Switching Portfolios

Yoram Singer
AT&T Labs, Room A277
180 Park Avenue, Florham Park, NJ 07932
singer@research.att.com

Abstract

Recently, there has been work on on-line portfolio selection algorithms which are competitive with the *best* constant rebalanced portfolio determined in *hindsight* [2, 6, 3]. By their nature, these algorithms employ the assumption that high yield returns can be achieved using a *fixed* asset allocation strategy. However, stock markets are far from being stationary and in many cases the return of a constant rebalanced portfolio is much smaller than the return of an ad-hoc investment strategy that adapts to changes in the market. In this paper we present an efficient portfolio selection algorithm that is able to track a changing market. We also describe a simple extension of the algorithm for the case of a general transaction cost, including a fixed percentage transaction cost which was recently investigated [1]. We provide a simple analysis of the competitiveness of the algorithm and check its performance on real stock data from the New York Stock Exchange accumulated during a 22-year period. On this data, our algorithm outperforms all the algorithms referenced above, with and without transaction costs.

1 Introduction

A constant rebalanced portfolio (CRP) is an asset allocation algorithm which keeps the same distribution of wealth among a set of assets along a period of time. Recently, there has been work on on-line portfolio selection algorithms which are competitive with the *best* constant rebalanced portfolio determined in *hindsight* [2, 6, 3]. However, these algorithms do not perform well when the market is changing as the following simple example shows. Assume that two hypothetical highly volatile stocks are available. The value of the first stock increases by a factor of $\frac{3}{2}$ on each of the first n trading periods. Then it changes behavior and its value falls by a factor of 4 on each of the consecutive n trading periods. The second stock behaves in an opposite manner. Its value falls by a factor of 4 on each of the first n trading periods and increases by a factor of $\frac{3}{2}$ on each of the second n trading days. The relative price change of the first stock can be described by the sequence $\frac{3}{2}, \frac{3}{2}, \ldots, \frac{3}{2}, \frac{1}{4}, \frac{1}{4}, \ldots, \frac{1}{4}$ and of the second by the sequence $\frac{1}{4}, \frac{1}{4}, \ldots, \frac{1}{4}\frac{3}{2}, \frac{3}{2}, \ldots, \frac{3}{2}$. Investing all of the initial wealth in any of the two stocks results in a fatal loss of almost all of the initial investment. Furthermore, it is easy to verify that the best constant rebalanced portfolio would redistribute the wealth equally after each trading day, resulting in an exponentially fast wealth decay. Thus, any competitive rebalancing portfolio selection algorithm, such as Cover's universal portfolio algorithm [2], would result in a similar miserable performance. In contrast, an prescient investor who puts all of hers money on the first stock for n days and then switches to the second stock would enjoy an enormous profit even in the presence of hefty transaction costs. The algorithm presented in this paper achieves almost the same wealth as the prescient investor without any prior knowledge.

The algorithm proposed here copes with changing markets by considering the possibility that the market changed its behavior after *each* trading day. The algorithm is provided with a set of investment strategies. The first version we present assumes that the apriori *duration* of using a given investment strategy is geometrically distributed with known parameters. If, in addition, the set of investment strategies includes only pure strategies, i.e., holding a single asset/stock, then a constant time per asset is required by the algorithm in order to compute the wealth redistribution after each trading day. We also give a more general version that does not employ a fixed value of parameters. This version achieves better competitiveness bounds but it is more computationally expensive. Both versions can be easily modified to deal with various models of transaction costs. Throughout the paper, we give examples using data from the New York Stock Exchange accumulated during a 22-year period and compare the results to the portfolio selection algorithms described in previous works [2, 6, 1]. On this data, the suggested portfolio selection algorithm outperforms the algorithms referenced above, with and without transaction costs.

2 Constant rebalanced portfolios

Consider a portfolio containing N stocks. On each trading day the performance of the stocks can be described by a vector of *price relatives*, denoted by $\mathbf{x} = (x_1, x_2, \ldots, x_N)$ where x_i is the next day's opening price of the ith 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 $\mathbf{w} = (w_1, w_2, \ldots, w_N)$ such that $w_i \geq 0$ and $\sum_{i=1}^{N} w_i = 1$. The ith entry of a portfolio $\mathbf{w}$ is the proportion of the total portfolio value invested in the ith asset. Therefore, if the total wealth is S and the worth of the ith asset S^i then $w_i = S^i/S$. 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 $\mathbf{w} \cdot \mathbf{x} = \sum_{i=1}^{N} w_i x_i$. Recent work in on-line portfolio selection algorithms has focused on changing an ensemble of portfolio vectors based on past performance. That is, at the start of each day t, the portfolio selection algorithm gets the previous price relatives of the stock market $\mathbf{x}^1, \ldots, \mathbf{x}^{t-1}$. From this information, the algorithm 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 price relatives $\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^T$ is observed and a sequence of portfolios $\mathbf{w}^1, \mathbf{w}^2, \ldots, \mathbf{w}^T$ is selected. From the beginning of day 1 through the beginning of day $T+1$, the wealth will have increased by a factor of $S_T(\{\mathbf{w}^t\}, \{\mathbf{x}^t\}) \stackrel{\mathrm{def}}{=} \prod_{t=1}^{T} \mathbf{w}^t \cdot \mathbf{x}^t$ or, alternatively, the logarithm of the increase is $LS_T(\{\mathbf{w}^t\}, \{\mathbf{x}^t\}) \stackrel{\mathrm{def}}{=} \sum_{t=1}^{T} \log(\mathbf{w}^t \cdot \mathbf{x}^t)$.

Cover [2] has defined a restricted class of investment strategies called constant rebalanced portfolios. As noted before, a CRP 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 $S_T(\mathbf{w}) = \prod_{t=1}^{T} \mathbf{w} \cdot \mathbf{x}^t$, $LS_T(\mathbf{w}) = \sum_{t=1}^{T} \log(\mathbf{w} \cdot \mathbf{x}^t)$. Note that such a strategy might require vast amounts of trading, since at the beginning of each trading day the investment proportions are rebalanced back to the vector $\mathbf{w}$. Given a sequence of daily price relatives 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. This portfolio vector is denoted by $\mathbf{w}^\star$. That is, $\mathbf{w}^\star \stackrel{\text{def}}{=} \arg\max_{\mathbf{w}} S_T(\mathbf{w}) = \arg\max_{\mathbf{w}} LS_T(\mathbf{w})$, 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). Recent work on portfolio selection has focused on on-line weight allocation algorithms that achieve the same asymptotic growth in normalized logarithmic wealth. An on-line algorithm that achieves such an asymptotic behavior is called *universal* portfolio. Recently, Blum and Kalai [1] extended the notion of universal portfolio to the case of fixed percentage transaction costs and gave an efficient randomized implementation of the algorithm. Throughout the paper we use Cover's (or Blum and Kalai's extension in the presence of transaction costs) universal portfolio algorithm as our straw-man for comparison.

3 Switching portfolios

In this section we provide two versions of our portfolio selection algorithm while ignoring transaction costs. A simple modification of the algorithm that takes transaction costs into account is discussed in the next section.

In contrast to previous work which has focused on finding a good portfolio vector, we assume instead that we are *provided* with a set of possible investment strategies which we term basic strategies. A basic strategy need not be complex. In fact, with the benefit of hindsight, on each day one can invest all of one's wealth in the single best-performing asset for that day. We thus use pure investment strategies, i.e., strategies that invest all the wealth in a single asset, as our basic investment strategies. Clearly, we do not have the luxury of foreseeing future behavior of stock markets. However, as we show, it is possible to track an investment regime that switches from one investment strategy to another as the market changes its behavior. We do so by employing a mixture of all possible switching regimes. The mixture technique enables us to hedge our bets against the individual switching regimes. We associate a prior probability with each switching regime. The weights are distributed among the different possible switching regimes such that more complicated regimes, that frequently switches from one strategy to another, are *apriori* less favorable. We then let the evidence, i.e., the actual returns, dictate which investment strategy to use.

In summary, our approach to portfolio selection is as follows. We first decide upon a set of investment strategies. We then choose a prior distribution over the possible switching sequences from one investment strategy to another. This prior distribution is recursive in order to enable an efficient evaluation of the portfolio vector. Last, we combine the actual return of each strategy on each day with the prior distribution to design a new portfolio vector before each trading day. The wealth achieved by the algorithm is no worse than the wealth of any specific switching regime times the prior probability of that regime. We therefore can give a simple lower bound on the minimal return of our algorithm compared to any available switching regime.

The first version assumes that the duration of using one strategy is geometrically distributed with a given parameter γ. Thus, if we started using the ith investment strategy at time t_0, then the *apriori* probability of using this strategy through time t_1 (and then switching to a new strategy) is $(1-\gamma)^{t_1-t_0}\gamma$. An investment switching regime $\mathcal{Q}$ for T trading days is described in terms of two lists, $(t_1 t_2 \dots t_l)$ and $(i_1 i_2 \dots i_l i_{l+1})$, where the t_j are indices of the trading days *after* which we switched to a new investment strategy and i_j are the indices of the basic investment strategies used. Defining $t_0 = 0$ and assuming that a new strategy is picked uniformly at random, the *apriori* probability of a switching regime $\mathcal{Q}$ from $t = 1$ through $t = T$ is,

$$
\begin{aligned}
P_0(\mathcal{Q}) &= N^{-1}(N-1)^{-l}\left[\prod_{i=1}^{l}(1-\gamma)^{t_i-t_{i-1}-1}\gamma\right](1-\gamma)^{T-t_l-1} && (1)\\
&= N^{-1}(N-1)^{-l}\gamma^{l}(1-\gamma)^{T-l-1}\quad, && (2)
\end{aligned}
$$

where N is the number of investment strategies which in the case of pure strategies is also the number of different assets/stocks. The wealth achieved by a switching regime $\mathcal{Q}$ after T trading days, $S_T(\mathcal{Q})$, is the product the returns of the investment strategies used by the regime, where for pure strategies is simply,

$$
S_T(\mathcal{Q}) = \prod_{j=1}^{l+1}\prod_{t=t_{j-1}+1}^{t_j} x_{i_j}^{t}\ . \tag{3}
$$

Therefore, the sum of the wealth achieved by each switching regimes, weighted by its prior probability, is simply $\sum_{\mathcal{Q}'} P_0(\mathcal{Q}')S_T(\mathcal{Q}')$. Evaluating this sum directly is clearly infeasible since the number of different switching regimes grows exponentially fast. However, since the geometric distribution is memoryless, we can calculate the sum in constant time per asset for each trading day as follows. Let S_t^i be the worth of the ith asset after t trading days. Then, at trading day $t+1$, we either stay with the current (pure) strategy with probability $1-\gamma$, and therefore keep holding the ith stock, or switch to a new strategy (with probability γ) by redistributing the wealth among all other assets. Put another way, the value of the ith asset is the sum of two terms: the first is the previous value of the asset times the probability that we kept using the ith investment strategy on the current trading day; the second is value of all other assets time the probability that we switched to the ith pure investment strategy. Formally, the value of the ith asset after trading day $t+1$ is,

$$
\begin{aligned}
S_{t+1}^{i} &= \left((1-\gamma)S_t^i + \frac{\gamma}{N-1}\sum_{j\neq i} S_t^j\right)x_i^{t+1} && (4)\\
&= \left(\left(1-\gamma-\frac{\gamma}{N-1}\right)S_t^i + \frac{\gamma}{N-1}\sum_{j=1}^{N} S_t^j\right)x_i^{t+1}\quad . && (5)
\end{aligned}
$$

The above equations give a simple procedure to incorporate the prior probability over switching regimes with the actual return. That is, the above scheme can be directly described as

Stocks	Best Stock	BCRP	EG(η) ($\eta = 0.05$)	Universal Portfolio	Switching ($\gamma = 1/4$)	Switching (Adaptive γ)
Iroquois & Kin Ark	8.92	73.70	70.85	39.97	72.58	**143.71**
Com. Met. & Kin Ark	52.02	144.00	117.15	80.54	118.68	**206.74**
Com. Met. & Mei. Corp.	52.02	102.96	97.93	74.08	98.89	**107.74**
IBM & Coca-Cola	13.36	**15.07**	14.90	14.24	15.02	15.05

Table 1: Comparison of the wealths achieved by various on-line portfolio selection algorithms. 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(η)-update rule (Helmbold et. al.), the wealth achieved by Cover's universal portfolio algorithm, and the switching portfolios algorithm with a fixed and adaptive switching probability.

a portfolio weight update. Writing $w_t^i = S_t^i / \sum_j S_t^j$, the weight w_{t+1}^i *before* trading day $t+1$ can be described in terms of w_t^i as follows,

$$w_{t+1}^i = \left(1 - \gamma - \frac{\gamma}{N-1}\right) w_t^i + \frac{\gamma}{N-1} \ . \tag{6}$$

This portfolio weight update scheme resembles the *fixed-share* weight update used by Herbster and Warmuth [7] for tracking the best expert in a binary prediction setting. The analysis presented in [7] does not simply carry over to our setting of unbounded returns. Since the next version of the algorithm includes this version as a special case, we defer the analysis of the first version to the end of the section.

We compared the performance of the switching portfolios algorithm to the returns achieved by Cover's universal portfolio algorithm [2], the multiplicative weight update algorithm [6] (denoted by EG(η)), the best constituent stock, and the best constant rebalanced portfolio. We compared the results for all subsets of stocks considered by Cover [2] in his experiments. The results are summarized in Table 1. In the absence of transactions costs, the average time to hold to an investment strategy can be relatively short. For the NYSE stock data we found that a reasonable time to hold a single stock ranges from a few days to a few weeks. To have a fair comparison we set $1/\gamma = 4$ in the all the experiments reported in Table 1, regardless of the stocks constituting the portfolios. In all the experiments reported, the wealth is calculated assuming an initial investment of one unit before the first trading day.

The problem with the first version described above is that the switching distribution is specified through a parameter γ which we need to set in advance. We now give a second version of the switching portfolios algorithm that does not need such prior knowledge. Instead of setting γ to a predefined value, we let γ vary in time and define $\hat{\gamma}(\Delta t) = \frac{1/2}{\Delta t + 1}$, to be the switching portability after using the same investment strategy for Δt consecutive trading days. The rationale behind this choice switching probability is beyond the scope of this short paper and is based on a universal coding approach (see [8] and the references therein for a motivating discussion and detailed analysis). We later show that this choice of switching probability lets us derive a simple competitiveness bound for the switching portfolios algorithm. Denote by S_{t,t_0}^i the wealth accumulated in the ith asset after trading day t given that we started holding the asset on trading day t_0. We now need to update the wealth accumulated for each asset based on the start date of the corresponding (pure) investment strategy. The wealth update scheme becomes,

$$\begin{aligned} S^i_{t+1,t_0} &= (1-\hat{\gamma}(t-t_0))\, S^i_{t,t_0}\, x^{t+1}_i \;=\; \frac{t-t_0+1/2}{t-t_0+1}\, S^i_{t,t_0}\, x^{t+1}_i \\ S^i_{t+1,t+1} &= \left(\sum_{j\neq i}\sum_{t_0=1}^{t} \hat{\gamma}(t-t_0)\, S^j_{t,t_0}\right) x^{t+1}_i \;=\; \left(\sum_{j\neq i}\sum_{t_0=1}^{t} \frac{1}{2(t-t_0+1)}\, S^j_{t,t_0}\right) x^{t+1}_i \quad . \end{aligned} \tag{7}$$

Note that this more general version has a price. The above wealth update can no longer be computed in a constant time per stock. Since at trading day t we perform $O(t)$ operations, the overall time complexity of the algorithm is $O(t^2)$ (as opposed to $O(t)$ for the first version). Furthermore, there is no equivalent update that directly manipulates a portfolio weight vector. Nonetheless, since the actual switching rate γ is induced by the actual returns, and thus may vary in time, we achieve significantly higher yields as shown in Table 1. The first version achieves results similar to the *best* CRP and second version significantly outperforms the best CRP in three out of the four cases. The fourth subset includes IBM and Coca-Cola. These stocks showed a 'lock-step' performance during the period 1963-1985. Thus, it is not surprising that all the portfolio selection algorithms achieved similar results. We also would like to note that we found that our the second version of the switching portfolios algorithm consistently outperforms the Best CRP on different (and larger) subsets of the historical NYSE stock market data, although the switching portfolios does not have the luxury of hindsight.

We now discuss the competitiveness properties of the switching portfolios algorithm. We compare the performance of the on-line algorithm to *any* switching regime that can be determined in hindsight. To derive a lower bound on wealth achieved by the switching portfolios algorithm we need the following lemma (proof omitted).

Lemma 1 $\log\left(\prod_{i=0}^{l-1}\frac{i+1/2}{i+1}\right) \;=\; \sum_{i=1}^{l-1}\log\left(\frac{i-1/2}{i}\right) \geq -\frac{1}{2}\log(l) - log(2)$.

Based on the above lemma we now give our main competitiveness result.

Theorem 2 *Let* $\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^T (T \geq 2)$ *be any sequence of price relatives for* N *assets. Let* $S_T(\mathcal{Q})$ *be the wealth achieved by a switching regime* $\mathcal{Q}$ *that uses the* N *pure investment strategies (as defined by Equ. (3)) and let* $l(\mathcal{Q})$ *be the number of times* $\mathcal{Q}$ *switches from one strategy to another. Then, the logarithmic wealth achieved by the switching portfolios algorithm is at least*

$$\log(S_T(\mathcal{Q})) - \frac{3}{2} l(\mathcal{Q}) \log\left(\frac{T}{l(\mathcal{Q})}\right) - 1/2\log(T) - (l(\mathcal{Q})+1)\log(4N) \quad .$$

Proof sketch: The wealth achieved by the switching portfolios algorithm given by Equation (7) is an efficient way to calculate the sum $\sum_{\mathcal{Q}'} P_0(\mathcal{Q}')S_T(\mathcal{Q}')$. Thus, the logarithmic wealth achieved by the switching portfolios algorithm is

$$\log(\sum_{\mathcal{Q}'} P_0(\mathcal{Q}')S_T(\mathcal{Q}')) \geq \log(P_0(\mathcal{Q})S(\mathcal{Q})) = \log(P_0(\mathcal{Q})) + \log(S(\mathcal{Q})) \quad .$$

Now, from the definition of $\hat{\gamma}(\Delta t)$ we get that,

$$P_0(\mathcal{Q}) = \frac{1}{N}(N-1)^{-l}\left(\prod_{i=1}^{l}\left(\prod_{j=1}^{t_i-t_{i-1}-1}\frac{j-1/2}{j}\right)\frac{1/2}{t_i-t_{i-1}}\right)\prod_{j=1}^{T-t_l-1}\frac{j-1/2}{j} \quad .$$

Let l be a shorthand for $l(\mathcal{Q})$. Using Lemma 1 and simple algebraic manipulations, we get that,

$$\log(P_0(\mathcal{Q})) = -\log(N) - l\log(N-1) - 3/2\sum_{i=1}^{l}\log(t_i-t_{i-1}) - 2l\log(2) - 1/2\log(T-t_l) - \log(2) \quad .$$

Now, using the log-sum inequality (see [4]) we can bound the above expression as follows,

$$\begin{aligned}\log(P_0(\mathcal{Q})) &\geq -\log(N) - l\log(N-1) - 3/2l\log(T/l) - (2l+1)\log(2) - 1/2\log(T) \\ &\geq -(l+1)\log(4N) - 3/2l\log(T/l) - 1/2\log(T) \quad . \qquad \square\end{aligned}$$

Deriving a competitiveness bound for the first version of the algorithm is a much easier task since γ is fixed. In short, we get that logarithmic wealth achieved by the switching portfolios algorithm with a fixed γ compared to any switching regime is smaller by at most, $(l+1)\log(N) + l\log(1/\gamma) + (T-l)\log(1/(1-\gamma))$.

The bound of Thm. 2 has a nice intuitive interpretation. Although a switching regimes that frequently switches from one investment strategy to another has the potential of achieving high yields the gap between the wealth achieved by the switching portfolios algorithm and such a frequently switching regime is large. Furthermore, since the bound holds for any switching regime we in fact get that the wealth achieved by the algorithm is at least,

$$\max_{\mathcal{Q}}\left\{\log(S_T(\mathcal{Q})) - \frac{3}{2}l(\mathcal{Q})\log\left(\frac{T}{l(\mathcal{Q})}\right) - 1/2\log(T) - l(\mathcal{Q})\log(4N)\right\} \quad .$$

Therefore, the switching portfolios algorithm encompasses a natural tradeoff between the yield of a switching regime and its complexity (in terms of the number of time it switches).

4 Transaction costs

Coping with a transaction costs model when using the switching portfolios algorithm with pure investment strategies is a simple task. Since on each trading day we take a portion of the wealth from each asset and redistribute it among the rest of the assets, we simply need to deduct the cost of selling asset i, and then further deduct the cost of buying asset j. Clearly, this is not the least expensive scheme to redistribute the wealth. We might end up selling and buying the same asset and thus pay more commission than would have been needed had we pre-calculated the amount we need to sell/buy per each asset. However, it is enough to use this scheme in order to achieve the same competitiveness bounds. For the case of a fixed percentage $\mathbf{c}$ of the amount we trade (either buy or sell), the first equation of the wealth update Equation (7) remains the same (no trading is taking place), while the second equation becomes,

$$S^i_{t+1,t+1} \;=\; (1-\mathbf{c})^2 \left(\sum_{j\neq i} \sum_{t_0=1}^{t} \frac{1}{2(t-t_0+1)} S^j_{t,t_0} \right) x_i^{t+1} \quad . \tag{8}$$

Although simplistic, this approach results in the same competitiveness bounds when using only pure strategies. When we switch from one strategy to another we have to sell one asset entirely and buy a new one. Therefore, the wealth achieved by a switching regime $\mathcal{Q}$ in the presence of fixed percentage transaction costs is now *exactly*, $S_T(\mathcal{Q}) = \prod_{j=1}^{l+1} (1-\mathbf{c})^2 \prod_{t=t_{j-1}+1}^{t_j} x^t_{i_j}$. Thus, the bound on gap between the wealth achieved by the switching portfolios algorithm and any specific switching regime does not change.

It is also simple to derive a wealth update scheme for other transaction costs models. For instance, Blum and Kalai [1] suggested a transaction costs model where first the entire portfolio vector is updated in *parallel* for all the assets, then the cost is subtracted from the total wealth (some from each stock) such that the wealth proportions would not change. Formally, the transaction cost of changing the wealth distribution from $\{S_i\}$ to $\{S'_i\}$ in that model is, $c\sum_i |S_i - S'_i|$. The analysis for this transaction cost model follows similar lines. For a switching regime $\mathcal{Q}$, we either keep the same asset or sell it entirely and buy a new asset. Thus, the former case incur a zero transaction cost while the latter decreases the wealth by a factor of $1-2\mathbf{c}$. Therefore, the wealth in case of a switch, as described by the second part of Equation (7), now becomes,

$$S^i_{t+1,t+1} \;=\; (1-2\mathbf{c}) \left(\sum_{j\neq i} \sum_{t_0=1}^{t} \frac{1}{2(t-t_0+1)} S^j_{t,t_0} \right) x_i^{t+1} \quad . \tag{9}$$

Again, the bound on wealth achieved by the switching portfolios algorithm remains the same. Blum and Kalai used this transaction costs model and described a modification to Cover's universal portfolio algorithm that is still competitive. In Table 2 we compare the wealths achieved by Blum and Kalai's algorithm and the switching portfolios algorithm for the same subsets of stocks appearing in Table 1. The results are given for a modest 0.1% transaction cost and a hefty 2% transaction cost. In all cases, the switching portfolios algorithm achieves much better results than the modified universal portfolio algorithm. Furthermore, in several cases the switching portfolios algorithm outperforms the best CRP. We also checked the performance on larger sets of stocks that were not reported in previous papers. In all cases checked, we found that the wealths achieved by the universal portfolio algorithm were significantly smaller than the wealths achieved by the switching portfolios algorithm, and, often smaller than the wealth achieved by the best constituent stocks. Moreover, we found that the universal portfolio scales rather poorly with the size of the portfolio (i.e., the number of stocks) while the switching portfolios is almost always able to take advantage of additional stocks.

To conclude this section, we give an illustrative example of a portfolio with three stocks: Dow Chemicals, Espey Manufacturing, and Kin Ark. These stocks are rather volatile and, as shown at the top right hand side of Figure 1, they exhibit different behavior during relatively long periods. The left and the middle graphs in the figure show the daily wealths achieved

	0.1% Transaction Costs			2% Transaction Costs		
Stocks	BCRP	Universal	Switching	BCRP	Universal	Switching
Iroquois & Kin Ark	68.02	33.92	**115.3**	22.05	9.22	**23.32**
Com. Met. & Kin Ark	136.6	65.09	**169.0**	56.20	21.53	**42.88**
Com. Met. & Mei. Corp.	**99.07**	64.33	97.41	52.76	29.81	**67.59**
IBM & Coca-Cola	**14.83**	13.85	14.79	13.02	10.24	**13.81**

Table 2: Comparison of the wealths achieved by the universal portfolio algorithm and the switching portfolios algorithm with 0.1% and 2% transaction costs. For comparison the wealth achieved by the best CRP is also provided.

by the switching portfolios algorithm and the universal portfolio for transaction costs of 0.5% and 3%. For comparison, we also give the daily value of the best constituent stock. The bottom right hand side plot describes, using a color coded scheme, the largest asset in the portfolio (for a 3% transaction cost). That is, if $\forall j \neq i : S_t^i > S_t^j$ then we draw a bar using the ith color at the location corresponding to the tth trading day. Clearly, there are long periods where a significant portion of the wealth is invested in a single asset, which is the asset that *locally* gives the highest returns. This behavior confirms empirically our basic assumption that high yield returns can be achieved using pure investment strategies combined with a switching regime.

We also implemented and evaluated more complex investment strategies in addition to pure strategies. For instance, we used Blum and Kalai's universal portfolio algorithm itself as a basic investment strategy and 'fed' it into the switching portfolios algorithm. In all the experiments we performed, we found a modest improvement (due to the lack of space the actual results are omitted) over using pure strategies only as our basic investment strategies. These findings are not surprising since, as argued earlier, with the benefit of hindsight, we can use pure investment and still achieve enormous returns. Roughly speaking, the switching portfolios algorithm tries to imitate an prescient observer and thus the improvement using more complex, yet oblivious, investment strategies is relatively small.

5 Conclusions

A simple and efficient portfolio selection algorithm was presented in this paper. The algorithm is competitive with any switching regime determined in hindsight. Similar approaches were investigated for different settings [7, 5]. One of the contributions of this paper is the distillation of key elements of previously known methods, and the synthesis with other learning-theory and information-theory results leading to an algorithm that outperforms previously published *competitive* portfolio selection algorithms. Furthermore, our algorithm achieves a slightly better competitiveness bound than the bounds reported in [7, 5] and in addition can also use *non-fixed* investment strategies as discussed above. The price that we pay for the improved bounds and flexibility is a more complex and time consuming algorithm. An important question is whether our more general algorithm can be implemented using a constant time per asset for each trading day.

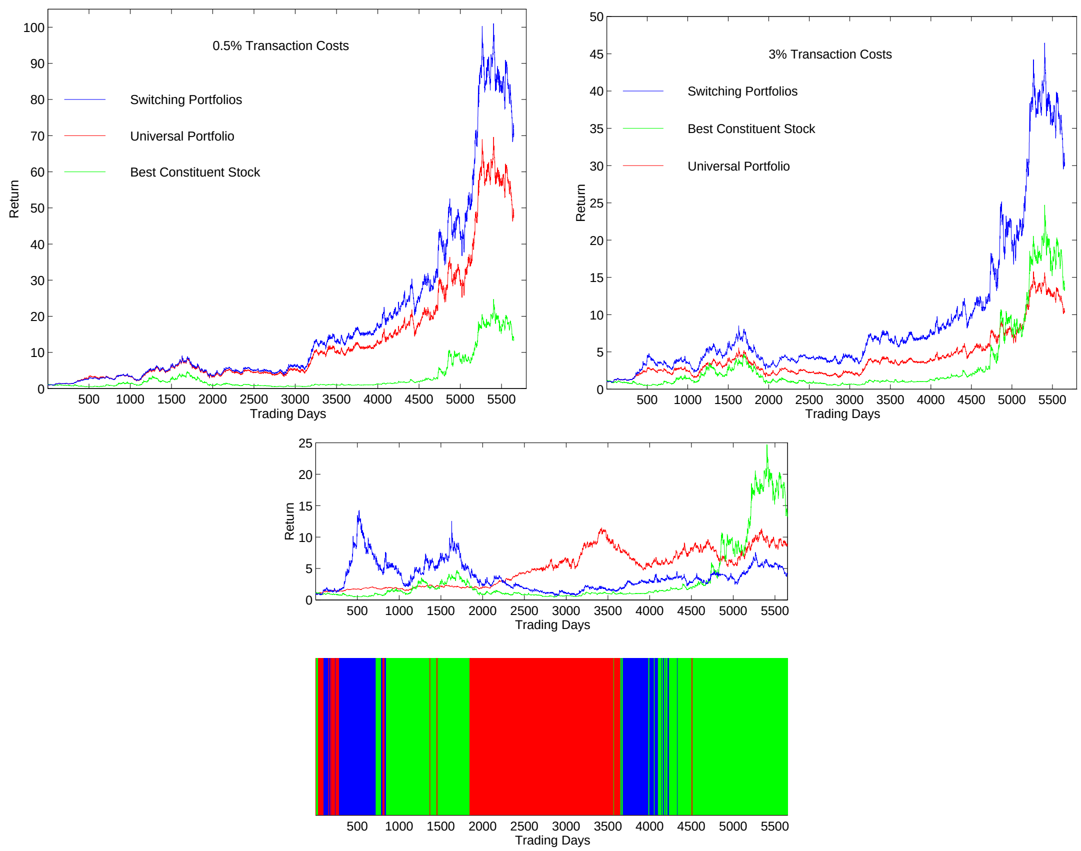

Figure 1: Top figures: the wealths achieved by the switching portfolios algorithm (with pure strategies), the modified universal portfolio algorithm and the best constituent stock in the presence of 0.5% and 3% transaction costs. Bottom figure: the performance of the individual stocks and a color coded plot of the largest invested stock held by the switching portfolios algorithm on each trading day.

There is still some room for improvement within the current framework. First, we use the least informative scheme, namely the uniform distribution, to redistribute the wealth upon a switch. More informative redistribution techniques that use past performance might yield higher returns as well as better competitiveness bounds. Second, our switching model employs a tacit assumption that the *prior* probability of switching is *stationary*. However, real stock markets alternate between periods of rapid price change and relatively calm periods. Such a behavior is not captured by the current *prior* probability model for switching. It might be possible to model time dependent prior distributions with an even more complex algorithm. Lastly, although we observed only a modest improvement in performance when we used complex investment strategies, this might simply reflect a limitation of the universal portfolio algorithm. Investigating alternative basic investment strategies that may result in higher returns is one of the future research goals.

Acknowledgments

Thanks to Y. Bengio, M. Herbster, R. Schapire, and M. Warmuth for helpful discussions.

References

[1] A. Blum and A. Kalai. Universal portfolios with and without transaction costs. In *Proc. of Tenth Ann. Conf. on Comp. Learning Theory*, 1997.

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

[3] T.M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2), 1996.

[4] T.M. Cover and J.A. Thomas. *Elements of information theory.* Wiley, 1991.

[5] Y. Freund, R.E. Schapire, Y. Singer, and M.K. Warmuth. Using and combining predictors that specialize. In *Proc. of the Twenty-Ninth Annual ACM Symp. on the Theory of Computing*, pages 334–343, 1997.

[6] D.P. Helmbold, R.E. Schapire, Y. Singer, and M.K. Warmuth. On-line portfolio selection algorithms using multiplicative updates. In *Machine Learning: Proceedings of the Thirteenth Intl. Conference*, 1996.

[7] M. Herbster and M.K. Warmuth. Tracking the best expert. In *Proc. of the Twelfth Intl. Conf. on Machine Learning*, pages 286–294, 1995.

[8] R.E. Krichevsky and V.K. Trofimov. The performance of universal coding. *IEEE Transactions on Information Theory*, 27:199–207, 1981.