

PAMR: Passive Aggressive Mean Reversion Strategy for Portfolio Selection

Bin Li · Peilin Zhao · Steven C. H. Hoi* · Vivekanand Gopalkrishnan

Received: date / Accepted: date

Abstract This article proposes a novel online portfolio selection strategy named “Passive Aggressive Mean Reversion” (PAMR). Unlike traditional trend following approaches, the proposed approach relies upon the mean reversion relation of financial markets. Equipped with online passive aggressive learning technique from machine learning, the proposed portfolio selection strategy can effectively exploit the mean reversion property of markets. By analyzing PAMR’s update scheme, we find that it nicely trades off between portfolio return and volatility risk and reflects the mean reversion trading principle. We also present several variants of PAMR algorithm, including a mixture algorithm which mixes PAMR and other strategies. We conduct extensive numerical experiments to evaluate the empirical performance of the proposed algorithms on various real datasets. The encouraging results show that in most cases the proposed PAMR strategy outperforms all benchmarks and almost all state-of-the-art portfolio selection strategies under various performance metrics. In addition to its superior performance, the proposed PAMR runs extremely fast and thus is very suitable for real-life online trading applications. The experimental testbed including source code and data sets is available at <http://www.cais.ntu.edu.sg/~chhoi/PAMR/>.

Keywords Portfolio selection · Mean reversion · Passive aggressive learning · Online learning

* Corresponding author: Prof. Steven C.H. Hoi, School of Computer Engineering, Nanyang Technological University, Singapore 639798, Tel: (+65) 6513-8040, Fax: (+65) 6792-6559.

Steven C.H. Hoi
E-mail: chhoi@ntu.edu.sg

Bin Li
E-mail: s080061@ntu.edu.sg

Peilin Zhao
E-mail: zhao0106@ntu.edu.sg

Vivekanand Gopalkrishnan
E-mail: vivekanand@deloitte.com

1 Introduction

Portfolio Selection (PS) is a practical financial engineering problem that requires determining a strategy of investing wealth among a set of assets in order to achieve certain objectives, such as maximizing cumulative wealth or risk-adjusted return, in the long run. In this article, we investigate *sequential portfolio selection* (also termed *online portfolio selection*) strategies, which sequentially determine portfolios based on publicly available information.

Traditionally in finance, portfolios are often selected according to *mean-variance* theory (Markowitz 1952, 1959) or its variants, to trade off between return and risk. In recent years, this problem has also been actively studied from a *learning to select portfolio* perspective, with roots in the fields of machine learning, data mining, information theory and statistics. Rather than trading with a single stock using computational intelligence techniques, learning to select portfolio approach focuses on a portfolio, which consists of multiple assets/stocks. Several approaches for online portfolio selection, often characterized by machine learning formulations and effective optimization solutions, have been proposed in literature (Kelly 1956; Breiman 1961; Cover 1991; Ordentlich and Cover 1996; Helmbold et al. 1996; Borodin and El-Yaniv 1998; Borodin et al. 2000, 2004; Stoltz and Lugosi 2005; Hazan 2006; Györfi et al. 2006; Blum and Mansour 2007; Levina and Shafer 2008; Györfi et al. 2008). Despite being studied extensively, most approaches are limited in some aspects or the other.

Our goal of this work is to investigate a new online portfolio selection strategy that employs online learning techniques to exploit the financial markets. Some existing strategies adopt the *trend following* approach, that is, they assume that price relative will follow its historical trading days. However, this philosophy fails when price relatives do not go in any particular direction, but rather actively move within a range. So in this study, we exploit another well-known principle in finance, viz., mean reversion (Jegadeesh 1990), through an online machine learning framework. To this end, we propose a novel portfolio selection strategy named “Passive Aggressive Mean Reversion” (PAMR), which exploits the mean reversion property of financial markets by online passive aggressive learning (Crammer et al. 2006). PAMR’s key idea is to formulate a new loss function that can effectively exploit the mean reversion property, and then adopt passive aggressive online learning to search for optimal portfolio among the asset pool to maximize the cumulative return.

Under different scenarios, the proposed PAMR strategy either passively keeps last portfolio or aggressively approaches a new portfolio by following the mean reversion principle. By solving three well formulated optimization problems, we arrive at three simple portfolio update rules. It is interesting to find that the final portfolio update scheme reaches certain trade-offs between portfolio return and volatility risk, and explicitly reflects the mean reversion trading rule. Moreover, we propose a mixture algorithm, which mixes PAMR and other strategies, and show that the mixture can be universal if one universal strategy is included. The key advantages of PAMR are its highly competitive performance and fairly attractive computation time efficiency. Our extensive numerical experiments on various real datasets show that in most cases the proposed PAMR strategy is quite performance efficient in comparison to a number of state-of-the-art portfolio selection strategies under a variety of performance metrics. At the same time, the proposed strategy costs linear time with respect to the product of the number of stocks and trading days, and its computational time in back tests is orders of magnitude less than its competitors, showing its applicability to real-world large scale online applications.

As a summary, our contributions in this article include:

1. We propose a new algorithm for online portfolio selection, named “Passive Aggressive Mean Reversion” (PAMR). To the best of our knowledge, it is the first portfolio selection strategy that exploits both the mean reversion property in finance and the powerful online passive aggressive learning technique in machine learning.
2. We propose a mixture algorithm to mix the proposed PAMR algorithms and other universal strategies, resulting in a theoretically guaranteed universal mixture strategy.
3. We analyze the final portfolio update scheme of PAMR and show that it is essentially related to certain trade-offs between portfolio return and volatility risk.
4. We conduct an extensive set of numerical experiments on a number of up-to-date datasets from various markets. The results show that in most cases the proposed PAMR strategy not only outperforms the benchmarks (including market index, best stock and challenging best constant rebalanced portfolio (Cover 1991) in hindsight), but also outperforms various state-of-the-art strategies under various performance metrics tested.
5. We also extend the proposed strategy to handle some practical issues for a real-life portfolio selection task, viz., transaction cost and margin buying, and show its practical viability through the extensive empirical study.
6. We show that the time complexity of the proposed algorithm is linear with respect to the number of stocks per trading day, and its empirical computational time in the back tests is quite competitive compared with the state of the arts, indicating the proposed strategy is suitable for online large-scale real applications.

The rest of the article is organized as follows. Section 2 formally states online portfolio selection problem. Section 3 reviews existing state-of-the-art approaches tackling this problem, and highlights their limitations. Section 4 presents our proposed PAMR strategy and analyzes the algorithm. Section 5 validates the effectiveness of PAMR by extensive empirical studies on historical financial markets. Finally, Section 6 summarizes this article and indicates future directions.

2 Problem Setting

Let us consider a financial market with m assets, over which we wish to invest. The changes of asset prices for n trading periods are represented by a sequence of non-negative, non-zero *price relative vectors* $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}_+^m$. Let us use \mathbf{x}^n to denote such a sequence of vectors. The i^{th} component of the t^{th} vector x_{ti} denotes the ratio of closing price to last closing price of the i^{th} asset on the t^{th} trading day, thus an investment in asset i on the t^{th} trading day increases by a factor of x_{ti} .

An investment in the market is specified by a *portfolio vector* $\mathbf{b}_t = (b_{t1}, \dots, b_{tm})$, where b_{ti} represents the proportion of wealth invested in the i^{th} asset. Typically, we assume portfolio is self-financed and no margin/short 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 $\Delta_m = \{\mathbf{b} : \mathbf{b} \in \mathbb{R}_+^m, \sum_{i=1}^m b_i = 1\}$. The investment procedure is represented by a *portfolio strategy*, that is, a sequence of mappings $\mathbf{b}_1 = (\frac{1}{m}, \dots, \frac{1}{m})$, $\mathbf{b}_t : \mathbb{R}_+^{m(t-1)} \rightarrow \Delta_m, t = 2, 3, \dots$, where $\mathbf{b}_t = \mathbf{b}_t(\mathbf{x}_1, \dots, \mathbf{x}_{t-1})$ is the portfolio used on the t^{th} trading period given past market price relatives $\mathbf{x}^{t-1} = \{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}\}$. Let us denote by \mathbf{b}^n the portfolio strategy for n trading periods.

For the t^{th} trading day, an investment according to portfolio \mathbf{b}_t results in a *portfolio daily return* s_t , that is, the wealth increases by a factor of $s_t = \mathbf{b}_t^\top \mathbf{x}_t = \sum_{i=1}^m b_{ti} x_{ti}$. Since we use price relative, the investment results in multiplicative cumulative return. Thus, after n trading days, the investment according to a portfolio strategy \mathbf{b}^n results in *portfolio*

cumulative wealth S_n , which is increased by a factor of $\prod_{t=1}^n \mathbf{b}_t^\top \mathbf{x}_t$, that is,

$$S_n(\mathbf{b}^n, \mathbf{x}^n) = S_0 \prod_{t=1}^n \mathbf{b}_t^\top \mathbf{x}_t,$$

where S_0 denotes the initial wealth, and is set to \$1 in this article for convenience.

Finally, we formulate the online portfolio selection problem as a sequential decision problem. The portfolio manager is a decision maker whose goal is to make a portfolio strategy on financial markets to satisfy certain requirements. In this study, his target is to maximize the portfolio cumulative wealth. He computes his portfolios in a sequential fashion. On each trading day t , the portfolio manager has access to all previous sequences of price relative vectors $\mathbf{x}^{t-1} = \{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}\}$, and previous sequences of portfolio vectors $\mathbf{b}^{t-1} = \{\mathbf{b}_1, \dots, \mathbf{b}_{t-1}\}$. On the basis of these historical information, the portfolio manager computes a new portfolio vector \mathbf{b}_t for coming price relative vector \mathbf{x}_t . Note that without historical information, the initial portfolio is set to uniform. The resulting portfolio is evaluated by its portfolio daily return. This procedure is repeated until the end of a trading period, and the portfolio is finally evaluated according to the portfolio cumulative wealth achieved. Figure 1 models the portfolio selection problem as a sequential decision problem.

```

Initialize  $S_0 = 1, \mathbf{b}_0 = (\frac{1}{m}, \dots, \frac{1}{m})$ 
for each trading day  $t = 1, 2, \dots, n$  do
  (1) Portfolio manager learns the portfolio  $\mathbf{b}_t$  based on historical information  $(\mathbf{b}^{t-1}, \mathbf{x}^{t-1})$ 
  (2) Market reveals the market price relative  $\mathbf{x}_t$ 
  (3) Portfolio incurs a portfolio daily return  $s_t = \mathbf{b}_t^\top \mathbf{x}_t$ 
end for

```

Fig. 1: Portfolio Selection as a Sequential Decision problem.

In the above portfolio selection model, we make several general assumptions:

1. Transaction cost: we assume no transaction cost or taxes exists in this portfolio selection model;
2. Market liquidity: we assume that one can buy and sell required quantities at last closing price of any given trading period;
3. Impact cost: we assume that market behavior is not affected by a portfolio selection strategy in our study.

3 Related work

In this section, we review some popular portfolio selection approaches, and some machine learning and trading philosophies that inspire the proposed approach.

3.1 Benchmark Approaches

The most common baseline is *Buy-And-Hold* (BAH) strategy, that is, one invests his/her wealth among a pool of assets with an initial portfolio and holds the portfolio all the time. The BAH strategy with a *uniform* initial portfolio is referred to as uniform BAH strategy,

which is adopted as market strategy producing the market index in our study. Contrary to the static BAH strategy, active trading strategies usually change portfolios regularly during the entire trading periods. A classical active strategy is *Constant Rebalanced Portfolios* (CRP) (Cover and Gluss 1986), which keeps a fixed fraction of an investor's wealth in each underlying asset every trading day. The best possible CRP strategy is often called *Best CRP* (BCRP), which apparently is only a hindsight strategy. The CRP strategy can take advantage of market fluctuations for active trading, and its underlying idea is based on the mean reversion principle, or known as "*Buy Low, Sell High*". To handle transaction cost issue for CRP strategy, Blum and Kalai (1999) proposed *semi-CRP* that partially balances between potential return and potential transaction cost and rebalances to initial portfolio at the end of any subset of the trading periods rather than every trading day.

3.2 Online Learning

In this section, we briefly introduce the related work on online machine learning (Rosenblatt 1958; Crammer and Singer 2003; Cesa-Bianchi et al. 2004; Crammer et al. 2006; Fink et al. 2006) to have the learning inspiration for our work. *Perceptron* algorithm (Rosenblatt 1958; Freund and Schapire 1999) is one important online approach which updates the learning function by adding a new example with a constant weight when it is misclassified. Recently a number of online learning algorithms have been proposed based on the criterion of maximum margin (Li and Long 1999; Gentile 2001; Kivinen et al. 2001; Crammer and Singer 2003; Crammer et al. 2006; Zhao et al. 2011). For example, *Relaxed Online Maximum Margin* (ROMMA) (Li and Long 1999) algorithm repeatedly chooses the hyper-planes that correctly classify the existing training example with the maximum margin. *Passive Aggressive* (PA) (Crammer et al. 2006) algorithm updates the classification function when a new example is misclassified or its classification score does not exceed some predefined thresholds. As empirical studies show, the maximum margin based online learning algorithms are generally more effectively than the Perceptron algorithm. In this article, we mainly adopt the idea of Passive Aggressive learning since it is suitable for our motivations as further illustrated in Section 4.1.

3.3 Learning to Select Portfolio

Learning to select portfolio has been extensively studied in information theory and machine learning. Generally, a strategy selects one optimal strategy (it can be market strategy, challenging BCRP strategy, or even Oracle strategy which chooses the best stock every trading day) and tries to obtain the same cumulative return. The *regret* of a strategy is defined as the gap between its logarithmic cumulative wealth achieved and that of the optimal strategy.

One important type of learning to select portfolio is regret minimization approach, which chooses BCRP strategy as the optimal strategy. Cover (1991) proposed *Universal Portfolio* (UP) strategy, where the portfolio is historical performance weighted average of all constant rebalanced portfolio experts. The regret achieved by Cover's UP is $O(m \log n)$, and its run time complexity is $O(n^m)$, where m denotes the number of stocks and n denotes the number of trading days. The implementation is exponential in the number of stocks and thus restricts the number of assets used in experiments and real applications. Kalai and Vempala (2002) presented a time-efficient implementation of Cover's UP based on non-uniform random walks that are rapidly mixing, which requires poly running time $O(m^7 n^8)$. Following

their work, Cover and Ordentlich (1996) developed universal procedures when side information¹ is taken into account as a finite number of values. Cross and Barron (2003) proposed a new universal portfolio strategy tracking the best in-hindsight wealth achievable within target classes of linearly parameterized portfolio sequences, which are more general than the standard CRP class and permit the portfolio to display a continuous form of dependence on past prices or other side information. Belentepe (2005) presented a statistical view of Cover's UP, showing that it is approximately equivalent to a constrained sequential portfolio optimization, which connects Cover's UP with traditional mean-variance portfolio theory.

Another famous strategy is *Exponential Gradient* (EG) strategy (Helmbold et al. 1997, 1996) for online portfolio selection using multiplicative updates. In general, EG strategy tries to maximize the expected logarithmic portfolio daily return (approximated using the last price relative), and minimize the deviation between next portfolio and last portfolio. The regret achieved by EG is $O(\sqrt{n \log m})$ with $O(mn)$ running time. While its regret is not as tight as Cover's UP, its linear time complexity is substantially less than the latter.

Recently, convex optimization has been applied to resolve the portfolio selection problem (Agarwal et al. 2006; Agarwal and Hazan 2005; Hazan 2006; Hazan et al. 2007). Examples include *Online Newton Step* (ONS) strategy (Agarwal et al. 2006), which aims to maximize the expected logarithmic cumulative wealth (approximated using historical price relatives) and to minimize the variation of the expected portfolio. ONS exploits the second order information of the log wealth function and applies it to the online scenario. It theoretically achieves a regret of $O(m \log n)$ which is the same as Cover's UP, and has running time complexity of $O(m^3 n)$. Following ONS, Hazan and Seshadhri (2009) recently proposed a new adaptive-regret approach with more decent theoretical results, which essentially is an ONS based strategy.

Another promising direction for portfolio selection is wealth maximization approach, which is based on the notion of approaching the Oracle as the optimal strategy. This idea was followed by Borodin et al. (2004) in their proposal of a non-universal portfolio strategy named *Anti-Correlation* (Anticor). Unlike the regret minimization approaches, Anticor strategy takes advantage of the statistical properties of financial market. The underlying motivation is to bet on the consistency of positive lagged cross-correlation and negative autocorrelation. It exploits the statistical information from the historical stock price relatives and adopts the classical mean reversion trading idea to transfer the wealth in the portfolio. Although it does not provide any theoretical guarantee, its empirical results (Borodin et al. 2004) showed that Anticor can outperform all existing strategies in most cases. Unlike the greedy algorithm by the Anticor strategy, Li, Hoi, Zhao and Gopalkrishnan (2011) very recently proposed *Confidence Weighted Mean Reversion* (CWMR) strategy to actively exploit the mean reversion property and the second order information of a portfolio, which produces better performance than Anticor.

In addition, Györfi et al. (2006) recently introduced a framework of *Nonparametric Kernel-based Moving Window* (B^K) learning strategies for portfolio selection based on non-parametric prediction techniques (Györfi and Schäfer 2003). Their algorithm first identifies a list of similar historical price relative sequences whose Euclidean distances with recent market windows are smaller than a threshold, then optimizes the portfolio with respect to the list of similar sequences. Under the same framework, Györfi et al. (2007) proposed another variant called *Nonparametric Kernel-based Semi-log-optimal* strategy, which is actually an approximation of the B^K strategy, mainly to improve the computational efficiency. Replacing log utility function by Markowitz-type utility function, Ottucsák and Vajda (2007)

¹ Side information includes interest rates, consumer confidence figures, etc.

proposed *Nonparametric Kernel-based Markowitz-type* strategy, which connects return and risk (or mean and variance) with the online portfolio selection strategy. Following the same framework as B^K strategy, *Nonparametric Nearest Neighbor learning* (B^{NN}) strategy proposed in Györfi et al. (2008) aims to search for the ℓ nearest neighbors in historical price relative sequences rather than search price relatives within a specified Euclidean ball. This method has been empirically shown to be a robust trading strategy. Along this direction, Li, Hoi and Gopalkrishnan (2011) recently proposed *Correlation-driven Nonparametric learning* (CORN) strategy to search for similar price relatives via correlation coefficient and considerably boosted the empirical performance of nonparametric learning approach.

Besides the main stream of learning to selection portfolio, another type of trading strategy is based on switching between various strategies, that is, maintaining a probability distribution among the strategies. Singer (1997) proposed *Switching Portfolios* (SP), which aims to deal with changing markets by taking into account the possibility that the market changes its behavior after each trading day. It switches among a set of basic investment strategies and assumes the a priori duration of using one basic strategy is geometrically distributed. Levina and Shafer (2008) proposed *Gaussian Random Walk* (GRW) strategy, which is a Markov switching strategy. GRW switches among the basic investment strategies as a Gaussian random walk in the simplex of portfolios.

Last, we note that our work is very different from another great body of existing work in literature (Kimoto et al. 1993; Tay and Cao 2001; Cao and Tay 2003; Tsang et al. 2004; Lu et al. 2009), which attempted to make financial time series forecasting and stock price predictions by applying machine learning techniques, such as neural networks (Kimoto et al. 1993), decision trees (Tsang et al. 2004), and support vector machines (SVM) (Tay and Cao 2001; Cao and Tay 2003; Lu et al. 2009), etc. The key difference between these work and ours is that their learning goal is to make explicit predictions of future prices/trends while our learning goal is to directly optimize portfolio without predicting prices explicitly.

3.4 Analysis of Existing Work

One popular trading idea in reality is *trend following* or *momentum* strategy, which assumes that historically better-performing stocks would still perform better than others in future. Some existing algorithms, such as EG and ONS, approximate the expected logarithmic daily return and logarithmic cumulative return respectively using historical price relatives. Though this idea is easy to understand and makes fortunes to many of the best traders and investors in the world, trend following is very hard to implement effectively. In addition, in the short-term, the stock price relatives may not follow previous trends as empirically evidenced by Jegadeesh (1990) and Lo and MacKinlay (1990).

Besides the trend following approach, another widely adopted approach in the learning community is *mean reversion* (Cover and Gluss 1986; Cover 1991; Borodin et al. 2004), which is also termed as *contrarian* approach. This approach stems from the CRP strategy (Cover and Gluss 1986), which rebalances to the initial portfolio every trading day. The idea behind this approach is that if one stock performs worse than others, it tends to perform better than others in the next trading day. As a result, the defining characteristic of a contrarian strategy is the purchase of securities that have performed poorly in the past and the sale of securities that have performed well, or quite simply, “*Sell the Winner, Buy the Loser*”. According to Lo and MacKinlay (1990), the effectiveness of mean reversion is a consequence of positive cross-autocovariances across securities. Among existing algorithms, CRP, UP, and Anticor adopt this trading idea. However, CRP and UP passively reverse to the mean,

while empirical evidence from Anticor algorithm (Borodin et al. 2004) shows that active reversion to the mean may better exploit the fluctuation of financial markets and is likely to obtain a much higher profit. On the other hand, although Anticor actively reverses to the mean, it is a heuristic method based on statistical correlations to transfer the wealth within the portfolio. In other words, it may not effectively exploit the mean reversion property.

In between, pattern matching based nonparametric learning algorithms (B^K and B^{NN} , etc.) can identify many market conditions including both mean reversion and trend following. However, when locating similar price relatives, the nonparametric learning approaches may locate both mean reversion and trend following price relatives, whose patterns are essentially opposite, thus weakening the maximization of the expected cumulative wealth.

In a word, both trend following and mean reversion can generate profit in the financial markets, if appropriately used. In the following, we will propose an active mean reversion based portfolio selection method. Though simple in update rules, it empirically outperforms the above existing portfolio selection strategies in most cases. The success of the proposed portfolio selection strategy indicates that it appropriately takes advantage of the mean reversion trading idea and generates significantly high profits in the back tests with real market data.

4 Passive Aggressive Mean Reversion Approach for Portfolio Selection

4.1 Intuition and Overview

The proposed approach is motivated by Constant Rebalanced Portfolios (Cover and Gluss 1986), which adopts the mean reversion trading idea. A simple but convincing example showing the mean reversion idea is illustrated in Table 1. Consider a fluctuating market with two stocks (A,B), and the stock price relative sequence is $(\frac{1}{2}, 2), (2, \frac{1}{2}), \dots$, where each stock is not going anywhere but actively moving within a range. Obviously, in a long-term period, market strategy cannot achieve any abnormal return from this sequence since the cumulative wealth of each stock remains the same after $2n$ trading days. However, Best CRP in hindsight can achieve a growth rate of $(\frac{5}{4})^n$ for a n -trading period. Now let us analyze the BCRP strategy on the stock price relative sequence to show the underlying mean reversion trading idea. Suppose the initial portfolio is $(\frac{1}{2}, \frac{1}{2})$ and at the end of the 1st trading day, the closing price adjusted wealth distribution becomes $(\frac{1}{5}, \frac{4}{5})$ with corresponding cumulative wealth increasing by a factor of $\frac{5}{4}$. At the beginning of the 2nd trading day, portfolio manager rebalances the portfolio to initial portfolio $(\frac{1}{2}, \frac{1}{2})$ by transferring the wealth from better-performing stock (B) to worse-performing stock (A) in the previous trading day. At the beginning of the 3rd trading day, the wealth transfer with the mean reversion trading idea continues. Although the market strategy gains nothing, BCRP can achieve a growth rate of $\frac{5}{4}$ per trading day using the mean reversion trading idea, which assumes that if one stock price performs worse, it tends to perform better in the subsequent trading day.

Another motivation of the proposed PAMR algorithm is inspired by the fact that in financial crisis, all stocks drop synchronously or certain stocks drop significantly. Under these situations, actively rebalancing may not be appropriate since it puts too much wealth on “mine” stocks, such as Bear Stearns during the recent financial crisis. To avoid the potential risk concerning such “mine” stocks, it is a good choice to stick to the previous portfolio, which constitutes the CRP strategy. Here, the reason to choose the passive CRP strategy is that identifying these “mine” stocks a priori is almost impossible, which are usually known in hindsight. Thus, to avoid suffering too much from such situations, PAMR alternates the

# day	Relative (A,B)	BCRP	BCRP Return	Wealth Proportion	Notes
1	$(\frac{1}{2}, 2)$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{1}{5}, \frac{4}{5})$	$B \rightarrow A$
2	$(2, \frac{1}{2})$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{4}{5}, \frac{1}{5})$	$A \rightarrow B$
3	$(\frac{1}{2}, 2)$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{1}{5}, \frac{4}{5})$	$B \rightarrow A$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

Table 1: Motivating example of CRP to show the mean reversion trading idea.

strategy between “aggressive” and “passive” reversion depending on the market conditions. The passive mean reversion strategy avoids the high risk of the aggressive approach that would put almost all wealth on these “mine” stocks when they drop significantly.

In this article, we propose a novel trading strategy named “Passive Aggressive Mean Reversion”, or PAMR for short. On the one hand, the underlying assumption of our approach is that better-performing stocks would perform worse than others in the next trading day. On the other hand, if the market drops too much, we would stop actively rebalancing the portfolio to avoid certain “mine” stocks and their associated risk. In order to exploit these intuitions, we suggest to adopt Passive Aggressive (PA) online learning (Crammer et al. 2006), which was originally proposed for classification tasks. Loosely speaking, the basic idea of PA for classification is that it *passively* keeps previous solution if loss is zero, while it *aggressively* updates the solution whenever the suffering loss is nonzero.

Let us now describe the basic idea of the proposed strategy in detail. Firstly, if the portfolio daily return is below a certain threshold, we will try to keep the previous portfolio such that it *passively* reverts to the mean to avoid the potential “mine” stocks. Secondly, if the portfolio daily return is above the threshold, we will actively rebalance the portfolio to ensure that the *expected* portfolio daily return is below the threshold in the belief that the stock price relatives will reverse in the next trading day. This sounds a bit counter-intuitive, but it is indeed reasonable, because if the stock price relative reverses, keeping the expected portfolio daily return below the threshold is able to maintain a high portfolio daily return in the next trading day. Here, the expected portfolio return is calculated with respect to the historical price relatives, for example, in our study, the last price relative, which is consistent with EG algorithm (Helmbold et al. 1997, 1996).

To further illustrate why aggressive reversion to the mean can be more effective than a passive one, let us continue the example in Table 1 that has a market going to nowhere but actively fluctuating. We show that in such markets, the proposed strategy is much more powerful than BCRP in hindsight, a passive mean reversion trading strategy. Table 2 compares the two trading strategies. As the motivating example shows, the growth rate of BCRP is $(\frac{5}{4})^n$ for a n -trading period, while at the same time, the growth rate of the proposed PAMR strategy is $\frac{5}{4} \times (\frac{3}{2})^{n-1}$ (The details of the calculation/algorithm will be presented later). We intuitively explain the success of PAMR below.

Assume the threshold for PAMR update is set to 1, that is, if portfolio daily return is below 1, we do nothing but keep the existing portfolio. Our strategy begins with a portfolio $(\frac{1}{2}, \frac{1}{2})$. For the 1st trading day, the return is $\frac{5}{4} > 1$. Then at the beginning of the 2nd trading day, we rebalance the portfolio to satisfy the condition that *approximate* portfolio daily return based on last price relatives is below the threshold 1, and the resulting portfolio is $(\frac{2}{3}, \frac{1}{3})$. Although it seems that we build a portfolio such that the approximate portfolio return is below the threshold, in practice, as the reversion to the mean suggests, we are maximizing the portfolio return in the next trading day. As we can observe, the return for the 2nd trading day is $\frac{3}{2} > 1$. Then following the same rule, we will rebalance the portfolio to

$(\frac{1}{3}, \frac{2}{3})$. As a result, in such a market, the growth rate of the proposed strategy is $\frac{5}{4} \times (\frac{3}{2})^{n-1}$ for a n -trading period, which is much more superior to that of BCRP, that is, $(\frac{5}{4})^n$.

# day	Relatives	BCRP	BCRP Return	PAMR	PAMR Return	Notes
1	$(\frac{1}{2}, 2)$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	rebalance to $(\frac{2}{3}, \frac{1}{3})$
2	$(2, \frac{1}{2})$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{2}{3}, \frac{1}{3})$	$\frac{5}{4}$	rebalance to $(\frac{1}{3}, \frac{2}{3})$
3	$(\frac{1}{2}, 2)$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{1}{3}, \frac{2}{3})$	$\frac{5}{4}$	rebalance to $(\frac{2}{3}, \frac{1}{3})$
4	$(2, \frac{1}{2})$	$(\frac{1}{2}, \frac{1}{2})$	$\frac{5}{4}$	$(\frac{2}{3}, \frac{1}{3})$	$\frac{5}{4}$	rebalance to $(\frac{1}{3}, \frac{2}{3})$
⋮	⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Table 2: Motivating example of comparison between BCRP and PAMR strategy.

4.2 Formulations

Now we shall formally devise the proposed Passive Aggressive Mean Reversion (PAMR) strategy for portfolio selection problem. The PAMR strategy is based on the mean reversion idea as described in Section 4.1, and is equipped with Passive Aggressive (PA) online learning technique (Crammer et al. 2006).

First of all, given a portfolio vector \mathbf{b} and a price relative vector \mathbf{x}_t , we define a ϵ -insensitive loss function for the t^{th} trading day as

$$\ell_\epsilon(\mathbf{b}; \mathbf{x}_t) = \begin{cases} 0 & \mathbf{b} \cdot \mathbf{x}_t \leq \epsilon \\ \mathbf{b} \cdot \mathbf{x}_t - \epsilon & \text{otherwise} \end{cases}, \quad (1)$$

where $\epsilon \geq 0$ is the sensitivity parameter which controls the mean reversion threshold. Since typically portfolio daily return fluctuates around 1, we often empirically choose $\epsilon \leq 1$ in order to buy worse performing stocks. The ϵ -insensitive loss is zero when return is less than the reversion threshold ϵ , and otherwise grows linearly with respect to the daily return. For conciseness, let us use ℓ_ϵ^t to denote $\ell_\epsilon(\mathbf{b}; \mathbf{x}_t)$, that is, the ϵ -insensitive loss of the t^{th} trading day. By defining this loss function, we can distinguish the two motivating cases described in Section 4.1.

In the following parts, we will formulate three variants of the proposed strategy, and will propose specific algorithms to solve them in the subsequent section. Recalling that \mathbf{b}_t denotes the portfolio vector on the t^{th} trading day, the first proposed method for Passive Aggressive Mean Reversion (PAMR) is formulated as the constrained optimization below: **Optimization Problem 1** (PAMR):

$$\mathbf{b}_{t+1} = \arg \min_{\mathbf{b} \in \Delta_m} \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 \quad \text{s. t.} \quad \ell_\epsilon(\mathbf{b}; \mathbf{x}_t) = 0. \quad (2)$$

The above formulation attempts to find an optimal portfolio by minimizing the deviation from last portfolio \mathbf{b}_t under the condition of satisfying the constraint of zero loss. On the one hand, the above approach *passively* keeps the last portfolio, that is, $\mathbf{b}_{t+1} = \mathbf{b}_t$ whenever $\ell_\epsilon^t = 0$ that means the portfolio daily return is below the threshold ϵ . On the other hand, whenever the loss is nonzero, it *aggressively* updates the solution by forcing it to strictly satisfy the constraint $\ell_\epsilon(\mathbf{b}_{t+1}; \mathbf{x}_t) = 0$. It is clear that this formulation is able to address the two motivations.

Although the above formulation is reasonable to address our concerns, it may have some undesirable properties in situations with noisy price relatives, which are common in real-world financial markets. For example, a noisy price relative appearing in some trending sequences may suddenly change the portfolio in a wrong direction due to the aggressive update. To avoid such problems, we propose two variants of PAMR that are able to trade off between aggressiveness and passiveness. The idea of formulating the two PAMR variants is similar to soft margin support vector machines by introducing some non-negative slack variables into optimization. Specifically, for the first variant, we modify the objective function by introducing a term that scales linearly with respect to ξ , which results in the following optimization:

Optimization Problem 2 (PAMR-1):

$$\mathbf{b}_{t+1} = \arg \min_{\mathbf{b} \in \Delta_m} \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 + C\xi \quad \text{s. t.} \quad \ell_\epsilon(\mathbf{b}; \mathbf{x}_t) \leq \xi \text{ and } \xi \geq 0, \quad (3)$$

where C is a positive parameter to control the influence of the slack variable term on the objective function. We refer to this parameter as the aggressiveness parameter similar to PA learning (Crammer et al. 2006) and call this variant ‘‘PAMR-1’’.

Instead of using a linear term of slack variable, in the second variant, we modify the objective function by introducing a slack variable term that scales quadratically with respect to ξ , which results in the following optimization problem:

Optimization Problem 3 (PAMR-2):

$$\mathbf{b}_{t+1} = \arg \min_{\mathbf{b} \in \Delta_m} \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 + C\xi^2 \quad \text{s. t.} \quad \ell_\epsilon(\mathbf{b}; \mathbf{x}_t) \leq \xi. \quad (4)$$

Note that in the above formulation we do not need to enforce the constraint $\xi \geq 0$ as ξ^2 is always non-negative. We refer to this variant as ‘‘PAMR-2’’.

4.3 Algorithms

We now derive the approximate solutions for the above three PAMR formulations using standard techniques from convex analysis (Boyd and Vandenberghe 2004), and present the proposed PAMR algorithms for portfolio selection task. Specifically, the following three propositions summarize the solutions to the PAMR methods.

Proposition 1 *The solution to the optimization problem 1 (PAMR) without considering the non-negativity constraint ($\mathbf{b} \succeq 0$) is expressed as:*

$$\mathbf{b} = \mathbf{b}_t - \tau_t (\mathbf{x}_t - \bar{x}_t \mathbf{1}), \quad (5)$$

where $\bar{x}_t = \frac{\mathbf{x}_t \cdot \mathbf{1}}{m}$ denotes the market return, and τ_t is computed as:

$$\tau_t = \max \left\{ 0, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2} \right\}. \quad (6)$$

Proof The proof can be found in Appendix A.

Proposition 2 *The solution to the optimization problem 2 (PAMR-1) without considering the non-negativity constraint ($\mathbf{b} \succeq 0$) is expressed as:*

$$\mathbf{b} = \mathbf{b}_t - \tau_t (\mathbf{x}_t - \bar{x}_t \mathbf{1}),$$

where $\bar{x}_t = \frac{\mathbf{x}_t \cdot \mathbf{1}}{m}$ denotes the market return, and τ_t is computed as:

$$\tau_t = \max \left\{ 0, \min \left\{ C, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2} \right\} \right\}. \quad (7)$$

Proof The proof can be found in Appendix B.

Proposition 3 *The solution to the optimization problem 3 (PAMR-2) without considering the non-negativity constraint ($\mathbf{b} \succeq 0$) is expressed as:*

$$\mathbf{b} = \mathbf{b}_t - \tau_t (\mathbf{x}_t - \bar{x}_t \mathbf{1}),$$

where $\bar{x}_t = \frac{\mathbf{x}_t \cdot \mathbf{1}}{m}$ denotes the market return, and τ_t is computed as:

$$\tau_t = \max \left\{ 0, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \frac{1}{2C}} \right\}. \quad (8)$$

Proof The proof can be found in Appendix C.

Figure 2 summarizes the details of the proposed PAMR algorithms. Firstly, with no historical information, the initial portfolio is set to uniform portfolio $\mathbf{b}_1 = (\frac{1}{m}, \dots, \frac{1}{m})$. At the beginning of t^{th} trading day, we rebalance according to the portfolio determined at the end of last trading day. At the end of t^{th} trading day, the market reveals a stock price relative vector, which represents the stock price movements. Since both the portfolio and the stock price relatives are already known, portfolio manager is able to measure the portfolio daily return $\mathbf{b}_t \cdot \mathbf{x}_t$ and the suffering loss $\ell_\epsilon(\mathbf{b}_t; \mathbf{x}_t)$ as defined in Eq. (1). Then, we calculate an optimal step size τ_t based on last portfolio and stock price relatives. Given the optimal step size τ_t , we can update the portfolio for next trading day. Finally, we perform a normalization step to obtain the final portfolio by projecting the updated portfolio into the simplex domain.

4.4 Analysis and Interpretation

To reflect the mean reversion trading idea, we are interested in analyzing the resulting update rules of the proposed PAMR algorithms, which mainly involve the portfolio \mathbf{b}_{t+1} and the step size τ_t . In particular, we want to examine how the update rules are related to return and risk – the two most important concerns in a portfolio selection task.

First of all, we analyze the resulting portfolio update rule in Eq. (5) for the three PAMR algorithms, that is, $\mathbf{b}_{t+1} = \mathbf{b}_t - \tau_t (\mathbf{x}_t - \bar{x}_t \mathbf{1})$. In the update rule, the step size τ_t is non-negative, and \bar{x}_t is the mean return or market return. For term $\mathbf{x}_t - \bar{x}_t \mathbf{1}$, we can see it represents stock abnormal returns with respect to the market on the t^{th} trading day. More precisely, we can interpret it as the directional vector for the weight transfer. The negative sign before the term indicates that the resulting update scheme is consistent with the motivation, that is, the weights shall be transferred from better performing stocks (with positive abnormal returns) to worse performing stocks (with negative abnormal returns) at the beginning of next day.

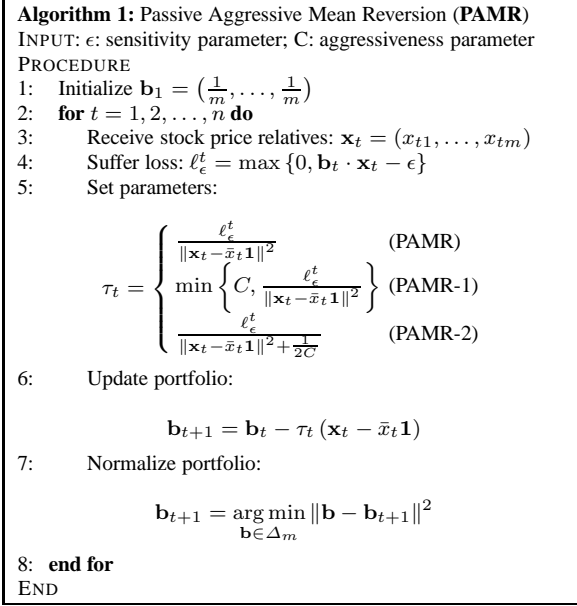


Fig. 2: The proposed Passive Aggressive Mean Reversion (PAMR) strategies.

Besides, another important update is the step size τ_t calculated as Eq. (6), Eq. (7), and Eq. (8), for three PAMR methods, respectively. The step size τ_t adaptively controls the weights to be transferred by taking effect on the directional vector. One interesting term in common for the three updates of τ_t is $\frac{\ell_\epsilon^t}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2}$. The numerator of the term equals to the t^{th} portfolio daily return minus the mean reversion threshold. Assuming other variables are constant, if the return is high (low), it leads to a large (small) value of τ_t , which would more (less) aggressively transfer the wealth from better performing stocks to worse performing stocks. The denominator is essentially the market quadratic variability, that is, the number of stocks times the market variance of the t^{th} trading day. In modern portfolio theory, variance of stock return is typically regarded as a volatility risk term for a portfolio (Markowitz 1952). As indicated by the denominator, if the risk is high (low), the step size τ_t would become small (large). As a result of small (large) step size, the weight transfer made by the update scheme will be weakened (strengthened), which is consistent with our intuition that prediction would be not accurate in drastically dropping markets, and we opt to make relatively less transfer in order to reduce risk. Moreover, PAMR-1 caps the step size by a constant C , while PAMR-2 decreases the step size by adding a constant $\frac{1}{2C}$ to its denominator. Both measures can prevent drastic weight transfer in case of noisy price relatives, which is consistent with their motivations.

From the above analysis on the updates of direction and step size, we can conclude that PAMR nicely balances between return and risk and clearly reflects the mean reversion trading idea. To the best of our knowledge, this important trade-off between return and risk has been considered by only one existing approach, that is, nonparametric kernel-based Markowitz-type strategy (Ottucsák and Vajda 2007). While the kernel-based Markowitz-type strategy trades off the return and risk with respect to similar historical price relatives, the proposed PAMR explicitly trades off the return and risk with respect to last price rela-

tives. This nice property distinguishes the proposed approach from most existing approaches that often cater to return, but ignore the risk concern, and are therefore undesirable according to modern portfolio theory (Markowitz 1952).

Now let us briefly analyze the time complexity of the proposed PAMR algorithms. From Figure 2, we can see that besides the normalization step, PAMR strategy takes $O(m)$ per trading day, where m denotes the number of assets. Moreover, the normalization or projection step (Step 7 in Figure 2) can be efficiently implemented (Michelot 1986; Duchi et al. 2008). In our implementation, we adopt the projection² according to Duchi et al. (2008), which takes linear time with respect to m . Thus, the total time complexity is $O(mn)$, where n is the total number of trading days. Such time complexity is the same as that of EG algorithm and is much superior to other existing methods. Linear time complexity enables the proposed algorithm to handle transactions in certain scenarios where low latency is of crucial importance, such as high frequency trading (Aldridge 2009).

4.5 Discussions

4.5.1 Discussion on Intuitions

Although the motivating example in Section 4.1 demonstrates the effectiveness of PAMR over BCRP strategy, PAMR may not always outperform BCRP. In general, PAMR is an online algorithm while BCRP is offline optimal for an i.i.d. market (see Cover and Thomas (1991), Theorem 15.3.1). Next, we discuss some possible situations where PAMR may fail to outperform BCRP.

Consider a special case where one stock crashes and the other explodes, e.g., a market sequence of two stocks as $(\frac{1}{2}, 2), (\frac{1}{2}, 2), \dots$. Assuming the same parameter settings as the motivating example, BCRP will increase at an exponential rate 2^n as it wholly invests in the 2nd asset, while PAMR will keep a fixed wealth on $\frac{5}{4}$ over the trading period. Obviously, in such a situation, PAMR performs much worse than BCRP does, i.e., PAMR produces a cumulative wealth of $\frac{5}{4}$ against 2^n achieved by BCRP over a n trading period. Though not shiny in such situations, PAMR still bounds its losses. Moreover, such a market, which violates the mean reversion assumption, is occasional, at least from the view point of our empirical studies.

4.5.2 Discussion on Loss Function

In our definition of loss function, that is, Eq. (1), we use the original portfolio expected return $\mathbf{b} \cdot \mathbf{x}_t$, while it is possible to use log utility (Latané 1959) on the return, that is, $\log(\mathbf{b} \cdot \mathbf{x}_t)$. With this log utility, the optimization problems Eq. (2), Eq. (3), and Eq. (4) are all non-convex and nonlinear, and thus difficult to solve. One way to solve these non-convex optimization problems is to use log's first-order Taylor expansion at last portfolio and ignore the higher order terms, that is, $\log(\mathbf{b} \cdot \mathbf{x}_t) \approx \log(\mathbf{b}_t \cdot \mathbf{x}_t) + \frac{\mathbf{x}_t}{\mathbf{b}_t \cdot \mathbf{x}_t} (\mathbf{b} - \mathbf{b}_t)$. After linear approximation, the optimization problems can be solved using the same techniques used in our derivation. However, such linear approximation of loss function may have some drawbacks. First of all, linear approximation yields an upper bound on regret in terms of a log utility loss function. There is no way to justify the goodness of the linear approximation.

² The precise matlab routine *ProjectOntoSimplex* can be found on <http://www.cs.berkeley.edu/~jduchi/projects/DuchiShSiCh08/>.

Moreover, if we use log utility, then the loss function is flat, then sharply rises and finally flattens out. While linear approximation is good in the two flat regimes, it is typically terrible at the point of non-differentiability and sub-par in the sharply rising region.

On the other hand, for the loss function in form of Eq. (1) without log utility or with linear approximation of log utility, the best possible regret in a minimax sense is at most $O(\sqrt{n})$ (Abernethy et al. 2009), while true log loss minimization algorithm can routinely achieve $O(\log n)$. However, although our loss function is non-differentiable and it would achieve a potential regret of $O(\sqrt{n})$, it is not a traditional loss function maximizing return (like traditional loss function, $-\log(\mathbf{b} \cdot \mathbf{x}_t)$), but only a tool to realize mean reversion. Thus the regret achieved using our loss function does not represent a regret about return, which may not be meaningful as traditional regret bound is.

Anyway, the potential worse bound may have unknown weaknesses, which may not be elicited by the following empirical evaluations. Though on our experiments PAMR works well, anyone who cares about its theoretical aspects should be notified about the possible worse bound.

4.5.3 Discussion on Formulation

Although our formulations mainly focus on the portfolio daily return without explicitly dealing with risk (e.g., volatility of daily returns), the final derived algorithms can be nicely interpreted as certain trade-offs between risk and return, as discussed in Section 4.4. Such interesting observation is further verified by our empirical evaluation in Section 5.4.2, which shows that the proposed PAMR algorithms achieve good risk-adjusted return in terms of two risk-related metrics (i.e., volatility risk and drawdown risk, respectively).

Similar to previous studies, we avoid incorporating transaction cost in the original formulations, which simplifies the formulations and clearly highlights PAMR's key ingredients. To further show the impact of transaction costs, it is not difficult to evaluate the effect of transaction costs, as shown in Section 5.2.2. In the following empirical study, we present results on both cases: with and without transaction costs. From the empirical results in Section 5.4.5, we find that in most markets, the proposed PAMR algorithms work well without or even with moderate transaction costs.

Besides, it is important to note that there are two key parameters in the proposed PAMR algorithm and its variants, viz., the sensitivity parameter ϵ and the aggressiveness parameter C . In practice, the choice of these parameters could affect the performance of the proposed algorithms. To achieve a good performance in a specific market, these parameters have to be finely tuned. We will thoroughly examine the effects of the two parameters on real-life datasets in Section 5.4.4, and make suggestions for the empirical selection of their values.

4.5.4 Discussion on PAMR Variants

In this section, we will show an example to illustrate different behaviors of the three update rules, viz., PAMR, PAMR-1, and PAMR-2. As discussed in Section 4.2, one objective for PAMR-1 and PAMR-2 is to prevent the portfolio being affected too much from noisy price relatives, which might drastically change the portfolio. Let us assume the environments and parameter settings as follows. Let the t^{th} price relative $\mathbf{x}_t = (1.00, 0.01)$, which represents the situations that the 2nd price relative is a noise, and the t^{th} portfolio $\mathbf{b}_t = (1, 0)$. Setting the parameters $\epsilon = 0.30$ and $C = 1.00$, let us calculate next portfolio \mathbf{b}_{t+1} . This market environment describes the situations where certain price relatives drop significantly, which is similar to some stocks during recent financial crisis. Without tuning, the original PAMR

algorithm would transfer a large proportion of wealth to the 2^{nd} asset in the next trading day. This can be verified by examining the portfolio calculated by PAMR, viz., PAMR calculates the update step size $\tau_t = 1.43$ and obtains the subsequent portfolio $\mathbf{b}_{t+1} = (0.29, 0.71)$. However, a natural choice of avoiding such noisy price relatives is to put less proportion of wealth to the second asset. Now, when calculating the next portfolios by PAMR-1 and PAMR-2, we obtain the update step size $\tau_t = 1.00$ and $\tau_t = 0.71$, respectively, which are smaller than the update step size of the original PAMR, that is, $\tau_t = 1.43$. Accordingly, we obtain the next portfolios $\mathbf{b}_{t+1} = (0.50, 0.50)$ and $\mathbf{b}_{t+1} = (0.65, 0.35)$ for PAMR-1 and PAMR-2, respectively. Clearly, PAMR-1 and PAMR-2 transfer less wealth to the 2^{nd} asset than the original PAMR does. Thus, PAMR-1 and PAMR-2 in general suffer relatively less from noisy price relatives, though we cannot completely avoid such suffering situation.

4.6 Mixture Algorithm

One theoretical result desired by existing online portfolio selection algorithms is *universal* property (Cover 1991). Since mean reversion trading idea is counter-intuitive (Borodin et al. 2004), we find it is hard to prove the universality of PAMR. Alternatively, we present a general mixture algorithm, which guarantees worst-case performance, not for PAMR itself but for the mixture algorithm.

Briefly speaking, the proposed mixture algorithm frames PAMR as one “expert” in a mixture-of-experts setting, while at least one universal algorithm serves as other “experts”. Then, the proposed mixture adopts no-regret expert learning (Cesa-Bianchi and Lugosi 2006) to bound the regret of the overall system with respect to the best of these experts. If the mixture algorithm contains at least one universal algorithm³, then the universality of the mixture algorithm can be straightforwardly proved according to Cesa-Bianchi and Lugosi (2006) (see example 10.3 and Theorem 10.3 for rigorous proofs). In our implementation, we adopt uniform buy and hold (BAH) mixture strategy, that is, we give equal proportion of portfolio wealth to each expert, let them run, and finally pool them again. We denote the BAH mixture algorithm as “MIX”. Other expert learning methods, such as exponential weighted, can also replace the buy and hold strategy, and they can also provide provable guarantees and get potentially stronger empirical performance. Though MIX seems trivial since it has a more involved mixing rule, one can make it nontrivial by extending the setting in a more general setting, such as the framework proposed by Akcoglu et al. (2002) and Das and Banerjee (2011). Obviously, such a mixture algorithm can be applied to any portfolio selection algorithm, either universal or not.

Though it is convenient to propose a mixture model consisting of PAMR such that the mixture model can achieve universality, PAMR’s universal consistency is still an open question and deserves further exploration.

5 Numerical Experiments

To examine the empirical efficacy of the proposed PAMR strategy, we conduct an extensive set of numerical experiments on a variety of real datasets. In our experiments, we adopt six real datasets, which were collected from several diverse financial markets. The performance metrics include cumulative wealth and risk-adjusted returns (volatility risk and drawdown

³ Such statement is also appeared in footnote 1 of Borodin et al. (2004).

risk). We also compare the proposed PAMR algorithms with all existing algorithms stated in the related work section.

5.1 Experimental Testbed on Real Data

In this study, we focus on historical daily prices in stock markets which are easy to obtain from public domains (such as Yahoo Finance and Google Finance), and thus publicly available to other researchers. Data from other types of market, such as high frequency intra-day quotes and Forex markets, are either too expensive or hard to obtain and process, and thus may reduce the experimental reproducibility. In general, we employ six real and diverse datasets from several types of financial markets⁴, which are summarized in Table 3.

Dataset	Market	Region	Time frame	# Trading days	# Assets
NYSE (O)	Stock	US	Jul. 3 rd 1962 - Dec. 31 st 1984	5651	36
NYSE (N)	Stock	US	Jan. 1 st 1985 - Jun. 30 th 2010	6431	23
TSE	Stock	CA	Jan. 4 th 1994 - Dec. 31 st 1998	1259	88
SP500	Stocks	US	Jan. 2 nd 1998 - Jan. 31 st 2003	1276	25
MSCI	Index	Global	Apr. 1 st 2006 - Mar. 31 st 2010	1043	24
DJIA	Stocks	US	Jan. 14 th 2001 - Jan. 14 th 2003	507	30

Table 3: Summary of the six real datasets in our numerical experiments.

The first one is NYSE dataset, one “standard” dataset pioneered by Cover (1991) and followed by several other researchers (Singer 1997; Helmbold et al. 1996; Borodin et al. 2004; Agarwal et al. 2006; Györfi et al. 2006, 2008). This dataset contains 5651 daily price relatives of 36 stocks⁵ in New York Stock Exchange (NYSE) for a 22-year period from Jul. 3rd 1962 to Dec. 31st 1984. We denote this dataset by “NYSE (O)” for short.

The second dataset is the extended version of the above NYSE dataset. For consistency, we collected the latest data in New York Stock Exchange (NYSE) from Jan. 1st 1985 to Jun. 30th 2010, which consists of 6431 trading days. We denote this new dataset as “NYSE (N)”⁶. It is worth noting that this new dataset consists of 23 stocks rather than the previous 36 stocks owing to amalgamations and bankruptcies. All self-collected price relatives are adjusted for splits and dividends, which is consistent with the previous “NYSE (O)” dataset.

The third dataset “TSE” is collected by Borodin et al. (2004), which consists of 88 stocks from Toronto Stock Exchange (TSE) containing price relatives of 1259 trading days, ranging from Jan. 4th 1994 to Dec. 31st 1998. The fourth dataset “SP500” is collected by Borodin et al. (2004), which consists of 25 stocks with the largest market capitalizations in the 500 SP500 components. It ranges from Jan. 2nd, 1998 to Jan. 31st 2003, containing 1276 trading days.

⁴ All the datasets and their compositions can be downloaded from <http://www.cais.ntu.edu.sg/~libin/portfolios>. Borodin et al. (2004)’s datasets can also be downloaded from <http://www.cs.technion.ac.il/~rani/portfolios/>.

⁵ According to Helmbold et al. (1996), the dataset was originally collected by Hal Stern. The stocks are mainly large cap stocks in NYSE, however, we do not know the criteria of choosing these 36 stocks.

⁶ The dataset before 2007 was collected by Gábor Gelencsér (<http://www.cs.bme.hu/~oti/portfolio>), we collected the remaining data from 2007 to 2010 via Yahoo Finance.

The fifth dataset is “MSCI”, a collection of global equity indices which are the constituents of MSCI World Index⁷. It contains 24 indices which represent the equity markets of 24 countries around the world, and consists of a total of 1043 trading days, ranging from Apr. 1st 2006 to Mar. 31st 2010. The final dataset is the “DJIA” dataset collected by Borodin et al. (2004), which consists of Dow Jones 30 composite stocks. DJIA contains 507 trading days, ranging from Jan. 14th 2001 to Jan. 14th 2003.

Besides the above six real market data, in the experiments, we also ran each dataset in their reverses (Borodin et al. 2004). For each dataset, we created a reversed dataset, which reverses the original order and inverts the price relatives. We denote these reverse datasets using a ‘-1’ superscript on the original dataset names. In nature, these reverse datasets are quite different from the original datasets, and we are interested in the behaviors of the proposed algorithm on these artificial datasets.

Unlike the previous studies, the above testbed covers much longer trading periods from 1962 to 2010 and much more diversified markets, which enables us to examine how the proposed PAMR strategy performs under different events and crises. For example, it covers several well-known events in the stock markets, such as dot-com bubble from 1995 to 2000 and subprime mortgage crisis from 2007 to 2009. The five stocks datasets are mainly chosen to test the capability of the proposed PAMR on regional stock markets, while the “MSCI” dataset aims to test PAMR’s capability on global indices, which may be potentially applicable to “Fund on Fund” (FOF)⁸. As a remark, although we numerically test the PAMR algorithm on stock markets, we note that the proposed strategy could be generally applied to any type of financial markets.

5.2 Experimental Setup and Metrics

Regarding the parameter settings, there are two key parameters in the proposed PAMR algorithms. One is the sensitivity parameter ϵ and the other is the aggressiveness parameter C . Roughly speaking, the best values for these parameters are often dataset dependent. In the experiments, we simply set these parameters empirically without tuning for each dataset separately. Specifically, for all datasets and experiments, we set the sensitivity parameter ϵ to 0.5 in the three algorithms, and set the aggressiveness parameter C to 500 in both PAMR-1 and PAMR-2, with which the cumulative wealth achieved tends to be stable for the proposed PAMR on most datasets. It is worth noting that these choices for parameters are not always the best. Our experiments on the parameter sensitivity in Section 5.4.4 show that the proposed PAMR algorithms are quite robust with respect to different parameter settings.

For the proposed mixture algorithm (MIX), we set the expert pool⁹ as initial uniform combination of PAMR, ONS, Anticor, and B^{NN}, and individual experts are set according to their respective studies.

We adopt the most common metric, *cumulative wealth*, to primarily compare different trading strategies. In addition to the cumulative wealth, we also adopt *annualized Sharpe Ratio* (SR) to compare the performance of different trading algorithms. In general, the higher the values of the cumulative wealth, and the annualized Sharpe Ratio, the better the performance of the compared algorithm. Besides, we also adopt *Maximum Drawdown* (MDD)

⁷ The constituents of MSCI World Index can be found from MSCI Barra (<http://www.msci Barra.com>), accessed on 28 May 2010.

⁸ It is worth noting that not every index is tradable through exchange traded funds (ETFs).

⁹ One can arbitrarily select experts, however, at least one universal algorithm should be included in order to guarantee the worst-case performance of the mixture algorithm.

and *Calmar Ratio* (CR) for analyzing the downside risk of the PAMR strategy. The lower the MDD value, the more preferable the trading algorithm concerning the downside risk. The higher the CR value, the more performance efficient the trading algorithm concerning the downside risk. The performance criteria are detailed in the following section.

5.2.1 Performance Criteria

One of the standard criteria to evaluate the performance of a strategy is *portfolio cumulative wealth* achieved by the strategy until the end of the whole trading period. In our study, we simply set the initial wealth $S_0 = 1$ and thus the notation S_n also denotes *portfolio cumulative return* at the end of the n^{th} trading day, which is the ratio of the portfolio cumulative wealth divided by the initial wealth. Another equivalent criterion is *annualized percentage yield* (APY) which takes the compounding effect into account, that is, $APY = \sqrt[y]{S_n} - 1$, where y is the number of years corresponding to n trading days. APY measures the average wealth increment that one strategy could achieve compounded in a year. Typically, the higher the value of portfolio cumulative wealth or annualized percentage yield, the more performance preferable the trading strategy is.

For some process-dependent investors (Moody et al. 1998), it is important to evaluate *risk and risk-adjusted return* of portfolios (Sharpe 1963, 1994). One common way to achieve this is to use *annualized standard deviation* of daily returns to measure the volatility risk and *annualized Sharpe Ratio* (SR) to evaluate the risk-adjusted return. For portfolio risk, we calculate the standard deviation of daily returns, and multiply by $\sqrt{252}$ (here 252 is the average number of annual trading days) to obtain annualized standard deviation. For risk-adjusted return, we calculate *annualized Sharpe Ratio* according to, $SR = \frac{APY - R_f}{\sigma_p}$, where R_f is the risk-free return (typically the return of Treasury bills, fixed at 4% in this work), and σ_p is the annualized standard deviation of daily returns. Basically, higher annualized Sharpe Ratios indicate better performance of a trading strategy concerning the volatility risk.

The investment community often analyzes *DrawDown* (DD) (Magdon-Ismael and Atiya 2004) to measure the decline from a historical peak in the cumulative wealth achieved by a financial trading strategy. Formally, let $\mathbf{S}(\cdot)$ denote the process of cumulative wealth achieved by a trading strategy, that is, $\{S_1, \dots, S_t, \dots, S_n\}$. The *DrawDown* at any time t , is defined as $DD(t) = \max[0, \max_{i \in (0, t)} S(i) - S(t)]$. The *Maximum DrawDown* for a horizon n , $MDD(n)$ is defined as, $MDD(n) = \max_{t \in (0, n)} [DD(t)]$, which is an excellent way to measure the downside risk of different strategies. Moreover, we also adopt *Calmar Ratio* (CR) to measure the return relative of the drawdown risk of a portfolio, calculated as $CR = \frac{APY}{MDD}$. Generally speaking, the smaller the Maximum DrawDown, the more downside risk tolerable the financial trading strategy. Higher Calmar Ratios indicate better performance of a trading strategy concerning the drawdown risk.

To test whether simple luck can generate the return of the proposed strategy, we can also conduct a statistical test to measure the probability of this situation, as is popularly done in the fund management industry (Grinold and Kahn 1999). First, we separate the portfolio daily returns into two components: one benchmark-related and the other non-benchmark-related by regressing the portfolio excess returns¹⁰ against the benchmark excess returns. Formally, $s_t - s_t(F) = \alpha + \beta(s_t(B) - s_t(F)) + \epsilon(t)$, where s_t stands for the portfolio daily returns, $s_t(B)$ denotes the daily returns of the benchmark (market index) and $s_t(F)$ is the daily returns of the risk-free assets (here we simply choose Treasury bill and set it

¹⁰ Excess return is daily return less risk-free return.

to 1.000156, or equivalently, annual interest of 4%). This regression estimates the portfolio's alpha (α), which indicates the performance of the investment after accounting for the involved risk. Then we conduct a statistical- t test to evaluate whether alpha is significantly different from zero, by using the t statistic $\frac{\alpha}{\text{SE}(\alpha)}$, where $\text{SE}(\alpha)$ is the standard error for the estimated alpha. Thus, by assuming the alpha is normally distributed, we can obtain the probability that the returns of the proposed strategy are generated by simple luck. Generally speaking, the smaller the probability, the higher confidence the trading strategy.

5.2.2 Practical Issues

While our model described in Section 2 is concise and not complicate to understand, it omits some practical issues in the portfolio management industry. We shall now relax some constraints in our model to address these issues.

In reality, an important and unavoidable issue is *transaction cost*. Generally, there are two ways to handle the transaction costs. The first, commonly adopted by learning to select portfolio strategies, is that the portfolio selection process doesn't take into account the transaction cost while the following rebalancing incurs transaction costs. The second is that the transaction cost is directly involved in the portfolio selection process (Györfi and Vajda 2008). In this work, we take the first way and adopt *proportional transaction cost* model proposed in Blum and Kalai (1999) and Borodin et al. (2004). To be specific, rebalancing the portfolio incurs a transaction cost on every buy and sell operation, based upon a transaction cost rate $\gamma \in (0, 1)$. At the beginning of the t^{th} trading day, the portfolio manager rebalances the portfolio from the previous closing price adjusted portfolio $\hat{\mathbf{b}}_{t-1}$ to a new portfolio \mathbf{b}_t , incurring a transaction cost of $\frac{\gamma}{2} \times \sum_i |b_{(t,i)} - \hat{b}_{(t-1,i)}|$, where the initial portfolio is set to $(0, \dots, 0)$. Thus, the cumulative wealth achieved by the end of the n^{th} trading day can be expressed as:

$$\mathbf{S}_n^{c(\gamma)} = \mathbf{S}_0 \prod_{t=1}^n \left[(\mathbf{b}_t \cdot \mathbf{x}_t) \times \left(1 - \frac{\gamma}{2} \times \sum_i |b_{(t,i)} - \hat{b}_{(t-1,i)}| \right) \right].$$

Another practical issue in portfolio selection is *margin buying*, which allows the portfolio managers to buy securities with cash borrowed from security brokers. Following previous studies (Cover 1991; Helmbold et al. 1996; Agarwal et al. 2006), we relax this constraint in the model and evaluate it empirically in Section 5.4.5. In this study, the margin setting is assumed to be 50% down and 50% loan, at an annual interest rate of 6%, so the interest rate of the borrowed money, c is set to 0.000238. Thus, for each security in the asset pool, a new asset named "Margin Component" is generated. Following the down and loan percentage, the price relative for the "Margin Component" of asset i would be $2 * x_{ti} - 1 - c$, where x_{ti} is the price relative of the i^{th} asset for the t^{th} trading day. In cases of $x_{ti} \leq \frac{1+c}{2}$, that is, certain stocks drop more than half, we simply set "Margin Component" to 0. By adding this "Margin Component", we magnify both the potential profit and loss of the trading strategy on the i^{th} asset.

5.3 Comparison Approaches

In our experiments, we implement the proposed PAMR strategy and its two variants, viz., PAMR-1 and PAMR-2. We compare them with a number of benchmarks and existing strate-

gies as described in Section 3. Below we summarize the list of compared algorithms, whose parameters are set according to the recommendations from their respective studies.

1. Market: Market strategy, that is, uniform Buy-And-Hold (BAH) strategy;
2. Best-Stock: Best stock in the market, which is a strategy in hindsight;
3. BCRP: Best Constant Rebalanced Portfolios strategy in hindsight;
4. UP: Cover’s Universal Portfolios implemented according to Kalai and Vempala (2002), where the parameters are set as $\delta_0 = 0.004$, $\delta = 0.005$, $m = 100$, and $S = 500$;
5. EG: Exponential Gradient (EG) algorithm with the best parameter $\eta = 0.05$ as suggested by Helmbold et al. (1996);
6. ONS: Online Newton Step (ONS) with the parameters suggested by Agarwal et al. (2006), that is, $\eta = 0$, $\beta = 1$, $\gamma = \frac{1}{8}$;
7. SP: Switching Portfolios with parameter $\gamma = \frac{1}{4}$ as suggested by Singer (1997);
8. GRW: Gaussian Random Walk strategy with parameter $\sigma = 0.00005$ recommended by Levina and Shafer (2008);
9. M0: Prediction based algorithm M0 with parameter $\beta = 0.5$ as suggested by Borodin et al. (2000);
10. Anticor: $\text{BAH}_{30}(\text{Anticor}(\text{Anticor}))$ as a variant of Anticor to smooth the performance, which achieves the best performance among the three solutions proposed by Borodin et al. (2004);
11. B^{K} : Nonparametric kernel-based moving window (B^{K}) strategy with $W = 5$, $L = 10$ and threshold $c = 1.0$ which has the best empirical performance according to Györfi et al. (2006);
12. B^{NN} : Nonparametric nearest neighbor based strategy (B^{NN}) with parameters $W = 5$, $L = 10$ and $p_\ell = 0.02 + 0.5 \frac{\ell-1}{L-1}$ as the authors suggested (Györfi et al. 2008).

5.4 Experimental Results

5.4.1 Experiment 1: Evaluation of Cumulative Wealth

We first compare the performance of the competing approaches based on their cumulative wealth. From the experimental results shown in Table 4, we can draw several observations below.

First of all, we observe that learning to select portfolio strategies generally perform better than three common benchmarks, which shows that it is promising to investigate learning algorithms for portfolio selection. Second, we find that although the cumulative wealth achieved by the regret minimization approaches (UP, EG and ONS) is higher than market strategy, their performance is significantly lower than that achieved by the wealth maximization approaches (Anticor, B^{K} and B^{NN}). This shows that to achieve better investment return, it is more powerful and promising to exploit the wealth maximization approaches for portfolio selection. Third, from the top two results indicated on each original dataset, it is clear that the proposed PAMR strategy (PAMR, PAMR-1, and PAMR-2) significantly outperforms most (except DJIA datasets) competitors including Anticor, B^{K} and B^{NN} , which are the state of the arts. The encouraging results in cumulative wealth validate the importance of exploiting the mean reversion property in the financial markets by an effective online learning strategy. On the other hand, though MIX beats the benchmarks on the DJIA dataset, PAMR algorithms perform bad on the DJIA dataset. This may be attributed to the reason that the motivating mean reversion does not exist in this dataset. This raises an im-

Methods	NYSE(O)	NYSE(N)	TSE	SP500	MSCI	DJIA
Market	14.50	18.06	1.61	1.34	0.91	0.76
Best-stock	54.14	83.51	6.28	3.78	1.50	1.19
BCRP	250.60	120.32	6.78	4.07	1.51	1.24
UP	26.68	31.49	1.60	1.62	0.92	0.81
EG	27.09	31.00	1.59	1.63	0.93	0.81
ONS	109.19	21.59	1.62	3.34	0.86	1.53
SP	27.08	31.55	1.60	1.65	0.93	0.81
GRW	27.73	30.45	1.61	1.64	0.93	0.81
M0	113.50	40.94	1.26	1.74	0.92	0.77
Anticor	2.41E+08	6.21E+06	39.36	5.89	3.22	2.29
B ^K	1.08E+09	4.64E+03	1.62	2.24	2.64	0.68
B ^{NN}	3.35E+11	6.80E+04	2.27	3.07	13.47	0.88
PAMR	5.14E+15	1.25E+06	264.86	5.09	15.23	0.68
PAMR-1	5.13E+15	1.26E+06	260.26	5.08	15.51	0.69
PAMR-2	4.88E+15	1.36E+06	249.95	5.00	16.87	0.71
MIX	1.28E+15	1.84E+06	78.58	4.36	8.16	1.35
Methods	NYSE(O) ⁻¹	NYSE(N) ⁻¹	TSE ⁻¹	SP500 ⁻¹	MSCI ⁻¹	DJIA ⁻¹
Market	0.12	1.27	1.67	0.88	1.26	1.44
Best-stock	0.33	24.59	37.65	1.65	3.45	2.77
BCRP	2.86	56.60	58.61	1.91	3.45	2.98
UP	0.23	0.3	1.18	1.10	1.26	1.54
EG	0.22	0.38	1.21	1.08	1.27	1.53
ONS	0.84	1.01	1.62	2.97	1.73	2.35
SP	0.23	0.35	1.19	1.10	1.27	1.54
GRW	0.24	0.34	1.18	1.09	1.26	1.55
M0	0.88	2.16	4.80	1.17	1.56	1.83
Anticor	1.38E+03	4.26E+04	7.24	9.64	6.31	4.58
B ^K	2.77E+07	162.74	8.81	1.01	4.47	1.43
B ^{NN}	4.60E+09	3.57E+04	66.09	1.89	30.06	1.85
PAMR	2.03E+04	3.07E+04	2.67	7.42	40.33	6.61
PAMR-1	2.02E+04	3.09E+04	2.68	7.43	39.82	6.62
PAMR-2	2.11E+04	3.21E+04	2.75	7.32	39.83	6.65
MIX	1.18E+09	2.70E+04	19.40	5.50	19.62	3.85

Table 4: Cumulative wealth achieved by various trading strategies on the six datasets and their reversed datasets. The top two best results in each dataset are highlighted in bold font.

portant question, “How to select the portfolio pool such that the motivating mean reversion exists on target portfolio?” Section 5.5.2 provides some discussions on this question.

Further examining the details, we find that the most impressive performance is achieved by PAMR on the standard NYSE (O) dataset, where its initial wealth grows by a factor of more than 5 quadrillion at the end of the 22-year period. We note that the main reason PAMR achieved such exceptional results is that it is powerful to exploit highly volatile price relatives. To verify this, we examine the detailed performance of PAMR in Table 4 by looking into individual stocks, and we find that it relies considerably on one single stock (“Kin Ark”) which has the highest volatility in terms of standard deviation. After removing this stock from the portfolio, we find that the cumulative wealth significantly reduces to 1.27E+08. We will investigate the volatility issue in more details by another experiment on dataset sensitivity in Section 5.4.3.

On the reverse datasets, though not performing as shiny as the original datasets, PAMR also performs well. Though some algorithms fail badly, in all cases, PAMR beats the benchmarks, including the market and BCRP strategies. In certain cases, it beats all competitors.

It is worth noting these reverse datasets are artificial datasets, which never exist in real markets. PAMR's performance on these datasets provides strong evidences that mean reversion does exist in even reverse market datasets and PAMR can successfully exploit it.

In addition to the final cumulative wealth, we are also interested in examining how the cumulative wealth changes over different trading periods. Figure 3 shows the trends of the cumulative wealth by the proposed PAMR algorithm and four algorithms (two benchmarks and two state-of-the-art algorithms). From the results, we can see that the proposed PAMR strategy consistently surpasses the benchmarks and the competing strategies over the entire trading period on most datasets (except DJIA datasets), which again validates the efficacy of the proposed technique.




Fig. 3: Trends of cumulative wealth achieved by various strategies during the entire trading period on the stock datasets.

Finally, to measure whether the excess return can be simply obtained by luck, we conduct a statistical t -test as described in Section 5.2.1. Table 5 shows the statistical results, which clearly show that the observed excess return is impossible to obtain by simple luck in most datasets. To be specific, the probabilities for achieving the excess returns by luck are almost 0 on datasets except DJIA. However, the statistics on DJIA dataset show that in this dataset, the assumption of mean reversion may not exist. Nevertheless, the results show that the PAMR strategy is a promising and reliable portfolio selection technique to achieve high return with high confidence.

5.4.2 Experiment 2: Evaluation of Risk and Risk-adjusted Return

We now evaluate the risk in terms of volatility risk and drawdown risk, and the risk-adjusted return in terms of annualized Sharpe ratio and Calmar ratio. Figure 4 shows the evaluation results on the six datasets. In addition to the proposed PAMR, we also plot two benchmarks (Market and BCRP) and two state-of-the-art algorithms (Anticor and B^{NN}) for comparison.

Statistics	NYSE (O)	NYSE (N)	TSE	SP500	MSCI	DJIA
Size	5651	6431	1259	1276	1043	507
Mean Excess Return (PAMR)	0.0069	0.0026	0.0054	0.0017	0.0029	-0.0003
Mean Excess Return (Market)	0.0005	0.0005	0.0004	0.0003	0.0000	-0.0004
Winning Ratio	0.5587	0.5175	0.5687	0.5337	0.5925	0.5187
α	0.0063	0.0021	0.0049	0.0013	0.0029	0.0002
β	1.2095	1.1241	1.4982	1.2375	1.1177	1.2393
t -statistics	15.7829	5.9979	3.9241	2.0020	6.1358	0.2195
p -value	0.0000	0.0000	0.0000	0.0227	0.0000	0.4132

Table 5: Statistical t -test of the performance of the PAMR on the stock datasets.

As shown in Figure 4, Figure 4a and 4b depict the volatility risk (standard deviation of daily returns) and the drawdown risk (maximum drawdown) on the six stock datasets. Figure 4c and Figure 4d compare the corresponding Sharpe ratio and Calmar ratio.

In previous cumulative wealth results, we find that PAMR achieved the highest cumulative return on most original datasets. Of course, high return is associated with high risk, which is commonly acceptable in finance, as no real financial instrument can guarantee a high return without risk. The volatility risk in Figure 4a shows that PAMR almost achieves the highest risk in terms of volatility risk. On the other hand, the drawdown risk in Figure 4b shows that PAMR achieves modest drawdown risk in most datasets. These results validate the above notion that high return is often associated with high risk.

To further evaluate the return and risk, we examine the risk-adjusted return in terms of annualized Sharpe ratio and Calmar ratio. The results shown in Figure 4c and Figure 4d clearly show that PAMR achieves excellent performance in most cases, except DJIA dataset. These encouraging results show that PAMR is able to reach a good trade-off between return and risk, even though we do not explicitly consider risk in our problem formulation.

5.4.3 Experiment 3: Dataset Sensitivity

As observed in Section 5.4.1, it is interesting that PAMR gained the excess return from the stock markets. In this section, we aim to examine how the dataset sensitivity affects the proposed PAMR strategy by evaluating performance on datasets of different volatilities.

To examine the effect of the dataset volatility, we create two datasets each consisting of 5 stocks, chosen from NYSE (N) dataset according to their volatility values. To be specific, we ranked the 23 stocks based on their daily volatility values measured by standard deviation of the logarithm of the price relatives (Hull 2008). Then we created two datasets of different volatility: NYSE (H) and NYSE (L), each consisting of 5 stocks of the highest and lowest volatility values, respectively. Table 6 shows the results achieved by various strategies on these two datasets.

From the results, we find that different strategies perform diversely on these two datasets. The regret minimization approaches (UP, EG and ONS), perform well regardless of the market volatilities as the theoretical universal property shows, while the wealth maximization approaches (Anticor, B^K and B^{NN}) and the proposed PAMR strategy achieved significantly higher cumulative wealth on NYSE (H), the high-volatility dataset. These results show that the volatility of datasets does considerably affect some algorithms, including the wealth maximization approaches and the proposed PAMR strategy. Specifically, we find that the proposed PAMR strategy could benefit much from a high-volatility dataset. For example, on the NYSE (L) dataset, the cumulative wealth achieved by PAMR algorithm is about 132,

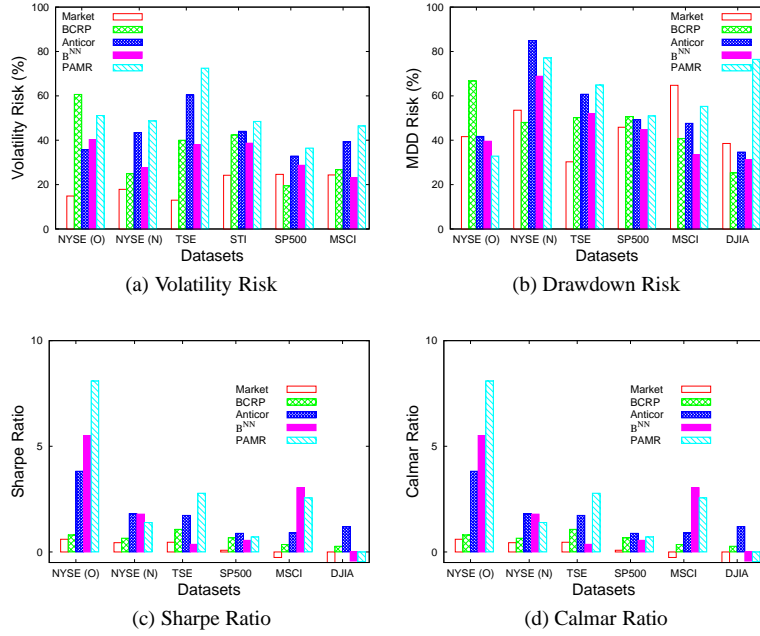


Fig. 4: Risk and risk-adjusted performance of various strategies on the six different datasets. In each diagram, the rightmost bars represent the results achieved by PAMR.

which is significantly boosted to $1.35\text{E}+05$ on the NYSE (H) dataset. To further examine which algorithm can benefit most from high-volatility dataset, we calculate the “H/L ratio” value, which is the ratio of cumulative wealth achieved on the high-volatility dataset over that achieved on the low-volatility dataset. From the ratios, we can observe that the PAMR strategy obtained the highest H/L ratio, indicating that PAMR can benefit most from the high-volatility dataset among all the competing methods.

Portfolio	NYSE (L)	NYSE (H)	H/L ratio
Market	24.69	9.15	0.37
Best-stock	43.87	17.46	0.40
BCRP	48.16	62.52	1.30
UP	32.89	26.12	0.79
EG	32.33	25.84	0.80
ONS	38.96	30.67	0.53
Anticor	$1.79\text{E}+03$	$1.43\text{E}+05$	79.89
B ^K	19.49	$3.82\text{E}+03$	196.00
B ^{NN}	180.85	$1.99\text{E}+05$	$1.10\text{E}+03$
PAMR	132.25	$1.39\text{E}+05$	$1.05+03$
PAMR-1	132.23	$1.66\text{E}+05$	$1.26\text{E}+03$
PAMR-2	142.53	$2.05\text{E}+05$	$1.44\text{E}+03$

Table 6: Cumulative wealth achieved by various strategies on portfolios of extreme volatilities. The “H/L ratio” column shows the ratio between the cumulative wealth achieved on the high-volatility dataset and that achieved on the low-volatility dataset.

5.4.4 Experiment 4: Parameter Sensitivity

We now evaluate how different choices of parameters affect the performance of the proposed PAMR strategy. All three PAMR algorithms require to set sensitivity parameter ϵ , while aggressiveness parameter C is needed for PAMR-1 and PAMR-2.

First, we examine the effect of the sensitivity parameter ϵ on the cumulative wealth achieved by PAMR. As ϵ becomes greater than 1, PAMR degrades to uniform CRP strategy and the wealth stabilizes at the wealth achieved by uniform CRP. Thus, we evaluate the effect of ϵ in the range of $[0, 1.5]$. Figure 5 shows the cumulative wealth achieved by PAMR with varying ϵ and those of the two benchmarks, that is, Market and BCRP strategies. Most results, besides DJIA dataset, show that the cumulative wealth achieved by PAMR grows as ϵ approaches 0, that is, the more sensitive the higher the wealth, which validates that the motivating mean reversion does exist on the stock markets. Moreover, in most cases, the cumulative wealth achieved by PAMR tends to stabilize as ϵ crosses certain dataset dependent thresholds. As stated before, we choose $\epsilon = 0.5$ in the experiments, with which the cumulative wealth becomes stabilized in most cases. We also note that on some datasets PAMR with $\epsilon = 0$ achieves the best. Though $\epsilon = 0$ means moving more weights to the worse performing stocks, it may not mean moving everything to the worst stock. On the one hand, the objectives in the formulations would prevent next portfolio far from last portfolio. On the other hand, PAMR-1 and PAMR-2 are designed to alleviate the huge changes. In a word, this experimental results clearly show that the proposed algorithm is robust with respect to the mean reversion sensitivity parameter. On the other side, for the failing case, DJIA, the mean reversion effect is different. As ϵ approaches 0, the cumulative wealth achieved by PAMR drops. This phenomena can be interpreted as that the motivating mean reversion does not exist in the DJIA dataset, at least in the sense of our motivation.

Second, we evaluate the other important parameter for both PAMR-1 and PAMR-2 algorithms, that is, aggressiveness parameter C . Figure 6 and Figure 7 show the effects on the cumulative wealth with varying sensitivity parameter ϵ from 0 to 1.5 and aggressiveness parameter C from 50 to 5000, on PAMR-1 and PAMR-2, respectively. Each heat map indicates the cumulative wealth achieved by PAMR with different C and ϵ combination. The indication bar on the right side of each heat map illustrates that each color represents a level of cumulative wealth achieved. It is clear that in most cases, except DJIA, we observe that as ϵ decreases and C increases, the cumulative wealth increases and then stabilizes as ϵ and C cross certain data-dependent thresholds. Moreover, we find C does not have a significant effect on the cumulative wealth achieved. We also find that the proposed PAMR algorithms are not so parameter sensitive, since a wide range of values correspond to the highest cumulative wealth. This again exhibits that the proposed PAMR strategy is robust with respect to its parameters. Similarly, the heat map on DJIA again shows that the mean reversion effect does not exist on the dataset, in the sense of our motivation.

5.4.5 Experiment 5: Evaluation of Practical Issues

For a real-world application, there are some important practical issues for portfolio selection, including the issues of transaction cost and margin buying. This experiment aims to examine how these practical issues affect the proposed PAMR strategy.

First, transaction cost is an important and unavoidable issue that should be addressed in practice. In our experiment, we adopt *proportional transaction cost* model stated in Section 5.2.2 to test the effect of the transaction cost on the proposed PAMR strategy. Figure 8 depicts the effect of proportional transaction cost when PAMR is applied on the six datasets,




Fig. 5: Parameter sensitivity of the cumulative wealth achieved by PAMR with respect to sensitivity parameter ϵ .




Fig. 6: Parameter sensitivity of the cumulative wealth achieved by PAMR-1 with respect to sensitivity parameter ϵ and aggressive parameter C .

where the transaction cost rate γ varies from 0 to 1%. We only present the results achieved by PAMR since the effect of its variants, that is, PAMR-1 and PAMR-2, is quite similar to that of PAMR. For comparison, we also plot the results achieved by two state-of-the-art strategies (Anticor and B^{NN}) and the cumulative wealth achieved by the two benchmarks (BCRP




Fig. 7: Parameter sensitivity of the cumulative wealth achieved by PAMR-2 with respect to sensitivity parameter ϵ and aggressive parameter C .

and Market). Since BCRP is the target strategy for regret minimization approaches (UP, EG and ONS) and for consistency, we do not plot the results achieved by these approaches.

From the results shown in the figure, we can observe that PAMR can withstand reasonable transaction cost rates. For example, with a transaction cost rate of 0.2%, PAMR can beat the BCRP strategy on the four datasets. The break-even transaction cost rates with respect to the market index ranges from 0.1% to 0.7% on the datasets, except DJIA. Since PAMR more actively reverses to the mean and thus results in more drastic portfolio changes, it surpasses Anticor with low or medium transaction costs while it underperforms Anticor with high transaction costs. On the other hand, it outperforms B^{NN} in most cases. Note that the transaction cost rate in real market is low¹¹. This experiment clearly shows the practical applicability of the proposed PAMR strategy when we take transaction cost into consideration.

Second, margin buying is another practical concern for a real-world portfolio selection task. In the following, we evaluate the performance of the approaches when margin buying is allowed with the model described in Section 5.2.2. Table 7 presents the cumulative wealth achieved by the competing approaches without/with margin loans on the six stock datasets. As we can observe, when margin buying is allowed, the profitability of PAMR increases, and in most cases, it achieves higher cumulative wealth than other competing approaches. These results clearly demonstrate that the proposed PAMR strategy can be extended to handle margin buying issue and benefit from margin buying, and thus has a better practical applicability.

¹¹ For example, without consideration taxes and bid-ask, Interactive Broker charges 0.005\$ per share traded. Considering the average price of Dow Jones Composite is around 50\$ (Accessed on June 2011), the percentage is about 0.01%.

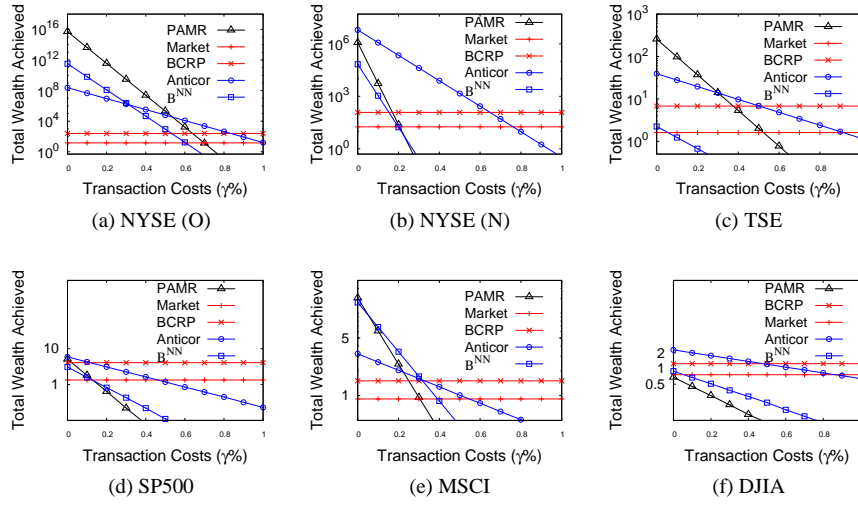


Fig. 8: Scalability of the cumulative wealth achieved by PAMR with respect to transaction cost rate (γ). The break-even transaction cost rates to the market index are about 0.7%, 0.4%, 0.1%, 0.3% and 0% on the six datasets, respectively.

Algorithm	NYSE(O)		NYSE(N)		TSE	
	No ML	with ML	No ML	with ML	No ML	with ML
Market	14.5	15.75	18.06	17.68	1.61	1.71
Best-stock	54.14	54.14	83.51	173.18	6.28	10.53
BCRP	250.6	3755.09	120.32	893.63	6.78	21.23
UP	27.41	62.99	31.49	57.03	1.60	1.69
EG	27.09	63.28	31.00	55.55	1.59	1.68
ONS	109.19	517.21	21.59	228.37	1.62	0.88
Anticor	2.41E+08	1.05E+15	6.21E+06	5.41E+09	39.36	18.69
B ^K	1.08E+09	6.29E+15	4.64E+03	3.72E+06	1.62	1.53
B ^{NN}	3.35E+11	3.17E+20	6.80E+04	5.58E+07	2.27	2.17
PAMR	5.14+15	5.57E+25	1.25E+06	1.12E+09	264.86	720.42
PAMR-1	5.13+15	5.55E+25	1.26E+06	1.13E+09	260.26	720.96
PAMR-2	4.88E+15	5.10E+25	1.36E+06	1.27E+09	249.95	711.90
Algorithm	SP500		MSCI		DJA	
	No ML	with ML	No ML	with ML	No ML	with ML
Market	1.34	1.03	0.91	0.69	0.76	0.59
Best-stock	3.78	3.78	1.50	1.50	1.19	1.19
BCRP	4.07	6.48	1.51	1.54	1.24	1.24
UP	1.62	1.75	0.92	0.71	0.81	0.66
EG	1.63	1.70	0.93	0.72	0.81	0.65
ONS	3.34	7.76	0.86	0.33	1.53	2.21
Anticor	5.89	10.73	3.22	3.40	2.29	2.89
B ^K	2.24	1.88	2.64	6.56	0.68	0.56
B ^{NN}	3.07	3.29	14.47	150.49	0.88	0.67
PAMR	5.09	15.91	15.23	68.83	0.68	0.84
PAMR-1	5.08	15.90	15.51	68.62	0.69	0.83
PAMR-2	5.00	16.26	16.87	70.08	0.71	0.86

Table 7: Cumulative wealth achieved by various strategies on the stock datasets with/without margin loans (ML). Top two achievements on each dataset are highlighted.

5.4.6 Experiment 6: Evaluation of Computational Time Cost

Our last experiment is to evaluate the computational time costs of different approaches, which is also an important issue in developing a practical online trading strategy. As stated in Section 4.3, the proposed PAMR algorithm enjoys linear time complexity per iteration, which is comparable to EG algorithm. Table 8 presents the computational time cost (in seconds) of the performance comparable approaches (Anticor, B^K and B^{NN}) on the six stock datasets. All the experiments were conducted on an Intel Core 2 Quad 2.66GHz processor with 4GB RAM, using Matlab 2009b on Windows XP.

From the results, we can clearly see that in all cases the proposed PAMR takes significant less computational time than the three performance comparable strategies. Even though the computational time in the back tests, especially per trading day, is small, it is important in certain scenarios such as high frequency trading (Aldridge 2009), where transactions may occur in a fraction of a second. Nevertheless, the results clearly demonstrate the computational efficiency of the proposed PAMR strategy, which is also an important concern for real-world large-scale applications.

Methods	NYSE (O)	NYSE (N)	TSE	SP500	MSCI	DJIA
Anticor	2.57E+03	1.93E+03	2.15E+03	387	306	175
B^K	7.89E+04	5.78E+04	6.35E+03	1.95E+03	2.60E+03	802
B^{NN}	4.93E+04	3.39E+04	1.32E+03	2.91E+03	2.55E+03	1.28E+03
PAMR	8	7	2	1	1	0.3

Table 8: Computational time cost on the real datasets (in seconds).

5.5 Discussions and Threads to Validity

5.5.1 Discussion on Model Assumption

Any statement about such encouraging empirical results would be incomplete without acknowledging the simplified assumptions made in Section 2. To recall, we had made several assumptions regarding transaction cost, market liquidity and market impact, which would affect the practical deployment of the proposed algorithm.

The first assumption is that no transaction cost exists. In Section 5.4.5 we have already examined the effect of varying transaction costs, and the results show that the proposed algorithm can withstand moderate transaction costs. Currently, with the wide-spread adoption of electronic communication networks (ECNs) and multilateral trading facilities (MTFs) on financial markets, various online trading brokers charge very small transaction cost rates, especially for large institutional investors. They also use a flat-rate¹², based on the volume threshold one reaches. Such measures can facilitate the portfolio managers to lower their transaction cost rates.

The second assumption is that the market is liquid and one can buy and sell any quantity at the quoted price. In practice, low market liquidity results in a large *bid-ask spread* – the gap between prices quoted for an immediate bid and an immediate ask. As a result, the

¹² For example, for US equities and options, E*Trade (<https://global.etrade.com/gl/home>, accessed on 16 March 2011.) charges only \$9.99 for \$50000+ or 30+ stocks per quarter.

execution of orders may incur a discrepancy between the prices sent by the algorithm and the prices actually executed. Moreover, stocks are often traded in multiples of *lot*, which is the standard trading unit containing certain number of stock shares. In this situation, the quantity of the stocks may not be arbitrary divisible. In the experiments, we have tried to minimize the effect of market liquidity by choosing the stocks that have large market capitalization, which usually have small bid-ask spreads and discrepancy, and thus have a high market liquidity.

The other assumption is that the portfolio strategy would have no impact on the market, that is, the stock market will not be affected by the trading algorithm. In practice, the impact can be neglected if the market capitalization of the portfolio is not too large. However, as the experimental results show, the portfolio wealth generated by PAMR increases astronomically, which would inevitably impact the market. One simple way to handle this issue is to scale down the portfolio, as done by many quantitative funds. Moreover, the development of algorithmic trading, which slices a big order into multiple smaller orders and schedules these orders to minimize the market impact, can significantly decrease the potential market impact of the proposed algorithm.

Here, we emphasize again that this study assumes a “perfect market”, which is consistent with previous studies in literature. It is important to note that even in such a perfect financial market, no algorithm has ever claimed such high performance, especially on the standard NYSE (O) dataset. Though it is common investment knowledge that past performance may not be reliable indicator of future performance, such high performance does provide us confidence that the proposed PAMR algorithm may work well in future unseen markets.

5.5.2 Discussion on PAMR Assumption

Though the proposed algorithm performs well on most datasets, we can not claim that PAMR can perform well on arbitrary portfolio pools. It is worth noting that PAMR relies on the assumption that mean reversion exists in a portfolio pool, that is, buying worse performing stocks is profitable. Preceding experiments seem to show that in most cases mean reversion does exist in the market. However, it is still possible that this assumption fails to exist in certain cases, especially when portfolio components are wrongly selected. PAMR’s performance on DJIA dataset indicates that mean reversion may not exist in its portfolio components. Though both based on mean reversion, PAMR and Anticor are formulated with different time periods of mean reversion, which may interpret why Anticor achieves a good performance on DJIA. Thus before investing in real market, it is of crucial importance to ensure that the motivating mean reversion does exist among the portfolio pools. In academic, mean reversion property in single stock has been extensively studied (Poterba and Summers 1988; Hillebrand 2003; Exley et al. 2004), one natural way is to calculate the sign of auto-correlation (Poterba and Summers 1988). On the contrary, the mean reversion property among a portfolio lacks academic attention. Compared with mean reversion in single stock, for a portfolio, not only the mean reversion of single stock matters, but rather the interaction among stocks matters.

On the other hand, the mixture algorithm, that is, MIX, performs well on the DJIA dataset, beating three benchmarks. As we discussed in Section 4.6, the mixture algorithm can provide a worst-case guarantee, which is lacked for the original PAMR algorithms. This can somehow solve the problem that PAMR itself does not have a worst-case guarantee. Moreover, it is worth noting that even with worst-case guarantee, some existing universal algorithms also perform poorly on the dataset.

Now let us briefly analyze the reason that PAMR failed on DJIA. To test whether mean reversion exists in the DJIA dataset, we propose a naïve trading strategy to test our motivating mean reversion in the dataset. The test strategy sets the weights proportional to differences between assets' returns and that of last best stock, that is, last best stock will be given zero weight, while the worst performing stock will be given a maximum weight. We are interested in whether this simple algorithm produces positive return among existing datasets. If it produces positive daily return, then the assumption that buying worse stocks may work well. Otherwise, our motivating assumption fails. The test is conducted on all six datasets. We calculated their arithmetic average daily returns and their standard deviations of daily returns. Since we are interested in absolute return, we compare their average values with 1. From the statistics in Table 9, we can find that the five successful datasets release average profit (> 1.0), while DJIA releases average loss (< 1.0). Thus, on DJIA dataset, it is expected to produce losses by purchasing worse performing stocks in the portfolio. Though expected daily loss is small, it would produce huge cumulative loss with a long trading period.

Statistics	NYSE (O)	NYSE (N)	TSE	SP500	MSCI	DJIA
Mean	1.000940	1.000835	1.000431	1.000544	1.000819	0.999848
Std. dev.	0.008920	0.013162	0.008562	0.014879	0.016423	0.016682

Table 9: Average daily return and standard deviation of the test strategy.

It is interesting to observe above results, however, we cannot claim that this method can definitely identify successful portfolio pools. Analyzing the mean reversion property in portfolio scenario and selecting portfolio components such that the portfolio satisfies mean reversion deserve further attention.

5.5.3 Discussion on Back Tests

Back tests in historical markets may suffer from “data-snooping bias” issue. One common “data-snooping bias” is dataset selection issue. On the one hand, we selected four datasets, that is, NYSE (O), TSE, SP500, and DJIA datasets, based on previous studies without consideration to the proposed approach. On the other hand, we developed the PAMR algorithm based solely on NYSE (O) dataset, while other five datasets (NYSE (N), TSE, SP500, MSCI and DJIA datasets) were obtained after the algorithm was fully developed. However, even we are cautious about the dataset selection issue, it may still appear in the experiments, especially for the datasets with relatively long history, that is, NYSE (O) and NYSE (N). The NYSE (O) dataset, pioneered by Cover (1991) and followed by other researchers, becomes one “standard” dataset in the learning community. Since it contains 36 large cap NYSE stocks that survived in hindsight for 22 years, thus it suffers from extreme survival bias. Nevertheless, it still has the merit to compare the performance among algorithms as done in all previous work. The NYSE (N) dataset, as a continuation of NYSE (O), contains 23 assets survived from previous 36 stocks for another 25 years. Therefore, it becomes even worse than the previous NYSE (O) dataset in terms of survival bias. In a word, even the experiment results on these datasets clearly show the effectiveness of the proposed PAMR algorithm, one can not make claims without noticing the deficiencies of these datasets.

Another common bias is asset selection issue. Four of the six datasets (NYSE (O), TSE, SP500, and DJIA) are collected by others, and to the best of our knowledge, their assets

are mainly the largest blue chip stocks in their respective markets. As a continuation of NYSE (O) dataset, we self-collected NYSE (N), which again contains several largest survival stocks in NYSE (O). The remaining dataset (MSCI) is chosen according to the world indices. In a word, we try to avoid the asset selection bias via arbitrarily choosing the representative stocks in their respective markets, which usually have large capitalization and thus high liquidity. Moreover, investing in these largest assets may reduce the market impact caused by the proposed portfolio strategy. Finally, following existing model assumption and experimental setting, we do not consider the assets of low quality, such as the bankrupt stocks and penny stocks. On the one hand, the bankrupt stock data is difficult to acquire, thus we cannot observe their behaviors and predict the behaviors of PAMR on datasets with bankrupt stocks. In reality, the bankruptcy situation happens rarely for the blue chip stocks as typically a bankrupt stock would be removed from the list of blue chip stocks before it actually goes bankrupt. On the other hand, the penny stocks lack the required liquidity to support the trading frequency in current research. Besides, one could also explore many practical strategies to exclude the low quality stocks from the asset pool at some early stage, such as some financial methods via either technical or fundamental analysis.

6 Conclusion

In this article, we proposed a novel portfolio selection strategy, “Passive Aggressive Mean Reversion” (PAMR). Motivated by the idea of mean reversion and passive aggressive learning, PAMR outperforms all benchmarks and various existing strategies on a number of real datasets from different markets. PAMR can also be easily extended to handle certain practical issues, e.g., transaction cost and margin buying. At the same time, PAMR executes in much less time than existing approaches, making it suitable for online applications. We also find that the update scheme of PAMR is based on the trade-off between the return and volatility risk, which is ignored by most existing learning strategies. This interesting property connects the PAMR strategy with modern portfolio theory, which may provide further explanation from the aspect of finance.

Although in most cases the proposed PAMR strategy achieves encouraging empirical results, it is still far from perfect for a real investment task, and may be improved in the following aspects. First of all, though universality may not be required in real investment, PAMR’s universality is still an open question. Second, none of existing algorithms considers the bankrupt assets, which may happen in real investment. It is thus interesting to study the behaviors and bankrupt assets and design strategies to exploit them. Besides, we note that PAMR sometimes fails when the mean reversion property does not exist in the portfolio components. Then it is crucial to propose efficient methods to test mean reversion. Finally, though PAMR handles the issue of transaction costs well, it is not formally addressed in our problem formulation. It would be interesting to incorporate the transaction cost issue when formulating the problem in order to improve the performance in case of high transaction costs and gain higher break-even ratios with respect to the market index.

Acknowledgements

This paper was fully supported by Singapore MOE Tier-1 Research Grant (RG67/07).

Appendix

Appendix A: Proof of **Proposition 1**

Proof First, if $\ell_\epsilon^t = 0$ then \mathbf{b}_t satisfies the constraint in Eq. (2) and is clearly the optimal solution.

Now let us focus on the case where $\ell_\epsilon^t \neq 0$. To solve the problem, we define the Lagrangian of the optimization problem in Eq. (2) to be,

$$\mathcal{L}(\mathbf{b}, \tau, \lambda) = \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 + \tau (\mathbf{x}_t \cdot \mathbf{b} - \epsilon) + \lambda (\mathbf{1} \cdot \mathbf{b} - 1), \quad (9)$$

where $\tau \geq 0$ is a Lagrange multiplier related to the loss function, λ is the Lagrange multiplier associated with the simplex constraint, and $\mathbf{1}$ denotes the column vector of m 1s. Note that the non-negativity of portfolio \mathbf{b} is not considered here since introducing this term causes too much complexity, and alternatively we project the resulting portfolio into a simplex to enforce the non-negativity constraint.

Setting the partial derivatives of \mathcal{L} with respect to the elements of \mathbf{b} to zero gives,

$$0 = \frac{\partial \mathcal{L}}{\partial \mathbf{b}} = (\mathbf{b} - \mathbf{b}_t) + \tau \mathbf{x}_t + \lambda \mathbf{1}.$$

Multiplying both sides with $\mathbf{1}^\top$, and $\mathbf{b}^\top \mathbf{1} = 1$, we can get $\lambda = -\frac{\tau}{m} \mathbf{x}_t \cdot \mathbf{1}$. Moreover, since $\bar{x}_t = \frac{\mathbf{x}_t \cdot \mathbf{1}}{m}$, where \bar{x}_t is the mean of the t^{th} asset price relative, or the market return, we can rewrite λ in the following form,

$$\lambda = -\tau \bar{x}_t. \quad (10)$$

And the solution for \mathcal{L} is,

$$\mathbf{b} = \mathbf{b}_t - \tau (\mathbf{x}_t - \bar{x}_t \mathbf{1}). \quad (11)$$

Plugging Eq. (10) and Eq. (11) to Eq. (9), we get,

$$\begin{aligned} \mathcal{L}(\tau) &= \frac{1}{2} \tau^2 \|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 - \tau^2 \mathbf{x}_t \cdot (\mathbf{x}_t - \bar{x}_t \mathbf{1}) + \tau (\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon) \\ &= -\frac{1}{2} \tau^2 \|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \tau (\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon). \end{aligned}$$

Note that in the derivation of the above formula, we used the following formula, that is,

$$\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 = \mathbf{x}_t \cdot \mathbf{x}_t - 2\bar{x}_t (\mathbf{x}_t \cdot \mathbf{1}) + \bar{x}_t^2 (\mathbf{1} \cdot \mathbf{1}) = \mathbf{x}_t \cdot \mathbf{x}_t - \bar{x}_t (\mathbf{x}_t \cdot \mathbf{1}) = \mathbf{x}_t \cdot (\mathbf{x}_t - \bar{x}_t \mathbf{1}).$$

Setting the derivative of $\mathcal{L}(\tau)$ with respect to τ to 0, we get,

$$0 = \frac{\partial \mathcal{L}}{\partial \tau} = -\tau \|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \mathbf{b}_t \cdot \mathbf{x}_t - \epsilon.$$

Then τ can be set to the following formula,

$$\tau = \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2}.$$

Since $\tau \geq 0$, we project τ to $[0, \infty)$, thus,

$$\tau = \max \left\{ 0, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2} \right\} = \frac{\ell_\epsilon^t}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2}.$$

Note that in case of zero market volatility, that is, $\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 = 0$, we just set $\tau = 0$. And we can state the update scheme for the case where $\ell_\epsilon^t = 0$ and the case where $\ell_\epsilon^t > 0$ by setting τ . Thus, we simplify the notation according to Eq. (1) and show the unified update scheme.

Appendix B: Proof of **Proposition 2**

Proof We derive the solution of PAMR-1 following the same procedure as the derivation of PAMR. Let us consider the situation when the loss is not 0 and we get the Lagrangian,

$$\mathcal{L}(\mathbf{b}, \xi, \tau, \mu, \lambda) = \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 + \tau (\mathbf{x}_t \cdot \mathbf{b} - \epsilon) + \xi (C - \tau - \mu) + \lambda (\mathbf{1} \cdot \mathbf{b} - 1).$$

Setting the partial derivatives of \mathcal{L} with respect to the elements of \mathbf{b} to zero gives,

$$0 = \frac{\partial \mathcal{L}}{\partial \mathbf{b}} = (\mathbf{b} - \mathbf{b}_t) + \tau \mathbf{x}_t + \lambda \mathbf{1},$$

Multiply both side with $\mathbf{1}^\top$, and $\mathbf{b}^\top \mathbf{1} = 1$, we can get, $\lambda = -\tau \frac{\mathbf{x}_t \cdot \mathbf{1}}{m} = -\tau \bar{x}_t$. And the approximation solution is

$$\mathbf{b} = \mathbf{b}_t - \tau (\mathbf{x}_t - \bar{x}_t \mathbf{1}).$$

Next, note that the minimum of the term $\xi(C - \tau - \mu)$ with respect to ξ is zero whenever $C - \tau - \mu = 0$. If $C - \tau - \mu \neq 0$ then the minimum can be made to approach $-\infty$. Since we need to maximize the dual we can rule out the latter case and pose the following constraint on the dual variables, $C - \tau - \mu = 0$. The KKT conditions confine μ to be non-negative so we conclude that $\tau \leq C$. We can project τ to the interval $[0, C]$ and get,

$$\tau = \max \left\{ 0, \min \left\{ C, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2} \right\} \right\} = \min \left\{ C, \frac{\ell_\epsilon^t}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2} \right\}.$$

Also note that we simplify the notation according to Eq. (1) and show the unified update scheme.

Appendix C: Proof of **Proposition 3**

Proof We derive the solution following the derivations of the PAMR and PAMR-1. Let us focus on the situation when the loss is not 0 and we can get the Lagrangian,

$$\mathcal{L}(\mathbf{b}, \xi, \tau, \mu, \lambda) = \frac{1}{2} \|\mathbf{b} - \mathbf{b}_t\|^2 + \tau (\mathbf{b} \cdot \mathbf{x}_t - \epsilon) + C\xi^2 - \tau\xi + \lambda (\mathbf{1} \cdot \mathbf{b} - 1).$$

Setting the partial derivatives of \mathcal{L} with respect to the elements of \mathbf{b} to zero gives,

$$0 = \frac{\partial \mathcal{L}}{\partial \mathbf{b}} = (\mathbf{b} - \mathbf{b}_t) + \tau \mathbf{x}_t + \lambda \mathbf{1},$$

Multiply both side with $\mathbf{1}^\top$, and $\mathbf{b}^\top \mathbf{1} = 1$, we can get, $\lambda = -\tau \frac{\mathbf{x}_t \cdot \mathbf{1}}{m} = -\tau \bar{x}$. And the approximation solution is,

$$\mathbf{b} = \mathbf{b}_t - \tau (\mathbf{x}_t - \bar{x} \mathbf{1}).$$

Setting the partial derivatives of \mathcal{L} with respect to the elements of ξ to zero gives,

$$0 = \frac{\partial \mathcal{L}}{\partial \xi} = 2C\xi - \tau \implies \xi = \frac{\tau}{2C}.$$

Expressing ξ as above and replacing \mathbf{b} , we rewrite the Lagrangian as

$$\tilde{\mathcal{L}}(\tau) = -\frac{\tau^2}{2} \left(\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \frac{1}{2C} \right) + \tau (\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon).$$

Take the derivative with respect to τ and set it to zero, we can get,

$$0 = \frac{\partial \tilde{\mathcal{L}}}{\partial \tau} = -\tau \left(\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \frac{1}{2C} \right) + (\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon).$$

Then we get the update scheme of τ , and project it to $[0, \infty)$

$$\tau = \max \left\{ 0, \frac{\mathbf{b}_t \cdot \mathbf{x}_t - \epsilon}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \frac{1}{2C}} \right\} = \frac{\ell_\epsilon^t}{\|\mathbf{x}_t - \bar{x}_t \mathbf{1}\|^2 + \frac{1}{2C}}.$$

References

- Abernethy, J., Agarwal, A., Barlett, P. L. and Rakhlin, A. (2009). A stochastic view of optimal regret through minimax duality, *Proceedings of Annual Conference on Learning Theory*.
- Agarwal, A. and Hazan, E. (2005). New algorithms for repeated play and universal portfolio management, *Technical report*, Princeton University.
- Agarwal, A., Hazan, E., Kale, S. and Schapire, R. E. (2006). Algorithms for portfolio management based on the newton method, *Proceedings of the International Conference on Machine Learning*, pp. 9–16.
- Akcoglu, K., Drineas, P. and Kao, M.-Y. (2002). Fast universalization of investment strategies with provably good relative returns, *Automata, Languages and Programming*, Vol. 2380, pp. 782–782.
- Aldridge, I. (2009). *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, Wiley, Hoboken, N.J.
- Belentepe, C. Y. (2005). *A Statistical View of Universal Portfolios*, PhD thesis, University of Pennsylvania.
- Blum, A. and Kalai, A. (1999). Universal portfolios with and without transaction costs, *Machine Learning* **35**(3): 193–205.
- Blum, A. and Mansour, Y. (2007). From external to internal regret, *Journal of Machine Learning Research* **8**: 1307–1324.
- Borodin, A. and El-Yaniv, R. (1998). *Online computation and competitive analysis*, Cambridge University Press, Cambridge.
- Borodin, A., El-Yaniv, R. and Gogan, V. (2000). On the competitive theory and practice of portfolio selection (extended abstract), *Proceedings of the Latin American Symposium on Theoretical Informatics*, pp. 173–196.
- Borodin, A., El-Yaniv, R. and Gogan, V. (2004). Can we learn to beat the best stock, *Journal of Artificial Intelligence Research* **21**: 579–594.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*, Cambridge University Press, New York.
- Breiman, L. (1961). Optimal gambling systems for favorable games, *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability* **1**: 65–78.
- Cao, L. J. and Tay, F. E. H. (2003). Support vector machine with adaptive parameters in financial time series forecasting, *IEEE Transactions on Neural Networks* **14**(6): 1506–1518.
- Cesa-Bianchi, N., Conconi, A. and Gentile, C. (2004). On the generalization ability of online learning algorithms, *IEEE Transactions on Information Theory* **50**(9): 2050–2057.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*, Cambridge University Press, New York.

- Cover, T. M. (1991). Universal portfolios, *Mathematical Finance* **1**(1): 1–29.
- Cover, T. M. and Gluss, D. H. (1986). Empirical bayes stock market portfolios, *Advances in applied mathematics* **7**(2): 170–181.
- Cover, T. M. and Ordentlich, E. (1996). Universal portfolios with side information, *IEEE Transactions on Information Theory* **42**(2): 348–363.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of information theory*, Wiley-Interscience, New York.
- Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S. and Singer, Y. (2006). Online passive-aggressive algorithms, *Journal of Machine Learning Research* **7**: 551–585.
- Crammer, K. and Singer, Y. (2003). Ultraconservative online algorithms for multiclass problems, *Journal of Machine Learning Research* **3**: 951–991.
- Cross, J. E. and Barron, A. R. (2003). Efficient universal portfolios for past-dependent target classes, *Mathematical Finance* **13**(2): 245–276.
- Das, P. and Banerjee, A. (2011). Meta optimization and its application to portfolio selection, *Proceedings of International Conference on Knowledge Discovery and Data Mining*.
- Duchi, J., Shalev-Shwartz, S., Singer, Y. and Chandra, T. (2008). Efficient projections onto the l_1 -ball for learning in high dimensions, *Proceedings of the International Conference on Machine Learning*, pp. 272–279.
- Exley, J., Mehta, S. and Smith, A. (2004). Mean reversion, *Technical report*, Faculty & Institute of Actuaries, Finance and Investment Conference, Brussels.
- Fink, M., Shalev-Shwartz, S., Singer, Y. and Ullman, S. (2006). Online multiclass learning by interclass hypothesis sharing, *Proceedings of the International Conference on Machine Learning*, pp. 313–320.
- Freund, Y. and Schapire, R. E. (1999). Large margin classification using the perceptron algorithm, *Machine Learning* **37**(3): 277–296.
- Gentile, C. (2001). A new approximate maximal margin classification algorithm, *Journal of Machine Learning Research* **2**: 213–242.
- Grinold, R. and Kahn, R. (1999). *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*, McGraw-Hill, New York.
- Györfi, L., Lugosi, G. and Udina, F. (2006). Nonparametric kernel-based sequential investment strategies, *Mathematical Finance* **16**(2): 337–357.
- Györfi, L. and Schäfer, D. (2003). Nonparametric prediction, in J. A. K. Suykens, G. Horváth, S. Basu, C. Micchelli and J. Vandevale (eds), *Advances in Learning Theory: Methods, Models and Applications*, IOS Press, Amsterdam, Netherlands, pp. 339–354.
- Györfi, L., Udina, F. and Walk, H. (2008). Nonparametric nearest neighbor based empirical portfolio selection strategies, *Statistics and Decisions* **26**(2): 145–157.
- Györfi, L., Urbán, A. and Vajda, I. (2007). Kernel-based semi-log-optimal empirical portfolio selection strategies, *International Journal of Theoretical and Applied Finance* **10**(3): 505–516.
- Györfi, L. and Vajda, I. (2008). Growth optimal investment with transaction costs, *Proceedings of the International Conference on Algorithmic Learning Theory*, pp. 108–122.
- Hazan, E. (2006). *Efficient Algorithms For Online Convex Optimization And Their Applications*, PhD thesis, Princeton University.
- Hazan, E., Agarwal, A. and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization, *Machine Learning* **69**(2-3): 169–192.
- Hazan, E. and Seshadhri, C. (2009). Efficient learning algorithms for changing environments, *Proceedings of the International Conference on Machine Learning*, pp. 393–400.
- Helmbold, D. P., Schapire, R. E., Singer, Y. and Warmuth, M. K. (1996). On-line portfolio selection using multiplicative updates, *Proceedings of the International Conference on*

- Machine Learning*, pp. 243–251.
- Helmbold, D. P., Schapire, R. E., Singer, Y. and Warmuth, M. K. (1997). A comparison of new and old algorithms for a mixture estimation problem, *Machine Learning* **27**(1): 97–119.
- Hillebrand, E. (2003). *Mean Reversion Models of Financial Markets*, PhD thesis, University of Bremen.
- Hull, J. C. (2008). *Options, Futures, and Other Derivatives*, Prentice Hall, Upper Saddle River, NJ.
- Jegadeesh, N. (1990). Evidence of predictable behavior of security returns, *Journal of Finance* **45**(3): 881–898.
- Kalai, A. and Vempala, S. (2002). Efficient algorithms for universal portfolios, *Journal of Machine Learning Research* **3**: 423–440.
- Kelly, J., J. (1956). A new interpretation of information rate, *Bell Systems Technical Journal* **35**: 917–926.
- Kimoto, T., Asakawa, K., Yoda, M. and Takeoka, M. (1993). Stock market prediction system with modular neural networks, *Neural Networks in Finance and Investing* pp. 343–357.
- Kivinen, J., Smola, A. J. and Williamson, R. C. (2001). Online learning with kernels, *Proceedings of the Annual Conference on Neural Information Processing Systems*.
- Latané, H. A. (1959). Criteria for choice among risky ventures, *The Journal of Political Economy* **67**(2): 144–155.
- Levina, T. and Shafer, G. (2008). Portfolio selection and online learning, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **16**(4): 437–473.
- Li, B., Hoi, S. C. H. and Gopalkrishnan, V. (2011). Corn: Correlation-driven nonparametric learning approach for portfolio selection, *ACM Transactions on Intelligent Systems and Technology* **2**(3).
- Li, B., Hoi, S. C. H., Zhao, P. and Gopalkrishnan, V. (2011). Confidence weighted mean reversion strategy for on-line portfolio selection, *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 434–442.
- Li, Y. and Long, P. M. (1999). The relaxed online maximum margin algorithm, *Proceedings of the Annual Conference on Neural Information Processing Systems*.
- Lo, A. W. and MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction?, *Review of Financial Studies* **3**(2): 175–205.
- Lu, C.-J., Lee, T.-S. and Chiu, C.-C. (2009). Financial time series forecasting using independent component analysis and support vector regression, *Decision Support Systems* **47**: 115–125.
- Magdon-Ismail, M. and Atiya, A. (2004). Maximum drawdown, *Risk Magazine* **10**: 99–102.
- Markowitz, H. (1952). Portfolio selection, *The Journal of Finance* **7**(1): 77–91.
- Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments*, Wiley, New York.
- Michelot, C. (1986). A finite algorithm for finding the projection of a point onto the canonical simplex of m , *Journal of Optimization Theory Applications* **50**: 195–200.
- Moody, J., Wu, L., Liao, Y. and Saffell, M. (1998). Performance functions and reinforcement learning for trading systems and portfolios, *Journal of Forecasting* **17**: 441–471.
- Ordentlich, E. and Cover, T. M. (1996). On-line portfolio selection, *Proceedings of the Annual Conference on Learning Theory*.
- Ottucsák, G. and Vajda, I. (2007). An asymptotic analysis of the mean-variance portfolio selection, *Statistics and Decisions* **25**: 63–88.
- Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices : Evidence and implications, *Journal of Financial Economics* **22**(1): 27–59.

-
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain, *Psychological Review* **65**: 386–407.
- Sharpe, W. F. (1963). A simplified model for portfolio analysis, *Management Science* **9**: 277–293.
- Sharpe, W. F. (1994). The sharpe ratio, *Journal of Portfolio Management* **21**(1): 49–58.
- Singer, Y. (1997). Switching portfolios, *International Journal of Neural Systems* **8**(4): 488–495.
- Stoltz, G. and Lugosi, G. (2005). Internal regret in on-line portfolio selection, *Machine Learning* **59**(1-2): 125–159.
- Tay, F. E. H. and Cao, L. (2001). Application of support vector machines in financial time series forecasting, *Omega* **29**(4): 309–317.
- Tsang, E., Yung, P. and Li, J. (2004). Eddie-automation, a decision support tool for financial forecasting, *Decision Support Systems* **37**: 559–565.
- Zhao, P., Hoi, S. C. H. and Jin, R. (2011). Double updating online learning, *J. Mach. Learn. Res.* **12**: 1587–1615.