, , 1–13 ()

Universal Portfolios With and Without Transaction Costs

AVRIM BLUM avrim+@cs.cmu.edu
ADAM KALAI akalai+@cs.cmu.edu
School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract. A constant rebalanced portfolio is an investment strategy which keeps the same distribution of wealth among a set of stocks from period to period. Recently there has been work on on-line investment strategies that are competitive with the best constant rebalanced portfolio determined in hindsight (Cover, 1991; Helmbold et al., 1996; Cover and Ordentlich, 1996a; Cover and Ordentlich, 1996b; Ordentlich and Cover, 1996; Cover, 1996). For the universal algorithm of Cover (Cover, 1991), we provide a simple analysis which naturally extends to the case of a fixed percentage transaction cost (commission), answering a question raised in (Cover, 1991; Helmbold et al., 1996; Cover and Ordentlich, 1996a; Cover and Ordentlich, 1996b; Ordentlich and Cover, 1996; Cover, 1996). In addition, we present a simple randomized implementation that is significantly faster in practice. We conclude by explaining how these algorithms can be applied to other problems, such as combining the predictions of statistical language models, where the resulting guarantees are more striking.

1. Introduction

A constant rebalanced portfolio (CRP) is an investment strategy which keeps the same distribution of wealth among a set of stocks from period to period. That is, the proportion of total wealth in a given stock is the same at the beginning of each period. Recently there has been work on on-line investment strategies which are competitive with the best CRP determined in hindsight (Cover, 1991; Helmbold et al., 1996; Cover and Ordentlich, 1996a; Cover and Ordentlich, 1996b; Ordentlich and Cover, 1996; Cover, 1996). Specifically, the daily performance of these algorithms on a market approaches that of the best CRP for that market, chosen in hindsight, as the lengths of these markets increase without bound.

As an example of a useful CRP, consider the following market with just two stocks (Helmbold et al., 1996; Ordentlich and Cover, 1996). The price of one stock remains constant, and the price of the other stock alternately halves and doubles. Investing in a single stock will not increase the wealth by more than a factor of two. However, a $(\frac{1}{2}, \frac{1}{2})$ CRP will increase its wealth exponentially. At the end of each period it trades stock so that it has an equal worth in each stock. On alternate periods the total value will change by a factor of $\frac{1}{2}(1) + \frac{1}{2}(\frac{1}{2}) = \frac{3}{4}$ and $\frac{1}{2}(1) + \frac{1}{2}(2) = \frac{3}{2}$, thus increasing total worth by a factor of $9/8$ every two periods.

We extend this model by adding a fixed percentage commission cost $c < 1$ to each transaction, as is common in financial modeling (Davis and Norman, 1990). To fully define the commission model, in addition to specifying the cost of each transaction, we must also specify how a CRP pays this overhead. In Section 2, we give three natural properties of such a specification which we will use for our analysis. These properties are satisfied by the optimal investor who must simultaneously rebalance her portfolio and pay for these transaction costs from the sale of stock.

In Section 3 we present a new, simpler analysis and implementation for Cover's universal algorithm (Cover, 1991). It is along the same lines of reasoning as an argument in (Foster and Vohra, 1995). We show how our argument extends to the case of commission in Section 4. The previous bound for a market with m assets and n periods was a performance ratio of at least

$$\frac{\text{wealth of UNIVERSAL}}{\text{wealth of best CRP}} \geq \frac{1}{(n+1)^{m-1}}.$$

In the presence of a commission c, this becomes

$$\frac{\text{wealth of UNIVERSAL}_c}{\text{wealth of best CRP}} \geq \frac{1}{((1+c)n+1)^{m-1}}.$$

The above ratio is a decreasing function of n. However, the average per-period ratio, $(1/(n+1)^{m-1})^{1/n}$, increases to 1 as n increases without bound. For example, if the best CRP makes one and a half times as much as we do over a period of 22 years, it is only making a factor of $1.5^{1/22} \approx 1.02$ as much as we do per year. Cover raises the question of of whether it is even possible to achieve the exponential growth rate of the best CRP in the presence of commission (Cover and Ordentlich, 1996a), and our analysis answers this question. However, we do not consider the Dirichelet$(1/2, \ldots, 1/2)$ Universal algorithm (Cover and Ordentlich, 1996a) which has the better guaranteed ratio of $2\sqrt{1/(n+1)^{m-1}}$.

Semi-constant-rebalanced portfolios, which may or may not rebalance during each period, have been suggested as an investment strategy in the presence of commission (Helmbold et al., 1996). In Section 5 we show that the exponential wealth of a semi-constant-rebalanced portfolio cannot, in general, be achieved without future knowledge of the market.

In Section 6, we present results of the universal algorithm with various commissions on some real-world stock data. Finally, in Section 7, we explain how these stock market algorithms can be used on other problems, such as combining the predictions of language models almost as well as if we had prior knowledge of the optimal mixture.

2. Notation and Definitions

We use mainly the notation of (Cover, 1991). A price relative for a given asset is the nonnegative ratio of closing price to opening price during a given period. If the market has m assets and trading takes place during n periods, then the market's performance can be expressed by n price relative vectors, $\mathbf{x}^n = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$, $\mathbf{x}_i \in \Re_+^m$, where x_{ij} is the nonnegative price relative of the j^{th} stock for the i^{th} period.

A portfolio is simply a distribution of wealth among the stocks. We represent a distribution of wealth by $\mathbf{b} \in \Re_+^m$, where $\sum_j b_j = 1$. So $\mathbf{b}$ is an element of the $(m-1)$-dimensional simplex β,

$$\beta = \{\mathbf{b} = (b_1, b_2, \ldots, b_m) : \sum_j b_j = 1, \forall_j \ b_j \geq 0\}.$$

The CRP investment strategy for a particular portfolio $\mathbf{b}$, $\mathrm{CRP}_\mathbf{b}$, redistributes its wealth at the end of each period so that the proportion of money in the j^{th} stock is b_j. An investment using a portfolio $\mathbf{b}$ during a period with price relatives $\mathbf{x}$ increases one's wealth by a factor of $\mathbf{b} \cdot \mathbf{x} = \sum b_j x_j$. Therefore, over n periods, the wealth achieved by $\mathrm{CRP}_\mathbf{b}$ is,

$$S_n(\mathbf{b}, \mathbf{x}^n) = \prod_{i=1}^{n} \mathbf{b} \cdot \mathbf{x}_i.$$

2.1. *Commission*

We now consider an extension of the above model to the case of a fixed percentage commission $0 \leq c \leq 1$. For simplicity, we will assume that the commission is charged only for purchases and not for sales (we explain at the end of this section why this can be assumed without loss of generality.) We now need to specify how an investor, who has a target distribution of wealth, pays for these transaction costs, each period. In our model, the investor must pay for all transaction costs by selling stock. Since we are comparing ourselves to the best CRP, it is natural to assume that the CRP investor makes the optimal trades so as to rebalance her portfolio and pay for her transaction costs. For example, suppose there is a hefty 40% commission on each purchase. Say, at the end of a period, an investor has \$200 in stock A, and \$800 in stock B, and this investor wishes to rebalance to a $(1/2, 1/2)$ portfolio for the start of the next period. The optimal investor would first sell \$100 of stock B to cover the upcoming transaction costs. She now has \$200 in stock A, \$700 in stock B, and \$100 in cash. When she trades \$250 of stock B for stock A to get \$450 in each stock, she then pays her (40%)\$250 = \$100 in transaction costs. This is better than someone who naively tries to rebalance to \$500 in each stock and then must sell \$60 worth of each stock to pay his (40%)\$300 = \$120 in transaction costs, leaving \$440 in each stock.

For our analysis, we do not need to know the specifics of this optimal rebalancer. But for the sake of completeness, let's look at how to compute these optimal costs. Say we start with one dollar distributed according to $\mathbf{b}$ and we would like to rebalance to $\mathbf{b}'$, optimally with respect to commission. If we know the the largest amount that we can have after rebalancing, α, then it is easy. In order to achieve α dollars distributed according to $\mathbf{b}'$, we sell the difference of every stock for which $\alpha b'_j < b_j$ and buy the others. Optimality requires that α is the solution of

$$\alpha = 1 - c \sum_{\alpha b'_j > b_j} (\alpha b'_j - b_j).$$

Algorithmically, to solve the above equation, we first determine which interval α is in by checking the critical points, $\alpha = b_j / b'_j$. Once we've determined the interval, the above reduces to a linear equation. For reasonable commission costs, we can easily approximate the optimal rebalancer described above by paying for the transactions proportionally from each stock, i.e.

$$\alpha \approx 1 - c \sum_{b'_j > b_j} (b'_j - b_j).$$

This is not optimal rebalancing because some of the same stock is bought and then sold during a single rebalance. However, our on-line guarantees also hold among this class of CRP investors, who do not optimally rebalance with respect to commission.

In our analysis, we will only use the following natural properties of optimal rebalancing:

1. The costs paid changing from distribution $\mathbf{b}_1$ to $\mathbf{b}_3$ is no more than the costs paid changing from $\mathbf{b}_1$ to $\mathbf{b}_2$ and then from $\mathbf{b}_2$ to $\mathbf{b}_3$.

2. The cost, per dollar, of changing from a distribution $\mathbf{b}$ to a distribution $(1 - \alpha)\mathbf{b} + \alpha\mathbf{b}'$ is no more than αc, simply because we are moving *at most* an α fraction of our money.

3. An investment strategy I which invests an initial fraction α of its money according to investment strategy I_1 and an initial $1 - \alpha$ of its money according to I_2, will achieve at least α times the wealth of I_1 plus $1 - \alpha$ times the wealth of I_2. (In fact, I may do even better by occasionally saving in commission cost if, for instance, strategy I_1 says to sell stock A and strategy I_2 says to buy it.)

Our model assumes commission on buying but not selling. Alternatively, one can imagine having two commissions, c_{buy} and c_{sell}, for buying and selling, as we do in our experiments. Our theoretical results will still hold for $c = c_{\text{buy}} + c_{\text{sell}}$ because one dollar in a single stock can be transferred to $(1 - c_{\text{sell}})/(1 + c_{\text{buy}}) \geq (1 - c)$ dollars in a different stock, which is all that is required for property 2.

3. Analysis Without Commission

In this section, we give a simple analysis of the universal algorithm of Cover, without commission. To get our bearings, let's first consider an easier question. Suppose you just want a strategy that is competitive with respect to the best single stock. In other words, you want to maximize the worst-case ratio of your wealth to that of the best stock. In this case, a good strategy is simply to divide your money among the m stocks and let it sit. You will always have at least $\frac{1}{m}$ times as much money as the best stock. Note that this deterministic strategy achieves the expected wealth of the randomized strategy that just places all its money in a random stock.

Now consider the problem of competing with the best CRP. Cover's universal portfolio algorithm is similar to the above. It splits its money evenly among all CRPs and lets it sit in these CRP strategies. (It does not transfer money between the strategies.) Likewise, it always achieves the expected wealth of the randomized strategy which invests all its money in a random CRP. In particular, the bookkeeping works as follows:

Definition 1. **(UNIVERSAL)** The universal portfolio algorithm, UNIVERSAL, at time i is specified by

$$\hat{\mathbf{b}}_i = \frac{\int_\beta \mathbf{b} S_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}{\int_\beta S_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}, \quad i = 1, 2, \ldots$$

with μ equal to the uniform distribution over portfolios $\mathbf{b}$. (We do not consider the Dirichelet$(1/2, 1/2, \ldots)$ distribution.)

This is the form in which Cover defines the algorithm. He also notes ((Cover and Ordentlich, 1996a)) that

$$\text{wealth of UNIVERSAL} = \mathrm{E}_{\mathbf{b}\in\beta}\left[\text{wealth of } \mathrm{CRP}_{\mathbf{b}}\right] \tag{1}$$

Let the *holdings* of a portfolio represent the actual amount of wealth we have in each stock, so that it is a vector whose elements do not necessarily add to 1. Then equation (1) can be seen very easily if we rewrite Definition 1 as, during period i,

$$\text{holdings of UNIVERSAL} = \mathrm{E}_{\mathbf{b}\in\beta}\left[\text{holdings of } \mathrm{CRP}_{\mathbf{b}}\right].$$

THEOREM 1 *As in (Cover and Ordentlich, 1996a),*

$$\begin{aligned} \frac{\textit{wealth of UNIVERSAL}}{\textit{wealth of best CRP}} &\geq \binom{n+m-1}{m-1}^{-1} \\ &\geq \frac{1}{(n+1)^{m-1}}, \end{aligned}$$

for all markets with m stocks and n periods.

Proof. As mentioned above, the wealth of UNIVERSAL is simply the expected wealth of a random CRP. The idea behind our argument is that portfolios "near" to each other perform similarly and that a large fraction of portfolios are "near" the optimal portfolio.

Say, in hindsight, $\mathbf{b}^*$ is an optimal CRP for the particular market. Here is what we mean when we say that "near" portfolios perform nearly as well as $\mathbf{b}^*$. If $\mathbf{b} = (1-\alpha)\mathbf{b}^* + \alpha\mathbf{z}$ for some $\mathbf{z} \in \beta$, then a single period's gain of $\mathrm{CRP}_{\mathbf{b}}$ must be at least $(1-\alpha)$ times as large as a single period's gain of $\mathrm{CRP}_{\mathbf{b}^*}$. After all, a $(1-\alpha)$ fraction of $\mathbf{b}$ is distributed exactly like $\mathbf{b}^*$. Compiled over n periods,

$$\text{wealth of } \mathrm{CRP}_{\mathbf{b}} \geq (1-\alpha)^n(\text{wealth of } \mathrm{CRP}_{\mathbf{b}^*}). \tag{2}$$

So, to prove the theorem, we will show that a sufficiently large volume of portfolios is sufficiently near $\mathbf{b}^*$. This is easy to compute because the set of near portfolios is a shrunken simplex, $\alpha\beta$, translated from the origin to $(1-\alpha)\mathbf{b}^*$.

$$\begin{aligned} \mathrm{Vol}_{m-1}(\{(1-\alpha)\mathbf{b}^* + \alpha\mathbf{z} : \mathbf{z} \in \beta\}) &= \mathrm{Vol}_{m-1}(\{\alpha\mathbf{z} : \mathbf{z} \in \beta\}) \\ &= \alpha^{m-1}\mathrm{Vol}_{m-1}(\beta). \end{aligned} \tag{3}$$

In particular, for $\alpha = 1/(n+1)$, we get a ratio of $\alpha^{m-1}(1-\alpha)^n > e^{-1}/(n+1)^{m-1}$ for the wealth of universal compared to the best CRP, because at least an $\alpha^{m-1} = 1/(n+1)^{m-1}$ fraction of portfolios perform at least $(1-\alpha)^n > e^{-1}$ as well as the best CRP.

We can get a better bound if we consider α to be the random variable defined by $1-\alpha = \min_j\{b_j/b_j^*\}$. Then, from equations (1) and (2),

$$\frac{\text{wealth of UNIVERSAL}}{\text{wealth of best CRP}} \geq \mathrm{E}_{\mathbf{b}\in\beta}[(1-\alpha)^n].$$

Now we'll compute this quantity exactly, starting with the application of an identity for non-negative random variables,

$$\mathrm{E}_{\mathbf{b}\in\beta}[(1-\alpha)^n] = \int_0^1 \mathrm{Prob}_{\mathbf{b}\in\beta}\left[(1-\alpha)^n \geq x\right] dx.$$

Now, by (3), the probability in the above integral is exactly $(1-x^{1/n})^{m-1}$. Making the change of variable $y = x^{1/n}$, and repeating integration by parts, this yields,

$$\begin{aligned}
\int_0^1 (1-x^{1/n})^{m-1} dx &= \\
n\int_0^1 y^{n-1}(1-y)^{m-1} dy &= \\
n\left(\frac{y^n(1-y)^{m-1}}{n}\bigg|_0^1 + \frac{m-1}{n}\int_0^1 y^n(1-y)^{m-2} dy\right) &= \\
n\left(\frac{m-1}{n}\int_0^1 y^n(1-y)^{m-2} dy\right) &= \\
n\left(\frac{m-1}{n}\,\frac{m-2}{n+1}\int_0^1 y^{n+1}(1-y)^{m-3} dy\right) &= \\
\ldots &= \\
n\left(\frac{(m-1)!(n-1)!}{(n+m-2)!}\int_0^1 y^{n+m-2} dy\right) &= \\
\frac{1}{\binom{n+m-1}{m-1}}&.
\end{aligned}$$

□

3.1. *Randomized Approximation*

The implementation presented in (Cover and Ordentlich, 1996a) has space and time requirements which grow like n^{m-1}. However, the universal algorithm is simply a weighted average of all CRPs. This suggests an obvious randomized approximation. First choose N portfolios uniformly at random. Then invest a $1/N$ fraction of the money in each of the N CRPs and let it sit within these CRPs. (Do not transfer between them.) If the best constant rebalanced portfolio achieves a wealth R times as large as the universal algorithm, Chebyshev's inequality guarantees that, using $N = (R-1)/\epsilon^2\delta$ random portfolios, with probability at least $1-\delta$, the approximation achieves a wealth at least $1-\epsilon$ times as large as the universal algorithm. For a given market with no commission, one can determine, in hindsight, the optimal CRP (Helmbold et al., 1995) and then estimate R. In the worst case R grows like n^{m-1}. However, experiments on stock market data, presented in (Cover, 1991; Helmbold et al., 1996), all have a ratio $R < 2$ for various combinations of two stocks, making this approach especially efficient.

This same observation can be used to implement the Dirichelet$(1/2, \ldots, 1/2)$ universal algorithm (Cover and Ordentlich, 1996a). In this case, it is an improvement because R grows, at worst, like $n^{(m-1)/2}$.

4. Analysis With Commission

In this section, we introduce a slight modification of UNIVERSAL, UNIVERSAL$_c$, which is competitive in the presence of a fixed commission $0 \leq c \leq 1$. The importance of this section is not as much in introducing a new algorithm as it is in showing that a trivial extension of UNIVERSAL is theoretically competitive in the presence of commission.

At the start of the i^{th} period, UNIVERSAL computes a weighted average of CRPs, where the weight of a particular CRP$_{\mathbf{b}}$ is proportional to the wealth it has accumulated during the first $i-1$ periods, $S_{i-1}(\mathbf{b}, \mathbf{x}^{i-1})$. UNIVERSAL$_c$ is similar. At the start of the i^{th} period, UNIVERSAL$_c$ computes a weighted average of CRPs, where the weight of a particular CRP$_{\mathbf{b}}$ is proportional to the wealth it has accumulated during the first $i-1$ periods *including commission costs*, $S^c_{i-1}(\mathbf{b}, \mathbf{x}^{i-1})$. Formally,

Definition 2. (UNIVERSAL$_c$) The universal portfolio with commission algorithm at time i is specified by

$$\hat{\mathbf{b}}^c_i = \frac{\int_\beta \mathbf{b} S^c_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}{\int_\beta S^c_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}, \; i = 1, 2, \ldots$$

with μ equal to the uniform distribution over portfolios $\mathbf{b}$.

Notice here we are defining the distribution of wealth held each day. The exact amount is determined by rebalancing *optimally* with respect to transaction costs. Notice that the universal algorithm maintains, each period, the same distribution of wealth as another algorithm which splits its money evenly among N independent random CRP investors (and does not transfer money between them), in the limit as N goes to infinity. Like the zero-commission case, this other algorithm achieves the expected wealth of a random CRP. However, the universal algorithm may actually do better because of lower commission costs due to offsetting trades (in the same way that a collection of investors operating together can save on commission by occasionally trading amongst themselves). This is an application of the third property given in Section 2.1.

THEOREM 2 *In the presence of commission $0 \leq c \leq 1$,*

$$\frac{\textit{wealth of UNIVERSAL}_c}{\textit{wealth of best CRP}} \geq \binom{(1+c)n+m-1}{m-1}^{-1} \geq \frac{1}{((1+c)n+1)^{m-1}},$$

for all markets with m stocks and n periods.

Proof. We need only modify equations (1) and (2) of the previous proof. As discussed above, equation (1) changes from an equality to an inequality,

$$\text{wealth of UNIVERSAL}_c \geq \mathrm{E}_{\mathbf{b}\in\beta}\left[\text{wealth of CRP}_{\mathbf{b}}\right],$$

which only helps.

Based on the three properties given in Section 2.1, if $b_j \geq (1-\alpha)b_j^*$ for all j, then

$$\frac{(\text{single-period profit of CRP}_{\mathbf{b}})}{(\text{single-period profit of CRP}_{\mathbf{b}^*})} \geq (1-\alpha)(1-c\alpha). \tag{4}$$

To see this, note first that $\text{CRP}_{\mathbf{b}}$ starts with at least $(1-\alpha)$ of $\text{CRP}_{\mathbf{b}^*}$'s wealth, both pay some cost to rebalance to $\mathbf{b}^*$, and $\text{CRP}_{\mathbf{b}}$ pays at most an additional $c\alpha$ fraction to rebalance from $\mathbf{b}^*$ to $\text{CRP}_{\mathbf{b}}$. In terms of our stated properties, the third property implies that $\text{CRP}_{\mathbf{b}}$ earns at least as much as an investor who only has $(1-\alpha)$ times as much money, but begins with it distributed according to $\mathbf{b}^*$ and ends with it distributed according to $\mathbf{b}$. In other words, the extra stock beyond $(1-\alpha)b_j^*$ cannot hurt. Secondly, note that this investor pays no more in commission than a naive investor who first rebalances to $\mathbf{b}^*$ and then to $\mathbf{b}$, by the first property. Finally, the second property implies that the commission cost for changing from $\mathbf{b}^*$ to $\mathbf{b}$ is no more than $c\alpha$, since it involves moving an α fraction of your portfolio. Thus, $\text{CRP}_{\mathbf{b}}$ does no worse than $(1-\alpha)(1-c\alpha)$ times an investor who starts and ends the period with distribution $\mathbf{b}^*$. This is (4).

It is easy to establish that $1-c\alpha \geq (1-\alpha)^c$ for $0 \leq c, \alpha \leq 1$ (see, e.g. Lemma 3.4.1 of (Littlestone, 1989)). Combining this with (4),

$$\text{wealth of CRP}_{\mathbf{b}} \geq (1-\alpha)^{(1+c)n}(\text{wealth of CRP}_{\mathbf{b}^*}).$$

Since (2) is the only use of n in the previous proof, we can replace n by $(1+c)n$ in the final guarantee. □

This algorithm can be implemented in the same randomized way as UNIVERSAL.

5. Semi-constant-rebalanced Portfolios

In (Helmbold et al., 1996), the idea of a semi-constant-rebalanced portfolio (SCRP) is proposed as a good strategy in the presence of transaction costs. An SCRP is a portfolio which may be rebalanced on any subset of the periods. For instance, one may prefer not to rebalance if the transaction costs outweigh the benefits of rebalancing. We show that no strategy can guarantee the exponential growth rate of the best SCRP in hindsight, even without commission.

As in (Cover and Ordentlich, 1996b), consider the set of all market sequences of length n over two stocks, where each period one of the stock relatives is 1 and the other is ϵ for some small $\epsilon \geq 0$,

$$\mathcal{K} = \{\mathbf{x}^n : \mathbf{x}_i = (1, \epsilon) \text{ or } \mathbf{x}_i = (\epsilon, 1), \text{ for all } \ i \leq n\}.$$

In these markets, each period one of the stocks crashes and the other stays the same.

Clearly, if we choose a random element of $\mathcal{K}$, every non-anticipating investment strategy (a strategy which has no knowledge of the future) will achieve an expected wealth of $(\frac{1+\epsilon}{2})^n$.

However, a SCRP chosen in hindsight can do much better. A good SCRP$_{(\frac{1}{2},\frac{1}{2})}$ strategy would initially divide its money into two parts, and then rebalance only when the market is about to switch (when $\mathbf{x}_i \neq \mathbf{x}_{i+1}$). If the market switches k times, this strategy achieves wealth $\geq 1/2^{k+1}$. Thus, if we choose a random market from $\mathcal{K}$, in hindsight we can make at least

$$(1/2)^n \sum_{k=0}^{n-1} 2 \binom{n-1}{k} (1/2)^{k+1} = (1/2)^n (1 + 1/2)^{n-1}.$$

Thus the hindsight strategy's expected performance is at least $(\frac{1}{1+\epsilon})^n (3/2)^{n-1}$ times as large as the expected performance of any non-anticipating strategy.

6. Experiments with NYSE data

In this section, we present results that show examples of the performance of the universal algorithm as commission is varied, on 22 years of NYSE data. For more thorough experiments, see (Cover, 1991; Helmbold et al., 1996). We tried the universal algorithm with various transaction costs. As was the case without commission, we found that with commission it performed well on some sets of stocks and not as well on others. The performance, of course, was better with smaller transaction costs. Rebalancing less often was very beneficial, especially with larger commissions. We chose to rebalance monthly. In practice, one way to select how often to rebalance is to choose the period length which gives the best result on past data. As an aside, the EG(η) algorithm of (Helmbold et al., 1996) usually outperformed the universal algorithm, even though there are no theoretical guarantees of its performance with commission.

Figure 1 shows the wealths achieved by the universal algorithm on two different sets of stocks, where rebalancing occurs monthly. We compare this to the best CRP and to the average value of the stocks. With a 2% commission charged on *both* purchases and sales, we did not get the amazing gains in wealth that were

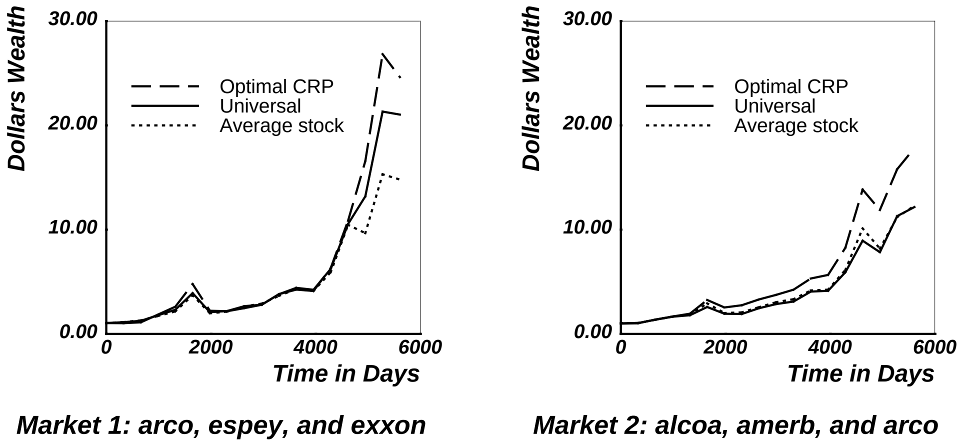

Figure 1. Performance of universal on two different sets of stocks. Compare with best CRP in hindsight and average stock price. 2% commission charged on all transactions, rebalancing monthly. This corresponds to $c = 0.04$ in our model.

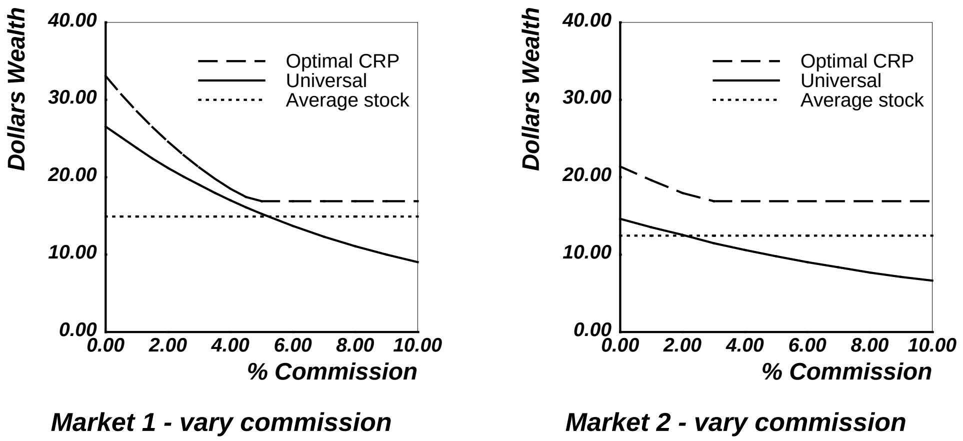

Figure 2. Performance of universal with varying commission rates. Commission charged on all transactions. Notice that with large enough commissions, we would be better off rebalancing less often than monthly.

found in the no-commission setting, but in the first set of stocks we were still doing significantly better than average.

Figure 2 shows the same stocks as we vary the commission costs. Since the strategy of investing in the best stock is a form of a CRP, the best CRP always does at least as well as the best stock and thus, of course, the average stock. And indeed, once the commission costs get prohibitively large, the best CRP is exactly the no-commission strategy of investing in a single stock. That is why the best CRP's performance plateaus. The universal algorithm, on the other hand, initially divides its money among all CRPs. As the commission increases, many of these CRPs do poorly. Over time, the investments in these worse CRPs become less valuable. As a result, the universal algorithm thus has an increasing fraction of its total holdings in CRPs that rebalance a smaller percentage of their holdings.

7. Application to Statistical Language Modeling

Analogies between universal compression and the UNIVERSAL algorithm have been made by Cover (Cover, 1996). The EG(η) portfolio algorithm of (Helmbold et al., 1996) has also been used for finding the best mixture of predictive models (Kivanen and Warmuth, 1994; Helmbold et al., 1995). We present the analogy between stocks and language models which shows how the UNIVERSAL algorithm can be applied to language models and other predictors. This analogy is more general than Kelly's racehorse analogy (Kelly, 1956) because it covers price relatives and probabilities other than just $\{0, 1\}$.

A statistical language model is a probability distribution over sequences of words. A language model is generally represented as a conditional probability distribution for the next word to be seen, given the previous words, i.e.,

$$P(w_i|h_{i-1}), h_{i-1} = w_1, w_2, \ldots, w_{i-1}.$$

The most common way to combine various language models is to linearly interpolate them. A mixture of three language models, for example a unigram model P_1, a bigram model P_2, and a trigram model P_3 might be,

$$P(w|h_i) = \lambda_1 P_1(w|h_i) + \lambda_2 P_2(w|h_i) + \lambda_3 P_3(w|h_i),$$

$$\text{where } \lambda_1 + \lambda_2 + \lambda_3 = 1, \lambda_i \geq 0.$$

This is similar to a CRP. Consider an analogy between language models and stocks. A price relative x_{ij} corresponds to a conditional probability $P_j(w_i|h_{i-1})$ and a portfolio $\mathbf{b}$ corresponds to a daily mixture of language models. In other words, feed the probabilities into a portfolio algorithm as if they were price relatives, and predict the next word as a linear combination of the predictions of the models where the coefficients are the weights recommended by the portfolio algorithm. Then we have the nice property that the probability that a combined language model algorithm assigns to a sequence $h_n = w_1, \ldots, w_n$ is simply the value of the holdings of the corresponding portfolio algorithm.

Now, the log of the probability assigned to the observed sequence is a common measure of a language model's performance. This leads us to the following algorithm. Given language models $P_1, P_2, \ldots, P_m$ and a sequence of words $h_n = w_1, \ldots, w_n$, if UNIVERSAL invests its money based on portfolios $\hat{\mathbf{b}}_i$ when observing price relatives $x_{ij} = P_j(w_i|h_{i-1})$, then the universal language model P predicts

$$P(w_i|h_{i-1}) = \sum_{j=1}^{m} \hat{b}_{ij} P_j(w_i|h_{i-1}).$$

THEOREM 3 *For any sequence of n words $h_n = w_1, \ldots, w_n$ and any m language models, the log probability assigned to h_n by the universal algorithm P is at least the log probability assigned to h_n by the best mixture, minus $(m-1)\log n$.*

This follows directly from the observation that the wealth of a portfolio algorithm is exactly the probability it assigns to the word sequence. It is impressive because it shows how these mixture parameters can be "learned" on the fly with on-line guarantees. Furthermore, the amortized cost of $(m-1)\log n/n$ bits per word is a small overhead as n gets large. It also helps to explain a relationship between portfolios and weighted-average-type algorithms for making predictions from expert advice (Cesa-Bianchi et al., 1993; DeSantis et al., 1988; Foster and Vohra, 1993; Haussler et al., 1994; Kivanen and Warmuth, 1994; Vovk, 1990; Vovk, 1995).

Acknowledgments

We thank the anonymous referees for their helpful comments. This research was supported in part by NSF National Young Investigator grant CCR-9357793, a grant from the AT&T Foundation, and an NSF Graduate Fellowship.

References

N. Cesa-Bianchi, Y. Freund, D.P. Helmbold, D. Haussler, R.E. Schapire, and M.K. Warmuth. How to use expert advice. In *Annual ACM Symposium on Theory of Computing*, pages 382–391, 1993.

T.M. Cover. Universal portfolios. *Math. Finance*, 1(1):1-29, January 1991.

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

T.M. Cover and E. Ordentlich. The Cost of Achieving the Best Portfolio in Hindsight. Department of Statistics Technical Report NSF-90, Stanford University, 1996.

T.M. Cover. Universal data compression and portfolio selection. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, pages 534-538, Oct 1996.

M.H.A. Davis and A.R. Norman. Portfolio Selection with Transaction Costs. *Mathematics of Operations Research*, 15(4), November 1990.

A. DeSantis, G. Markowsky, and M. Wegman. Learning probabilistic prediction functions. In *Proceedings of the 29th IEEE Symposium on Foundations of Computer Science*, pages 110–119, Oct 1988.

D.P. Foster and R.V. Vohra. A randomization rule for selecting forecasts. *Operations Research*, 41:704–709, 1993.

D.P. Foster and R.V. Vohra. Regret in the On-line Decision Problem. *'Something for Nothing' Workshop.*, May 1995.

D. Haussler, J. Kivinen, and M. Warmuth. Tight worst-case loss bounds for predicting with expert advice. Technical Report UCSC-CRL-94-36, University of California, Santa Cruz, November 1994.

D. Helmbold, R. Schapire, Y. Singer, and M. Warmuth. A comparison of new and old algorithms for a mixture estimation problem. *Proceedings of the Eighth Annual Workshop on Computational Learning Theory*, pages 69-78, 1995.

D. Helmbold, R. Schapire, Y. Singer, and M. Warmuth. On-line portfolio selection using multiplicative updates. *Machine Learning: Proceedings of the Thirteenth International Conference*, 1996.

J. Kelly. A new interpretation of information rate. Bell Sys. Tech. Journal, 35: 917-926, 1956.

J. Kivinen and M Warmuth. Exponentiated gradient versus gradient descent for linear predictors. Technical Report UCSC-CRL-94-16, University of California, Santa Cruz, June 1994.

N. Littlestone. Mistake Bounds and Logarithmic Linear-threshold Learning Algorithms. Technical Report UCSC-CRL-89-11, University of California, Santa Cruz 1989.

E. Ordentlich. and T.M. Cover. Online portfolio selection. In *Proceedings of the 9th Annual Conference on Computational Learning Theory*, Desenzano del Garda, Italy, 1996.

V. Vovk. Aggregating strategies. In *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–383. Morgan Kaufmann, 1990.

V. G. Vovk. A game of prediction with expert advice. In *Proceedings of the 8th Annual Conference on Computational Learning Theory*, pages 51–60. ACM Press, New York, NY, 1995.