# Growth-Optimal Portfolio Selection under CVaR Constraints

**Guy Uziel**
Computer Science Department
Technion - Israel Institute of Technology

**Ran El-Yaniv**
Computer Science Department
Technion - Israel Institute of Technology

## Abstract

Online portfolio selection research has so far focused mainly on minimizing regret defined in terms of wealth growth. Practical financial decision making, however, is deeply concerned with both wealth and risk. We consider online learning of portfolios of stocks whose prices are governed by arbitrary (unknown) stationary and ergodic processes, where the goal is to maximize wealth while keeping the conditional value at risk (CVaR) below a desired threshold. We characterize the asymptomatically optimal risk-adjusted performance and present an investment strategy whose portfolios are guaranteed to achieve the asymptotic optimal solution while fulfilling the desired risk constraint. We also numerically demonstrate and validate the viability of our method on standard datasets.

## 1 Introduction

It has long been recognized that the value of any financial investment should be quantified using both return and risk, where risk is traditionally measured by the variance of the return. A common quantification for risk-adjusted return is the Sharpe ratio [42], which is essentially the (annualized) mean return divided by the (annualized) standard deviation of the return. Nevertheless, in online *portfolio selection* [13], which has become a focal point in online learning research, risk is rarely considered and the primary quantity to be optimized is still the return alone. The creation of an online learning technique that optimizes risk-adjusted return is a longstanding goal and a major challenge [29].

In an adversarial (regret minimization) online learning setting, risk-adjusted portfolio selection with no regret is known to be an impossible goal [15, 39]. Recently, within an i.i.d. setting, Mahdavi et al. presented a framework that can be utilized for achieving this goal [37], and Haskell et al. considered risk-aware algorithms [22], but i.i.d. modeling has been criticized for being unsuitable for modeling the stock prices faithfully [33]. The problem with i.i.d. modeling is the lack of time dependencies between stock returns. A substantially richer family of stochastic models is the class of stationary and ergodic processes, which are sufficiently expressive to model arbitrary dependencies among stock prices. It is the place to note that in the literature there exist other papers dealing with long term constraints (e.g., [24, 36]), however they are not suitable for our goal since they assume that the constraint function is fully known a priory (i.e. do not depend on time) which is obviously not the case studied in this paper.

Many publications have considered stationary and ergodic markets [3, 21, 20, 18, 30], and all these works consider strategies that are oblivious to risk. All the learning strategies they consider rely on non-parametric estimation techniques (e.g., histogram, kernel, or nearest neighbors methods). Moreover, these strategies always use a countably infinite set of experts, and the guarantees provided for these strategies are always asymptotic. This is no coincidence, as it is well known that finite sample guarantees for these methods cannot be achieved without additional strong assumptions on the source distribution [14, 35]. Similarly, it is also known that non-parametric strategies in this context must rely on infinitely many experts [17].

Approximate implementations of non-parametric strategies (which apply only a finite set of experts), however, turn out to work exceptionally well and, despite the inevitable approximation, are reported [21, 20, 18, 28, 29] to significantly outperform strategies designed to work in an adversarial, no-regret setting. For example, the nearest-neighbor investment strategy of [21] is shown in [32, 29] to beat Cover's universal portfolios (UP) [13], the exponentiated gradient (EG) method [23], and the online Newton steps strategy of [1] on most of the common datasets. We also note that practical approximate use of asymptotic methods is prevalent in other areas of machine learning such as (deep) reinforcement learning with function approximation [6]).

For a market with $n$ stocks, and within a stochastic online

Proceedings of the $21^{st}$ International Conference on Artificial Intelligence and Statistics (AISTATS) 2018, Lanzarote, Spain. PMLR: Volume 84. 

learning framework, we develop a novel online portfolio selection strategy called *CVaR-Adjusted Nearest Neighbor* (CANN), which guarantees the best possible asymptotic performance while keeping the risk contained to a desired threshold. This is done using a novel mechanism that facilitates the handling of multiple objectives. Rather than using standard deviation to measure risk, we consider the well-known CVaR, a coherent and widely-accepted risk measure, which improves upon the traditional measure by appropriately capturing the downside risk [41]. We prove the asymptotic optimality of our strategy for general stationary and ergodic processes, thus allowing for arbitrary (unknown) dependencies among stock prices. We also present numerical examples where we apply an approximate application of our strategy (with a finite set of experts) that validates the method and beautifully demonstrates how risk can be controlled.

The problem of risk-adjusted prediction is, of course, not limited only to online portfolio selection but to other areas as well, and in recent years, there is a growing interest in the risk accompanied by the prediction. In [15] it was proved that in the expert setting one can not efficiently trade off between return and risk (measured by the variance) when the setting is adversarial. Other papers have tried to incorporate a coherent risk measure (see Section 2) in their predictive algorithms. For example, papers such as [11, 12] have discussed risk sensitive algorithms within a Markov decision process (MDP) framework and [22] incorporated coherent risk measures in the i.i.d. setting.

## 2 Online Portfolio Selection

We consider the following standard online portfolio selection game (OPS) with short selling and leverage, as defined by Györfi et al. [19]. The game is played through $T$ days over a market with $n$ stocks. On each day $t$, the market is represented by a *market vector* $\mathbf{X}_t$ of relative prices, $\mathbf{X}_t \triangleq (x_1^t, x_2^t, ..., x_n^t)$, where for each $i = 1, \dots, n$, $x_i^t \geq 0$ is the *relative price* of stock $i$, defined to be the ratio of its closing price on day $t$ relative to its closing price on day $t-1$. A *wealth allocation* vector or *portfolio* for day $t$ is $\mathbf{b}_t \triangleq (b_0^t, b_1^t, b_2^t, \dots, b_n^t)$, where $b_0^t$ is a cash allocation (not invested in any stock), and for $i > 0$, $b_i^t$ is the wealth allocation for stock $i$ (i.e. the fraction of total wealth allocated into the asset), where a positive component, $b_i^t > 0$, represents a *long position* in stock $i$, and a negative one, $b_i^t < 0$, is a *short position* in stock $i$. We also allow leverage; that is, the investor can borrow and invest additional cash, so as to amplify her profits. For the borrowed cash, the investor must pay a daily interest rate, $r > 0$, and we assume that the investor receives the same interest $r$ for deposited cash ($b_0^t$). Consider a portfolio $\mathbf{b}_t$ played at the start of day $t$. After the market vector $\mathbf{X}_t$ is revealed, the portfolio changes in response to changes in stock price, as follows. For each portfolio component $b_i$, if $b_i^t > 0$ is a long position, its revised value is $b_i^t x_i^t$. However, if $b_i^t < 0$ is a short position, then, after we take into account the interest owed on borrowing the stock for the short sale, the revised value of this position is $b_i^t(x_i^t - 1 + r)$ (note that in this case, the investor profits when the price drops and vice versa). Clearly, short selling and leveraging are risky; for example, a short position has unbounded potential loss that is further amplified by leveraging. Following [19], we assume that no stock can lose or gain more than $B \times 100\%$ of its value from one day to another, where $B \in (0, 1)$. In other words, for each $i, t$,

$$1 - B \leq x_i^t \leq 1 + B. \tag{1}$$

Thus in order to preclude the possibility of bankruptcy (see, e.g., [19], Chapter 4), the allowed leverage is

$$L_{B,r} \triangleq \frac{B+1}{r+1}.$$

Note that in order to avoid a degenerate market where the best investment is the risk-free asset, it is reasonable to assume that $B > r$, which results in $L_{B,r} > 1$. Using the notation

$$(\mathbf{b})^+ \triangleq (\max\{b_1, 0\}, \dots, \max\{b_n, 0\})$$

and

$$(\mathbf{b})^- \triangleq (\min\{b_1, 0\}, \dots, \min\{b_n, 0\}),$$

and considering the interest accredited for deposited cash, the interest debited for borrowed stocks (short positions), and the interest paid for leveraged wealth, we obtain, by the end of the day, an overall daily return of

$$b_0(1+r) + \langle (\mathbf{b}_t)^+, \mathbf{X}_t \rangle + \langle (\mathbf{b}_t)^-, \mathbf{X}_t - 1 + r \rangle \\ -(L_{B,r} - 1)(1+r). \tag{2}$$

The investor chooses a portfolio from the following set,

$$\left\{ (b_0, \dots, b_n) \in \mathbb{R}^{n+1} \quad | \quad \sum_{i=0}^{n} |b_i| = L_{B,r} \right\}, \tag{3}$$

which is, unfortunately, not convex. We thus apply a simple transformation proposed by Györfi et al. [19]: transform the market vector $\mathbf{X}_t$ into a vector with $2n+1$ entries (one entry for cash, $n$ entries for the long components, and $n$ for the short ones). Formally, we define the transformed market vector as

$$\mathbf{X}'_t \triangleq (1 + r, x_1^t, 2 - x_1^t + r, \dots, x_n^t, 2 - x_n^t + r),$$

which is uniquely defined as a function of the original market vector. The transformed portfolio set is now defined as

$$\mathcal{B}' \triangleq \\ \{(b_0, \dots, b_{2n}) \in \mathbb{R}^{2n+1} \quad | \quad b_i \geq 0, \sum_{i=0}^{2n} b_i = L_{B,r}\}, \tag{4}$$

which is an unnormalized simplex. With this transformed market vector and portfolio set, at the start of each trading day $t$, the player chooses a portfolio $\mathbf{b}_t \in \mathcal{B}'$ based on the previous market sequences. It can easily be shown [19] that by the end of day $t$, the player's daily multiplicative return is simplified to

$$\langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1 + r). \quad (5)$$

With respect to a fixed stationary and ergodic process, we denote by $\mathbf{X} \triangleq \{\mathbf{X}_t\}_{-\infty}^{\infty}$ [1] the induced sequence of stationary and ergodic market vectors, and define the player's *investment strategy*, denoted by $\mathbf{S}$, as a sequence of portfolios $\mathbf{b}_1, \mathbf{b}_2, \ldots$. Then, assuming initial wealth of \$1, we obtain after $T$ days the following cumulative wealth,

$$R_T(\mathbf{S}, \mathbf{X}) \triangleq \prod_{t=1}^{T} \left( \langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1 + r) \right). \quad (6)$$

Defining the *average growth rate*,

$$W_T(\mathbf{S}) \triangleq \frac{1}{T} \sum_{t=1}^{T} \log \left( \langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1 + r) \right), \quad (7)$$

we have

$$R_T(\mathbf{S}, \mathbf{X}) = \prod_{t=1}^{T} \left( \langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1 + r) \right) = e^{\sum_{t=1}^{T} \log\left( \langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1+r) \right)} = e^{T W_T(\mathbf{S})}.$$

Notice that maximizing $W_T(\mathbf{S})$ is equivalent to maximizing $R_T(\mathbf{S}, \mathbf{X})$. In Section 4 , we denote the summand of $W_T(\mathbf{S})$ (7) by

$$\omega(\mathbf{b}_t, \mathbf{X}_t) \triangleq - \log \left( \langle \mathbf{b}_t, \mathbf{X}'_t \rangle - (L_{B,r} - 1)(1 + r) \right).$$

## 3 Introducing Risk

The traditional quantity for measuring financial risk is the variance (standard deviation) of the return. This measure, however, is criticized for being inadequate to measure risk. One of the reasons is its inability to distinguish between downside risk and upside risk (which corresponds to a desirable behavior). Various alternative measures have been proposed, such as the maximum drawdown, and value at risk (VaR). An axiomatic approach proposed by Artzner et al. [4] identifies *coherent* risk measures, which satisfy the proposed axioms. Accordingly, the most popular coherent risk measure is *conditional value at risk* (CVaR). For any parameter $\alpha \in (0, 1)$, $\text{CVaR}_\alpha$ is essentially the average loss that the investor suffers on the $(1 - \alpha)\%$ worst returns. For a continuous, bounded mean random variable $Z$ the $\text{CVaR}_\alpha$ is defined as

**Definition 1** ($\text{CVaR}_\alpha$)**.** *Let $Z$ be a continuous random variable representing loss. Given a parameter $0 < \alpha < 1$, the CVaR$_\alpha$ of $Z$ is*

$$\mathit{CVaR}_\alpha(Z) = \mathbb{E}[Z \mid Z \geq \min\{c \mid \mathbb{P}_Z(Z \leq c) \geq \alpha\}].$$

Assuming that we already know the distribution of returns, a direct calculation of CVaR from the above formula requires a calculation of the $\alpha\%$ quantile followed by averaging over the right tail of the loss distribution. Alternatively, it was shown in [41] that $\text{CVaR}_\alpha$ can be computed by solving the following convex optimization problem. Define

$$\phi(\mathbf{b}, c) \triangleq c + \frac{1}{1 - \alpha} \mathbb{E} \left[ (\omega(\mathbf{b}, \mathbf{X}) - c)^+ \right], \quad (8)$$

where we overload the previously defined $(\cdot)^+$ for vectors, and define for any scalar $x$, $(x)^+ \triangleq \max\{0, x\}$.

**Theorem 1** ([41])**.** *The function $\phi(\mathbf{b}, c)$ is convex and continuously differentiable. Moreover, the CVaR$_\alpha$ of the loss associated with any portfolio $\mathbf{b}$ is*

$$\mathit{CVaR}_\alpha(\mathbf{b}) = \min_{c \in \mathbb{R}} \phi(\mathbf{b}, c). \quad (9)$$

Theorem 1 is essential to the development and analysis of our strategy. By our market boundedness assumption (1), it follows that $\omega(\mathbf{b}, X)$ is contained in $[-M, M]$ for some $M > 0$. Thus, any $c$ that minimizes Equation (9) must reside in $[-M, M]$. For a complete proof of this simple fact, see [22]. In Section 4, we require the following definition,

$$\mathcal{B} \triangleq \mathcal{B}' \times [-M, M].$$

## 4 Optimality of $\mathcal{W}^*$

Let $\mathcal{F}_\infty$ be the $\sigma$-algebra generated by the infinite past $X_{-1}, X_{-2}, \ldots$, and let $\mathbb{P}_\infty$, be the induced regular conditional probability distribution of $X_0$ given the infinite past. Thus, all expectations w.r.t. $X_0$ are conditional given the infinite past. A well-known result appearing in [3, 2] proves the following upper bound on the asymptotic average growth rate of any investment strategy $\mathbf{S}$ under stationary and ergodic markets:

$$\limsup_{T \to \infty} W_T(\mathbf{S}) \leq \mathbb{E} \left[ \max_{\mathbf{b} \in \mathcal{B}'()} \mathbb{E}_{\mathbb{P}_\infty} \left[ -\omega(\mathbf{b}, \mathbf{X}_0) \right] \right], \quad (10)$$

where $\mathcal{B}'()$ denotes the $\mathcal{F}_\infty$-measurable functions. Over the years, several algorithms achieving this asymptotic bound were proposed [20, 18, 21] (for the case of long-only portfolios).

[1] By Kolmogorov's extension theorem [10], the stationary and ergodic process $(X_n)_1^\infty$ can be extended to $(X_n)_{-\infty}^{\infty}$ such that the ergodicity holds for both $n \to \infty$ and $n \to -\infty$.

Our goal is to achieve the optimal asymptotic average growth rate while keeping the CVaR bounded. By Theorem 1, the desired growth for a given risk threshold $\gamma > 0$ rate is given by the solution to the following minimization problem,

$$\begin{aligned} \underset{(\mathbf{b},c)\in\mathcal{B}()}{\text{minimize}} \quad & \mathbb{E}[\omega(\mathbf{b},\mathbf{X}_0)] \\ \text{subject to} \quad & \phi(\mathbf{b},c) \le \gamma, \end{aligned} \tag{11}$$

Optimization problem (11) motivates a definition of a $\gamma$-*bounded strategy*, whose long-term average CVaR, calculated according to the available information at the beginning of each round, is bounded by $\gamma$.

**Definition 2** ($\gamma$-bounded strategy). *An investment strategy* $\mathbf{S}$ *will be called* $\gamma$*-bounded if, almost surely,*

$$\limsup_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T} \min_{c\in\mathbb{R}} \left( c + \frac{1}{1-\alpha}\mathbb{E}_{\mathbb{P}_{X_i|X_0^{i-1}}}\left[(\omega(\mathbf{b}_t,\mathbf{X}_t)-c)^+\right]\right) \le \gamma,$$

*where* $b_i$ *is portfolio chosen by* $\mathbf{S}$ *at day* $i$*. The set of all* $\gamma$*-bounded strategies is denoted* $\mathcal{S}_\gamma$*.*

Clearly, there is always a solution to optimization problem (11), and therefore, $\mathcal{S}_\gamma \neq \emptyset$. For example, the vacuous strategy that always invests everything in cash is $\gamma$-bounded for any $\gamma > 0$. Let $(\mathbf{b}^*_\infty, c^*_\infty)$ be a solution to (11). Define the $\gamma$-*feasible optimal value* as

$$\mathcal{W}^* \triangleq \mathbb{E}\left[\omega(\mathbf{b}^*_\infty, X_0)\right] \quad a.s.$$

Optimization problem (11) is equivalent to finding the saddle-point of the Lagrangian function [34], namely,

$$\min_{(\mathbf{b},c)\in\mathcal{B}()} \max_{\lambda\in\mathbb{R}^+} \mathcal{L}((\mathbf{b},c),\lambda), \tag{12}$$

where the Lagrangian is

$$\mathcal{L}((\mathbf{b},c),\lambda) \triangleq \mathbb{E}\left[\omega(\mathbf{b},X_0)\right] + \lambda\left(\phi(\mathbf{b},c)-\gamma\right). \tag{13}$$

Let $\lambda^*_\infty$ be the value of $\gamma$ optimizing (12), and assume it is unique and that it is possible to identify a constant $\lambda_{\max}$ such that $\lambda_{\max} > \lambda^*_\infty$ [37]. With this constant available, we set $\Lambda \triangleq [0, \lambda_{\max}]$.

Our first result is that $\mathcal{W}^*$ bounds the performance of any strategy in $\mathcal{S}_\gamma$. This result, as stated in Theorem 2, is a generalization of the well-known result of [2] regarding the best possible performance for wealth alone (without constraints).

**Theorem 2** (Optimality of $\mathcal{W}^*$). *For any investment strategy* $\mathbf{S} \in \mathcal{S}_\gamma$ *whose portfolios are* $\mathbf{b}_1, \mathbf{b}_2, \ldots$*, the following holds a.s.*

$$\liminf_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T} \omega(\mathbf{b}_i, X_i) \ge \mathcal{W}^*.$$

From Theorem 2 it follows that an investment strategy, $\mathbf{S} \in \mathcal{S}_\gamma$, is optimal if, for any bounded, stationary and ergodic process $\{X_i\}_{-\infty}^{\infty}$,

$$\lim_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T} \omega(\mathbf{b}_i, X_i) = \mathcal{W}^* \quad a.s. \tag{14}$$

A $\gamma$-bounded investment strategy whose asymptotic average growth rate is $\mathcal{W}^*$ will be called $\gamma$-*universal* , By the above theorem its asymptotic average growth rate is not worse than any other $\gamma$-bounded strategy. We find just such a strategy in Section 5.

# 5 CVaR-Adjusted Nearest Neighbor Investment Strategy

In this section we present an investment strategy in $\mathbf{S} \in \mathcal{S}_\gamma$ that satisfies (14). The strategy, which we call *CVaR-Adjusted Nearest Neighbor*, henceforth CANN, is summarized in the pseudo-code in Algorithm 1. To define the strategy we require the following definition of the *instantaneous Lagrangian*:

$$l(\mathbf{b},c,\lambda,x) \triangleq \omega(\mathbf{b},x) + \lambda\left(c + \frac{1}{1-\alpha}\left(\omega(\mathbf{b},x)-c\right)^+ - \gamma\right). \tag{15}$$

The strategy maintains a countable array of experts $\{H_{k,l}\}$, where on each day $t$ an expert $H_{k,l}$ outputs a triplet $(\mathbf{b}^t_{k,l}, c^t_{k,l}, \lambda^t_{k,l}) \in \mathcal{B}\times\Lambda$, defined to be the minimax solution corresponding to an empirical distribution using nearest neighbor estimates (see details below). We prove that, as $t$ grows, those empirical estimates converge (weakly) to $\mathbb{P}_\infty$ and thus converge to $\mathcal{W}^*$. Each day $t$, CANN outputs a prediction $(\mathbf{b}_t, c_t, \lambda_t) \in \mathcal{B}\times\Lambda$. The sequence of predictions $(\mathbf{b}_1, c_1), (\mathbf{b}_2, c_2), \ldots$ output by CANN is designed to minimize the average loss, $\frac{1}{T}\sum_{i=1}^{T} l(\mathbf{b}, c, \lambda_i, x_i)$. Similarly, the sequence of predictions $\lambda_1, \lambda_2, \ldots$ is designed to maximize the average loss, $\frac{1}{T}\sum_{i=1}^{T} l(\mathbf{b}_i, c_i, \lambda, x_i)$. Each of $(\mathbf{b}_i, c_i)$ and $\lambda_i$ is generated by aggregating the experts' predictions $(\mathbf{b}, c)^i_{k,l}$ and $\lambda^i_{k,l}$, $k, l = 1, 2, \ldots$, respectively. In order to ensure that CANN will perform as well as any other expert for both the $(\mathbf{b}, c)$ and $\lambda$ predictions, we apply, twice alternately, the Weak Aggregating Algorithm (WAA) of [43], and [26].[2] It will also ensure that the average loss of the strategy will converge a.s. to $\mathcal{W}^*$.

We now turn to defining the countable set of experts $\{H_{k,h}\}$: For each $h = 1, 2, \ldots$, we choose $p_h \in (0,1)$ such that for the sequence $\{p_h\}_{h=1}^{\infty}$, $\lim_{h\to\infty} p_h = 0$. Setting $\hat{h} = \lfloor tp_h \rfloor$, for expert $H_{k,h}$ we define, for a fixed $(k \times n)$-

[2] In [43] it was proved that the WAA is capable of dealing with countable number of experts.

**Algorithm 1** CVaR-Adjusted Nearest Neighbor Investment Strategy (CANN)

**Input:** Countable set of experts $\{H_{k,h}\}$, $\alpha > 0$ $(\mathbf{b}_0, c_0) \in \mathcal{B}$ $\lambda_0 \in \Lambda$, initial probability $\{\beta_{k,h}\}$,

**For** $t = 0$ **to** $\infty$

Play $\mathbf{b}_t, c_t, \lambda_t$.

Nature reveals market vector $\mathbf{X}_t$

Suffer loss $l(\mathbf{b}_t, c_t, \lambda_t, x_t)$.

Update the cumulative loss of the experts

$$l^{k,h}_{(\mathbf{b},c),t} \triangleq \sum_{i=0}^{t} l(\mathbf{b}^i_{k,h}, c^i_{k,h}, \lambda_i, x_i)$$

$$l^{k,h}_{\lambda,t} \triangleq \sum_{i=0}^{t} l(\mathbf{b}_i, c_i, \lambda^i_{k,h}, x_i)$$

Update experts' weights

$$w^{(k,h)}_{t+1,(\mathbf{b},c)} \triangleq \beta_{k,h} \exp\left(-\frac{1}{\sqrt{t}} l^{k,h}_{(\mathbf{b},c),t}\right)$$

$$p^{(k,h)}_{t+1,(\mathbf{b},c)} \triangleq \frac{w^{(k,h)}_{t+1,(\mathbf{b},c)}}{\sum_{h=1}^{\infty}\sum_{k=1}^{\infty} w^{(k,h)}_{t+1,(\mathbf{b},c)}}$$

Update experts' weights $w^{\lambda,(k,h)}_{n+1}$

$$w^{(k,h)}_{t+1,\lambda} \triangleq \beta_{k,h} \exp\left(\frac{1}{\sqrt{t}} l^{k,h}_{\lambda,t}\right)$$

$$p^{(k,h)}_{t+1,\lambda} \triangleq \frac{w^{(k,h)}_{t+1,\lambda}}{\sum_{h=1}^{\infty}\sum_{k=1}^{\infty} w^{(k,h)}_{t+1,\lambda}}$$

Choose $\mathbf{b}_{t+1}, c_{t+1}$ and $\lambda_{t+1}$ as follows

$$\mathbf{b}_{t+1} = \sum_{k,h} p^{(k,h)}_{t+1,(\mathbf{b},c)} \mathbf{b}^{t+1}_{k,h}$$

$$c_{t+1} = \sum_{k,h} p^{(k,h)}_{t+1,(\mathbf{b},c)} c^{t+1}_{k,h}$$

$$\lambda_{t+1} = \sum_{k,h} p^{(k,h)}_{t+1,\lambda} \lambda^{t+1}_{k,h}$$

**End For**

dimensional vector, denoted $w$, the following set,

$$B^{w,(1,t)}_{k,h} \triangleq \{x_i \mid k+1 \le i \le t, X^{i-1}_{i-k} \text{ is among the } \hat{h} \text{ nearest neighbors of } w \text{ among } X^k_1, \ldots, X^{t-1}_{t-k}\},$$

where

$$X^{j+k}_j \triangleq (\mathbf{X}_j, \ldots, \mathbf{X}_{j+k}) \in \mathbb{R}^{k\times n}.$$

Thus, expert $H_{k,h}$ has a window of length $k$ and it looks for the $\hat{h}$ euclidean nearest-neighbors of $w$ in the past. This results in a set of market vectors $B^{w,(1,t)}_{k,h}$. This set can also be seen as a conditional probability over the space of possible market vectors. Then, each expert recommends the actions $(\mathbf{b}, c) \in \mathcal{B}$ and $\lambda \in \Lambda$, which are the corresponding minimax solution. More formally, we define

$$h^{(\mathbf{b},c)}_{k,h}(X^{t-1}_1, w) \triangleq \arg\min_{(\mathbf{b},c)\in\mathcal{B}} \left(\max_{\lambda\in\Lambda} \frac{1}{|B^{w,(1,t)}_{k,h}|} \sum_{x_i \in B^{w,(1,t)}_{k,h}} l_{k,l,t}(\mathbf{b}, c, \lambda, x_i)\right)$$

and

$$h^{\lambda}_{k,h}(X^{t-1}_1, w) \triangleq \arg\max_{\lambda\in\Lambda} \left(\min_{(\mathbf{b},c)\in\mathcal{B}} \frac{1}{|B^{w,(1,t)}_{k,h}|} \sum_{x_i \in B^{w,(1,t)}_{k,h}} l_{k,l,t}(\mathbf{b}, c, \lambda, x_i)\right)$$

for

$$l_{k,h,t}(\mathbf{b}, c, \lambda, x_i) \triangleq l(\mathbf{b}, c, \lambda, x_i) + \left(||(\mathbf{b}, c)||^2 - ||\lambda||^2\right)\left(\frac{1}{t} + \frac{1}{h} + \frac{1}{k}\right),$$

Using the above, we define the predictions of $H_{k,h}$ to be:

$$H^{(\mathbf{b},c)}_{k,h}(X^{t-1}_1) = h^{(\mathbf{b},c)}_{k,h}(X^{t-1}_1, X^{t-1}_{t-k}), \quad t = 1, 2, 3, \ldots \tag{16}$$

$$H^{\lambda}_{k,h}(X^{t-1}_1) = h^{\lambda}_{k,h}(X^{t-1}_1, X^{t-1}_{t-k}), \quad t = 1, 2, 3, \ldots \tag{17}$$

Note that $l_{k,h,t}(\mathbf{b}, c, \lambda, x)$ is an approximation of $l(\mathbf{b}, c, \lambda, x)$, which guarantees that the minimax solution of every expert is unique. This technicality is used in the proof of Theorem 3.

Theorem 3 below states that the CANN strategy, applied on the experts defined above, is $\gamma$-universal. We note that the theorem utilizes a standard assumption (see, e.g., [7, 21]).

**Theorem 3** ($\gamma$-universality)**.** *Assume that for any vector $w \in \mathbb{R}^{n\times k}$ the random variable $||X^k_1 - w||$ has a continuous distribution. Then, for any $\gamma > 0$ and for any bounded process $\{X_i\}^{\infty}_{-\infty}$, CANN is $\gamma$-universal.*

The full proof of this theorem appears in the supplementary material. The main idea is to show first that the minimax (12) value of the Lagrangian (13) is continuous with respect to the probability measure. Then, we prove that the minimax measurable selection (which gives the optimal actions) is also continuous and every accumulation point of induced sequence of optimal actions is optimal. This is formulated in the following lemmas:

**Lemma 1** (Continuity and Minimax)**.** *Let $\mathcal{B}, \Lambda, \mathcal{X}$ be compact real spaces. $l : \mathcal{B} \times \Lambda \times \mathcal{X} \to \mathbb{R}$ be a continuous function. Denote by $\mathbb{P}(\mathcal{X})$ the space of all probability measures* *on $\mathcal{X}$ (equipped with the topology of weak-convergence). Then the following function $L^* : \mathbb{P}(\mathcal{X}) \to \mathbb{R}$ is continuous*

$$L^*(\mathbb{Q}) = \inf_{(\mathbf{b},c)\in\mathcal{B}} \sup_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}}\left[l((\mathbf{b},c),\lambda,x)\right]. \qquad (18)$$

*Moreover, for any $\mathbb{Q} \in \mathbb{P}(\mathcal{X})$,*

$$\inf_{(\mathbf{b},c)\in\mathcal{B}} \sup_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}}\left[l((\mathbf{b},c),\lambda,x)\right] = \sup_{\lambda\in\Lambda} \inf_{(\mathbf{b},c)\in\mathcal{B}} \mathbb{E}_{\mathbb{Q}}\left[l((\mathbf{b},c),\lambda,x)\right].$$

**Lemma 2** (Continuity of the optimal selection)**.** *Let $\mathcal{B}, \Lambda, \mathcal{X}$ be compact real spaces. Then, there exist two measurable selection functions $h^X$,$h^\lambda$ such that*

$$h^{(\mathbf{b},c)}(\mathbb{Q}) \in \arg\min_{(\mathbf{b},c)\in\mathcal{B}} \left(\max_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}}\left[l((\mathbf{b},c),\lambda,x)\right]\right),$$

$$h^{\lambda}(\mathbb{Q}) \in \arg\max_{\lambda\in\Lambda} \left(\min_{(\mathbf{b},c)\in\mathcal{B}} \mathbb{E}_{\mathbb{Q}}\left[l((\mathbf{b},c),\lambda,x)\right]\right)$$

*for any $\mathbb{Q} \in \mathbb{P}(\mathcal{X})$. Moreover, let $L^*$ be as defined in Equation (18). Then, the set*

$$Gr(L^*) \triangleq \{(u^*, v^*, \mathbb{Q}) \mid u^* \in h^{(\mathbf{b},c)}(\mathbb{Q}), v^* \in h^{\lambda}(\mathbb{Q}), \mathbb{Q} \in \mathbb{P}(\mathcal{X})\},$$

*is closed in $\mathcal{B} \times \Lambda \times \mathbb{P}(\mathcal{X})$.*

Those lemmas are helpful since the empirical measures generated by the experts will converge weakly, using the ergodic theorem [9] and the properties of the nearest neighbors estimates [14], to $\mathbb{P}_\infty$. Thus, we will get that the asymptotic average growth rate of the experts is $\mathcal{W}^*$, and by the guarantee of WAA, the asymptotic average growth rate of CANN will also be $\mathcal{W}^*$.

Table 1: Datasets

| DATASET | STARTING DAY | # DAYS | # STOCKS |
|---|---|---|---|
| NYSE-N | 1/1/1983 | 2250 | 23 |
| MSCI | 4/1/2006 | 1043 | 24 |

## 6 Empirical results

To apply the CANN strategy, we implemented it with a finite set of experts, and in this section we present our empirical results on some standard datasets. One objective of our experiments is to examine how well CANN maintains the CVaR constraints. Another objective is to compare it to several well-known adversarial no-regret portfolio selection algorithms and to stochastically universal strategies. The benchmark algorithms we tested are:

- Best Constant Rebalancing Portfolio (BCRP) [13]: Constant rebalancing portfolio (CRP) is a strategy which uses the same wealth allocation at each round. The BCRP is the optimal strategy among the class of CRP strategies. Clearly, the BCRP is calculated in hindsight and it only stands as a benchmark strategy. It was shown in [13] that the BCRP is the optimal strategy whenever market sequences are i.i.d.

- Cover's Universal Portfolios (UP) [13]: This algorithm invests a fraction of the wealth in every CRP, where the updates ensures that well preforming CRPs will get an higher fraction. The algorithm guarantees logarithmic regret w.r.t the wealth achieved by the BCRP.

- Exponentiated Gradient (EG) [23]: This well known no-regret algorithm can be applied to the problem of OPS by looking at every stock as an individual expert. This algorithm ensures a square-root regret.

- Online Newton Steps (ONS) [1]: This algorithm exploits the exp-concavity property of the loss function in order to guarantee logarithmic regret w.r.t. the wealth achieved by BCRP.

- The nearest-neighbor based strategy (long-only and non-leveraged) of Györfi et al. ($\mathcal{B}_{NN}$) [21]: $\mathcal{B}_{NN}$, which is a (stochastically) universal strategy whose asymptotic growth rate is optimal when the market follows a stationary and ergodic process.

- The nearest-neighbor based strategy (with short and leveraged): $\mathcal{B}^L_{NN}$: A variant of $\mathcal{B}_{NN}$ which allows short and leveraged investment. This variant was described in [19]

The experiments were conducted on two datasets that were used in many previous works (see, e.g., [28, 29, 8]). The first is the NYSE dataset, which consists of 23 stocks between the years 1985-1995. The second is the MSCI dataset, which consists of 24 stocks between the years 2006-2010. The statistics of these two datasets are summarized in Table 1

Following [19, 25], for both datasets we used a daily interest rate of $r = 0.000245$ and set $B = 0.4$, which implies that $L_{B,r} = 2.49$. This interest rate is suitable for the NYSE dataset and is higher than the true rate that was at the time of the MSCI dataset. However, this high choice of interest only reduces the returns of our algorithm, which rarely deposits cash and must pay a lot for short selling and loans. Similarly to the implementation of $\mathcal{B}_{NN}$ [21], our implementation of CANN took the following experts, $k = 1, \ldots, 5$ $h = 1, \ldots, 10$, for a total of 50 experts, and we set $p_l = \frac{1}{20} + \frac{h-1}{18}$. The initial expert prior was set to be uniform and we chose the typical value of $\alpha = 0.95$ for the calculation of CVaR. The hyper-parameters for the benchmark algorithms were according to [31].

Table 2: Wealth of CANN and benchmark algorithms.

| Dataset | BCRP | UP | EG | ONS | $\mathcal{B}_{NN}$ | $\mathcal{B}^{L}_{NN}$ | CANN$_{.05}$ |
|---|---|---|---|---|---|---|---|
| NYSE | 12.53 | 5.05 | 5.03 | 5.83 | 39.56 | 1054 | 58.8 |
| MSCI | 1.51 | 0.92 | 0.93 | 0.86 | 13.47 | 6.32E+05 | 6.06E+03 |

Table 3: CVaR$_{0.95}$ of CANN with different values of $\gamma$.

| Dataset | $\mathcal{B}^{L}_{NN}$ | CANN$_{.07}$ | CANN$_{.06}$ | CANN$_{.05}$ | CANN$_{.04}$ | CANN$_{.03}$ | CANN$_{.02}$ | CANN$_{.01}$ |
|---|---|---|---|---|---|---|---|---|
| NYSE | 6.3% | 4.9% | 3.82% | 3.2% | 2.9% | 2.46% | 1.86% | 1.24% |
| MSCI | 7.76% | 5.49% | 5.15% | 4.44% | 3.81% | 2.98% | 2.27% | 1.59% |

Table 2 presents the total wealth of all the algorithms, where CANN was applied with $\gamma = 0.05$, denoted by CANN$_{.05}$. It is evident that the stochastically universal algorithms are superior to all the worst-case universal algorithms. It is also evident that allowing margin and short trading is beneficial in terms of final wealth. Another phenomenon we can observe is that the final wealth of the our risk-adjusted algorithm is less compare to the risk-oblivious algorithm $\mathcal{B}^{L}_{NN}$, this of-course should not be surprising.

We conducted another experiment where we applied CANN with different choices of $\gamma$ in the range $[0.01, 0.07]$. The results are presented in Table 3, where the CVaR$_{0.95}$ is presented. it is evident that CANN indeed enjoys lower CVaR$_{0.95}$ rates compares to $\mathcal{B}^{L}_{NN}$. We can also observe that by lowering the choice of $\gamma$ we can adjust the CVaR resulting in less risky strategies. In Figure 1, we present the mean-CVaR tradoff where the $y$-axis shows the average return and the $x$-axis shows the CVaR$_{0.95}$. The concave shape suggests that by choosing an appropriate $\gamma$, one may achieve a better mean-CVaR trade-off.

In Figure 2 we present the smoothed PDF of the returns of the daily returns of both $\mathcal{B}^{L}_{NN}$ (red curve) and CANN (grenn curve) on the NYSE and the MSCI datasets. The left tails of these PDFs show that our algorithm effectively decreases the risk accompanied with the investment. Another interesting aspect of our strategy is its lower variance.

In Figure 3 we plot the log-returns of the above instances on both datasets where it can be seen that the performance of our algorithm is consistent during different investment periods.

## 7 Concluding Remarks

In this paper we introduced the CVaR-adjusted nearest-neighbor portfolio selection strategy, which is the first CVaR-adjusted universal portfolio selection strategy when the underlying market process is stationary and ergodic. It should be noted that it is possible to revise our method to work with other modern measures of risk such as the *optimized certainty equivalent* [5], *distortion risk measures* (mixture of CVaR) [16, 27], and *law-invariant coherent risk measures* [27], in a straight forward manner since all of the above can be calculated by solving a convex minimization problem as was done for CVaR.

Early works in modern finance assumed that markets are stochastic and very simple (e.g., the returns are normally distributed) [40, 38]. This modeling assumption was later found to be too simplistic [33]. At the other extreme, Cover initiated the study of adversarial portfolio selection whereby stock prices are controlled by an adversary. Neither extreme led to overly effective strategies. It appears that a more sophisticated stochastic modeling, as we pursue here, can lead to effective strategies. However, despite the empirical success of these methods, there are two caveats, the first one is that the obtained bounds are asymptotic and the second is the computational deficiency of the method. To overcome those barriers, additional, and possibly strong, assumptions on the market process will be required. Moreover, it is not entirely clear whether such goals can be pursued without harming the empirical success of the methods. In the future, we wish to pursue those goals while not over-committing to dubious assumptions.

Another interesting future direction is designing a metric which is suitable for the problem of potfolio selection. In our paper and in [21] the euclidean metric was used. However, it was demonstrated in [30] that the Euclidean metric may not be suitable for OPS since it ignores profitable phenomenas such as mean-reversion. A careful design of similarity measure might lead to significant improvement in terms of final wealth.

A final caveat would be that the OPS framework ignores several factors which are essential for real paper trading (e.g., transaction costs, implementation shortfall and market impact). Incorporating and considering those factors is a difficult task and still considered to be an open problem. We leave those issues for future work.

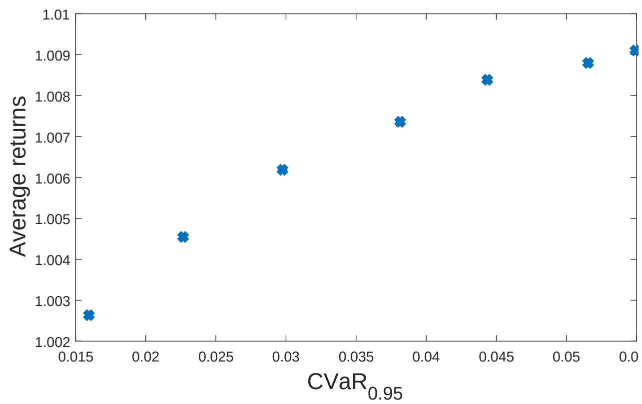


(a) MSCI dataset

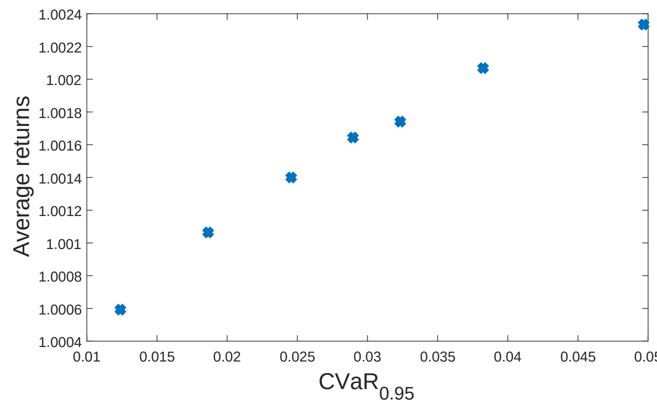


(b) NYSE dataset

Figure 1: Mean-CVaR trade-off

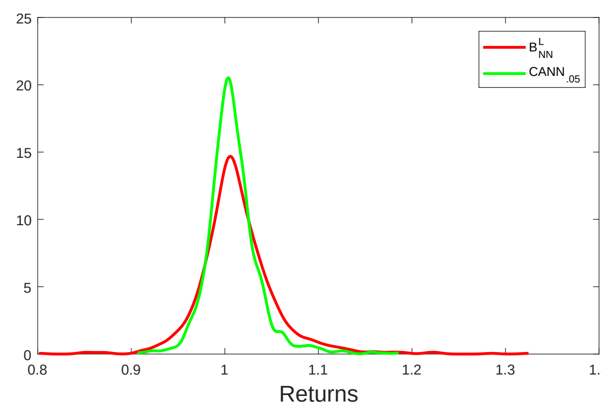


(a) MSCI dataset

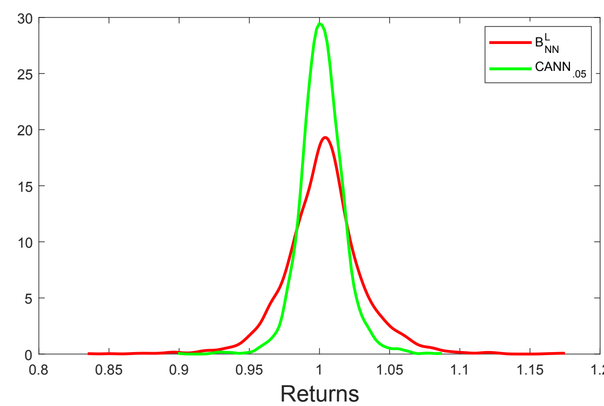


(b) NYSE dataset

Figure 2: The smoothed PDF

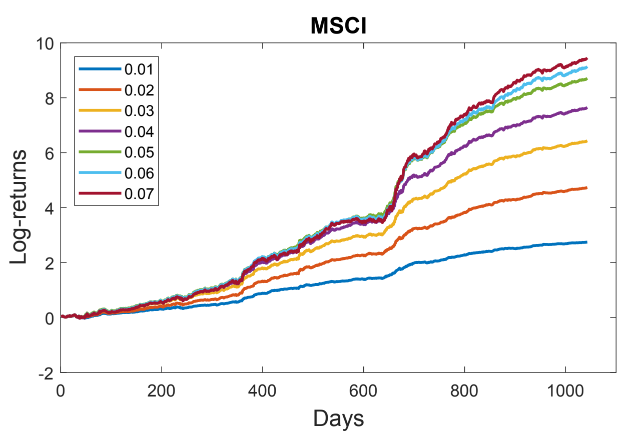


(a) MSCI dataset

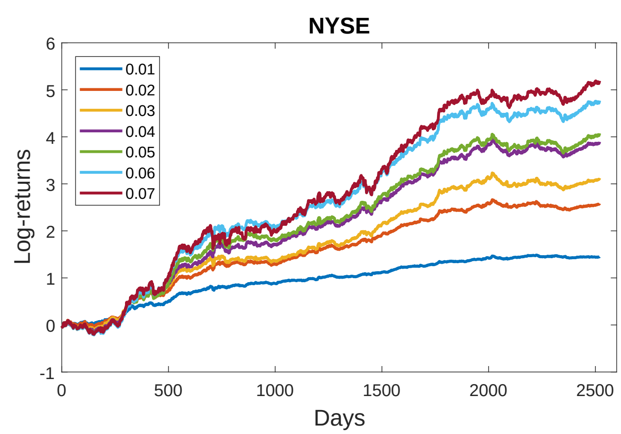


(b) NYSE dataset

Figure 3: The log-returns of CANN applied with different choices of $\gamma$.

## Acknowledgments

We would like to thank the anonymous reviewers for providing helpful comments. This research was supported by The Israel Science Foundation (grant No. 1890/14)

## References


[1] A. Agarwal, E. Hazan, S. Kale, and R.E. Schapire. Algorithms for portfolio management based on the newton method. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 9–16. ACM, 2006.

[2] P.H. Algoet. The strong law of large numbers for sequential decisions under uncertainty. *IEEE Transactions on Information Theory*, 40(3):609–633, 1994.

[3] P.H. Algoet and T.M. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *The Annals of Probability*, pages 876–898, 1988.

[4] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.

[5] A. Ben-Tal and M. Teboulle. An old-new concept of convex risk measures: the optimized cerrtainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007.

[6] Shalabh Bhatnagar, Doina Precup, David Silver, Richard S Sutton, Hamid R Maei, and Csaba Szepesvári. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Advances in Neural Information Processing Systems*, pages 1204–1212, 2009.

[7] G. Biau and B. Patra. Sequential quantile prediction of time series. *IEEE Transactions on Information Theory*, 57(3):1664–1674, 2011.

[8] A. Borodin, R. El-Yaniv, and V. Gogan. Can we learn to beat the best stock? *Journal of Artificial Intelligence Research*, pages 579–594, 2004.

[9] L. Breiman. The individual ergodic theorem of information theory. *The Annals of Mathematical Statistics*, 28(3):809–811, 1957.

[10] Leo Breiman. Probability, volume 7 of classics in applied mathematics. *Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA*, 1992.

[11] Y. Chow and M. Ghavamzadeh. Algorithms for cvar optimization in mdps. In *Advances in neural information processing systems*, pages 3509–3517, 2014.

[12] Y. Chow, A. Tamar, S. Mannor, and M. Pavone. Risk-sensitive and robust decision-making: a cvar optimization approach. In *Advances in Neural Information Processing Systems*, pages 1522–1530, 2015.

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

[14] L. Devroye, L. Györfi, and G. Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.

[15] E. Even-Dar, M. Kearns, and J. Wortman. Risk-sensitive online learning. In *Algorithmic Learning Theory*, pages 199–213. Springer, 2006.

[16] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and stochastics*, 6(4):429–447, 2002.

[17] L. Györfi, G. Lugosi, and G. Morvai. A simple randomized algorithm for sequential prediction of ergodic time series. *IEEE Transactions on Information Theory*, 45(7):2642–2650, 1999.

[18] L. Györfi, G. Lugosi, and F. Udina. Nonparametric kernel-based sequential investment strategies. *Mathematical Finance*, 16(2):337–357, 2006.

[19] L. Györfi, G. Ottucsák, and H. Walk. *Machine Learning for Financial Engineering*, volume 8. World Scientific, 2012.

[20] L. Györfi and D. Schäfer. Nonparametric prediction. *Advances in Learning Theory: Methods, Models and Applications*, 339:354, 2003.

[21] L. Györfi, F. Udina, and H. Walk. Nonparametric nearest neighbor based empirical portfolio selection strategies. *Statistics & Decisions, International Mathematical Journal for Stochastic Methods and Models*, 26(2):145–157, 2008.

[22] W. Haskell, H. Xu, Q. Chao, and Y. Zhiyue. Online risk-aware optimization. 2016.

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

[24] R. Jenatton, J. Huang, and C. Archambeau. Adaptive algorithms for online convex optimization with long-term constraints. In *International Conference on Machine Learning*, pages 402–411, 2016.

[25] N. Johnson and A. Banerjee. Structured hedging for resource allocations with leverage. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 477–486. ACM, 2015.

[26] Y. Kalnishkan and M. Vyugin. The weak aggregating algorithm and weak mixability. In *International Conference on Computational Learning Theory*, pages 188–203. Springer, 2005.

[27] S. Kusuoka. On law invariant coherent risk measures. In *Advances in Mathematical Economics*, pages 83–95. Springer, 2001.

[28] B. Li and S.C.H. Hoi. On-line portfolio selection with moving average reversion. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 273–280, 2012.

[29] B. Li and S.C.H. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3):35, 2014.

[30] B. Li, S.C.H Hoi, and V. Gopalkrishnan. Corn: Correlation-driven nonparametric learning approach for portfolio selection. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):21, 2011.

[31] B. Li, D. Sahoo, and S.C.H. Hoi. Olps: A toolbox for online portfolio selection. *Journal of Machine Learning Research (JMLR)*, 2015.

[32] Bin Li and Steven Chu Hong Hoi. *Online Portfolio Selection: Principles and Algorithms*. CRC Press, 2015.

[33] A. Lo and A. MacKinlay. *A non-random walk down Wall Street*. Princeton University Press, 2002.

[34] D. Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.

[35] U.V. Luxburg and B. Schölkopf. Statistical learning theory: Models, concepts, and results. *arXiv preprint arXiv:0810.4752*, 2008.

[36] M. Mahdavi, R. Jin, and T. Yang. Trading regret for efficiency: online convex optimization with long term constraints. *Journal of Machine Learning Research*, 13(Sep):2503–2528, 2012.

[37] M. Mahdavi, T. Yang, and R. Jin. Stochastic convex optimization with multiple objectives. In *Advances in Neural Information Processing Systems*, pages 1115–1123, 2013.

[38] B.G. Malkiel. *A random walk down Wall Street*. WW Norton & Company, 1999.

[39] S. Mannor, J. Tsitsiklis, and J.Y. Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10(Mar):569–590, 2009.

[40] H. Markowitz. Portfolio selection*. *The journal of finance*, 7(1):77–91, 1952.

[41] R. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

[42] William F Sharpe. Adjusting for risk in portfolio performance measurement. *The Journal of Portfolio Management*, 1(2):29–34, 1975.

[43] V. Vovk. Competing with stationary prediction strategies. In *International Conference on Computational Learning Theory*, pages 439–453. Springer, 2007.