Artificial Intelligence 222 (2015) 104–123

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Moving average reversion strategy for on-line portfolio selection ☆

Bin Li [a], Steven C.H. Hoi [b,*], Doyen Sahoo [b], Zhi-Yong Liu [c]

[a] Economics and Management School, Wuhan University, Wuhan 430072, PR China
[b] School of Information Systems, Singapore Management University, 178902, Singapore
[c] Institute of Automation, Chinese Academy of Sciences, Beijing 100080, PR China

ARTICLE INFO

Article history:
Received 17 December 2012
Received in revised form 24 January 2015
Accepted 28 January 2015
Available online 2 February 2015

Keywords:
Portfolio selection
On-line learning
Mean reversion
Moving average reversion

ABSTRACT

On-line portfolio selection, a fundamental problem in computational finance, has attracted increasing interest from artificial intelligence and machine learning communities in recent years. Empirical evidence shows that stock's high and low prices are temporary and stock prices are likely to follow the mean reversion phenomenon. While existing mean reversion strategies are shown to achieve good empirical performance on many real datasets, they often make the *single-period mean reversion* assumption, which is not always satisfied, leading to poor performance in certain real datasets. To overcome this limitation, this article proposes a *multiple-period mean reversion*, or so-called "Moving Average Reversion" (MAR), and a new on-line portfolio selection strategy named "On-Line Moving Average Reversion" (OLMAR), which exploits MAR via efficient and scalable online machine learning techniques. From our empirical results on real markets, we found that OLMAR can overcome the drawbacks of existing mean reversion algorithms and achieve significantly better results, especially on the datasets where existing mean reversion algorithms failed. In addition to its superior empirical performance, OLMAR also runs extremely fast, further supporting its practical applicability to a wide range of applications. Finally, we have made all the datasets and source codes of this work publicly available at our project website: http://OLPS.stevenhoi.org/.

1. Introduction

Portfolio selection is a fundamental problem of computational finance extensively explored across several fields, ranging from traditional finance theory and quantitative finance, to machine learning and artificial intelligence [43]. It generally aims to achieve some targets in the long run by sequentially allocating wealth across a set of assets. Two major schools of principles and theories for portfolio selection include: (i) Mean-variance theory [48] that trades off between expected return (mean) and risk (variance) of a portfolio, which is suitable for single-period portfolio selection; and (ii) Kelly investment [41,10,23] that aims to maximize the expected log return of a portfolio and is naturally available to multiple-period portfolio selection. Due to the sequential nature of a real-world portfolio selection task, many recent on-line portfolio selection techniques often follow the second approach.

☆ The short version of this work [42] appeared at the 29th International Conference on Machine Learning (ICML 2012).

* Corresponding author.
E-mail addresses: binli.whu@whu.edu.cn (B. Li), chhoi@smu.edu.sg (S.C.H. Hoi), doyen.2014@phdis.smu.edu.sg (D. Sahoo), zhiyong.liu@ia.ac.cn (Z.-Y. Liu).

http://dx.doi.org/10.1016/j.artint.2015.01.006

One important property exploited by many existing studies [8,46,45] is the *mean reversion* property, which assumes poor performing stocks will perform well in the subsequent periods and vice versa. Although some recently proposed mean reversion algorithms [46,45] have achieved promising results on many real datasets, they might perform extremely poorly on certain datasets, such as the DJIA dataset [8]. Comparing with Borodin et al. [8], which exploits multi-period correlation, we found that the assumption of single-period prediction may attribute to the performance degradation. On the other hand, as illustrated in existing studies [46,45], Borodin et al. [8] cannot fully exploit the potential of (multi-period) mean reversion for its heuristical nature.

To address the above drawbacks, we present a new approach for on-line portfolio selection, named "On-Line Moving Average Reversion" (OLMAR). The basic idea is to represent multi-period mean reversion as "Moving Average Reversion" (MAR), which explicitly predicts next price relatives using moving averages, and then learns portfolios via online learning techniques [38]. To the best of our knowledge, OLMAR is the first algorithm that exploits moving average in the setting of on-line portfolio selection. Though simple in nature, OLMAR has a reasonable updating strategy and has been empirically validated via a set of extensive experiments on real markets. The experimental results show that OLMAR not only can achieve better performance (in terms of *cumulative wealth*) than existing algorithms, but also can avoid the significant performance degradation on certain datasets, such as the DJIA dataset [8,46]. Finally, OLMAR runs highly efficiently, and thus is suitable for large-scale trading applications.

The rest of the paper is organized as follows. Section 2 formulates the on-line portfolio selection problem, and Section 3 reviews the background and analyzes related work. Section 4 presents the proposed OLMAR approach, and its effectiveness is validated by extensive empirical studies on real stock markets in Section 5. Section 6 discusses our empirical findings and some threats to validity. Section 7 summarizes the paper and provides directions for future work.

2. Problem setting

Before we formulate the online portfolio selection problem, we first introduce some notations to be used throughout the article. Symbols in bold are vectors, e.g., we define a as a scalar, and $\mathbf{b}$ and $\mathbf{c}$ as vectors. One typical operation is the product or division between a scalar and a vector, which means each element of the vector will multiply or divide by the scalar, e.g., $[a\mathbf{b}]_i = a * b_i$ and $[\frac{a}{\mathbf{b}}]_i = \frac{a}{b_i}$. For the element-wise operations between two vectors, we denote by $\frac{\mathbf{a}}{\mathbf{b}}$ the element-wise division and $\mathbf{a} \otimes \mathbf{b}$ the element-wise product, which means $[\frac{\mathbf{a}}{\mathbf{b}}]_i = a_i / b_i$ and $[\mathbf{a} \otimes \mathbf{b}]_i = a_i * b_i$, respectively. The dot product between two vectors is defined as $\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^T \mathbf{b} = \sum_{i=1}^n a_i * b_i$.

Now let us consider an investment task over a financial market with m assets for n periods. On the tth period, the asset prices are represented by a *close price vector* $\mathbf{p}_t \in \mathbb{R}_+^m$, and each element $p_{t,i}$ represents the close price of asset i. The price changes are represented by a *price relative vector* $\mathbf{x}_t \in \mathbb{R}_+^m$, and $x_{t,i} = \frac{p_{t,i}}{p_{t-1,i}}$. Thus, an investment in asset i on the tth period increases by a factor of $x_{t,i}$. Let us denote $\mathbf{x}_{t1}^{t2} = \{\mathbf{x}_{t1}, \mathbf{x}_{t2}, \ldots, \mathbf{x}_{t2}\}$ as a sequence of price relative vectors ranging from period $t1$ to $t2$. Therefore, $\mathbf{x}_1^n = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ represents the sequence of price relative vectors over the entire n periods.

An investment on the tth period is specified by a *portfolio vector* $\mathbf{b}_t = (b_{t,1}, \ldots, b_{t,m})$, where $b_{t,i}$ represents the proportion of wealth invested in asset i. Typically, we assume the portfolio is self-financed and no margin/short sale is allowed, therefore each entry of a portfolio is non-negative and adds up to one, that is, $\mathbf{b}_t \in \Delta_m$, where $\Delta_m = \{\mathbf{b}_t : \mathbf{b}_t \in \mathbb{R}_+^m, \sum_{i=1}^m b_{t,i} = 1\}$. The investment procedure is represented by a *portfolio strategy*, that is, $\mathbf{b}_1 = \frac{1}{m}\mathbf{1}$ and the following sequence of mappings $\mathbf{b}_t : \mathbb{R}_+^{m(t-1)} \to \Delta_m$, $t = 2, 3, \ldots$, where $\mathbf{b}_t = \mathbf{b}_t(\mathbf{x}_1^{t-1})$ is the tth portfolio given past market sequence of $\mathbf{x}_1^{t-1} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{t-1}\}$. We denote by $\mathbf{b}_1^n = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ the strategy for n periods.

On the tth period, a portfolio $\mathbf{b}_t$ produces a *portfolio period return* s_t, that is, the wealth increases by a factor of $s_t = \mathbf{b}_t^\top \mathbf{x}_t = \sum_{i=1}^m b_{t,i} x_{t,i}$. Since we reinvest and adopt price relative, the portfolio wealth would multiplicatively grow. Thus, after n periods, a portfolio strategy $\mathbf{b}_1^n$ produces a *portfolio cumulative wealth* of S_n, which increases the initial wealth by a factor of $\prod_{t=1}^n \mathbf{b}_t^\top \mathbf{x}_t$, that is, $S_n(\mathbf{b}_1^n, \mathbf{x}_1^n) = S_0 \prod_{t=1}^n \mathbf{b}_t^\top \mathbf{x}_t$, where S_0 is set to \$1 for convenience.

Finally, we formally formulate the on-line portfolio selection procedure, and outline the algorithmic framework in Algorithm 1. In this task, a portfolio manager is a decision maker, whose goal is to produce a portfolio strategy $\mathbf{b}_1^n$, aiming to maximize the cumulative wealth S_n. He/she computes the portfolios sequentially. On each period t, the manager has access to the sequence of previous price relative vectors $\mathbf{x}_1^{t-1}$. Then, he/she computes a new portfolio $\mathbf{b}_t$ for next price relative vector $\mathbf{x}_t$, where the decision criterion varies among different managers. The portfolio $\mathbf{b}_t$ is scored based on portfolio period return s_t. This procedure is repeated until the end, and the portfolio strategy is finally scored according to portfolio cumulative wealth S_n.

It is important to note that we have made several general and common assumptions in the above model:

1. Transaction cost: no commission fees or taxes exists;
2. Market liquidity: one can buy and sell any desired amount, even fractional, at last closing price of any given trading period;
3. Market impact: any portfolio selection strategy shall not influence the market, or other stocks' prices.

Note that although these assumptions are commonly made in many literature studies, they are nontrivial in practice. We will further analyze and discuss their implications and effects in the empirical studies.