Robust median reversion strategy for on-line portfolio selection

Dingjiang HUANG
East China University of Science and Technology

Junlong ZHOU
East China University of Science and Technology

Bin LI
Nanyang Technological University

Steven HOI
Singapore Management University, chhoi@smu.edu.sg

Shuigeng ZHOU
Fudan University

Robust Median Reversion Strategy for On-Line Portfolio Selection*

Dingjiang Huang[1,2,3]**, Junlong Zhou**[1]**, Bin Li**[4]**, Steven C.H. Hoi**[4]**, Shuigeng Zhou**[2,3†]

[1]Department of Mathematics, East China University of Science and Technology,
Shanghai 200237, China
[2]School of Computer Science, Fudan University, Shanghai 200433, China
[3]Shanghai Key Lab of Intelligent Information Processing, Fudan University,
Shanghai 200433, China
[4]School of Computer Engineering, Nanyang Technological University, Singapore 639798
{djhuang,sgzhou}@fudan.edu.cn; jlzhou8812@gmail.com; {binli,chhoi}@ntu.edu.sg

Abstract

On-line portfolio selection has been attracting increasing interests from artificial intelligence community in recent decades. Mean reversion, as one most frequent pattern in financial markets, plays an important role in some state-of-the-art strategies. Though successful in certain datasets, existing mean reversion strategies do not fully consider noises and outliers in the data, leading to estimation error and thus non-optimal portfolios, which results in poor performance in practice. To overcome the limitation, we propose to exploit the reversion phenomenon by robust L_1-median estimator, and design a novel on-line portfolio selection strategy named "Robust Median Reversion" (RMR), which makes optimal portfolios based on the improved reversion estimation. Empirical results on various real markets show that RMR can overcome the drawbacks of existing mean reversion algorithms and achieve significantly better results. Finally, RMR runs in linear time, and thus is suitable for large-scale trading applications.

1 Introduction

Portfolio Selection (PS) problem is concerned with determining a portfolio for allocating the wealth among a set of assets to achieve some financial objectives in the long run. There are two main mathematical models for this problem: the mean-variance model [Markowitz, 1952] and the Kelly investment [Kelly, 1956]. In general, mean-variance theory, which trades off between the expected return (mean) and risk (variance) of a portfolio, is suitable for single-period (batch) PS. While Kelly investment, which tends to maximize the expected log

return, focuses on multiple-period sequential PS. Although these two theories, initially, generated little interest, they are now mainstream theories whose principles are constantly visited and re-invented. One popular research is On-line PS, which often designs algorithms following the Kelly investment because of its sequential nature, and has been actively explored in AI [Cover, 1991; Ordentlich and Cover, 1996] and machine learning communities [Agarwal *et al.*, 2006; Borodin *et al.*, 2004; Li and Hoi, 2012].

Some state-of-the-art on-line PS strategies [Gyorfi *et al.*, 2006; 2008] assume that the current best performing stocks would also perform well in the next trading day, but empirical evidence [Jegadeesh., 1990] indicates that such trends may be often violated, especially in the short term. This observation leads to the strategy of buying poor performing stocks and selling those with good performance. This trading principle, known as "mean reversion", is a valid investment principle.

Recently, on-line PS [Cover, 1991; Borodin *et al.*, 2004; Li *et al.*, 2013; 2012a; Li and Hoi, 2012b] has exploiting the mean reversion principle. Though these mean reversion algorithms achieve encouraging results on many datasets, they perform poor on certain datasets, such as DJA dataset [Li *et al.*, 2013; 2012a]. This is because all existing mean reversion strategies do not fully consider the noisy data and outliers, which often leads to estimation error, and thus makes the portfolio non-optimal (see [Merton, 1980]). Furthermore, the assumption of single-period prediction [Li *et al.*, 2013; 2012a] also leads to estimation error, which thus makes the performance extremely poor [Li and Hoi, 2012b].

To address the above drawbacks, we present a new multi-period on-line PS strategy named "Robust Median Reversion"(RMR). The basic idea is to exploit the reversion phenomenon via robust L_1-median estimator [Weber, 1909; Weiszfeld, 1937; Vardi and Zhang, 2000], which explicitly estimates next price relative and is more accurate than simple mean estimator. Then we learn optimal portfolios based on the improved reversion estimation and state of the art online learning techniques.

To the best of our knowledge, RMR is the first algorithm

that exploits the reversion phenomenon by robust L_1-median estimator. Though simple in nature, RMR can release better estimation than existing algorithms and has been empirically validated via extensive experiments on real markets. The experimental results show that RMR significantly surpasses a number of state-of-the-art strategies in terms of long-term compound return. Moreover, it is robust to different parameter settings and it can withstand small transaction costs. Finally, RMR has linear time complexity, which is suitable for large-scale applications.

The rest of the paper is organized as follows. Section 2 formulates the on-line PS problem and Section 3 reviews some related work. Section 4 presents the proposed algorithm and Section 5 empirically evaluates its efficacy on real markets. Section 6 finally summarizes the article.

2 Problem Setting

Consider a financial market with d assets for n trading periods to be invested. On the t^{th} period, the asset prices are represented by a *close price vector* $\mathbf{p}_t \in \mathbb{R}^d_+$, and each element p_t^i represents the close price of asset i. The changes of asset prices are represented by a *price relative vector* $\mathbf{x}_t = (x_t^1, \ldots, x_t^d) \in \mathbb{R}^d_+$, where x_t^j expresses the ratio of close price to last close price of asset j at the t^{th} period, i.e., $x_t^j = p_t^j/p_{t-1}^j$. We denote $\mathbf{x}_1^n = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ as the sequence of price relative vectors for n periods.

At the beginning of the t^{th} period, we diversify the capital among the d assets according to a *portfolio vector* $\mathbf{b}_t = (b_t^1, \ldots, b_t^d) \in \mathbb{R}^d_+$, where b_t^j represents the proportion of wealth invested in the j^{th} asset. Typically, we assume the portfolio is self-financed and no margin/short is allowed, which means $\mathbf{b}_t \in \Delta_d$, where $\Delta_d = \{\mathbf{b}_t : \mathbf{b}_t \in \mathbb{R}^d_+, \sum_{j=1}^d b_t^j = 1\}$. The investment procedure is represented by a *portfolio strategy*, that is, $\mathbf{b}_1 = \frac{1}{d}\mathbf{1}$ and following sequence of mappings $\mathbf{b}_t : (\mathbb{R}^d_+)^{t-1} \to \Delta_d, t = 1, 2, \ldots$, where $\mathbf{b}_t = \mathbf{b}_t(\mathbf{x}_1^{t-1})$ is the portfolio used on the t^{th} trading period given past market sequence $\mathbf{x}_1^{t-1} = (\mathbf{x}_1, \ldots, \mathbf{x}_{t-1})$. We denote by $\mathbf{b}_1^n = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ the strategy for n periods.

On the t^{th} trading period, a portfolio $\mathbf{b}_t$ achieves a *portfolio period return* s_t, that is, the wealth increases by a factor of $s_t = \mathbf{b}_t^T \mathbf{x}_t = \sum_{j=1}^d b_t^j x_t^j$. Since we reinvest and adopt price relative, the portfolio wealth would multiplicatively grow. Thus, after n trading periods, a portfolio strategy $\mathbf{b}_1^n$ produces a *portfolio cumulative wealth* S_n, which increases the initial wealth by a factor of $\prod_{t=1}^n \mathbf{b}_t^T \mathbf{x}_t$, that is, $S_n(\mathbf{b}_1^n, \mathbf{x}_1^n) = S_0 \prod_{t=1}^n (\mathbf{b}_t^T \mathbf{x}_t)$, where S_0 is the initial wealth, which is set to 1 in this paper.

Finally, we formulate the on-line PS problem as a sequential decision task. The portfolio manager aims to design a strategy $\mathbf{b}_1^n$ to maximize the portfolio cumulative wealth S_n. The portfolios are selected in a sequential fashion. On each period t, given the historical information, the manager learns to select a new portfolio vector $\mathbf{b}_t$ for next price relative vector $\mathbf{x}_t$, where the decision criterion varies among different managers. The resulting portfolio $\mathbf{b}_t$ is scored based on the portfolio period return of s_t. Such procedure repeats until

the end of trading periods and the portfolio strategy is finally scored by the cumulative wealth S_n.

In the above model, we ideally assume no transaction cost, perfect market liquidity and zero impact cost. These assumptions are not trivial, which has been explained in all existing work (refer to Section 3 for detail). We will empirically analyze the effects of transaction costs in Section 5.2.

3 Related Work

On-line portfolio selection has been extensively explored following the principle of Kelly investment [Kelly, 1956]. One classical strategy is *Constantly Rebalanced Portfolios* (CRP), which keeps fixed weight on each asset for every period. *Best CRP* (BCRP) [Cover, 1991], the best CRP strategy over a whole market sequence, is an optimal strategy if the market is i.i.d. [Cover and Thomas, 1991]. *Successive Constantly Rebalanced Portfolios* (SCRP) [Gaivoronski and Stella, 2000] and *Online Newton Step* (ONS) [Agarwal *et al.*, 2006] implicitly estimate next price relative via all historical price relatives with a uniform probability[1].

Besides estimation via all historical price relatives, some strategies predict next price relatives by selecting a set of similar price relatives. *Nonparametric kernel based moving window* (B^K) [Gyorfi *et al.*, 2006] measures the similarity by kernel method. Following the same framework, *Nonparametric Nearest Neighbor* (B^{NN}) [Gyorfi *et al.*, 2008] locates the set of price relatives via nearest neighbor methods. Li et.al [2011] proposed *Correlation-driven Nonparametric learning* (CORN), which measures the similarity via correlation.

Moreover, another category of estimation is to predict next price relative via a single-value prediction. *Exponential Gradient* (EG)[Helmbold *et al.*, 1998] estimates next price relative as last price relative. *Passive Aggressive Mean Reversion* (PAMR) [Li *et al.*, 2012a] and *Confidence Weighted Mean Reversion* (CWMR) [Li *et al.*, 2013] estimate next price as the inverse of last price relative, which is in essence the mean reversion principle[2]. Recently, Li and Hoi [2012b] proposed *On-Line Moving Average Reversion* (OLMAR), which predicts the next price relative using moving averages and explores the multi-period mean reversion.

Finally, some algorithms do not focus on estimation, either explicitly or implicitly. *Universal portfolios* (UP) [Cover, 1991] is the historical performance weighted average of all CRPs. *Anti-Correlation* (Anticor) [Borodin *et al.*, 2004] adopts the consistency of positive lagged cross-correlation and negative autocorrelation to adjust the portfolio.

3.1 Analysis of Existing Work

Now, let us focus on the estimation methods of existing work. In practice, a Kelly portfolio manager [Kelly, 1956; Thorp, 1971] firstly predicts $\hat{\mathbf{x}}_{t+1}$ in terms of k possible values $\hat{\mathbf{x}}_{t+1}^1, \ldots, \hat{\mathbf{x}}_{t+1}^k$ and their corresponding probabilities $p_1, \ldots, p_k$, where each $\hat{\mathbf{x}}_{t+1}^i$ denotes one possible combination vector of individual price relative predictions. Then

---

[1]SCRP and ONS's formulations are similar, while they use different techniques to solve the formulations.

[2]PAMR and CWMR adopt the same estimation, while they exploit the principle via different techniques.

he/she can figure out a portfolio by maximizing the expected log return on the possible combinations,

$$\mathbf{b}_{t+1} = \arg\max_{\mathbf{b} \in \Delta_d} \sum_{i=1}^{k} p_i \log \left(\mathbf{b} \cdot \hat{\mathbf{x}}_{t+1}^i \right).$$

As different estimation methods result in different $\hat{\mathbf{x}}_{t+1}^i$ and p_i and lead to different portfolio, an accurate estimation method is crucial to the success of a strategy.

Below, we focus on the algorithms PAMR, CWMR and OLMAR, which estimate next price relative via a single value prediction based on mean reversion or moving average reversion. PAMR and CWMR implicitly assume $\hat{\mathbf{x}}_{t+1}^1 = \frac{1}{\mathbf{x}_t}$ with $p_1 = 100\%$, i.e., they estimate next price relative as the inverse of last price relative, which is in essence the mean reversion principle. From the price perspective [Li and Hoi, 2012b], they implicitly assume that next price $\hat{\mathbf{p}}_{t+1}$ will revert to last price $\mathbf{p}_{t-1}$,

$$\hat{\mathbf{x}}_{t+1} = \frac{1}{\mathbf{x}_t} \Rightarrow \frac{\hat{\mathbf{p}}_{t+1}}{\mathbf{p}_t} = \frac{\mathbf{p}_{t-1}}{\mathbf{p}_t} \Rightarrow \hat{\mathbf{p}}_{t+1} = \mathbf{p}_{t-1},$$

where $\mathbf{x}$ and $\mathbf{p}$ are all vectors and the above operations are element-wise. Rather than $\hat{\mathbf{p}}_{t+1} = \mathbf{p}_{t-1}$, OLMAR estimates the next price as a moving average at the end of t^{th} period, that is, $\hat{\mathbf{p}}_{t+1} = MA_t(w) = \frac{1}{w} \sum_{i=t-w+1}^{i=t} \mathbf{p}_i$ where $MA_t(w)$ denotes the moving average with a w-window. Though empirically effective on most datasets, current estimation methods in PAMR/CWMR and OLMAR cause two potential problems. Firstly, the single-period mean reversion assumption of PAMR and CWMR may not be satisfied in the real world. One real example [Li *et al.*, 2012a] is DJA dataset [Borodin *et al.*, 2004], on which PAMR performs the worst among the state of the art. Secondly, all three algorithms suffer from the frequently fluctuating raw prices, which often contain a lot of noises and outliers. The two drawbacks thus motivate the proposed method.

4 Robust Median Reversion

4.1 Motivation

Empirical results [Li *et al.*, 2013; 2012a] show that if asset price follows the normal distribution, the mean of historical prices may better explain the markets. OLMAR, which estimates next price via moving average, also achieves good results on most datasets. However, due to noises and outliers in the data, the price distribution often has longtail, thus previous estimation methods are sub-optimal on the real markets.

To illustrate the drawbacks of mean and moving average, let us see a toy example. The toy market consists of one volatile stock, and $t_i(i \geq 0)$ denotes the period that requires estimation. Several sequences of market prices are listed in Table 1. A_0, A_1 are single-period price sequences and their prices change by sequent factor of $2, \frac{1}{2}, 2, \frac{1}{2}, \ldots$. For example, let P_{t_i} be the price of the i^{th} period, then $P_{t_1} = P_{t_0} \times 2 = 1 \times 2 = 2$, $P_{t_2} = P_{t_1} \times \frac{1}{2} = 2 \times \frac{1}{2} = 1$, $P_{t_3} = P_{t_2} \times 2 = 1 \times 2 = 2, \ldots$. B_0, B_1 are two-period price sequences and their prices change by sequent factor of $2, 2, \frac{1}{2}, \frac{1}{2}, 2, 2, \ldots$. C_0, C_1 are the three-period price sequences, and the price changes

Price:$t_0 \to t_1 \to \ldots$	Acc	PAMR/ CWMR	OLMAR	RMR
$A_0 : 1, 2, 1, 2, ?$	1	1	1.5	1.5
$A_1 : 1, 2, (10), 2, ?$	1	10	3.75	2
$B_0 : 1, 2, 4, 2, ?$	1	4	2.25	2
$B_1 : 1, 2, (10), 2, ?$	1	10	3.75	2
$C_0 : 1, 2, 4, 8, 4, 2, ?$	1	4	3.5	3
$C_1 : 1, 2, 4, 8, (10), 2, ?$	1	10	4.5	3

Table 1: Illustration of different price estimation methods on toy markets. A_0, A_1; B_0, B_1 and C_0, C_1 represent single-period, two-period and three-period price sequence respectively. A_0, B_0, C_0 are exact price sequence, and A_1, B_1, C_1 are price sequence contaminated by an outlier of 10. "Acc" is the accurate price. Other three items represent three estimators based on three different methods.

by sequent factor of $2, 2, 2, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 2, 2, 2, \ldots$. Moreover, A_0, B_0, C_0 are exact price sequences, while A_1, B_1, C_1 are the sequences contaminated by a outlier of 10. "?" denotes the price to be estimated and "Acc" is the accurate price. The estimated prices clear show that the next prices estimated by PAMR/CWMA and OLMAR are far away from the accurate values, which thus leads to inaccurate price relatives and suboptimal portfolios. Contrarily, the proposed methods release much better estimations (the calculation will be detailed later) and subsequent better portfolios. Note that although the toy example is on a single asset, such estimation goodness can be easily extended to the scenario of multiple assets.

To better exploit (multiple period) reversion property, we proposed a new type of algorithms for on-line PS, named "*Robust Median Reversion*" (RMR). The essential idea is to exploit multiple period reversion via robust L_1-median estimator [Weber, 1909; Weiszfeld, 1937; Vardi and Zhang, 2000] and power online machine learning. Rather than $\hat{\mathbf{p}}_{t+1} = \mathbf{p}_{t-1}$ or $\hat{\mathbf{p}}_{t+1} = MA_t(w)$, RMR estimates next price by robust L_1-median estimator at the end of t^{th} period, that is, $\hat{\mathbf{p}}_{t+1} = L_1med_{t+1}(w) = \boldsymbol{\mu}$, where w is the window size, $\boldsymbol{\mu}$ denotes the value in L_1-median estimator that satisfied the Optimization Problem 1 below (see section 4.2). Then the expected price relative with L_1-median estimator is

$$\hat{\mathbf{x}}_{t+1}(w) = \frac{L_1med_{t+1}(w)}{\mathbf{p}_t} = \frac{\boldsymbol{\mu}}{\mathbf{p}_t}. \tag{1}$$

Without detailing the calculation, we list the estimated next price of RMR in different toy markets in Table 1. Clearly, for the multiple period price sequences B_0, B_1 and C_0, C_1, RMR estimator is much closer to the Accurate estimator than PAMR/CWMR, showing that RMR method can deal with multiple period price sequence. For the contaminated price sequences A_1, B_1, C_1, RMR is also closer to the Accurate estimator than OLMAR and PAMR/CWMR estimators, which shows that RMR is a robust method.

Moreover, L_1-median estimator is much better than mean estimators statistically. In fact, the L_1-median has an attractive statistical properties, that is, its breakdown point is 0.5 [Lopuhaa and Rousseeuw, 1991], i.e., only if more than 50% of the data points are contaminated, the L_1-median can take values beyond all bounds. Note that breakdown point, the proportion of incorrect observations an estimator can handle,

is a statistical metric of robustness. The higher the breakdown point of an estimator is, the more robust it is. However, the breakdown point of mean is 0, which means that the mean estimator is sensitive to the noisy data and outliers.

Based on the expected price relative vector in Eq. (1), RMR further adopts the idea of an effective online learning algorithm, that is, Passive Aggressive (PA) [Grammer *et al.*, 2006] learning, to exploit median reversion. Generally proposed for classification, PA passively keeps the previous solutions if the classification is correct, while aggressively approaches a new solution if the classification is incorrect. After formulating the proposed RMR optimization, we solve its closed form update and design the proposed algorithm.

4.2 Formulation

The proposed formulation, RMR, is to exploit median reversion via robust L_1-estimator and PA online learning. The basic idea is to obtain next price relative $\hat{\mathbf{x}}_{t+1}$ via multivariate L_1-median, and then maximize the expected return $\mathbf{b} \cdot \hat{\mathbf{x}}_{t+1}$ while keeping last portfolio information via regularization.

To estimate the next price relative $\hat{\mathbf{x}}_{t+1}$, we calculate the multivariate L_1-median of historical prices. In statistics, the L_1-median (also named spatial median) of a k-historical price window, $\boldsymbol{\mu}$, is the solution of following optimization,

Optimization problem 1: L_1-median

$$\boldsymbol{\mu} = \operatorname*{argmin}_{\boldsymbol{\mu}} \sum_{i=0}^{k-1} \|\mathbf{p}_{t-i} - \boldsymbol{\mu}\|, \tag{2}$$

where $\|\cdot\|$ denotes the Euclidean norm. In a word, L_1-median is the point with minimal sum of Euclidean distances to the k given price data points. This problem is formulated in an even more general form by Weber [Weber, 1909] (Fermat-Weber problem), as he refers to location issues in industrial contexts. If the data points are not collinear, the solution to problem (2) is unique [Weiszfeld, 1937].

After obtaining the L_1-median estimator of price, we can calculate next price relative $\hat{\mathbf{x}}_{t+1}$ by Eq. (1). Based on the obtained price relative $\hat{\mathbf{x}}_{t+1}$, we can select next portfolio via PA online learning technique. Thus, following the similar idea PAMR and OLMAR [Li *et al.*, 2012a; Li and Hoi, 2012b], we can formulate the following optimization,

Optimization problem 2: RMR

$$\mathbf{b}_{t+1} = \operatorname*{argmin}_{\mathbf{b} \in \Delta_d} \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 \quad s.t. \ \mathbf{b} \cdot \hat{\mathbf{x}}_{t+1} \geq \varepsilon. \tag{3}$$

The above formulation attempts to find an optimal portfolio by minimizing the deviation from last portfolio $\mathbf{b}_t$ under the condition of $\mathbf{b} \cdot \hat{\mathbf{x}}_{t+1} \geq \varepsilon$. Such formulation explicitly reflects the reversion idea underlying the proposed RMR. In fact, $\hat{\mathbf{x}}_{t+1}$ is the price relative estimated via L_1-median estimator, while the constraint $\mathbf{b} \cdot \hat{\mathbf{x}}_{t+1} \geq \varepsilon$ means that next price will revert to the L_1-median.

4.3 Algorithm

To obtain the L_1-median of historical prices, we apply the Modified Weiszfeld Algorithm [Vardi and Zhang, 2000],

which converges monotonically to the L_1-median. The solution of L_1-median described in Eq. (2) is illustrated below.

Proposition 1 *The solution of L_1-median optimization problem 1 is calculated through iteration, and the iteration process is described as:*

$$\boldsymbol{\mu} \to T(\boldsymbol{\mu}) = \left(1 - \frac{\eta(\boldsymbol{\mu})}{\gamma(\boldsymbol{\mu})}\right)^+ \tilde{T}(\boldsymbol{\mu}) + \min\left(1, \frac{\eta(\boldsymbol{\mu})}{\gamma(\boldsymbol{\mu})}\right)\boldsymbol{\mu},$$

where

$$\eta(\boldsymbol{\mu}) = \begin{cases} 1 & if \ \boldsymbol{\mu} = \mathbf{p}_{t-i}, \ i = 0, \ldots, k-1 \\ 0 & otherwise \end{cases},$$

$$\gamma(\boldsymbol{\mu}) = \|\tilde{R}(\boldsymbol{\mu})\|, \ \ \tilde{R}(\boldsymbol{\mu}) = \sum_{\mathbf{p}_{t-i} \neq \boldsymbol{\mu}} \frac{\mathbf{p}_{t-i} - \boldsymbol{\mu}}{\|\mathbf{p}_{t-i} - \boldsymbol{\mu}\|},$$

$$\tilde{T}(\boldsymbol{\mu}) = \left\{\sum_{\mathbf{p}_{t-i} \neq \boldsymbol{\mu}} \frac{1}{\|\mathbf{p}_{t-i} - \boldsymbol{\mu}\|}\right\}^{-1} \sum_{\mathbf{p}_{t-i} \neq \boldsymbol{\mu}} \frac{\mathbf{p}_{t-i}}{\|\mathbf{p}_{t-i} - \boldsymbol{\mu}\|}.$$

In general, we often practically set the convergence criterion during the iteration. Here, once the constraint of $\|\boldsymbol{\mu}_{\iota-1} - \boldsymbol{\mu}_\iota\|_1 \leq \tau\|\boldsymbol{\mu}_\iota\|_1$ is satisfied, we terminate the iteration. Note that τ represents a toleration level.

We now obtain the final portfolio selection formula by solving the Optimization problem 2, which is convex and thus straightforward to solve via the Lagrange multiplier method [Boyd and Vandenberghe, 2004].

Proposition 2 *The solution of the Optimization problem 2 without considering the non-negativity constraint is*

$$\mathbf{b}_{t+1} = \mathbf{b}_t - \alpha_{t+1}(\hat{\mathbf{x}}_{t+1} - \overline{x}_{t+1} \cdot \mathbf{1}),$$

where $\overline{x}_{t+1} = \frac{1}{d}(\mathbf{1} \cdot \hat{\mathbf{x}}_{t+1})$ denotes the average predicted price relative and α_{t+1} is the Lagrangian multiplier calculated as,

$$\alpha_{t+1} = \min\left\{0, \frac{\hat{\mathbf{x}}_{t+1}\mathbf{b}_t - \epsilon}{\|\hat{\mathbf{x}}_{t+1} - \overline{x}_{t+1} \cdot \mathbf{1}\|^2}\right\}.$$

Note that it is possible that the resulting portfolio in Proposition 2 goes out of the simplex domain since we do not consider the non-negativity constraint. Thus, to ensure that the portfolio is non-negative, we finally project the above portfolio to the simplex domain [Duchi *et al.*, 2008].

To this end, we can design the proposed algorithms based on the above Propositions. The estimated process of price relative $\hat{x}_{t+1}$, mainly based on Proposition 1, is illustrated in Algorithm 1. The proposed RMR procedure, which is designed up on Proposition 2, is shown in Algorithm 2. Finally, Algorithm 3 presents the on-line portfolio selection following the problem setting in Section 2.

4.4 Analysis

It is widely known that computational time is important to certain trading environments, such as high frequency trading [Aldridge, 2010], which occurs in fractions of a second. RMR's time complexity is linear with respect to d and n, where n is much larger than d. In the RMR implementation, the max number of loop (Line 6 in Algorithm 1) can be implemented in $O(m)$. Thus, Algorithm 1 take $O(m)$ time

Algorithm 1 L_1 median $(\mathbf{p}_t, \mathbf{p}_{t-1}, \ldots, \mathbf{p}_{t-w+1}, m, \tau)$

1: **Input:** data $\mathbf{p}_t, \mathbf{p}_{t-1}, \ldots, \mathbf{p}_{t-w+1}$; iteration maximum m; toleration level τ
2: **Output:** estimated $\hat{\mathbf{x}}_{t+1}$
3: **Procedure:**
4: Initialize $\boldsymbol{\mu}_1 = median(\mathbf{p}_t, \mathbf{p}_{t-1}, \ldots, \mathbf{p}_{t-w+1})$.
5: **for** $i = 2$ **to** m **do**
6: $\quad \boldsymbol{\mu}_i = T(\boldsymbol{\mu}_{i-1})$
7: $\quad$ **if** $\|\boldsymbol{\mu}_{i-1} - \boldsymbol{\mu}_i\|_1 \leq \tau \|\boldsymbol{\mu}_i\|_1$ **then**
8: $\quad\quad$ break
9: $\quad$ **end if**
10: **end for**
11: $\hat{\mathbf{p}}_{t+1} = \boldsymbol{\mu}_i$
12: $\hat{\mathbf{x}}_{t+1} = \hat{\mathbf{p}}_{t+1}/\mathbf{p}_t$

Algorithm 2 RMR$(\epsilon, \hat{x}_{t+1}, \mathbf{b}_t)$

1: **Input:** reversion threshold $\epsilon > 1$; predicted next price relative vector $\hat{\mathbf{x}}_{t+1}$; current portfolio $\mathbf{b}_t$;
2: **Output:** next portfolio $\mathbf{b}_{t+1}$
3: **Procedure:**
4: Calculate the following variable:

$$\alpha_{t+1} = \min\left\{0, \frac{\hat{\mathbf{x}}_{t+1}\mathbf{b}_t - \epsilon}{\|\hat{\mathbf{x}}_{t+1} - \overline{x}_{t+1} \cdot \mathbf{1}\|^2}\right\}$$

5: Update the portfolio:

$$\mathbf{b}_{t+1} = \mathbf{b}_t - \alpha_{t+1}(\hat{\mathbf{x}}_{t+1} - \overline{x}_{t+1} \cdot \mathbf{1})$$

6: Normalize $\mathbf{b}_{t+1}$: $\mathbf{b}_{t+1} = \operatorname{argmin}_{\mathbf{b} \in \Delta_d} \|\mathbf{b} - \mathbf{b}_{t+1}\|^2$

---

per period. Moreover, Algorithm 2 takes $O(d)$ per period. In total, the whole time complexity is $O(dn) + O(mn)$. Table 2 compares the time complexity of RMR with that of existing strategies. Clearly, the proposed RMR algorithm takes no more time than any others.

Methods	Time Complexity	Methods	Time Complexity
UP	$O\left(n^d\right)/O\left(d^7 n^8\right)$	ONS	$O\left(d^3 n\right)$
EG	$O\left(dn\right)$	Anticor	$O\left(N^3 d^2 n\right)$
PAMR/CWMR	$O\left(dn\right)$	B^K/B^{NN}	$O\left(N^2 dn^2\right)$
/OLMAR		/CORN	$+O\left(Ndn^2\right)$
RMR	$O\left(dn\right) + O\left(mn\right)$		

Table 2: Summary of time complexity analysis. d denotes the number of stocks; n is the number of trading periods; N denotes the number of experts; m denotes the number of loop in Algorithms 1.

5 Experiments

5.1 Datasets

We test the portfolio strategies on four public datasets from real markets [3], which are summarized in Table 3.

Pioneered by Cover [1991], NYSE(O) is one benchmark dataset, which consists of 36 stocks from the New York Stock

---

Algorithm 3 Portfolio selection with RMR

1: **Input:** reversion threshold $\epsilon > 1$; iteration maximum m; window size $w \geq 2$; toleration level τ; market sequence $\mathbf{x}_1^n$
2: **Output:** S_n: Cumulative wealth after n^{th} periods
3: **Procedure:**
4: Initialization: $b_1 = \frac{1}{d}\mathbf{1}, S_0 = 1, \mathbf{p}_0 = \mathbf{1}$
5: **for** $i = 1$ **to** n **do**
6: $\quad \mathbf{p}_i = \mathbf{x}_i \cdot \mathbf{p}_{i-1}$
7: **end for**
8: **for** $t = 1, 2, \ldots, n$ **do**
9: $\quad$ Receive stock price: $\mathbf{x}_t$
10: $\quad$ Update cumulative return: $S_t = S_{t-1} \times (\mathbf{b}_t \cdot \mathbf{x}_t)$
11: $\quad$ Predict next price relative vector:

$$\hat{\mathbf{x}}_{t+1} = L_1 median(\mathbf{p}_t, \mathbf{p}_{t-1}, \ldots, \mathbf{p}_{t-w+1}, m, \tau)$$

12: $\quad$ Update the portfolio:

$$\mathbf{b}_{t+1} = RMR(\epsilon, \hat{\mathbf{x}}_{t+1}, \mathbf{b}_t)$$

13: **end for**

Exchange. Extended from NYSE(O), NYSE(N) [Li *et al.*, 2013] contains 23 stocks. The dataset DJA is collected by Borodin *et al.* [2004] and consists of Dow Jones 30 composite stocks. The final dataset MSCI is a collection of 25 global equity indices that are the constituents of MSCI World Index.

DATA SET	REGION	TIME FRAME	#DAYS	#ASSETS
NYSE(O)	US	3/7/1962-31/12/1984	5651	36
NYSE(N)	US	1/1/1985-30/6/2010	6431	23
DJA	US	1/1/2001-14/1/2003	507	30
MSCI	GLOBAL	1/4/2006-31/3/2010	1043	24

Table 3: Summary of 4 real datasets.

RMR has two possible parameters, i.e., w and ε, which can be tuned to obtain optimal results. To consistently compare different methods, we empirically set the parameter w=5, ε=5 and $m = 200$ on all settings. Our experiments on parameter sensitivity clearly show that our empirical choice is not the best in hindsight.

5.2 Experimental Results

Cumulative wealth

Table 4 summarizes the cumulative wealth achieved by various methods. From the figure, we can see that RMR outperform the state of the art on NYSE(O), NYSE(N) and DJA. On MSCI, RMR beats most existing algorithms, including OL-MAR. Besides, the maximum cumulative wealth is even better than RMR, showing the potential of the proposed method. Finally, Table 5 shows some statistics [Grinold and Kahn, 1999] of RMR. From the results, we can observe small p-values, which means that RMR's excellent performance is not due to luck but owed to the strategy principle.

Parameter sensitivity

Firstly, we examine the effect of sensitivity parameter ϵ on cumulative wealth, in Figure 1. The cumulative wealth sharply

---

[3]All datasets and their compositions can be downloaded from http://www.cais.ntu.edu.sg/ chhoi/olps/datasets.html.

Methods	NYSE(O)	NYSE(N)	DJA	MSCI
Market	14.50	18.06	0.76	0.91
Best-stock	54.14	83.51	1.19	1.50
BCRP	250.60	120.32	1.24	1.51
UP	26.68	31.49	0.81	0.92
EG	27.09	31.00	0.81	0.93
ONS	109.91	21.59	1.53	0.86
B^k	1.08E+09	4.64E+03	0.68	2.64
B^{NN}	3.35E+11	6.80E+04	0.88	13.47
CORN	1.48E+13	5.37E+05	0.84	**26.19**
Anticor	2.41E+08	6.21E+06	2.29	3.22
PAMR	5.14E+15	1.25E+06	0.68	15.23
CWMR	6.49E+15	1.41E+06	0.68	17.28
OLMAR	4.04E+16	2.24E+08	2.05	16.33
RMR	**1.64E + 17**	**3.25E + 08**	**2.67**	16.76
RMR(max)	2.81E+17	4.73E+08	3.47	19.07

Table 4: Cumulative wealth achieved by various strategies on the four datasets. The best results (excluding the best experts at the bottom, which is in hindsight) in each dataset are highlighted in **bold**.

Stat.	NYSE(O)	NYSE(N)	DJA	MSCI
Size	5651	6431	507	1043
MER(RMR)	0.0077	0.0036	0.0024	0.0030
MER(Market)	0.0005	0.0005	-0.0004	0.0000
α	0.0071	0.0031	0.0030	0.0030
β	1.2718	1.1628	1.2427	1.1885
t-statistics	15.7325	7.4222	2.5217	5.8380
p-value	0.0000	0.0000	0.0060	0.0000

Table 5: Statistical Test of RMR.

grows as ϵ increases and then flattens when ϵ crosses a threshold. Secondly, we examine the effect of window size w, in Figure 2. The cumulative wealth decreases with increasing w. Obviously, RMR's performance with a large number of ϵ and w is better than the two benchmarks. All above observations clearly show that RMR is robust w.r.t. different parameters and it is convenient to choose satisfying parameters.

Transaction costs

In practice, transaction cost is an important and unavoidable issue that should be addressed. We thus test the effect of proportional transaction cost when the transaction cost rate γ varies from 0 to 1%. We also plot the results achieved by two benchmarkes of BCRP and Market and the cumulative wealth achieved by PAMR[Li *et al.*, 2012a], OLMAR[Li and Hoi, 2012b]. From Figure 3, we can observe that RMR can withstand reasonable transaction cost rates, and can beat the two benchmarks and PAMR, OLMAR.

6 Conclusions

In this paper, we propose a novel multiple period on-line portfolio selection strategy named "robust median reversion" (RMR), which exploits the reversion phenomenon of stock prices by robust L_1-median estimator and online learning technologies. The proposed approach can solve the problems of the state of the art caused by the noisy data and outliers as well as the single-period mean reversion. Extensive experiments on real markets show that the proposed RMR can achieve satisfying performance. In future, we will study other robust estimation methods and adaptive parameter methods and theoretically analyze the median reversion.


(a) NYSE(O) (b) NYSE(N)

(c) DJA (d) MSCI

Figure 1: Parameter sensitivity of RMR w.r.t. ϵ with fixed w (w=5)

References

[Agarwal *et al.*, 2006] A. Agarwal, E. Hazan, S. kale, and R.E. Schapire. Algorithms for portfolio management based on the newton method. In *Proceedings of the International Conference on Machine Learning (ICML 2006)*, pages 9–16, 2006.

[Aldridge, 2010] Irene Aldridge. *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Hoboken, NJ : Wiley, Hoboken, N.J., 2010.

[Borodin *et al.*, 2004] A. Borodin, R. El-Yaniv, and V. Gogan. Can we learn to beat the best stock. *JAIR*, 21:579–594, 2004.

[Boyd and Vandenberghe, 2004] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, New York, 2004.

[Cover and Thomas, 1991] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. New York : Wiley, New York, 1991.

[Cover, 1991] T.M. Cover. Universal portfolios. *Mathematical Finance*, 1:1–29, 1991.

[Duchi *et al.*, 2008] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the l_1-ball for learning in high dimensions. In *Proceedings of the International Conference on Machine Learning (ICML 2008)*, 2008.

[Gaivoronski and Stella, 2000] Alexei A. Gaivoronski and Fabio Stella. Stochastic nonstationary optimization for finding universal portfolios. *Annals of Operations Research*, 100:165–188, 2000.

[Grammer *et al.*, 2006] K. Grammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer. Online passive-aggressive algorithms. *JMLR*, 7:551–585, 2006.

[Grinold and Kahn, 1999] R. Grinold and R. Kahn. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*. McGraw-Hill, New York, 1999.
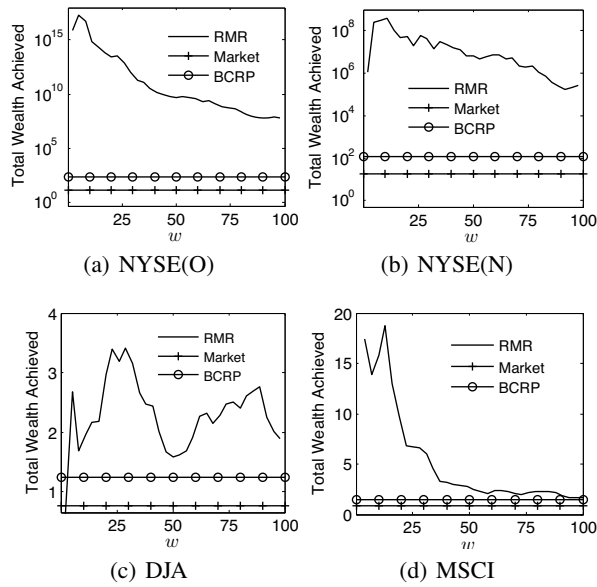
(a) NYSE(O) (b) NYSE(N)

(c) DJA (d) MSCI

Figure 2: Parameter sensitivity of RMR w.r.t. w with fixed ϵ ($\epsilon = 5$)

(a) NYSE(O) (b) NYSE(N)

(c) DJA (d) MSCI

Figure 3: Scalability of the total wealth achieved by RMR with respect to transaction cost rate $\gamma\%$

[Gyorfi *et al.*, 2006] L. Gyorfi, G. Lugosi, and F. Udina. Nonparametric kernel-based sequential investment strategies. *Math. Finan.*, 16(2):337–357, 2006.

[Gyorfi *et al.*, 2008] L. Gyorfi, F. Udina, and H. Walk. Nonparametric nearest neighbor based empirical portfolio selection strategies. *Statistics and Decisions*, 26(2):145–157, 2008.

[Helmbold *et al.*, 1998] D.P. Helmbold, R.E. Schapire, Y. Singer, and M.K. Warmuth. On-line portfolios portfolio selection using multiplicative updates. *Math.Finan*, 8(4):325–347, 1998.

[Jegadeesh., 1990] N. Jegadeesh. Evidence of predictable behavior of security returns. *Journal of Finance*, 45(3):881–898, 1990.

[Kelly, 1956] J. Jr. Kelly. A new interpretation of information rate. *AT&T Tech. J.*, 35:917–926, 1956.

[Li and Hoi, 2012] B. Li and S.C.H. Hoi. On-line portfolio selection: A survey. *CoRR*, abs/1212.2129, 2012.

[Li and Hoi, 2012b] B. Li and S. C. H. Hoi. On-line portfolio selection with moving average reversion. In *Proceedings of the International Conference on Machine Learning (ICML 2012)*, 2012b.

[Li *et al.*, 2011] B. Li, S. C. H. Hoi, and V. Gopalkrishnan. Corn: correlation-driven nonparametric learning approach for portfolio selection. *ACM TIST*, 2(3):21:1–21:29, 2011.

[Li *et al.*, 2012a] B. Li, P. Zhao, S. C.H. Hoi, and V. Gopalkrishnan. Pamr: Passive aggressive mean reversion strategy for portfolio selection. *Mach. Learn.*, 87(2):221–258, 2012a.

[Li *et al.*, 2013] B. Li, S. C. H. Hoi, P.L. Zhao, and V. Gopalkrishnan. Confidence weighted mean reversion strategy for online portfolio selection. *ACM Transactions on Knowledge Discovery from Data*, 7(1):4:1–4:38, 2013.

[Lopuhaa and Rousseeuw, 1991] H. Lopuhaa and P. Rousseeuw. Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *The Annals of Statistics*, 19(1):229–248, 1991.

[Markowitz, 1952] H. Markowitz. Mean-variance analysis in portfolio choice and capital markets. *J. Finance*, 7(1):77–91, 1952.

[Merton, 1980] R. C. Merton. On estimating the expected return on a market: An exploratory investigation. *J. Financial Econ.*, 8:323–361, 1980.

[Ordentlich and Cover, 1996] E. Ordentlich and T.M. Cover. On-line portfolio selection. In *COLT*, pages 310–313, 1996.

[Thorp, 1971] E. O. Thorp. Portfolio choice and the kelly criterion. In *Business and Economics Section of the American Statistical Association*, 1971.

[Vardi and Zhang, 2000] Yehuda Vardi and Cun Hui Zhang. The multivariate l_1-median and associated data depth. *PNAS*, 97(4):1423–1426, 2000.

[Weber, 1909] A. Weber. *Uber den Standort der Industrien*. Tubingen : Mohr, 1909.

[Weiszfeld, 1937] E. Weiszfeld. Sur le point pour lequel la somme des distances de n points donnes est minimum. *Tohoku Mathematical Journal*, 43:355–386, 1937.