are independent and identically distributed (see, e.g., Algoet & Cover, 1988). Due to this fact, a number of studies (see, e.g., Fleten, Høyland, & Wallace, 2002; Maranas, Androulakis, Floudas, Berger, & Mulvey, 1997; Takano & Gotoh, 2011; Takano & Sotirov, 2012) have dealt with multi-period portfolio optimization with the constant rebalancing strategy.

However, it has been demonstrated, e.g., in Jegadeesh and Titman (1993), Lo and MacKinlay (1990), that stock returns are serially dependent; therefore, it is probably effective to dynamically rebalance the portfolio in view of the observed asset returns. For instance, DeMiguel, Nogales, and Uppal (2013) improve the out-of-sample investment performance of single-period models by predicting future stock returns through the use of a vector autoregressive (VAR) model. More importantly, Fleten et al. (2002) have shown by means of an out-of-sample simulation test that the stochastic dynamic approach dominates the constant rebalancing strategy. These observations motivated us to develop a rebalancing strategy for exploiting the time-series dependence of stock returns.

Multi-period portfolio selection was first framed as a stochastic control problem (see Infanger, 2006 for detailed references). In general, however, it is very difficult to handle a stochastic control problem of a practical size because it requires one to solve a large-scale dynamic programing problem or partial differential equations. The most popular framework for solving such problems is provided by multi-stage stochastic programing models. Among them the simulated path model (see, e.g., Hibiki, 2006) describes multi-period scenarios of asset returns using a number of simulated paths. The actual market behavior can be simulated in detail by this model, but there is no room for conditional investment decisions in this model due to what is called the "non-anticipativity condition," which requires one to prevent investment decisions from depending on future observations on each simulated path. By contrast, the scenario tree model (see Steinbach, 2001 for detailed references) enables one to make conditional investment decisions in each future state; however, this model is disadvantageous in that the size of the resultant optimization problem grows exponentially as the number of time periods increases. The hybrid model devised by Hibiki (2003) integrates the simulated path model and the scenario tree model; nevertheless, it is still computationally burdensome to make conditional investment decisions in the hybrid model as well as in the scenario tree model.

* Corresponding author. Tel.: +81 3 5734 2361.
E-mail addresses: takano.y.ad@m.titech.ac.jp (Y. Takano), jgoto@indsys.chuo-u.ac.jp (J. Gotoh).

0957-4174/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2013.11.043

Under several assumptions about stochastic process of asset returns, closed-form optimal solutions to stochastic control problems can be derived (see, e.g., DeMiguel, Martin-Utrera, & Nogales, 2013; DeMiguel, Mei, & Nogales, 2013; Garleanu & Pedersen, 2013). However, these assumptions often fail in simulating a complicated market behavior accurately. Additionally, the presence of portfolio composition constraints makes it harder to obtain a closed-form optimal solution. Meanwhile, a number of studies focus on numerical methods for approximately solving such a stochastic control problem. Among them are approximate dynamic programing (e.g., Boyd, Mueller, O'Donoghue, & Wang, 2012) and model predictive control (e.g., Boyd et al., 2012; Yamada & Primbs, 2012). Besides, control actions are frequently prescribed by control policies, which map past outcomes to the portfolio adjustments. Most studies (Barmish & Primbs, 2012; Calafiore, 2008; Calafiore, 2009; Calafiore & Campi, 2005; Fonseca & Rustem, 2012; Moallemi & Sağlam, 2012) only deal with linear control policies because determining the best control policy generally leads to a computationally intractable optimization. Although these approaches are computationally attractive, they can only attain a "sub-optimal" investment strategy.

On the other hand, the authors of the present paper build in Takano and Gotoh (2011) a computational framework based on the kernel method for finding the best nonlinear control policy. Bhat, Moallemi, and Farias (2012) also use the kernel method for a discrete time Markov decision process. Such a kernel method is often employed in machine learning for estimating nonlinear statistical models (see, e.g., Schölkopf & Smola, 2001), and it enables one to easily incorporate a nonlinear transformation in a non-parametric manner. Since the computational model developed by Takano and Gotoh (2011) is based on the stochastic programing, it can impose practical constraints on the portfolio composition. Moreover, this kernel-based non-parametric approach does not require any assumptions about the underlying stochastic process of asset returns. Note also that, in contrast to the sub-optimal strategies mentioned above, the kernel-based control policy method selects an investment strategy from a wider class of policies. With the kernel method, we can extract the complex time-series dependence of asset returns. Indeed, Bhat et al. (2012) and Takano and Gotoh (2011) empirically show that kernel-based methods perform better than other standard parametric methods (e.g., linear control policy).

In view of these facts, we shall employ the kernel-based control policy in the multi-period portfolio selection problem. However, we are yet confronted by two difficulties in using the kernel-based control policy: Long computation time and overfitting. Indeed, the experiments in Takano and Gotoh (2011) show that substantial time is required even for small-sized problems, despite the fact that the problem is formulated as a convex quadratic programing problem. In addition, since the kernel approach admits a highly nonlinear mapping, the resulting control policy may overfit the noisy financial data used in the optimization problem. It has been reported in Takano and Gotoh (2011) that such a overfitting weakens out-of-sample performance of the kernel-based control policy.

The purpose of this paper is to devise a new approach for efficiently solving the multi-period portfolio selection problem with a kernel-based control policy (Takano & Gotoh, 2011) and for further improving its investment performance. To this end, a method of problem reduction is posed based on a dimensionality reduction technique which is often used in principal component analysis (PCA). More precisely, our application is directly related to what is called kernel principal component analysis (kernel PCA), which is an extension of PCA into a feature space of (possibly, infinitely) high dimension (see, e.g., Schölkopf & Smola, 2001). Yajima, Ohi, and Mori (2003) also use a dimensionality reduction technique to reduce the problem size of a nonlinear support vector classification. Their results encouraged us to apply a similar reduction method to our multi-period portfolio selection problem. In addition, it has been demonstrated, e.g., in Mika et al. (1998), that kernel PCA has an effect of de-noising. This means that our dimensionality reduction technique has the potential of not just achieving a high degree of computation efficiency, but also improving investment performance.

We test the effectiveness of our approach through numerical experiments with historical data on actual stock returns. Contributions of this paper are summarized as follows:

An efficient solution algorithm.
We develop an efficient solution algorithm based on the dimensionality reduction technique. Numerical results show that the proposed method sharply lessened the computation time.

Improvement in investment performance.
We improve investment performance with the dimensionality reduction technique. Numerical experiments show that the dimensionality reduction method enhanced the out-of-sample investment performance by avoiding overfitting.

The rest of the paper is organized as follows: Section 2 formulates a multi-period portfolio selection model equipped with a kernel-based control policy. Section 3 develops a method for reducing the problem size by means of eigenvalue decomposition and formulate an optimization problem in a reduced form. Numerical results are reported in Section 4, and concluding remarks are given in Section 5.

2. Control policy for multi-period portfolio selection

In this section, after giving a mathematical description of portfolio dynamics, we formulate the multi-period portfolio selection problem with a kernel-based nonlinear control policy.

2.1. Preliminaries and portfolio dynamics

The terminology and notation used in this subsection are as follows:

Index sets		
$\mathcal{I} := \{1, 2, \ldots, I\}$	:	Index set of investable financial assets (where asset 1 is cash)
$\mathcal{S} := \{1, 2, \ldots, S\}$	:	Index set of given scenarios (or simulated paths)
$\mathcal{T} := \{1, 2, \ldots, T\}$	:	Index set of planning time periods
Decision variables		
$x_{i,s}(t)$	:	Investment amount in asset i at the end of period t in scenario s ($i \in \mathcal{I},\ s \in \mathcal{S},\ t \in \mathcal{T}$)
$u_i(t)$	:	Adjustment of asset i at the beginning of period t ($i \in \mathcal{I},\ t \in \mathcal{T}$)
$u_{i,s}(t)$	:	Adjustment of asset i at the beginning of period t in scenario s ($i \in \mathcal{I},\ s \in \mathcal{S},\ t \in \mathcal{T} \setminus \{1\}$)
$v_s(t)$	:	Portfolio value at the end of period t in scenario s ($s \in \mathcal{S},\ t \in \mathcal{T}$)
$a(t)$	:	The value-at-risk (VaR) in period t ($t \in \mathcal{T}$)
$z_s(t)$	:	Auxiliary decision variable for calculating the conditional value-at-risk (CVaR) in period t ($s \in \mathcal{S},\ t \in \mathcal{T}$)
Given constants		
$\bar{x}_i(0)$	:	The initial holdings of asset i ($i \in \mathcal{I}$)
$C(t)$	:	Net cash flow at the beginning of period