Journal of Machine Learning Research 1 (2002) 1-48 Submitted 8/01; Published 11/01

Efficient Algorithms for Universal Portfolios

Adam Kalai AKALAI@MIT.EDU
Department of Mathematics
Massachusetts Institute of Technology
Cambridge MA, 02139

Santosh Vempala VEMPALA@MATH.MIT.EDU
Department of Mathematics
Massachusetts Institute of Technology
Cambridge MA, 02139

Editor: Philip M. Long

Abstract

A constant rebalanced portfolio is an investment strategy that keeps the same distribution of wealth among a set of stocks from day to day. There has been much work on Cover's Universal algorithm, which is competitive with the best constant rebalanced portfolio determined in hindsight (3, 9, 2, 8, 16, 4, 5, 6). While this algorithm has good performance guarantees, all known implementations are exponential in the number of stocks, restricting the number of stocks used in experiments (9, 4, 2, 5, 6). We present an efficient implementation of the Universal algorithm that is based on non-uniform random walks that are rapidly mixing (1, 14, 7). This same implementation also works for non-financial applications of the Universal algorithm, such as data compression (6) and language modeling (11).

1. Introduction

A constant rebalanced portfolio (CRP) is an investment strategy which keeps the same distribution of wealth among a set of stocks from day to day. That is, the proportion of total wealth in a given stock is the same at the beginning of each day. Recently there has been work on on-line investment strategies which are competitive with the best CRP determined in hindsight (3, 9, 2, 8, 16, 4, 5, 6). 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 (9, 5). 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 day it trades stock so that it has an equal worth in each stock. On alternate days 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 days.

The main contribution of this paper is an efficient implementation of Cover's UNIVERSAL algorithm for portfolios (3). Cover and Ordentlich (4) show that, in a market with n

stocks, over t days,

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

By performance, we mean the return per dollar on an investment. The above ratio is a decreasing function of t. However, the average per-day ratio, $(1/(t+1)^{n-1})^{1/t}$, increases to 1 as t increases without bound. For example, if the best CRP makes 1.5 times as much as we do each day 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. In this paper, we do not consider the Dirichlet$(1/2, \ldots, 1/2)$ UNIVERSAL (4) which has the better guaranteed ratio of $2\sqrt{1/(t+1)^{n-1}}$.

All previous implementations of Cover's algorithm are exponential in the number of stocks with worst-case run times of $\Theta(t^{n-1})$. In some sense, Cover's algorithm divides its money evenly among all CRPs. Unfortunately, for some market sequences, the number of CRPs which perform near optimally can be as small as $1/\Theta(t^{n-1})$. In these cases, Blum and Kalai's randomized approximation based on sampling from the uniform distribution, requires $\Theta(t^{n-1})$ samples to perform nearly as well as UNIVERSAL, with high probability.

We show that by sampling portfolios from a non-uniform distribution, only polynomially many samples are required to have a high probability of performing nearly as well as UNIVERSAL. This non-uniform sampling can be achieved by random walks on the simplex of portfolios.

2. Notation and Definitions

A price relative for a given stock is the nonnegative ratio of closing price to opening price during a given day. If the market has n stocks and trading takes place during T days, then the market's performance can be expressed by T price relative vectors, $(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_T), \vec{x}_i \in \Re^n_+$, where x_i^j is the nonnegative price relative of the jth stock for the ith day.

A portfolio is simply a distribution of wealth among the stocks. As such, it is really an $n-1$ dimensional quantity where the last component can be determined from the other $n-1$. We view the set of portfolios as the $(n-1)$-dimensional simplex,

$$\Delta = \{\vec{b} \in \Re^{n-1} | \sum_{j=1}^{n-1} b^j \leq 1 \wedge b^j \geq 0\}.$$

We also consider a portfolio also as an n-dimensional vector, where

$$b^n = 1 - \sum_{1}^{n-1} b^j.$$

This abuse of notation allows us to view the portfolio as an n component vector and the set of portfolios as an $(n-1)$-dimensional set, wherever convenient. This is especially valuable, since the random walk result we will be using (7) is for a full-dimensional set such as Δ.

The CRP investment strategy for a particular portfolio $\vec{b}$, $\text{CRP}_{\vec{b}}$, redistributes its wealth at the end of each day so that the proportion of money in the jth stock is b^j. An investment using a portfolio $\vec{b}$ during a day with price relatives $\vec{x}$ increases one's wealth by a factor of $\vec{b} \cdot \vec{x} = \sum_1^n b^j x^j$. Therefore, over t days, the wealth achieved by $\text{CRP}_{\vec{b}}$ is,

$$P_t(\vec{b}) = \prod_{i=1}^{t} \vec{b} \cdot \vec{x}_i. \tag{1}$$

Finally, we let μ be the uniform distribution on Δ.

3. Universal portfolios

Before we define the universal portfolio, consider the problem of being 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 n stocks and let it sit. You will always have at least $\frac{1}{n}$ 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 at time t has portfolio $\vec{u}_t$, which for stock j is, on the first day $u_0^j = 1/n$, and on the end of the tth day,*

$$u_t^j = \frac{\int_\Delta v^j P_t(\vec{v}) d\mu(\vec{v})}{\int_\Delta P_t(\vec{v}) d\mu(\vec{v})}, \quad i = 1, 2, \ldots$$

(Recall that μ is the uniform distribution over the $(n-1)$-dimensional simplex of portfolios, Δ.)

This is the form in which Cover defines the algorithm. He also notes (4) that UNIVERSAL achieves the average performance of all CRPs, i.e.,

$$\begin{aligned} \text{Performance of UNIVERSAL} &= \prod_{t=1}^{T} \vec{u}_{t-1} \cdot \vec{x}_t \\ &= \int_\Delta P_T(\vec{v}) d\mu(\vec{v}) \end{aligned}$$

4. Efficient Approximation

Unfortunately, the straightforward method of evaluating the integral in the definition of UNIVERSAL takes time exponential in the number of stocks. Since UNIVERSAL is really just an average of CRP's, it is natural to approximate the portfolio by sampling (2). Simply imagine dividing the wealth into many random portfolios and see what distribution of wealth one would have. In particular, you would just take a weighted average of the portfolios you've chosen, with weights proportional to their performance. The problem is that there may be a very small set of portfolios that did well while most portfolios did very poorly. In order to get a good sample, one would need to get a draw from this set, which can require $\Omega(t^{n-1})$ samples in the worst case.

The key to our algorithm is sampling according to a biased distribution. Instead of sampling according to μ, the uniform distribution on Δ, we sample according to ρ_t, which weights portfolios in proportion to their performance, i.e.,

$$d\rho_t(\vec{b}) = \frac{P_t(\vec{b})d\mu(\vec{b})}{\int_\Delta P_t(\vec{v})d\mu(\vec{v})}$$

UNIVERSAL can be thought of as computing each component of the portfolio by taking the expectation of draws from ρ_t, i.e.,

$$u_t^j = \int_\Delta v^j d\rho_t(\vec{v}) = \mathrm{E}_{\vec{v}\in\rho_t}\left[v^j\right] \tag{2}$$

We use existing random walk theorems to show one can sample from ρ_t in time polynomial in t and n. The current best provable bounds are for sampling on a discretization of the simplex, although many other random walks might also work. For the purpose of this random walk, we will need the following modification Q_t of the wealth function P_t:

$$Q_t(\vec{b}) = P_t(\vec{b}) \min\left\{\exp\left(\frac{b^n - 2\delta_0}{n\delta}\right), 1\right\}. \tag{3}$$

The function Q_t is a slightly "damped" version of P_t: it is equal to P_t on a slightly smaller simplex, namely the set $\{\vec{b} \in \Delta | b^n \geq 2\delta_0\}$ and falls off rapidly outside this set. We introduce it for technical reasons; it is quite possible that the algorithm below works with P_t in place of Q_t (see Section 7). The parameter δ_0 will be specified in Section 5.3. We also only sample grid points whose coordinates are at least δ_0 each (again for technical reasons). We refer to the algorithm as R-UNIVERSAL because it's a randomized approximation.

R-UNIVERSAL(δ_0, δ, m, S)
// δ_0 = minimum coordinate
// δ = spacing of grid
// m = number of samples
// S = number of steps in random walk

On each day, we take the average of m samples obtained as follows:

1. Start each one at the point $\vec{r} = (\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n})$.
2. From each of them, take S steps of the following random walk:
   (a) Choose $1 \le j \le n-1$ at random
   // We'll try to increment or decrement r^j with equal probability.
   (b) Choose $X \in \{-1, +1\}$ randomly. If $\delta_0 \le r^i + X\delta$ and $\delta_0 \le r^n - X\delta$,
      i. Let $x := Q_t(r^1, r^2, \ldots, r^n)$.
      ii. Let $y := Q_t(r^1, r^2, \ldots, r^j + X\delta, \ldots, r^{n-1}, r^n - X\delta)$.
      iii. With probability $\mathrm{Min}(1, x/y)$,
         - $r^j := r^j + X\delta$
         - $r^n := r^n - X\delta$

The function Q_t defined by (3), can be evaluated in time $O(nt)$. Thus the runtime of the algorithm on day t is $O(mSnt)$. In the analysis section, we will show that m and S can be chosen to be polynomial in n and t.

5. Analysis

Here we show that, with non-uniform sampling, the algorithm approximates the portfolio efficiently. With high probability $(1-\eta)$, we can achieve performance of at least $(1-\epsilon)$ times the performance of UNIVERSAL, with S (number of steps) polynomial in $1/\epsilon$, $\log(1/\eta)$, n (the number of stocks), and T (the number of days). We will show the following:

Theorem 2 *There is a constant A such that, for all ϵ, η, δ, and δ_0 with*

$$\delta_0 \le \frac{\epsilon}{8nT(n+T)^2}, \quad \delta \log \frac{1}{\delta} = \frac{\epsilon\delta_0}{A(n+T)^2},$$

$m \ge 64T^2(n+T)\ln(nT/\eta)/\epsilon^2$ samples on each day, and $S \ge \frac{An}{\delta^2} \log \frac{n+T}{\epsilon\delta}$ steps of the random walk per sample, R-UNIVERSAL performs at least $(1-\epsilon)$ times as well as UNIVERSAL with probability at least $1-\eta$.

As will become clear in the next two sections, the random walk, which is the basis of our sampling algorithm, tends to a stationary distribution proportional to the function Q_t. The key observation that leads to a polynomial-time implementation is the fact that Q_t is a *log-concave* function (as is P_t). Such functions can be sampled in polynomial-time, albeit with several technical restrictions. It is the latter that force us to introduce the parameter

δ_0 and also to use Q_t instead of just P_t. Using Frieze and Kannan's theorem (our Theorem 3), we show that the random walk quickly converges to its stationary distribution (Theorem 11) in Section 6.

It then remains to show that the approximation provided by the random walk is sufficient for the algorithm. This is done in Section 5.3. The above theorem follows from Theorems 4 and 11. Let's begin by describing what has been analyzed.

5.1 Analysis Overview

To compute $\vec{u_t}$, We would like to compute the expectation of random points drawn according to the distribution P_t. As we will show, P_t is a log-concave function. This suggests using existing technology based on random walks for sampling from log-concave distributions (7) over an arbitrary convex set K. However, their algorithms work by discretizing space into cubes and sampling from cube centers. One important parameter of their algorithm is the total probability of cubes that are not completely contained in K. In our application, the shape of our set is a simplex, so it is impossible to discretize it in such a way that there are few such border cubes. However, we will show that the mean of P_t is not near the border of the simplex. Then, we use a damped function Q_t to reduce the probability of border cubes, in such a way that the mean of Q_t is near to the mean of P_t and Q_t is log-concave while sufficiently small near the borders of the simplex.

5.2 Summary of Frieze-Kannan '97 (7)

Suppose you have some nonnegative log-concave function f on $\Re^n$, meaning simply that $\log f$ is a concave function. The goal is to sample from f restricted to some convex set $K \subset \Re^n$. They first divide the space into cubes of side length δ. The spacing δ has to be

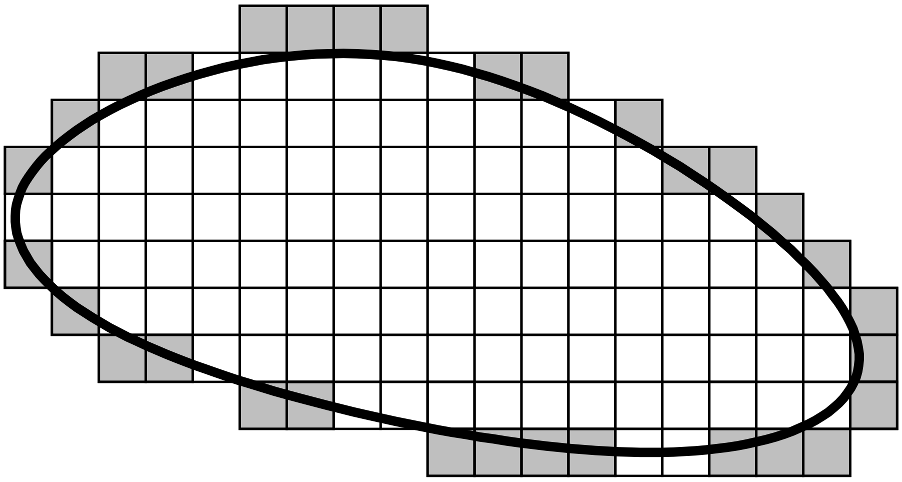

Figure 1: A convex set K and the set of all cubes that intersect K. The centers of these cubes form the set C, while $C_{1/2}$ is the set of centers of the cubes that have more than half of their volume in the set. These cubes are the unshaded cubes.

small enough so that the function f varies by a factor of at most a constant factor within a cube. Then they perform all future operations on the centers of the cubes which intersect K, a set called C, as shown in the figure. Let $C(\vec{x})$ be the cube centered at $\vec{x} \in C$ of width δ.

Consider the following walk.

- Start at some arbitrary cube center $\vec{x} \in C$.
- Choose a random coordinate and add or subtract δ in that coordinate with equal probability, to get another cube center $\vec{y}$.
- If $\vec{y} \in C$ (haven't stepped out of the set), then move to $\vec{y}$ with probability $\min(1, f(\vec{y})/f(\vec{x}))$.

They first observe that the stationary distribution of this walk is

$$\pi(\vec{x}) = \frac{f(\vec{x})}{\sum_{\vec{x}' \in C} f(\vec{x}')} \qquad \forall \vec{x} \in C$$

This follows from the fact that the walk is time reversible. That is, when in the above distribution, on a single step, the probability of being at $\vec{x}$ and moving to $\vec{y}$ is equal to the probability of being at $\vec{y}$ and going to $\vec{x}$, i.e.,

$$\begin{aligned} \frac{f(\vec{x})}{\sum_{\vec{x}' \in C} f(\vec{x}')} \frac{1}{2n} \min\left(1, \frac{f(\vec{y})}{f(\vec{x})}\right) &= \frac{1}{\sum_{\vec{x}' \in C} f(\vec{x}')} \frac{1}{2n} \min(f(\vec{y}), f(\vec{x})) \\ &= \frac{f(\vec{y})}{\sum_{\vec{x}' \in C} f(\vec{x}')} \frac{1}{2n} \min\left(1, \frac{f(\vec{x})}{f(\vec{y})}\right). \end{aligned}$$

Now, there are several parameters for the analysis of this walk. The diameter of K is d. The dimensionality is n. The stationary distribution is π, and $\pi_* = \min \pi$. Let p_s be the distribution obtained after s steps of the walk. For any $0 \leq \theta \leq 1$, let $C_\theta = \{\vec{x} \in C | \mathrm{vol}(C(\vec{x}) \cap K) \geq \theta \delta^n\}$ and $\pi_\theta = \sum_{\vec{x} \notin C_\theta} \pi(\vec{x})$. Finally, let $M = \max_{\vec{x} \in C} \frac{p_0(\vec{x})}{\pi(\vec{x})} \log \frac{p_0(\vec{x})}{\pi(\vec{x})}$.

Theorem 3 *(Theorem 1 of (7)) Assume $d \geq \delta n^{1/2}$. Then there is an absolute constant $\gamma > 0$ such that,*

$$2\left(\sum_{\vec{x} \in C} |p_s(\vec{x}) - \pi(\vec{x})|\right)^2 \leq e^{-\frac{\gamma s \delta^2}{n d^2}} \log\left(\frac{1}{\pi_*}\right) + \frac{M \pi_{1/2} n d^2}{\gamma \delta^2}.$$

In our problem, the dimensionality is actually $n - 1$ if there are n stocks, but this can be absorbed by γ. For us M can be very large, but we'll choose parameters that make $\pi_{1/2}$ small enough to compensate.

5.3 Approximation suffices

To apply Theorem 3, we first need to ensure that π varies by at most a constant factor within the δ-cube around any grid point. This is the technical reason why we use a slightly smaller simplex Δ'. In the full simplex Δ, the function P_t (and hence ρ_t) can vary a lot within a small distance. In this section, we show that this approximation is good enough for the algorithm.

The random walk actually samples from the following simplex:

$$\Delta' = \{\vec{v} \in \Delta | v^j \geq \delta_0 \text{ for } j = 1, 2, \ldots, n\}.$$

Let C' denote the set of cube centers contained in Δ'. The stationary distribution of the walk is proportional to Q_t; let it be called π_t. Note that for any $\vec{x} \in C'$,

$$\pi_t(\vec{x}) = \frac{Q_t(\vec{x})}{\sum_{\vec{y} \in C'} Q_t(\vec{y})}.$$

Let the actual distribution obtained on them (after some number of steps of the random walk) be $\tilde{\pi}_t$. The main theorem of this section shows that this suffices as long as $\tilde{\pi}_t$ is sufficiently close to π_t.

Theorem 4 *Suppose that* $\delta_0 \leq \frac{\epsilon}{8nT(n+T)^2}$, $\delta \leq \frac{\epsilon\delta_0}{8T(n+T)}$, *and we can sample grid points in the simplex* Δ' *according to a distribution* $\tilde{\pi}_t$ *where*

$$\sum_{x \in C'} |\tilde{\pi}_t(x) - \pi_t(x)| \leq \frac{\epsilon}{4T(n+T)}.$$

Then with $m \geq 64T^2(n+T)\ln(nT/\eta)/\epsilon^2$, *the algorithm R-UNIVERSAL performs at least* $(1-\epsilon)$ *as well as UNIVERSAL with probability at least* $1-\eta$.

To prove the theorem, we will need several lemmas. First, we observe that subsets of the simplex of sufficient size will have large volume under ρ_t.

Lemma 5 *Let* β *be a subset of the simplex, shrunken by a factor of* $1-z$ *for* $0 \leq z \leq 1$, *i.e., for some* $\vec{w} \in \Delta$,

$$\beta = z\vec{w} + (1-z)\Delta = \{\vec{v} \in \Delta | \vec{v} = z\vec{w} + (1-z)\vec{v'}, \vec{v'} \in \Delta\}. \tag{4}$$

The probability that a random portfolio selected in proportion to its performance is in β *is at least,*

$$\rho_t(\beta) \geq (1-z)^{t+n-1}$$

Proof. Geometrically, β is also a simplex, translated by $\vec{w}$ and shrunken by a factor of $(1-z)$. Since the simplex has dimension $n-1$, this set has volume at least,

$$\mu(\beta) = (1-z)^{n-1},$$

under the uniform measure μ on the simplex.

As indicated in (4), there is a simple bijection between $\vec{v} \in \beta$ and $\vec{v'} \in \Delta$. Furthermore, on any day the performance of $\vec{v}$ must be at least $(1-z)$ times as good as the performance of the corresponding $\vec{v'}$ since a $(1-z)$ fraction of its holdings are distributed exactly like $\vec{v'}$. Over t days, we see that

$$P_t(\vec{v}) \geq (1-z)^t P_t(\vec{v'})$$

Consequently, the performance of a uniformly random portfolio in β is at least $(1-z)^t$ as good as a uniformly random portfolio in Δ. Since a $(1-z)^{n-1}$ fraction of the portfolios are in β,

$$\rho_t(\beta) \geq (1-z)^t(1-z)^{n-1} = (1-z)^{t+n-1}$$

□

A corollary, which we will use later is,

Corollary 6 *For all $j \leq n$ and t,*

$$u_t^j \geq \frac{1}{n+t}$$

Proof. WLOG $j = 1$. Using the lemma, it is easy to see that $u_t^1 \geq 1/(e(n+t))$. This is because $u_t^1 = E_{\rho_t}[v^1]$ and the set of portfolios with $v^1 \geq 1/(n+t)$ has volume at least $1/e$, i.e.,

$$\rho_t\left(\frac{1}{n+t}(1,0,0,\ldots,0) + \left(1 - \frac{1}{n+t}\right)\Delta\right) \geq \left(1 - \frac{1}{n+t}\right)^{t+n-1} \geq \frac{1}{e}$$

To remove the $1/e$ factor, note that the expectation of a random variable $0 \leq X \leq 1$ is $E[X] = \int_0^1 \text{Prob}(X \geq z)dz$.

$$\begin{aligned}
u_t^1 &= E_{\vec{v}\in\rho_t}[v^1] \\
&= \int_0^1 \rho_t(\{\vec{v} \in \Delta | v^1 \geq z\})dz \\
&= \int_0^1 \rho_t(\{\vec{v} \in \Delta | \vec{v} = z(1,0,0,\ldots,0) + (1-z)\vec{v'}, \vec{v'} \in \Delta\})dz \\
&\geq \int_0^1 (1-z)^{t+n-1}dz \\
&= \frac{1}{n+t}
\end{aligned}$$

□

Lemma 7 *For any grid point $\vec{z}$ in C', and any point $\vec{v} \in C(z)$, we have*

$$(1-\delta/\delta_0)^T P_t(\vec{z}) \leq P_t(\vec{v}) \leq (1+\delta/\delta_0)^T P_t(\vec{z}).$$

Proof. Since $\vec{v} \in C(z)$ and $\vec{z} \in \Delta'$,

$$\begin{aligned}
v^j &\leq z^j + \delta \\
&\leq (1+\delta/\delta_0)z^j
\end{aligned}$$

(The above holds for $\delta/2$ as well.) Therefore, over t days,

$$\begin{aligned}
P_t(\vec{v}) &= \prod_{i=1}^t \vec{v}\cdot\vec{x_i} \\
&\leq \prod_{i=1}^t (1+\delta/\delta_0)\vec{z}\cdot\vec{x_i} \\
&= (1+\delta/\delta_0)^t P_t(\vec{z})
\end{aligned}$$

The RHS follows from $t \leq T$, and the other inequality is similar. $\square$

The main lemma of this section is the following. It says that the average over Δ' is close to the one computed by UNIVERSAL over Δ.

Lemma 8 *Let $\delta_0 \leq \frac{\epsilon}{8nT(n+T)^2}$, $\delta \leq \frac{\epsilon\delta_0}{8T(n+T)}$. Then for each stock j on any day t,*

$$E_{\pi_t}[v^j] \geq \left(1 - \frac{\epsilon}{2T}\right) u_t^j.$$

Proof. The universal portfolio is

$$\begin{aligned} u_j^t &= \int_\Delta v^j \, d\rho_t(\vec{v}) \\ &= \frac{\int_\Delta P_t(\vec{v}) v^j \, d\mu(\vec{v})}{\int_\Delta P_t(\vec{v}) d\mu(\vec{v})} \end{aligned} \tag{5}$$

We are sampling the function Q_t over cube centers contained in Δ' rather than the set Δ. What we have is,

$$\begin{aligned} E_{\pi_t}[v^j] &= \sum_{\vec{v} \in \Delta'} \pi_t(\vec{v}) v^j \\ &= \frac{\sum_{\vec{v} \in C'} Q_t(\vec{v}) v^j}{\sum_{\vec{v} \in C'} Q_t(\vec{v})}. \end{aligned} \tag{6}$$

So, in order to prove the lemma, we have to show that in (6), the numerator is large enough and the denominator is small enough compared to the quantities in (5).

Now, by Lemma 7, we can say that the integral over a cube is bounded by its value in the center $\vec{z}$, i.e.,

$$\begin{aligned} (1 - \delta/\delta_0)^T P_t(\vec{z}) &\leq \frac{1}{\delta^{n-1}} \int_{\vec{v} \in C(\vec{z})} P_t(\vec{v}) d\mu(\vec{v}), \text{ and} \\ (1 + \delta/\delta_0)^{T+1} P_t(\vec{z}) z^j &\geq \frac{1}{\delta^{n-1}} \int_{\vec{v} \in C(\vec{z})} P_t(\vec{v}) v^j \, d\mu(\vec{v}) \end{aligned}$$

Also notice that the union of all cubes with centers in Δ' is contained in Δ. To see this, take any grid point $\vec{z} \in C'$ and $\vec{v} \in C(\vec{z})$. Clearly the first $n-1$ coordinates of $\vec{v}$ are all at least $\delta_0 - \delta/2 > 0$ and the nth coordinate is,

$$\begin{aligned} v^n &= z^n + \sum_1^{n-1} z^j - v^j \\ &\geq \delta_0 - n\delta/2 \\ &\geq 0. \end{aligned}$$

So $\vec{v} \in \Delta$. Since $Q_t(\vec{v}) \leq P_t(\vec{v})$ for all points, the denominator in (6) is at most,

$$\sum_{\vec{v} \in C'} Q_t(\vec{v}) \leq \frac{1}{(1 - \delta/\delta_0)^T} \frac{1}{\delta^{n-1}} \int_{\vec{v} \in \Delta} P_t(\vec{v}) d\mu(\vec{v})$$

In order to lower bound the numerator of (6), we consider the set Δ'' which is entirely contained in the union of cubes with centers in Δ',

$$\Delta'' = \{\vec{v} \in \Delta' | v^n \geq 2\delta_0\}.$$

To see that Δ'' is contained in this set, take any point $\vec{v} \in \Delta''$ and let z be the center of the cube that contains it. Since $v^j \geq \delta_0$ for $1 \leq j \leq n-1$, we also have[1] $z^j \geq \delta_0 + \delta/2$. Finally, $v^n \geq 2\delta_0$ so

$$\begin{aligned} z^n &= v^n + \sum_1^{n-1} v^j - z^j \\ &\geq 2\delta_0 - n\delta/2 \\ &\geq \delta_0. \end{aligned}$$

This shows that $\vec{z} \in \Delta'$.

Since Δ'' is entirely contained in the union of cubes with centers C' and $P_t(\vec{v}) = Q_t(\vec{v})$ for $\vec{v} \in \Delta''$,

$$\sum_{\vec{v} \in C'} Q_t(\vec{v}) v^j \geq \frac{1}{(1+\delta/\delta_0)^T} \frac{1}{\delta^{n-1}} \int_{\vec{z} \in \Delta''} P_t(\vec{v}) v^j d\mu(\vec{v})$$

Putting these together with (6), we get

$$\begin{aligned} E_{\pi_t}[v^j] &= \frac{\sum_{\vec{v} \in C'} Q_t(\vec{v}) v^j}{\sum_{\vec{v} \in C'} Q_t(\vec{v})} \\ &\geq \left(\frac{1-\delta/\delta_0}{1+\delta/\delta_0}\right)^{T+1} \frac{\int_{\vec{v} \in \Delta''} P_t(\vec{v}) v^j d\mu(\vec{v})}{\int_{\vec{v} \in \Delta} P_t(\vec{v}) d\mu(\vec{v})} \\ &\geq (1 - 2\delta/\delta_0)^{T+1} \int_{\vec{v} \in \Delta''} v^j d\rho_t(\vec{v}) \\ &\geq (1 - 2(T+1)\delta/\delta_0) \int_{\vec{v} \in \Delta''} v^j d\rho_t(\vec{v}) \end{aligned}$$

We need to compare this to UNIVERSAL, which is an integral over all of Δ. Using Lemma 5 on the set,

$$\beta = (\delta_0, \delta_0, \ldots, \delta_0, 2\delta_0) + (1 - (n+1)\delta_0)\Delta,$$

we get,

$$\begin{aligned} \rho_t(\Delta'') &\geq (1 - (n+1)\delta_0)^{T+n-1} \\ &\geq 1 - (T+n-1)(n+1)\delta_0 \\ &\geq 1 - \frac{\epsilon}{4T(n+T)} \end{aligned}$$

Hence,

$$\begin{aligned} \int_{\Delta''} v^j d\rho_t(\vec{v}) &= \int_{\Delta} v^j d\rho_t(\vec{v}) - \int_{\Delta \setminus \Delta''} v^j d\rho_t(\vec{v}) \\ &\geq u_t^j - \rho_t(\Delta \setminus \Delta'') \\ &\geq u_t^j - \frac{\epsilon}{4T(n+T)} \geq (1 - \frac{\epsilon}{4T}) u_t^j. \end{aligned}$$

1. We have chosen δ_0 and placed the grid so that $(\delta_0, \delta_0, \ldots, \delta_0) \in \Re^{n-1}$ is a *corner* of a cube.

This finally implies that

$$\begin{aligned}
E_{\pi_t}[v^j] &\geq (1-2(T+1)\delta/\delta_0)(1-\frac{\epsilon}{4T})u_t^j \\
&\geq (1-2(T+1)\frac{\epsilon}{8T(n+T)})(1-\frac{\epsilon}{4T})u_t^j \\
&\geq (1-\frac{\epsilon}{4T})(1-\frac{\epsilon}{4T})u_t^j \\
&\geq (1-\frac{\epsilon}{2T})u_t^j.
\end{aligned}$$

□

Proof (of Theorem 4).

We will first show that on each day, the expected value of each stock j as computed by the algorithm is close to u_t^j.

$$\begin{aligned}
|E_{\tilde{\pi}_t}[v^j] - E_{\pi_t}[v^j]| &\leq \sum_{v\in C'} |\tilde{\pi}_t(v) - \pi_t(v)| v^j \\
&\leq \sum_{v\in C'} |\tilde{\pi}_t(v) - \pi_t(v)| \\
&\leq \frac{\epsilon}{4T(n+T)} \\
&\leq \frac{\epsilon}{4T} u_t^j.
\end{aligned}$$

Here we have used the assumption of the theorem that $\tilde{\pi}_t$ is close to π_t and then Corollary 6 which states that each u_t^j is at least $1/(n+t)$.

Next, using Lemma 8, we have that

$$\begin{aligned}
E_{\tilde{\pi}_t}[v^j] &\geq E_{\pi_t}[v^j] - \frac{\epsilon}{4T}u_t^j \\
&\geq (1-\frac{\epsilon}{2T})u_t^j - \frac{\epsilon}{4T}u_t^j \\
&\geq (1-\frac{3\epsilon}{4T})u_t^j.
\end{aligned}$$

Finally, we apply Chernoff bounds to show that with probability $1-\eta$, the value a_t^j of each stock j on each day t satisfies

$$a_t^j \geq (1-\frac{\epsilon}{T})u_t^j.$$

Then, on any individual day, the performance of the $\vec{a}_t$ is at least $(1-\epsilon/T)$ times as good as the performance of $\vec{u}_t$. Thus, over T days, our approximation's performance is at least $(1-\epsilon/T)^T \geq 1-\epsilon$ times the performance of UNIVERSAL.

The multiplicative Chernoff bound for approximating a random variable $0 \leq X \leq 1$, with mean $\bar{X}$, by the sum S of m independent draws is,

$$\Pr\left[S < (1-\alpha)\bar{X}m\right] \leq e^{-m\bar{X}\alpha^2/2}.$$

In our case, we are using m samples for each stock. In our case $\bar{X} \geq 1/2(n+T)$ and we want $\alpha = \epsilon/4T$. Since this must hold for nT different a_t^j's, it suffices for,

$$e^{-m\epsilon^2/(64T^2(n+T))} \leq \frac{\eta}{nT},$$

which holds for the number of samples m chosen in the theorem. □

6. The time per sample

In this section we show that the random walk quickly produces samples from a distribution $\tilde{\pi}(x)$ satisfying the requirements of Theorem 4. The random walk has stationary distribution proportional to Q_t. We begin with the observation that Q_t is log-concave.

Lemma 9 *The function $Q_t(\vec{b})$ is log-concave for nonnegative vectors.*

Proof. First, we prove that P_t is log-concave. This follows easily from the concavity of the log function.

$$\begin{aligned}
\log P_t\left(\frac{\vec{a}+\vec{b}}{2}\right) &= \log \prod_{i=1}^{t} \left(\frac{\vec{a}+\vec{b}}{2}\right) \cdot \vec{x}_i \\
&= \sum_{i=1}^{t} \log \frac{\vec{a}+\vec{b}}{2} \cdot \vec{x}_i \\
&= \sum_{i=1}^{t} \log \frac{\vec{a}\cdot\vec{x}_i + \vec{b}\cdot\vec{x}_i}{2} \\
&\geq \sum_{i=1}^{t} \frac{\log \vec{a}\cdot\vec{x}_i + \log \vec{b}\cdot\vec{x}_i}{2} \\
&= \frac{\log P_t(\vec{a}) + \log P_t(\vec{b})}{2}
\end{aligned}$$

Next we observe that $\exp(\frac{b^n - 2\delta_0}{n\delta})$ is a log-concave function (its *log* is linear in $b^n = 1 - b^1 - \cdots b^{n-1}$); finally we recall that the minimum of two log-concave functions is log-concave, and so is their product. □

The main issue is how fast the random walk approaches its stationary distribution, π_t, which is proportional to $Q_t(x)$. To analyze this we will use the theorem of Frieze and Kannan. To apply their theorem, we first need to ensure that π_t varies by at most a constant factor within the δ-cube around any grid point. In the full simplex Δ, the function P_t (and hence ρ_t) can vary a lot within a small distance[2]. With the smaller simplex however, the variation can be bounded using Lemma 7.

2. For example, one CRP may have 0 performance P_t while the center of its cube may have nonzero performance

Corollary 10 *With $\delta \le \frac{\delta_0}{T}$, for any grid point $\vec{z} \in C'$ and any $\vec{v} \in C(\vec{z})$,*

$$\frac{1}{5} Q_t(\vec{z}) \le Q_t(\vec{v}) \le 5 Q_t(\vec{z}).$$

Proof. From Lemma 7,

$$(1 - 1/T)^T P_t(\vec{z}) \le P_t(\vec{v}) \le (1 + 1/T)^T P_t(\vec{z}).$$

This gives bounds of $1/e$ and e on the LHS and RHS, respectively. Also, we can bound the maximum difference,

$$|z^n - v^n| = \left| \sum_1^{n-1} v^j - z^j \right| < \delta n/2$$

Therefore $Q_t(\vec{w})/P_t(\vec{w}) = \min(1, \exp(\frac{w^n - 2\delta_0}{n\delta}))$ differs by at most a factor of $e^{1/2}$ at $\vec{z}$ and $\vec{v}$. Finally, $e^{3/2} < 5$. □

Consider any particular day t, and let the distribution attained by the random walk after s steps be p_s, i.e. $p_s(x)$ is the probability that the walk is at the grid point x after s steps. The next theorem bounds the progress of the random walk towards its stationary distribution.

Theorem 11 *There is a constant A such that with $\delta_0 < \frac{1}{(n+T)^2}$, $\delta \log \frac{1}{\delta} = \frac{\epsilon \delta_0}{A(n+T)^2}$, and any $\epsilon > 0$, after $s \ge \frac{An}{\delta^2} \log \frac{n+T}{\epsilon \delta}$ steps,*

$$\sum_{\vec{x} \in C'} |p_s(\vec{x}) - \pi_t(\vec{x})| \le \frac{\epsilon}{4T(n+T)}.$$

Proof. The diameter of our set is $\sqrt{2}$. Applying the theorem of Frieze and Kannan we get that

$$2 \left(\sum_{\vec{x} \in C'} |p_s(\vec{x}) - \pi_t(\vec{x})| \right)^2 \le e^{-\frac{s\gamma\delta^2}{2n}} \log \left(\frac{1}{\pi_*} \right) + \frac{2M\pi_{1/2} n}{\gamma \delta^2} \tag{7}$$

where $\gamma > 0$ is a constant, π_* is $\min_{z \in C'} \pi(z)$, and

$$\pi_{\frac{1}{2}} = \sum_{x \in C' : \frac{vol(C(x) \cap \Delta')}{vol(C(x))} \le \frac{1}{2}} \pi(x).$$

In words, $\pi_{1/2}$ is the probability of the grid points whose cubes intersect the simplex in less than $\frac{1}{2}$ fraction of their volume. The parameter M is $\max_{\vec{z} \in C'} \frac{p_0(\vec{z})}{\pi_t(\vec{z})} \log \frac{p_0(\vec{z})}{\pi_t(\vec{z})}$.

There are two terms in (7). If the set K we're sampling from were a perfect cube, then we would have no cubes that were only partly in K, and $\pi_{1/2}$ in the second term would be 0. While it would probably be possible to reprove their theorem with simplexes rather than cubes, instead we chose to modify the walk to make $\pi_{1/2}$ very small. As we will show, the damping term in Q_t does exactly that, reducing border cubes by a factor of $e^{-\frac{\delta_0}{n\delta}}$. With the value of δ_0 we have chosen,

$$e^{-\frac{\delta_0}{n\delta}} = \delta^{\frac{A(n+T)^2}{n\epsilon}}, \tag{8}$$

a quantity that is smaller than basically any other quantity we are dealing with, for sufficiently large A.

First we bound π_*. For any $\vec{v}, \vec{w} \in \Delta$, we see that $v^j \geq \delta_0 w^j$, for $1 \leq j \leq n$, so that each day $\text{CRP}_{\vec{v}}$ does at least δ_0 as well as $\text{CRP}_{\vec{w}}$. This implies that

$$\begin{aligned} P_t(\vec{v}) &\geq \delta_0^T P_t(\vec{w}) \\ Q_t(\vec{v}) &\geq \delta_0^T e^{-\frac{\delta_0}{n\delta}} Q_t(\vec{w}) \\ \pi_t(\vec{v}) &= \frac{Q_t(\vec{v})}{\sum_{\vec{w}\in C'} Q_t(\vec{w})} \\ &\geq \delta_0^T e^{-\frac{\delta_0}{n\delta}} / (\#\text{ cubes}) \end{aligned}$$

Because δ^{-n} is an upper bound on the number of cubes (they all fit inside the unit cube) and (8),

$$\begin{aligned} \frac{1}{\pi_*} &\leq \delta^{-n} \delta_0^{-T} e^{\frac{\delta_0}{n\delta}} \\ &\leq \delta^{-2\frac{A(n+T)^2}{n\epsilon}} \end{aligned} \tag{9}$$

Next we bound M. Let $\vec{\sigma} = (1/n, 1/n, \ldots, 1/n)$. Since we are starting with $p_0(\vec{\sigma}) = 1$, we have $M = \log(1/\pi_t(\vec{\sigma}))/\pi_t(\vec{\sigma})$. Following logic similar to above, if $\vec{v}$ is any portfolio, then each day $\text{CRP}_{\vec{\sigma}}$ does at least $1/n$ as well as $\text{CRP}_{\vec{v}}$. So, over t days,

$$\begin{aligned} P_t(\vec{\sigma}) &\geq n^{-T} P_t(\vec{v}) \\ \pi_t(\vec{\sigma}) &\geq n^{-T}/(\#\text{ cubes}) \\ &\geq n^{-T}\delta^n \\ M &\leq n^T \delta^{-n} \log(n^T \delta^{-n}) \\ &\leq \delta^{-2n-2T} \end{aligned} \tag{10}$$

To bound $\pi_{\frac{1}{2}}$, we define K to the be the following set:

$$K = \{\vec{y} \in \Delta' | y^n \geq \delta_0 + (n-1)\delta/2\}.$$

With this definition of K, Frieze and Kannan's walk becomes exactly the walk in R-UNIVERSAL. To show this, we must show that the set of centers of cubes that intersect K is exactly C'. Recall that $(\delta_0, \delta_0, \ldots, \delta_0) \in \Re^{n-1}$ was chosen as a corner of a cube. First, suppose $\vec{v} \in K$ and $\vec{v} \in C(\vec{z})$ for grid point $\vec{z}$. Now, because of our grid position, $z^j \geq \delta_0$ for $j = 1, 2, \ldots, n-1$. Further, $z^n = v^n + \sum_1^{n-1} v^j - z^j$ so $z^n \geq v^n - (n-1)\delta/2$, which means that $z \in \Delta'$. Conversely, suppose $\vec{z} \in C'$, so $z^j \geq \delta_0 + \delta/2$ for $j < n$ and $z^n \geq \delta_0$. The point $\vec{v} = \vec{z} - (\delta/2, \delta, 2, \ldots, \delta/2) \in \Re^{n-1}$ has $v^j \geq \delta_0$ for $j < n$ and $v^n \geq \delta_0 + (n-1)\delta/2$, so $\vec{v} \in C(\vec{z}) \cap K$. Recall that the vectors here are $n-1$ dimensional, and v^n is shorthand for $1 - v^1 - \cdots v^{n-1}$.

Further, the only grid points whose cubes have less than half their volume inside K are those points $\vec{z}$ such that $\vec{z} \in C'$ but $\vec{z} \notin K$, i.e., $\delta_0 \leq z^n < \delta_0 + (n-1)\delta/2$. That is because any cut through the center of a cube divides it into two congruent pieces, so any cut that does not pass through the center divides it into a larger and smaller piece, with the center in the larger piece. Let this set of grid points be denoted by F. For any $\vec{z} \in F$,

$$\begin{aligned} Q_t(\vec{z}) &\leq \exp\left(\frac{-\delta_0 + n\delta/2}{n\delta}\right) P_t(\vec{z}) \\ &\leq e^{-\frac{\delta_0}{n\delta}} e^{1/2} P_t(\vec{z}). \end{aligned}$$

Moreover, if δ_0 were $1/(n+T)$ then at most a $1 - 1/e$ fraction of P_t would be in $\Delta' \setminus \Delta''$, by Lemma 5. The δ_0 we are using is much smaller[3], so we can easily say,

$$\sum_{\vec{v} \in F} P_t(\vec{v}) \leq \sum_{\vec{v} \in C''} P_t(\vec{v}),$$

and thus,

$$\sum_{\vec{v} \in F} Q_t(\vec{v}) \leq e^{-\frac{\delta_0}{n\delta}} e^{1/2} \sum_{\vec{v} \in C''} P_t(\vec{v}).$$

Since $P_t = Q_t$ for $\vec{v} \in C''$,

$$\begin{aligned} \pi_{1/2} &\leq e^{-\frac{\delta_0}{n\delta}} e^{1/2} \\ &= \delta^{\frac{A(n+t)^2}{n\epsilon}} e^{1/2} \end{aligned} \tag{11}$$

Substituting (9), (10), and (11) into (7) and using $s \geq \frac{An}{\delta^2} \log \frac{n+T}{\epsilon\delta}$,

$$2\left(\sum_{\vec{x} \in C} |p_s(\vec{x}) - \pi_t(\vec{x})|\right)^2 \leq \left(\frac{\epsilon\delta}{n+T}\right)^{A\gamma/2} \frac{2A(n+T)^2}{n\epsilon} \log\frac{1}{\delta} + \frac{2\delta^{-2n-2T}\delta^{\frac{A(n+t)^2}{n\epsilon}} e^{1/2} n}{\gamma\delta^2}.$$

Since $\delta^{1/\epsilon} < \epsilon$, both terms on the right hand side become smaller than $\epsilon/(n+T)$ to any constant power for sufficiently large A. □

7. Practical considerations

In this section, we give some suggestions which may help speed things up in practice. Our algorithm has an unnatural asymmetry in that we treat coordinate n differently than the rest. If it were not for the tapering function Q_t, then one could pick two arbitrary coordinates $i \neq j$, increase v^i, and decrease v^j. Alternatively, one could do tapering in any direction. The tapering seems to be an artifact of our analysis and the fact that we are trying to analyze a walk on a simplex by a grid. This is like fitting a triangular peg into a square hole. It seems more natural to implement it without any tapering.

Evaluating P_t can be costly for long market sequences. We can save some evaluations as follows. Imagine that the way we branch in step (iii) is by choosing a random $\alpha \in [0,1]$ and checking if $\alpha < x/y$. Now, by various means, one can calculate a lower bound on x/y for two neighboring portfolios. For small δ, this will be near 1. By doing this, one can then

3. This is the only reason we have required $\delta_0 < \frac{1}{(n+T)^2}$, which is probably smaller than necessary at this point. But the final δ_0 will be much smaller anyway.

avoid many evaluations by first choosing α in step (i). Only if α is larger than the lower bound, then compute x and y. Otherwise we know $\alpha < x/y$ without computing x and y. This saves many evaluations of P_t. Another natural idea is to vary the step size from 1 to larger numbers or use random step sizes. Theoretically, speaking, it is easy to show that the stationary distribution remains the same with varying step sizes, but we do not know how to show that the walk will converge faster.

We also believe that the algorithm may be sped up by starting near the maximum of the P_t function rather than at the center. Although we have no theoretical guarantee of this, the idea behind the random walk is to spend most of its time in places with high P_t. Thus, it could be helpful to start there, and there are known simple, practical techniques such as the EM algorithm for efficiently finding the best portfolio in hindsight (10).

The naive sampling approach described earlier, i.e. pick random portfolios and average them weighted by their performance, has been shown to require a number of samples that is on the order of the ratio of the performance of the average CRP to the best CRP. For many markets or for short periods of time, this ratio may be small. Thus, for calm markets with small changes, random walks are probably not necessary. In any case, it would be interesting to find the situations and improvements to the algorithm which make random walks do better than the naive sampling approach.

8. Conclusions

We have presented an efficient randomized approximation of the UNIVERSAL algorithm. With high probability $(1-\eta)$ it is within $(1-\epsilon)$ times the performance of universal, and runs in time polynomial in $\log(1/\eta)$, $1/\epsilon$, the number of days, and the number of stocks. With money, it is especially important to achieve this expectation. For example, a 50% chance at 10 million dollars may not be as valuable to most people as a guaranteed 5 million dollars.

While our implementation can be used for other applications of UNIVERSAL, such as data compression (6) and language modeling (11), we do not know an implementation for the case of transaction costs (2) or for the Dirichlet $(1/2, \ldots, 1/2)$ UNIVERSAL (4).

Acknowledgments

We are grateful to Philip Long and the anonymous referees for very many helpful suggestions and improvements. This research was supported by NSF CAREER award CCR-9875024 and an NSF postdoctoral research fellowship.

References

[1] D. Applegate and R. Kannan. Sampling and integration of near log-concave functions. *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, 156-163, 1991.

[2] A. Blum and A. Kalai. Universal Portfolios With and Without Transaction Costs. *Machine Learning*, 35(3):193-205, 1999.

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

[4] T.M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348-363, 1996.

[5] E. Ordentlich. and T.M. Cover. Online portfolio selection. *Proceedings of the 9th Annual Conference on Computational Learning Theory*, 310-313, 1996.

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

[7] A. Frieze and R. Kannan. Log-Sobolev inequalities and sampling from log-concave distributions. *Annals of Applied Probability*, 9:14-26, 1999.

[8] D.P. Foster and R.V. Vohra. Regret in the On-line Decision Problem. *Games and Economic Behavior*, 29(1/2):7-35, 1999.

[9] D. Helmbold, R. Schapire, Y. Singer, and M. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325-347, 1998.

[10] D. Helmbold, R. Schapire, Y. Singer, M. Warmuth. A comparison of new and old algorithms for a mixture estimation Problem. *Machine Learning* 27(1):97-119, 1997.

[11] A. Kalai, S. Chen, A. Blum, and R. Rosenfeld. On-line Algorithms for Combining Language Models. *Proceedings of the International Conference on Accoustics, Speech, and Signal Processing*, 2:745-748, 1999.

[12] R. Kannan, L. Lovasz and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. Random Structures and Algorithms, 11(1):1-50, 1997.

[13] A. Kalai and S. Vempala. Efficient algorithms for universal portfolios. *Proceedings of the 41st Annual Symposium on the Foundations of Computer Science*, 486-491, 2000.

[14] L. Lovasz and M. Simonovits. On the randomized complexity of volume and diameter. *Proceedings of the IEEE Symp. on Foundation of Computer Science*, 482-491, 1992.

[15] N. Metropolis, Rosenberg, Rosenbluth, Teller and Teller. Equation of state calculation by fast computing machines. In *Journal of Chemical Physics*, 21:1087-1092, 1953.

[16] V. Vovk and C.J.H.C. Watkins. Universal Portfolio Selection. *Proceedings of the 11th Annual Conference on Computational Learning Theory*, 12-23, 1998.