

---

## On-Line Portfolio Selection Using Multiplicative Updates

---

**David P. Helmbold**

Computer & Information Sciences  
University of California  
Santa Cruz, CA 95064  
dph@cse.ucsc.edu

**Robert E. Schapire**

AT&T Laboratories  
600 Mountain Avenue  
Murray Hill, NJ 07974  
schapire@research.att.com

**Yoram Singer**

AT&T Laboratories  
600 Mountain Avenue  
Murray Hill, NJ 07974  
singer@research.att.com

**Manfred K. Warmuth**

Computer & Information Sciences  
University of California  
Santa Cruz, CA 95064  
manfred@cse.ucsc.edu

### Abstract

We present an on-line investment algorithm which achieves almost the same wealth as the best constant-rebalanced portfolio determined in hindsight from the actual market outcomes. The algorithm employs a multiplicative update rule derived using a framework introduced by Kivinen and Warmuth. Our algorithm is very simple to implement and requires only constant storage and computing time per stock in each trading period. We tested the performance of our algorithm on real stock data from the New York Stock Exchange accumulated during a 22-year period. On this data, our algorithm clearly outperforms the best single stock as well as Cover's universal portfolio selection algorithm. We also present results for the situation in which the investor has access to additional "side information."

### 1 INTRODUCTION

We present an on-line investment algorithm which achieves almost the same wealth as the best constant-rebalanced portfolio investment strategy. The algorithm employs a multiplicative update rule derived using a framework introduced by Kivinen and Warmuth [17]. Our algorithm is very simple to implement and its time and storage requirements grow linearly in the number of stocks. Experiments on real New York Stock Exchange data indicate that our algorithm outperforms Cover's [8] universal portfolio algorithm.

The following simple example demonstrates the power of constant-rebalanced portfolio strategies. Assume that two investments are available. The first is a risk-free, no-growth investment stock whose value never changes. The second investment is a hypothetical highly volatile stock. On even days, the value of this stock doubles and on odd days its value is halved. The relative returns of the first stock can be described by the sequence  $1, 1, 1, \dots$  and of the second by the sequence  $\frac{1}{2}, 2, \frac{1}{2}, 2, \dots$ . Neither investment alone can increase in value by more than a factor of 2, but a strategy combining the two investments can grow

exponentially. One such strategy splits the investor's total wealth evenly between the two investments, and maintains this even split at the end of each day. On odd days the relative wealth decreases by a factor of  $\frac{1}{2} \times 1 + \frac{1}{2} \times \frac{1}{2} = \frac{3}{4}$ . However, on even days the relative wealth grows by  $\frac{1}{2} \times 1 + \frac{1}{2} \times 2 = \frac{3}{2}$ . Thus, after two consecutive trading days the investor's wealth grows by a factor of  $\frac{3}{4} \times \frac{3}{2} = \frac{9}{8}$ . It takes only six days to double the wealth and over  $2n$  trading days the wealth grows by a factor of  $(\frac{9}{8})^n$ .

Investment strategies which maintain a fixed fraction of the total wealth in each of the underlying investments, like the one described above, are called *constant-rebalanced portfolio strategies*. Previously, Cover [8] described a portfolio-selection algorithm that provably performs "almost as well" as the *best* constant-rebalanced portfolio. In this paper, we describe a new algorithm with similar properties. Like the results for Cover's algorithm, this performance property is proven without making any statistical assumptions on the nature of the stock market.

The theoretical bound we prove on the performance of our algorithm relative to the best constant-rebalanced portfolio is not as strong as the bound proved by Cover and Ordentlich [10]. However, the time and space required for our algorithm is linear in the number of stocks whereas Cover's algorithm is exponential in the number of stocks. Moreover, we tested our algorithm experimentally on historical data from the New York Stock Exchange (NYSE) accumulated over a 22-year period, and found that our algorithm clearly outperforms the algorithm of Cover and Ordentlich.

Following Cover and Ordentlich [10], we also present results for the situation in which the investor has some finite "side information," such as the current interest rate. Side information may provide hints to the investor that one or a set of stocks are likely to outperform the other stocks in the portfolio. Moreover, the side information may be dependent on the past and future behavior of the market. At the beginning of each trading day, the side information is presented to the investor as a single scalar representing the "state" of the finite side information; the significance of this information must be learned by the investor.

## 2 PRELIMINARIES

Consider a portfolio containing  $N$  stocks. Each trading day,<sup>1</sup> the performance of the stocks can be described by a vector of *price relatives*, denoted by  $\mathbf{x} = (x_1, x_2, \dots, x_N)$  where  $x_i$  is the next day's opening price of the  $i$ th stock divided by its opening price on the current day. Thus the value of an investment in stock  $i$  increases (or falls) to  $x_i$  times its previous value from one morning to the next. A portfolio is defined by a weight vector  $\mathbf{w} = (w_1, w_2, \dots, w_N)$  such that  $w_i \geq 0$  and  $\sum_{i=1}^N w_i = 1$ . The  $i$ th entry of a portfolio  $\mathbf{w}$  is the proportion of the total portfolio value invested in the  $i$ th stock. Given a portfolio  $\mathbf{w}$  and the price relatives  $\mathbf{x}$ , investors using this portfolio increase (or decrease) their wealth from one morning to the next by a factor of  $\mathbf{w} \cdot \mathbf{x} = \sum_{i=1}^N w_i x_i$ .

### 2.1 ON-LINE PORTFOLIO SELECTION

In this paper, we are interested in on-line portfolio selection strategies. At the start of each day  $t$ , the portfolio selection strategy gets the previous price relatives of the stock market  $\mathbf{x}^1, \dots, \mathbf{x}^{t-1}$ . From this information, the strategy immediately selects its portfolio  $\mathbf{w}^t$  for the day. At the beginning of the next day (day  $t + 1$ ), the price relatives for day  $t$  are observed and the investor's wealth increases by a factor of  $\mathbf{w}^t \cdot \mathbf{x}^t$ .

Over time, a sequence of daily price relatives  $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^T$  is observed and a sequence of portfolios  $\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^T$  is selected. From the beginning of day 1 through the beginning of day  $T + 1$ , the wealth will have increased by a factor of

$$S_T(\{\mathbf{w}^t\}, \{\mathbf{x}^t\}) \stackrel{\text{def}}{=} \prod_{t=1}^T \mathbf{w}^t \cdot \mathbf{x}^t.$$

Since in a typical market the wealth grows exponentially fast, the formal analysis of our algorithm will be presented in terms of the normalized logarithm of the wealth achieved. We denote this normalized logarithm of the wealth by

$$LS_T(\{\mathbf{w}^t\}, \{\mathbf{x}^t\}) \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=1}^T \log(\mathbf{w}^t \cdot \mathbf{x}^t).$$

### 2.2 CONSTANT-REBALANCED PORTFOLIOS

With the benefit of hindsight, on each day one can invest all of one's wealth in the single best-performing stock for that day. It is certainly absurd to hope to perform as well as a prescient agent with this level of information about the future. Instead, in this paper, we compete against a more restricted class of investment strategies called constant-rebalanced portfolios. As noted in the introduction, a constant-rebalanced portfolio is rebalanced each day so that a fixed fraction of the wealth is held in each of the underlying investments. Therefore, a constant-rebalanced portfolio strategy employs the same investment vector  $\mathbf{w}$  on

each trading day and the resulting wealth and normalized logarithmic wealth after  $T$  trading days are

$$S_T(\mathbf{w}) \stackrel{\text{def}}{=} S_T(\mathbf{w}, \{\mathbf{x}^t\}) = \prod_{t=1}^T \mathbf{w} \cdot \mathbf{x}^t,$$

$$LS_T(\mathbf{w}) \stackrel{\text{def}}{=} LS_T(\mathbf{w}, \{\mathbf{x}^t\}) = \frac{1}{T} \sum_{t=1}^T \log(\mathbf{w} \cdot \mathbf{x}^t).$$

Note that such a strategy might require vast amounts of trading, since at the beginning of each day  $t$  the investment proportions are rebalanced back to the vector  $\mathbf{w}$ . In this paper we ignore commission costs (however, see the discussion in Section 6).

Given a sequence of daily price relatives  $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^T$  we can define, in retrospect, the best rebalanced portfolio vector which would have achieved the maximum wealth  $S_T$ , and hence also the maximum logarithmic wealth,  $LS_T$ . We denote this portfolio by  $\mathbf{w}^*$ . That is,

$$\mathbf{w}^* \stackrel{\text{def}}{=} \arg \max_{\mathbf{w}} S_T(\mathbf{w}) = \arg \max_{\mathbf{w}} LS_T(\mathbf{w}),$$

where the maximum is taken over all possible portfolio vectors (i.e., vectors in  $\mathbb{R}^N$  with non-negative components that sum to one). Iterative methods for finding this vector using the entire sequence of price relatives  $\mathbf{x}^1, \dots, \mathbf{x}^T$  are discussed in our earlier paper [13] which gives several updates for solving a general mixture estimation problem, including multiplicative updates like those described in this paper. We denote the wealth and the logarithmic wealth achieved using the optimal constant-rebalanced portfolio  $\mathbf{w}^*$  by  $S_T^*(\mathbf{x}^t)$  and  $LS_T^*(\mathbf{x}^t)$ , respectively. Whenever it is clear from the context, we will omit the dependency on the price relatives and simply denote the above by  $S^*$  and  $LS^*$ . Clearly,  $\mathbf{w}^*$  depends on the entire sequence of price relatives  $\{\mathbf{x}^t\}$  and may be dramatically different for different market behaviors.

Obviously, the optimal vector  $\mathbf{w}^*$  can only be computed after the entire sequence of price relatives is known (at which point, it is no longer of value). However, the algorithm described in this paper (as well as Cover's [8] algorithm) performs almost as well as  $\mathbf{w}^*$  while using only the previously observed history of price relatives to make each day's investment decision.

### 2.3 UNIVERSAL PORTFOLIOS

Cover [8] introduced the notion of universal portfolio. An on-line portfolio selection algorithm that results in the sequence  $\{\mathbf{w}^t\}$  is said to be *universal* (relative to the set of all constant-rebalanced portfolios) if

$$\lim_{T \rightarrow \infty} \max_{\{\mathbf{x}^t\}} [LS^*(\{\mathbf{x}^t\}) - LS(\{\mathbf{w}^t\}, \{\mathbf{x}^t\})] = 0.$$

That is, a universal portfolio selection algorithm exhibits the same asymptotic growth rate in normalized logarithmic wealth as the best rebalanced portfolio for *any* sequence of price relatives  $\{\mathbf{x}^t\}$ .

<sup>1</sup>The unit of time "day" was chosen arbitrarily; we could equally well use minutes, hours, weeks, etc. as the time between actions.

In Section 3 we adapt a framework developed for supervised learning and give a simple update rule that selects a new portfolio vector from the previous one. We prove that this algorithm is universal in Section 5.

## 2.4 SIDE INFORMATION

In reality, the investor might have more information than just the price relatives observed so far. Side information such as prevailing interest rates or consumer-confidence figures can indicate which stocks are likely to outperform the other stocks in the portfolio. Following Cover and Ordentlich [10], we denote the side information by an integer  $y$  from a finite set  $\{1, 2, \dots, K\}$ . Thus, the behavior of the market including the side information is now denoted by the sequence  $\{\mathbf{x}^t, y^t\}$ .

Following Cover and Ordentlich [10], we allow the constant-rebalanced portfolio to exploit the side information by expanding the single portfolio into a set of portfolios, one for each possible value of the side information. Thus, a constant-rebalanced portfolio with side information consists of the vectors  $\mathbf{w}(1), \mathbf{w}(2), \dots, \mathbf{w}(K)$  and uses portfolio vector  $\mathbf{w}(y^t)$  on day  $t$ . The wealth and normalized logarithmic wealth resulting from using a set of constant-rebalanced portfolios based on side information are,

$$S_T(\mathbf{w}(\cdot), \{\mathbf{x}^t, y^t\}) \stackrel{\text{def}}{=} \prod_{t=1}^T \mathbf{w}(y^t) \cdot \mathbf{x}^t, \\ LS_T(\mathbf{w}(\cdot), \{\mathbf{x}^t, y^t\}) \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=1}^T \log(\mathbf{w}(y^t) \cdot \mathbf{x}^t).$$

Just like the definition of the best constant-rebalanced portfolio, we define the best side information dependent portfolio set  $\mathbf{w}^*(\cdot)$  as the maximizer of  $S_T(\mathbf{w}(\cdot), \{\mathbf{x}^t, y^t\})$ . Note that the dimension of a side information dependent portfolio selection problem is  $K$  times larger than the single portfolio selection problem.

The sequence of side information  $\{y^t\}$  could be meaningless random noise, neither a function of the past market nor a predictor of future markets. On the other hand, it might be a perfect indicator of the best investment. Extending the two-investment example given in Section 1, we might have side information  $y = 1$  on odd days (when the volatile stock loses half its value) and  $y = 2$  on even days (when the volatile stock doubles). This side information can be exploited by the constant-rebalanced portfolio set  $\mathbf{w}(1) = (1, 0)$  and  $\mathbf{w}(2) = (0, 1)$  to double its wealth every other trading day. However, the only side information communicated to the investor (at the beginning of day  $t$ ) is the single value  $y^t$  with no further “explanations,” and the sequence  $\{y^t\}$  may or may not contain any useful information. Hence, the importance of each side information value must be learned from the performance of the market during previous trading days.

An on-line investment algorithm in this setting has access on day  $t$  both to the past history of price relatives

(as before) and to the past and current side information values  $y^1, \dots, y^t$ . The goal of the algorithm now is to invest in a manner competitive with  $S_T(\mathbf{w}^*(\cdot), \{\mathbf{x}^t, y^t\})$ , the wealth of the best constant-rebalanced portfolio with side information. One can easily define a notion of universality analogous to the definition given in Section 2.3.

As noticed by Cover and Ordentlich [10], the investor can partition the trading days based on the side information, and treat each partition separately. Exploiting the side information is therefore no more difficult than running  $K$  copies of our algorithm, one for each possible value of the side information. Since the logarithm of the wealth is additive, the logarithm of the wealth on the entire sequence with side information is just the sum of the logarithms of the wealths generated by the  $K$  copies of the algorithm.

## 2.5 RELATED WORK

Distributional methods are probably the most common approach to adaptive investment strategies for rebalanced portfolios. Kelly [16] assumed the existence of an underlying distribution of the price relatives and used Bayes decision theory to specify the next portfolio vector. Under various conditions, it was demonstrated (e.g. [5, 7, 6, 4, 2]) that with probability one the Bayes decision approach achieves the same growth rate of the wealth as the best rebalanced portfolio. In this approach, the price relative sequences can be drawn from one of a known set of possible distributions. This approach was used by Algoet [1] who considered the set of all ergodic and stationary distributions on infinite sequences, and estimated the underlying distribution in order to choose the next portfolio vector. Cover and Gluss [9] considered the restricted case where the set of price relatives is *finite* and gave an investment scheme with universal properties.

The most closely related previous results are by Cover [8] and Cover and Ordentlich [10]. They prove that certain investment strategies are universal without making any (or almost any) statistical assumptions on the nature of the stock market. Cover [8] proved that the wealth achieved by his universal portfolio algorithm is “almost as large” as the best constant-rebalanced portfolio. His analysis depends on a sensitivity matrix that characterizes the behavior of the market and he assumes that there is an upper bound on the price relatives and that they are bounded away from zero. Cover and Ordentlich [10] introduced the notion of side information and generalized Cover’s universal portfolio algorithm by using the Dirichlet( $1/2, \dots, 1/2$ ) and the Dirichlet( $1, \dots, 1$ ) priors over the set of all possible portfolio vectors.

Cover and Ordentlich’s investment strategies use an averaging method to pick their portfolio vectors. The portfolio vector used on day  $t$  is the weighted average over *all* feasible portfolio vectors (all  $N$  dimensional vectors with non-negative components that sum to 1), where the weight of each possible portfolio vector is determined by its per-

formance in the past. That is,

$$\mathbf{w}^t = \frac{\int \mathbf{w} S_{t-1}(\mathbf{w}) d\mu(\mathbf{w})}{\int S_{t-1}(\mathbf{w}) d\mu(\mathbf{w})},$$

where  $d\mu$  is one of the Dirichlet distributions mentioned above. Note that the portfolio vectors are weighted according to their past performance,  $S_{t-1}(\mathbf{w})$ , as well as the prior  $\mu(\mathbf{w})$ . Discrete approximation [8] or recursive series expansion [10] are used to evaluate the above integrals. In both cases, however, the time and space required for finding the new portfolio vector appears to grow exponentially in the number of stocks. While the bounds achieved by the generalized universal portfolio algorithm of Cover and Orntlich are stronger than ours, we show that on historical stock data our algorithm performs better while requiring time and space linear in the number of stocks.

### 3 MULTIPLICATIVE PORTFOLIO SELECTION ALGORITHMS

Our framework for updating a portfolio vector is analogous to the framework developed by Kivinen and Warmuth [17] for on-line regression. In this on-line framework the portfolio vector itself encapsulates the necessary information from the previous price relatives. Thus, at the start of day  $t$ , the algorithm computes its new portfolio vector  $\mathbf{w}^{t+1}$  as a function of  $\mathbf{w}^t$  and the just observed price relatives  $\mathbf{x}^t$ . In the linear regression setting analyzed by Kivinen and Warmuth, they show that good performance can be achieved by choosing a vector  $\mathbf{w}^{t+1}$  that is “close” to  $\mathbf{w}^t$ . We adapt their method and find a new vector  $\mathbf{w}^{t+1}$  that (approximately) maximizes the following function:

$$F(\mathbf{w}^{t+1}) = \eta \log(\mathbf{w}^{t+1} \cdot \mathbf{x}^t) - d(\mathbf{w}^{t+1}, \mathbf{w}^t), \quad (1)$$

where  $\eta > 0$  is some parameter called the learning rate and  $d$  is a distance measure that serves as a penalty term. This penalty term,  $-d(\mathbf{w}^{t+1}, \mathbf{w}^t)$ , tends to keep  $\mathbf{w}^{t+1}$  close to  $\mathbf{w}^t$ . The purpose of the first term is to maximize the logarithmic wealth if the current price relative  $\mathbf{x}^t$  is repeated. The learning rate  $\eta$  controls the relative importance between the two terms. Intuitively, if  $\mathbf{w}^t$  is far from the best constant-rebalanced portfolio  $\mathbf{w}^*$  then a small learning rate means that  $\mathbf{w}^{t+1}$  will move only slowly toward  $\mathbf{w}^*$ . On the other hand, if  $\mathbf{w}^t$  is already close to  $\mathbf{w}^*$  then a large learning rate may cause the algorithm to be misled by day-to-day fluctuations.

Different distance functions lead to different update rules. One of the main contributions of this line of work is the use of the relative entropy as a distance function for motivating updates:

$$D_{RE}(\mathbf{u}||\mathbf{v}) \stackrel{\text{def}}{=} \sum_{i=1}^N u_i \log \frac{u_i}{v_i}.$$

Many other on-line algorithms with multiplicative weight updates [18, 3, 17, 12] are also motivated by this distance

function and are thus rooted in the minimum relative entropy principle of Kullback [15, 11].

We also use a second-order Taylor approximation (at  $\mathbf{u} = \mathbf{v}$ ) of the relative entropy called the  $\chi^2$ -distance, since it leads to updates that are computationally cheaper:

$$D_{\chi^2}(\mathbf{u}||\mathbf{v}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{i=1}^N \frac{(u_i - v_i)^2}{v_i}.$$

Note that both distance functions are non-negative and zero if and only if  $\mathbf{u} = \mathbf{v}$ .

It is hard to maximize  $F$  since both terms depend non-linearly on  $\mathbf{w}^{t+1}$ . Instead, we replace the first term with its first-order Taylor polynomial around  $\mathbf{w}^{t+1} = \mathbf{w}^t$ . We also use a Lagrange multiplier to handle the constraint that the components of  $\mathbf{w}^{t+1}$  must sum to one. This leads us to maximize  $\hat{F}$  instead of  $F$ :

$$\begin{aligned} \hat{F}(\mathbf{w}^{t+1}, \gamma) &= \eta \left( \log(\mathbf{w}^t \cdot \mathbf{x}^t) + \frac{\mathbf{x}^t \cdot (\mathbf{w}^{t+1} - \mathbf{w}^t)}{\mathbf{w}^t \cdot \mathbf{x}^t} \right) \\ &\quad - d(\mathbf{w}^{t+1}, \mathbf{w}^t) + \gamma \left( \sum_{i=1}^N w_i^{t+1} - 1 \right). \end{aligned}$$

This is done by setting the  $N$  partial derivatives to zero (for  $1 \leq i \leq N$ ):

$$\begin{aligned} \frac{\partial \hat{F}(\mathbf{w}^{t+1}, \gamma)}{\partial w_i^{t+1}} &= \eta \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} - \frac{\partial d(\mathbf{w}^{t+1}, \mathbf{w}^t)}{\partial w_i^{t+1}} + \gamma = 0. \quad (2) \end{aligned}$$

If the relative entropy is used as the distance function then Equation (2) becomes

$$\eta \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} - (\log \frac{w_i^{t+1}}{w_i^t} + 1) + \gamma = 0$$

or

$$w_i^{t+1} = w_i^t \exp \left( \eta \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} + \gamma - 1 \right).$$

Enforcing the additional constraint  $\sum_{i=1}^N w_i^{t+1} = 1$  gives a portfolio update which we call the *exponentiated gradient* (EG( $\eta$ )) update:

$$w_i^{t+1} = \frac{w_i^t \exp(\eta x_i^t / \mathbf{w}^t \cdot \mathbf{x}^t)}{\sum_{j=1}^N w_j^t \exp(\eta x_j^t / \mathbf{w}^t \cdot \mathbf{x}^t)}. \quad (3)$$

A similar update for the case of linear regression was first given by Kivinen and Warmuth [17]. If we use the  $\chi^2$ -distance measure in place of the relative entropy then Equation (2) becomes

$$\eta \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} - \left( \frac{w_i^{t+1}}{w_i^t} - 1 \right) + \gamma = 0$$

or

$$w_i^{t+1} = \eta w_i^t \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} + w_i^t (\gamma + 1).$$

Now we sum the latter  $N$  equalities and use the constraints that  $\sum_{i=1}^N w_i^t = 1$  and  $\sum_{i=1}^N w_i^{t+1} = 1$  plus the fact that  $\sum_{i=1}^N w_i^t \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} = 1$ . This gives  $\gamma = -\eta$  and we obtain the update

$$w_i^{t+1} = w_i^t \left( \eta \left( \frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} - 1 \right) + 1 \right). \quad (4)$$

We call Equation (4) the  $\chi^2(\eta)$ -update. The  $\chi^2(\eta)$ -update can be viewed as a first order approximation of the  $\text{EG}(\eta)$ -update and the approximation is accurate when the exponents  $\eta(\frac{x_i^t}{\mathbf{w}^t \cdot \mathbf{x}^t} - 1)$  are small. The advantage of the  $\chi^2(\eta)$ -update is that it is computationally cheaper as it avoids the exponentiation. However, the  $\text{EG}(\eta)$ -update is easier to analyze. Our experiments with stock data indicate that these two update rules tend to approximate each other well yielding about the same wealth. In the next section we compare the performance of the  $\text{EG}(\eta)$ -update and  $\chi^2(\eta)$ -update with other on-line portfolio selection algorithms for different settings. The analysis of the updates is presented later in Section 5.

In addition to the updates, we also need to choose an initial portfolio vector  $\mathbf{w}^1$ . When no prior information is given, a reasonable choice would be to start with an equal weight assigned to each of the stocks in the portfolio, that is,  $\mathbf{w}^1 = (1/N, \dots, 1/N)$ . When side information is presented, we employ a set of portfolio vectors. We use the  $\text{EG}(\eta)$  or the  $\chi^2(\eta)$  updates to change the portfolio vector indexed by the side information. Hence, the problem of portfolio selection with side information simply reduces to a parallel selection of  $K$  different portfolios. If the side information is indeed informative, the set of portfolios will achieve larger wealth than a sequence of portfolio vectors resulting from the entire sequence. We demonstrate this in the experimental section that follows.

## 4 EXPERIMENTS WITH NYSE DATA

We tested our update rules on historical stock market data from the New York Stock Exchange accumulated over a 22-year period. For each experiment we restricted our attention to a subset of the stocks and compared the  $\text{EG}(\eta)$ -update and  $\chi^2(\eta)$ -update with each selected stock and with the best constant-rebalanced portfolio for the subset. We found the best constant-rebalanced portfolio by applying a batch maximum-likelihood mixture estimation procedure as described in our earlier paper [13]. After determining the best constant-rebalanced portfolio we then computed its performance on the price relative sequence. We also compared the performance of our update rules to that of Cover’s universal portfolio algorithm. We compared the results for all subsets of stocks considered by Cover [8] in his experiments.

Surprisingly, the wealth achieved by the universal portfolio strategy using the  $\text{Dirichlet}(1, \dots, 1)$  prior performed better than the  $\text{Dirichlet}(1/2, \dots, 1/2)$  prior, despite the better theoretical bounds proved for the  $\text{Dirichlet}(1/2, \dots, 1/2)$  prior [10]. Furthermore, the wealth

achieved by the  $\text{EG}(\eta)$ -update and  $\chi^2(\eta)$ -update was larger than the wealth achieved by the universal portfolio algorithm — again, despite the superior worst-case bounds proved for the universal portfolio algorithm. The difference in performance was largest when the portfolio is composed of volatile stocks.

The first example given by Cover is a portfolio based on Iroquois Brands Ltd. and Kin Ark Corp., two NYSE stocks chosen for their volatility. During the 22-year period ending in 1985, Iroquois increased in price by a factor of 8.92, while Kin Ark increased in price by a factor of 4.13. The best constant-rebalanced portfolio achieves a factor of 73.70 and the universal portfolio a wealth of 39.97. Using the  $\text{EG}(\eta)$ -update with  $\eta = 0.05$  yields a factor of 70.85, which is almost as good as the best constant-rebalanced portfolio. The results of the wealth achieved over the 22 years are depicted for this subset of stocks, as well as other subsets, in Figure 1. We got quantitatively similar results for the different portfolios considered by Cover [8]: Commercial Metals and Kin Ark, Commercial Metals and Meicco Corp., and IBM and Coca-Cola. The results are summarized in Table 1. However, when the stocks considered are not volatile and show a lockstep performance, as in the case of IBM and Coca-Cola, the wealth achieved by the universal portfolios and the  $\text{EG}(\eta)$ -update as well as the best constant-rebalanced portfolio barely outperform the individual stocks.

Following Cover [8], we also tested the case when we can invest in stock when margin loans are allowed. This case can be modeled by adding an additional “margin component” for each stock to the vector of price relatives. We assumed that all margin purchases were made 50% down and with a 50% loan. Thus the margin price relative for a stock  $i$  on day  $t$  is  $2x_i^t - 1 - c$  where  $c$  is the daily interest rate (recall that  $x_i^t$  is the price relative of stock  $i$ ). We tested this case with  $c = 0.000233$  which corresponds to an annual interest rate of 6%. The results are given in Table 2. It is clear from the table that the four-investment market containing the same two stocks plus “buying on the margin” results in a greater wealth. The efficiency of our update rules enables us to test our updates on more than two stocks. Moreover, as shown by the analysis, the wealth “lost” by our algorithms compared to the best constant-rebalanced portfolio scales like  $O(\sqrt{N})$ , whereas for the bounds on Cover’s universal portfolio algorithms the loss in wealth is linear in the number of investment options  $N$ . Thus our algorithm is more likely to tolerate additional investment options, such as buying on margin.

We found that the wealths achieved by the  $\text{EG}(\eta)$ -update and the  $\chi^2(\eta)$ -update were comparable. It turns out that the performance is not too sensitive to particular choices of a learning rate  $\eta$ . Learning rates from 0.01 to 0.15 all achieved great wealth, greater than the wealth achieved by the universal portfolio algorithm and in many cases comparable to the wealth achieved by the constant-rebalanced portfolio. The wealth achieved for different learning rates for the four-investment portfolio discussed

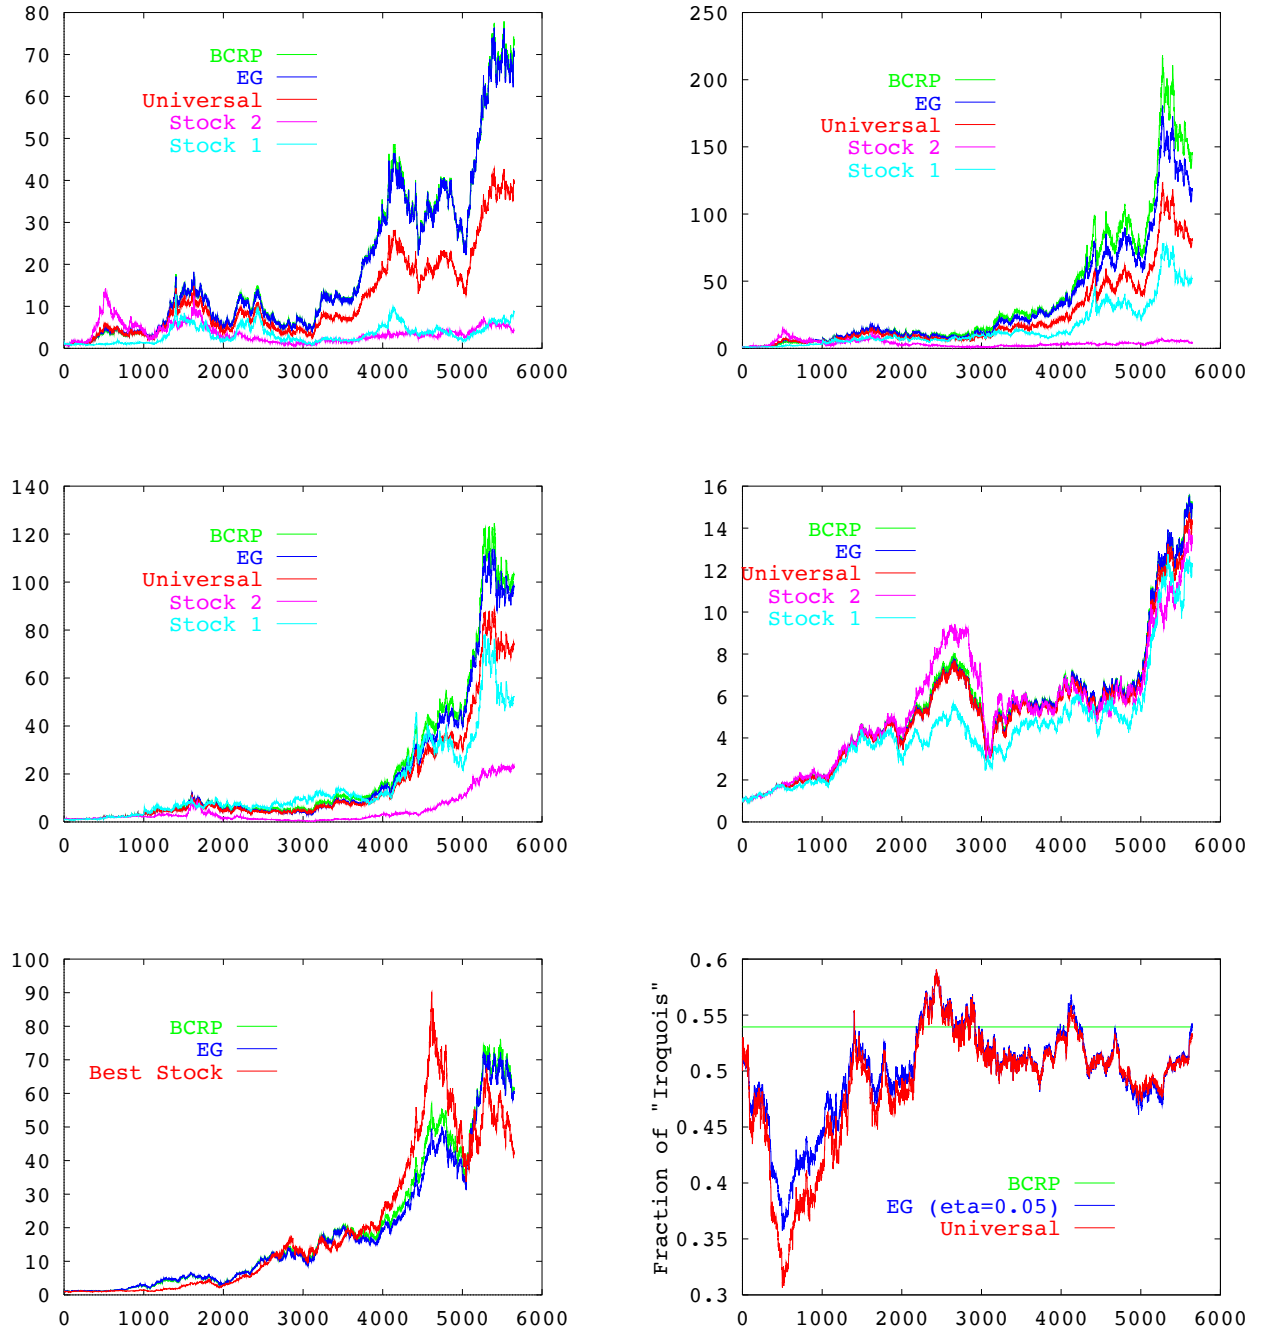


Figure 1: Comparison of wealths achieved by the best constant-rebalanced portfolio, the  $EG(\eta)$ -update, and the universal portfolio algorithm in certain markets. The markets consist of: Iroquois Brands and Kin Ark (top left), Commercial Metals and Kin Ark (top right), Commercial Metals and Meicco Corp. (middle left), IBM and Coca-Cola (middle right), and the three stocks Gulf, HP, and Schlum (bottom left). In all of these cases the wealth achieved by the  $EG(\eta)$ -update is close to the wealth of the best rebalanced portfolio and exceeds that achieved by the universal portfolio algorithm. It is interesting to note that after 4000 trading days in the three stock set a single stock (Schlum) achieves larger wealth than the best constant-rebalanced portfolio. However, the stock's value plummets around day 5000 and both the best constant-rebalanced portfolio and the  $EG(\eta)$ -update outperform Schlum over the 22 year period. At the bottom right we plot the fraction of the wealth invested in Iroquois by the strategies over time for the Iroquois/Kin Ark market.

Stocks	Best Stock	Const. Rebal. Port. (BCRP)	EG( $\eta$ ) ( $\eta = 0.05$ )	EG( $\eta$ ) / BCRP	Universal Portfolio	Universal / BCRP
Iroquois & Kin Ark	8.92	73.70	70.85	0.96	39.97	0.54
Comm. Metals & Kin Ark	52.02	144.00	117.15	0.81	80.54	0.56
Comm. Metals & Meicco Corp.	52.02	102.96	97.93	0.95	74.08	0.72
IBM & Coca-Cola	13.36	15.07	14.90	0.99	14.24	0.94
Gulf & HP & Morris & Schlum	54.14	69.94	65.64	0.94	—	—

Table 1: Comparison of the wealth achieved by the EG( $\eta$ )-update and the universal portfolio algorithm. For all the portfolios considered, we give the wealth achieved by the best constituent stock in the portfolio, the wealth achieved by the best constant-rebalanced portfolio (BCRP) computed in hindsight from the entire price relatives sequence, the wealth achieved by the EG( $\eta$ )-update rule, and the wealth achieved by the universal portfolio algorithm. We also give the proportion of the wealth achieved compared to the wealth of the BCRP. In all cases, the wealth achieved by EG( $\eta$ )-update is larger than the wealth of the universal portfolio algorithm. Moreover, in several cases the wealth of the EG( $\eta$ )-update is almost as good as the wealth of the best constant-rebalanced portfolio. We also tested portfolios consisting of more than two stocks and in most portfolios tested, the wealth achieved by the EG( $\eta$ )-update was almost as good as the wealth of the best constant-rebalanced portfolio.

	W/O Loans	Margin Loans
Commercial Metals	52.02	19.73
Kin Ark	4.13	0.00
BCRP	144.01	262.40
Universal portfolio	78.47	98.42
EG( $\eta$ ) ( $\eta = 0.05$ )	110.76	121.98

Table 2: Comparison of the portfolio selection algorithms when margin loans for each stock are available.

BCRP		Universal			
262.40		98.42			
EG ( $\eta =$ )					
0.01	0.02	0.05	0.10	0.15	0.20
119.87	121.49	121.94	113.33	103.06	91.31

Table 3: Comparison of the wealth achieved by the EG( $\eta$ )-update and for various learning rates and the universal portfolio algorithm, for the stocks considered in Table 2.

above (two stocks plus margin) are given in Table 3.

Finally, we tested the performance of our portfolio update algorithm when side information is presented. There are many possible forms of side information on which these algorithms might be tested. In our experiments, we chose to define the side information value to be the index of the stock with the best growth of wealth on the last 100 trading days — information that would certainly be available to an investor in a real trading situation. Thus, the possible set of values for the side information is  $1, \dots, K$  where  $K = N$ .

The results are summarized in Table 4. It is evident from the examples given in the table that using the side information (i.e., keeping  $N$  portfolio vectors) results in a significant improvement in the wealth achieved, even when using such simple and readily available side information. However, the gap between the best side information depen-

Stocks	With Side Information		
	BCRP	EG( $\eta$ )	Univ.
Iroq. & Kin Ark	307.9	99.4	86.6
Com. & Kin Ark	451.3	257.2	115.7
Com. & Meicco	436.2	186.1	110.9
IBM & Coke	118.5	89.9	21.1

Table 4: Comparison of the wealth achieved by the best constant-rebalanced portfolio (BCRP) and the EG( $\eta$ )-update when side information about the best stock in the last 100 trading days is presented. We have used the same learning rate ( $\eta = 0.05$ ) for both cases. (Refer to Table 1 for a comparison to the wealth achieved when no side information is presented.)

dent constant-rebalanced portfolio and the wealth achieved by the EG( $\eta$ )-update with side information is now much larger. One of the reasons is that we used the same learning rate regardless of the side information value. Large learning rates cause the update algorithms to quickly approach the best constant-rebalanced portfolio, but make it difficult for the algorithm to reach this portfolio exactly. On the other hand, small learning rates aid convergence to the best constant-rebalanced portfolio, but may cause the algorithm to spend a long time far away from this value. Therefore, when the side information splits the number of trading days unevenly, different learning rates for the different side information values may be required.

## 5 ANALYSIS

In this section, we analyze the logarithmic wealth obtained by the EG( $\eta$ ) portfolio update rule. We state worst-case bounds on the update which imply that the EG( $\eta$ ) update is almost as good as the best constant-rebalanced portfolio when certain assumptions hold on the relative volatility of the stocks in the portfolio. We also present a variant of EG( $\eta$ ) which requires no such assumptions. All proofs

have been omitted for lack of space.

Although the analysis is presented for a single portfolio vector, it can be generalized to the multiple vectors kept when side information is present by partitioning the trading days based on the side information and treating each partition separately.

Since  $x_i^t$  represents price relatives, we have that  $x_i^t \geq 0$  for all  $i$  and  $t$ . Furthermore, we assume that  $\max_i x_i^t = 1$  for all  $t$ . We can make this assumption without loss of generality since multiplying the price relatives  $\mathbf{x}^t$  by a constant  $c$  simply adds  $\log c$  to the logarithmic wealth, leaving the difference between the logarithmic wealth achieved by the  $\text{EG}(\eta)$ -update and the best achieved logarithmic wealth  $LS^*$  unchanged. Put another way, the assumed lower bound  $r$  on  $x_i^t$  used in Theorem 1 (below) can be viewed as a lower bound on the ratio of the worst to best price relatives for trading day  $t$ .

The following theorem characterizes a general property of the  $\text{EG}(\eta)$ -update.

**Theorem 1** *Let  $\mathbf{u} \in \mathbb{R}^N$  be a portfolio vector, and let  $\mathbf{x}^1, \dots, \mathbf{x}^T$  be a sequence of price relatives with  $x_i^t \geq r > 0$  for all  $i, t$  and  $\max_i x_i^t = 1$  for all  $t$ . For  $\eta > 0$  the logarithmic wealth due to the portfolio vectors produced by the  $\text{EG}(\eta)$ -update is bounded from below as follows:*

$$\sum_{t=1}^T \log(\mathbf{w}^t \cdot \mathbf{x}^t) \geq \sum_{t=1}^T \log(\mathbf{u} \cdot \mathbf{x}^t) - \frac{D_{RE}(\mathbf{u} \parallel \mathbf{w}^1)}{\eta} - \frac{\eta T}{8r^2}.$$

Furthermore, if  $\mathbf{w}^1$  is chosen to be the uniform proportion vector, and we set  $\eta = 2r\sqrt{2 \log N / T}$  then we have

$$\sum_{t=1}^T \log(\mathbf{w}^t \cdot \mathbf{x}^t) \geq \sum_{t=1}^T \log(\mathbf{u} \cdot \mathbf{x}^t) - \frac{\sqrt{2T \log N}}{2r}.$$

Since  $LS_T = \frac{1}{T} \sum_{t=1}^T \log(\mathbf{w}^t \cdot \mathbf{x}^t)$ , Theorem 1 immediately gives  $LS_T^* - LS_T \leq \sqrt{\ln N / (2r^2 T)}$  (under the conditions of Theorem 1). Thus, for an appropriate choice of  $\eta$ , when the number of days  $T$  becomes large, the difference between the logarithmic wealth achieved by  $\text{EG}(\eta)$  is guaranteed to converge to the logarithmic wealth of the best constant-rebalanced portfolio. However, Theorem 1 is not strong enough to show that  $\text{EG}(\eta)$  is a universal portfolio algorithm. This is because choosing the proper  $\eta$  requires knowledge of both the number of trading days and the ratio  $r$  in advance. We will deal with both of these difficulties, starting with the dependence of  $\eta$  on  $r$ .

When no lower bound  $r$  on  $x_i^t$  is known, we can use the following portfolio update algorithm which is parameterized by a real number  $\alpha \in [0, 1]$ . Let

$$\tilde{\mathbf{x}}^t = (1 - \alpha/N)\mathbf{x}^t + (\alpha/N)\mathbf{1}$$

where  $\mathbf{1}$  is the all 1's vector. As before, we maintain a portfolio vector  $\mathbf{w}^t$  which is updated using  $\tilde{\mathbf{x}}^t$  rather than  $\mathbf{x}^t$ :

$$w_i^{t+1} = \frac{w_i^t \exp(\eta \tilde{x}_i^t / \mathbf{w}^t \cdot \tilde{\mathbf{x}}^t)}{\sum_i w_i^t \exp(\eta \tilde{x}_i^t / \mathbf{w}^t \cdot \tilde{\mathbf{x}}^t)}.$$

Further, the portfolio vector that we invest with is also slightly modified. Specifically, the algorithm uses the portfolio vector  $\tilde{\mathbf{w}}^t = (1 - \alpha)\mathbf{w}^t + (\alpha/N)\mathbf{1}$  and so the logarithmic wealth achieved is  $\log(\tilde{\mathbf{w}}^t \cdot \mathbf{x}^t)$ .

We call this modified algorithm  $\widetilde{\text{EG}}(\alpha, \eta)$ .

**Theorem 2** *Let  $\mathbf{u} \in \mathbb{R}^N$  be a portfolio vector, and let  $\mathbf{x}^1, \dots, \mathbf{x}^T$  be a sequence of price relatives with  $x_i^t \geq 0$  for all  $i, t$  and  $\max_i x_i^t = 1$  for all  $t$ . For  $\alpha \in (0, 1/2]$  and  $\eta > 0$ , the logarithmic wealth due to the portfolio vectors produced by the  $\widetilde{\text{EG}}(\alpha, \eta)$ -update is bounded from below as follows:*

$$\sum_{t=1}^T \log(\tilde{\mathbf{w}}^t \cdot \mathbf{x}^t) \geq \sum_{t=1}^T \log(\mathbf{u} \cdot \mathbf{x}^t) - 2\alpha T - \frac{D_{RE}(\mathbf{u} \parallel \mathbf{w}^1)}{\eta} - \frac{\eta T}{8(\alpha/N)^2}.$$

Furthermore, if  $\mathbf{w}^1$  is chosen to be the uniform proportion vector,  $T \geq 2N^2 \log N$ , and we set  $\alpha = (N^2 \log N / (8T))^{1/4}$  and  $\eta = \sqrt{8\alpha^2 \log N / (N^2 T)}$  then we have

$$\sum_{t=1}^T \log(\tilde{\mathbf{w}}^t \cdot \mathbf{x}^t) \geq \sum_{t=1}^T \log(\mathbf{u} \cdot \mathbf{x}^t) - 2(2N^2 \log N)^{1/4} \cdot T^{3/4}. \quad (5)$$

Dividing inequality (5) of Theorem 2 by the number of trading days  $T$  shows that the logarithmic wealth achieved by the  $\widetilde{\text{EG}}(\alpha, \eta)$ -update converges to that of the best constant-rebalanced portfolio (for an appropriate choice of  $\eta$  dependent on  $T$ ). However, we still have the issue that the learning rate must be chosen in advance as a function of  $T$ . The following algorithm and corollary shows how a doubling trick can be used to obtain a universal portfolio algorithm.

The staged  $\widetilde{\text{EG}}(\alpha, \eta)$ -update runs in stages which are numbered from 0. The number of days in stage 0 is  $2N^2 \log N$ , and the number of days in each stage  $i > 0$  is  $2^i N^2 \log N$ . Thus if  $T > 2N^2 \log N$  is the total number of days, the last stage entered is numbered  $\lceil \log_2(\frac{T}{2N^2 \log N}) \rceil$ . At the start of each stage the portfolio vector is re-initialized to the uniform proportion vector and  $\alpha$  and  $\eta$  are set as in Theorem 2 using the number of days in the stage as the value for  $T$ .

**Corollary 3** *The staged  $\widetilde{\text{EG}}(\alpha, \eta)$ -update is a universal portfolio selection algorithm.*

In sum, the difference between the average daily logarithmic increase in wealth of the  $\widetilde{\text{EG}}(\alpha, \eta)$ -update and the best constant-rebalanced portfolio drops to zero at the rate  $O(((N^2 \log N)/T)^{1/4})$  for  $T \geq 2N^2 \log N$ . When the ratio between the best and worst stock on each day is bounded and relatively small (as can often be expected in practice),



the  $\text{EG}(\eta)$ -update can be used instead giving a convergence rate to zero of  $O(\sqrt{(\log N)/T})$ . In comparison, the bounds proved by Cover and Ordentlich [10] for their algorithm converge to zero at the rate  $O((N \log T)/T)$ . In terms of the number of trading days  $T$ , their bounds are much superior, especially compared to our bound for  $\widetilde{\text{EG}}(\alpha, \eta)$ . The only case in which our bounds have an advantage is when the number of stocks  $N$  included in the portfolio is relatively large and the market has bounded relative volatility so that  $\text{EG}(\eta)$  can be used. Despite the comparative inferiority of our theoretical bounds, in our experiments, we found that our algorithm did better, even though the number of trading days  $T$  was large (over 5,000) and the portfolios included only a few stocks.

## 6 DISCUSSION AND FUTURE RESEARCH

Although the experimental results presented in this paper are encouraging, we have ignored one important aspect of a real market — trading costs. Typically, there are two types of commissions imposed in a real market. In the first case, the investor needs to pay a percentage of the transaction to a broker. In this case, we can still write down a closed form expression for the wealth achieved at each time step while taking the trading costs into account. However, the wealth function we are trying to maximize becomes highly non-linear and it is hard to derive an update rule. The second type of commission is to pay a fixed amount per transaction, that is, per purchase or sale of a stock. Therefore, there might be days for which the wealth will be larger if no trading is performed, especially if the portfolio vector after the new trading day is close to the desired portfolio vector. We can define a semi-constant-rebalanced portfolio which is rebalanced only on a subset of the possible trading days. Now, in addition to the best constant-rebalanced portfolio, we need also to find the best subset of the sequence that results in the maximal wealth. We suspect that finding the best subset is computationally hard. Still, it is not clear whether finding a competitive approximation is hard as well.

This paper and most other work on investment strategies employ a tacit assumption that the market is stationary and seek a strategy that successfully competes against the best *single* constant-rebalanced portfolio. However, this assumption is far from being realistic. An interesting question is whether the techniques developed for tracking a drifting concept [3, 14] can be applied to the case of on-line portfolio selection in a changing market. Clearly, using a scheme that tracks a drifting portfolio vector might yield a more powerful investment strategy, both theoretically and empirically.

## Acknowledgments

Thanks to Tom Cover and Erik Ordentlich for providing us with the stock market data (originally generated by Hal Stern) used in

our experiments. We are also grateful to Erik Ordentlich for a careful reading and helpful comments on an earlier draft.

## References

- [1] P. H. Algoet. Universal schemes for prediction, gambling, and portfolio selection. *Annals of Probability*, 20:901–941, 1992.
- [2] P. H. Algoet and T. M. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *Annals of Probability*, 16(2):876–898, 1988.
- [3] P. Auer and M. K. Warmuth. Tracking the best disjunction. In *36th Annual Symposium on Foundations of Computer Science*, 1995.
- [4] A. Barron and T. M. Cover. A bound on the financial value of information. *IEEE Transactions on Information Theory*, 34:1097–1100, 1988.
- [5] R. Bell and T. M. Cover. Competitive optimality of logarithmic investment. *Mathematics of Operations Research*, 5:161–166, 1980.
- [6] R. Bell and T. M. Cover. Game-theoretic optimal portfolios. *Management Science*, 34:724–733, 1988.
- [7] T. M. Cover. An algorithm for maximizing expected log investment return. *IEEE Transactions on Information Theory*, 30:369–373, 1984.
- [8] T. M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.
- [9] T. M. Cover and D. Gluss. Empirical Bayes stock market portfolios. *Advances in Applied Mathematics*, 7, 1986.
- [10] T. M. Cover and E. Ordentlich. Universal portfolios with side information. Unpublished manuscript, 1995.
- [11] D. Haussler, N. Littlestone, and M. K. Warmuth. Predicting  $\{0, 1\}$ -functions on randomly drawn points. *Information and Computation*, 115(2):284–293, 1994.
- [12] D. P. Helmbold, J. Kivinen, and M. K. Warmuth. Worst-case loss bounds for sigmoided linear neurons. In *Advances in Neural Information Processing Systems* 8, 1996.
- [13] D. P. Helmbold, R. E. Schapire, Y. Singer, and M. K. Warmuth. A comparison of new and old algorithms for a mixture estimation problem. In *Proceedings of the Eighth Annual Workshop on Computational Learning Theory*, pages 69–78, 1995.
- [14] M. Herbster and M. K. Warmuth. Tracking the best expert. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 286–294, 1995.
- [15] G. Jumarie. *Relative information*. Springer-Verlag, 1990.
- [16] J. L. Kelly. A new interpretation of information rate. *Bell Systems Technical Journal*, 35:917–926, 1956.
- [17] J. Kivinen and M. K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. Technical Report UCSC-CRL-94-16, University of California, Santa Cruz, Computer Research Laboratory, June 1994. Revised December, 1995. An extended abstract appeared in the *Proceedings of the Twenty-Seventh Annual ACM Symposium on the Theory of Computing*, pages 209–218, 1995.
- [18] N. Littlestone. Redundant noisy attributes, attribute errors, and linear threshold learning using Winnow. In *Proceedings of the Fourth Annual Workshop on Computational Learning Theory*, pages 147–156, 1991.