

# An Adaptive News-Driven Method for CVaR-sensitive Online Portfolio Selection in Non-Stationary Financial Markets

Qianqiao Liang<sup>1</sup>, Mengying Zhu<sup>1</sup>, Xiaolin Zheng<sup>1\*</sup> and Yan Wang<sup>2</sup>

<sup>1</sup> College of Computer Science, Zhejiang University, Hangzhou, China

<sup>2</sup> Department of Computing, Macquarie University, Sydney, NSW, Australia

{liangqq, mengyingzhu, xlzheng}@zju.edu.cn, yan.wang@mq.edu.au

## Abstract

CVaR-sensitive online portfolio selection (CS-OLPS) becomes increasingly important for investors because of its effectiveness to minimize conditional value at risk (CVaR) and control extreme losses. However, the non-stationary nature of financial markets makes it very difficult to address the CS-OLPS problem effectively. To address the CS-OLPS problem in non-stationary markets, we propose an effective news-driven method, named CAND, which adaptively exploits news to determine the adjustment tendency and adjustment scale for tracking the dynamic optimal portfolio with minimal CVaR in each trading round. In addition, we devise a filtering mechanism to reduce the errors caused by the noisy news for further improving CAND’s effectiveness. We rigorously prove a sub-linear regret of CAND. Extensive experiments on three real-world datasets demonstrate CAND’s superiority over the state-of-the-art portfolio methods in terms of returns and risks.

## 1 Introduction

CVaR, an effective risk measure, stands for Conditional Value at Risk. In recent years, CVaR-sensitive online portfolio selection (CS-OLPS) has become increasingly important for risk-averse investors, which sequentially tracks the optimal portfolio, i.e., portfolio that has minimal CVaR, to control extreme losses and ensure stable returns. However, in the real world, financial markets are non-stationary, i.e., the distribution of assets’ returns constantly change. Such non-stationary nature of financial markets causes the constant changes of a portfolio’s CVaR, which makes it very difficult to address the CS-OLPS problem effectively. In this paper, we study the CS-OLPS problem in non-stationary markets, which aims to track the optimal portfolio dynamically.

In financial markets, the non-stationary changes of a portfolio’s CVaR may be the effect of related news. In Figure 1, we visualize such changes via a real-world example. Specifically, this example focuses on two stocks, i.e., Boeing and Zoom, and their portfolios from Feb 3, 2020 to Feb 19, 2020

\*Corresponding Author

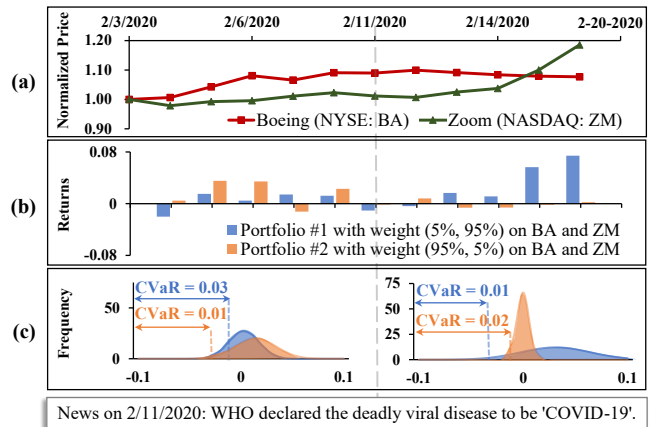


Figure 1: A motivating example. (a) Two stocks’ daily normalized prices. (b) Two portfolios’ daily returns. (c) CVaR estimated by the returns of the two portfolios before and after Feb 11, 2020.

when the COVID-19 epidemic has influenced the market. Suppose that portfolio #1 and portfolio #2 maintain two portfolio weights of (5%, 95%) and (95%, 5%) on Boeing and Zoom respectively on each day. In Figure 1, after the news “WHO declared the deadly viral disease to be ‘COVID-19’” was released after Feb 11, 2020, the aircraft manufacturer Boeing’s prices has responded negatively because of more restrictions on traveling, whereas the online meeting platform Zoom’s prices have responded positively (see Figure 1(a)) for the work-from-home policy. These different responses led to the different fluctuation changes of the two portfolios’ returns (see Figure 1 (b)), which indicated the significant changes of the two portfolios’ CVaR (see Figure 1 (c)). The CS-OLPS problem in non-stationary financial markets aims to select the optimal portfolio with minimal CVaR in each trading round.

The above real-world example motivates us to exploit news to capture the changes of portfolios’ CVaR for addressing the CS-OLPS problem in non-stationary financial markets. However, such idea faces three significant challenging problems. In each trading round, investors can only observe limited rewards of the selected portfolio, which are insufficient to estimate its CVaR precisely and further decide the adjustment tendency towards the low-CVaR portfolio. Therefore, (CH1) it is difficult to determine the adjustment tendency when the

*CVaR is unobservable.* In addition to determining the adjustment tendency, investors still need to determine an appropriate adjustment scale based on the non-stationary degree of the market [Zhang *et al.*, 2018]. Generally, the more violent the market changes within the entire investment horizon, the larger the adjustment scale are required in order to “keep up with” the potential changes. However, it is difficult to know in advance the non-stationary degree in hindsight, as a result, **(CH2)** *it is difficult to determine the adjustment scale when the non-stationary degree of the market is unobservable.* Moreover, even though news can provide hint about the non-stationary market changes to some extent, trusting all the obtained news may introduce bias in a news-driven portfolio selection strategy because some pieces of news may be noisy. In other words, **(CH3)** *it is difficult to determine which pieces of news to exploit when some of them are noisy.*

In order to address the above three challenging problems, we propose a novel CVaR-sensitive Adaptive News Driven method, named **CAND**, which constructs effective portfolios for the CS-OLPS problem in non-stationary markets. We devise three key modules in CAND, i.e., the adjustment tendency module, the adjustment scale module, and the filtering module. In each trading round, the *adjustment tendency module* estimates a CVaR-related function and decides the news-driven adjustment tendency based on this function (for addressing CH1). Then, the *adjustment scale module* tracks the model’s performance under several adjustment scales to keep up with CVaR changes (for addressing CH2). Under each adjustment scale, the *filtering module* filters the news-driven adjustment tendency based on a novel mechanism with confidence level to reduce the errors caused by the noisy news, which significantly improves the performance of the constructed portfolio (for addressing CH3). At the end of the current round, CAND constructs a portfolio based on the above three modules and updates itself based on the portfolio’s return, aiming to construct better portfolio in the next round.

To sum up, our main contributions are as follows: (1) We devise an effective news-driven method, named CAND, which adaptively leverages news to address the CS-OLPS problem in non-stationary financial markets; (2) We have rigorously proven a sub-linear upper regret bound of CAND, i.e., the sub-optimal portfolios from CAND reduces rapidly over time; (3) We have conducted extensive experiments on three real-world datasets to demonstrate CAND’s superiority over the state-of-the-art portfolio selection methods in terms of returns and risks.

## 2 Related Work

Studies have shown that minimizing portfolios’ CVaR contributes to safe and robust returns with relatively low losses [Uziel and El-Yaniv, 2018; Nakagawa *et al.*, 2020]. The CVaR-sensitive portfolio selection methods can be classified into three types. The first type replaces variance in the mean-variance optimization model [Markowitz, 1952] with CVaR, which is common for the offline portfolio selection problem [Iyengar and Ma, 2013]. The second type maximizes the return under a certain CVaR constraints, which appears in both offline [Sharma *et al.*, 2017] and online [Uziel and El-Yaniv,

2018] portfolio selection studies. The third type formulates the problem of minimizing CVaR as an online convex optimization problem [Cardoso and Xu, 2019], which can be applied to the CS-OLPS problem. However, all of them assume a stationary CVaR of the same portfolio during the trading horizon, which is too ideal in real non-stationary markets.

In the non-stationary financial markets, news-driven online portfolio selection model has become increasingly popular, which unveils the market dynamics and provides hints to adjust portfolios’ weights. Hu *et al.* [2018] proposed a deep neural network, which leverages news to capture assets’ prices changes. Ye *et al.* [2020] proposed a reinforcement learning method to exploit the news in non-stationary markets. Du *et al.* [2020] proposed a news embedding method to quantify the changing covariance of assets. In addition to the news-driven studies, some studies on dynamic online convex optimization are effective in the non-stationary environment [Yang *et al.*, 2016; Zhang *et al.*, 2018]. Their models choose appropriate adjustment scales to keep up with the market changes although they do not specifically focus on the portfolio selection problem. Nevertheless, the above studies do not consider minimizing CVaR, so they cannot be directly applied to the CS-OLPS problem in non-stationary markets.

## 3 Methodology

In this section, we present the problem settings and preliminaries of our proposed method CAND, an overview of CAND, and the details of three key modules in CAND.

### 3.1 Problem Settings and Preliminaries

We consider a CS-OLPS problem with  $M$  assets during  $T$  rounds. Let  $\mathbf{p}_t = (p_{t,1}, \dots, p_{t,M})$ , where  $p_{t,i}$  is the closing price of asset  $i$  in round  $t$ . Let  $\mathbf{r}_t = (r_{t,1}, \dots, r_{t,M})$  be the price growth rates, where  $r_{t,i} = \frac{p_{t,i} - p_{t-1,i}}{p_{t-1,i}} \in [C_1, C_2]$  with two constants  $C_1 \leq C_2$ . We define CVaR as follows.

**Definition 3.1** (CVaR). *Given a parameter  $\alpha \in [0, 1]$ , the CVaR of a portfolio  $\mathbf{w}$  in the trading round  $t$  is defined as*

$$\text{CVaR}_t(\mathbf{w}) = \min_z F(\mathbf{w}, z) = \min_z z + \frac{1}{\alpha} \mathbb{E}_{-\mathbf{w}^\top \mathbf{r}_t} [-\mathbf{w}^\top \mathbf{r}_t - z]_+, \quad (1)$$

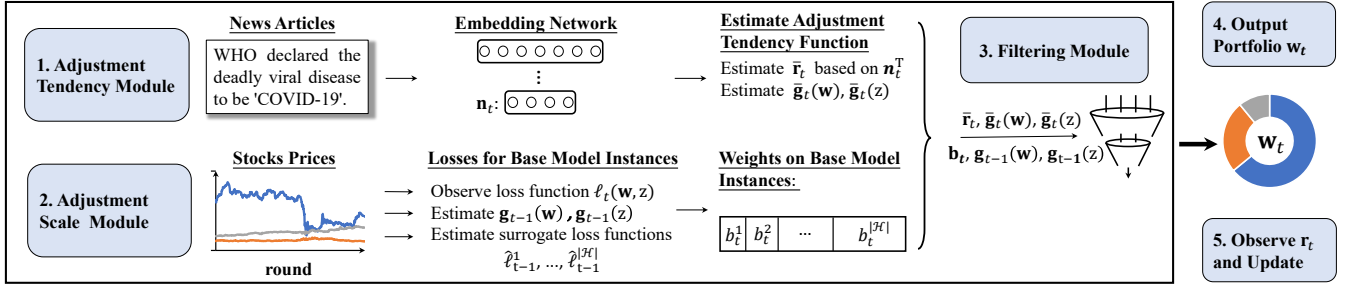
where the portfolio  $\mathbf{w} \in \mathcal{W} = \{\mathbf{w} | \mathbf{w} \geq 0, \mathbf{w}^\top \mathbf{1} = 1\} \subset \mathbb{R}^M$  and  $z \in \mathcal{Z} = [-C_2, -C_1]$  is an auxiliary parameter.

However, it is hard to observe or precisely estimate a portfolio  $\mathbf{w}_t$ ’s CVaR in round  $t$  because we only have limited observations of the portfolio’s loss, i.e.,  $-\mathbf{w}_t^\top \mathbf{r}_t$ , in Eq. (1). To address this issue, we introduce an observable CVaR-related loss function [Rockafellar *et al.*, 2000] in Definition 3.2.

**Definition 3.2** (CVaR-related loss function). *Using the same notations as those of Definition 3.1, the CVaR-related loss function in the trading round  $t$  is defined as:*

$$\ell_t(\mathbf{w}, z) = \frac{1}{\alpha} [-\mathbf{w}^\top \mathbf{r}_t - z]_+ + z. \quad (2)$$

With the above definitions, we formulate the CS-OLPS problem as an online convex optimization (OCO) problem. The base model for this OCO problem is defined as follows.


 Figure 2: Graphical representation of the CAND in round  $t$ .

**Definition 3.3** (Base model). Let  $\eta_h$  be the adjustment scale of a base model  $h$ . The decision process is based on online gradient descent, which is described as follows. At the beginning of trading round  $t$ , based on  $\eta_h$ ,  $\mathbf{r}_1, \dots, \mathbf{r}_{t-1}$ , and the previous CVaR-related loss functions  $\ell_1(\mathbf{w}, z), \dots, \ell_{t-1}(\mathbf{w}, z)$ , the base model constructs a portfolio  $\mathbf{w}_t$  and updates the auxiliary parameter  $z_t$ , aiming to minimize  $\sum_{\tau=1}^{t-1} \ell_{\tau}(\mathbf{w}, z)$ . At the end of round  $t$ ,  $\mathbf{r}_t$  and  $\ell_t(\mathbf{w}, z)$  are observed for the decision process in the next round.

Besides, we follow some widely used assumptions in online portfolio selection [Li and Hoi, 2014], i.e., no margin/short, unlimited market liquidity, and zero market impact.

### 3.2 Overview of CAND

To address the CS-OLPS problem in non-stationary markets, CAND adaptively maintains a set  $\mathcal{H}$  containing multiple instances of the base model, each of which selectively utilizes the news-driven adjustment tendency as a descent direction to construct its portfolio based on its unique adjustment scale. CAND consists of three key modules, i.e., the adjustment tendency module, the adjustment scale module, and the filtering module. The graphical representation of CAND is shown in Figure 2 and the pseudocode is presented in Algorithm 2.

Before the trading begins, CAND initializes its parameters and  $|\mathcal{H}|$  base model instances. At each round  $t$ , CAND processes in five step: (1) leverages news to determine the news-driven adjustment tendency, (2) leverages the previous performance of each base model instance to determine the weights on all instances in the adjustment scale module, (3) filters the news-driven adjustment tendency for each base model instance based on the the confidence level under its adjustment scale in the filtering module; (4) constructs and outputs the portfolio  $\mathbf{w}_t$ ; (5) observes assets' price growth rates for model's update.

In the following three subsections, we elaborate CAND's three key modules respectively in a trading round  $t$ .

### 3.3 Adjustment Tendency Module

In CAND, the adjustment tendency function  $-\mathbf{g}_t(\mathbf{w})$  is denoted by  $-\nabla_{\mathbf{w}} \ell_t(\mathbf{w}, z)$ , which is unobservable until the end of round  $t$ . Therefore, we need to estimate this function in round  $t$ . Because news can guide the adjustment tendency of a portfolio, CAND exploits news to estimate this function, which is beneficial to improve the estimation accuracy.

Specifically, CAND gets the news features in round  $t$  via a neural language processing model similar to the settings in [Ye *et al.*, 2020]. We will elaborate the news feature generation process in Section 5.1. Then, CAND estimates the assets' price growth rates  $\bar{\mathbf{r}}_t$  via an online multivariate ridge regression model [Abbasi-Yadkori *et al.*, 2011]. Based on  $\bar{\mathbf{r}}_t$ , CAND estimates the CVaR-related loss function as  $\bar{\ell}_t(\mathbf{w}, z) = \frac{1}{\alpha}[-\mathbf{w}^T \bar{\mathbf{r}}_t - z]_+ + z$ . Then, CAND estimates the gradients of  $\bar{\ell}_t(\mathbf{w}, z)$ , i.e.,  $\bar{\mathbf{g}}_t(\mathbf{w}) = \nabla_{\mathbf{w}} \bar{\ell}_t(\mathbf{w}, z)$  and  $\bar{\mathbf{g}}_t(z) = \nabla_z \bar{\ell}_t(\mathbf{w}, z)$ , in which  $-\bar{\mathbf{g}}_t(\mathbf{w})$  is the adjustment tendency function, which is applied in each base model instance.

### 3.4 Adjustment Scale Module

In real markets, we cannot know the non-stationary degree of the market in advance, which makes it very difficult to select an appropriate adjustment scale to keep up with potential market changes. Therefore, CAND maintains a pool of multiple candidate adjustment scales, and then invoke multiple instances of the base model simultaneously, where each instance  $h$  is associated with a candidate adjustment scales  $\eta_h$  and constructs  $\mathbf{w}_t^h$  and  $z_t^h$ . CAND adaptively integrates the portfolios from all instances based on a weight vector  $\mathbf{b}_t \in \mathbb{R}^{|\mathcal{H}|}$ .

Maintaining multiple base model instances needs to query  $\ell_{t-1}(\mathbf{w}_t^h, z_t^h)$  for every  $h \in \mathcal{H}$ , which is inefficient in the decision process. To reduce the number of queries to one, we introduce the surrogate loss. Let  $\mathbf{g}_{t-1}(\mathbf{w}) = \nabla_{\mathbf{w}} \ell_{t-1}(\mathbf{w}, z)$  and  $\mathbf{g}_{t-1}(z) = \nabla_z \ell_{t-1}(\mathbf{w}, z)$  denote the true gradient functions of the previous loss function. According to the first-order condition of Eq. (2)'s convexity [Cesa-Bianchi and Lugosi, 2006], we have  $\forall \mathbf{x}, \mathbf{y} \in \mathcal{W} \times \mathcal{Z}$ ,  $\ell_t(\mathbf{x}) \geq \ell_t(\mathbf{y}) + \langle \nabla \ell_t(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$ . Then, the surrogate loss for each instance  $h \in \mathcal{H}$  in round  $t-1$  is constructed as

$$\hat{\ell}_{t-1}^h = \mathbf{g}_{t-1}^T(\mathbf{w}_{t-1})(\mathbf{w}_{t-1}^h - \mathbf{w}_{t-1}) + \mathbf{g}_{t-1}(z_{t-1})(z_{t-1}^h - z_{t-1}), \quad (3)$$

where  $\mathbf{w}_{t-1}$  and  $z_{t-1}$  are the average portfolio and auxiliary parameter from all instances weighted by  $\mathbf{b}_t$ . With these surrogate losses, CAND updates the probability vector  $\mathbf{b}_t$  based on the exponential weighting model [Cesa-Bianchi and Lugosi, 2006] (Line 8, Algorithm 2).

### 3.5 Filtering Module

In this module, CAND utilizes the news-driven adjustment tendency as a descent direction to amend the update of each base model instance. However, fully trusting the news-driven

**Algorithm 1: Update and Filter (UF)**

**Input:** Index of the instance  $h$ , previous portfolio  $\mathbf{w}_{t-1}^h$ , previous auxiliary parameter  $z_{t-1}^h$ , adjustment scale  $\eta_h$ , estimation error  $\epsilon_t$ , predicted price growth rates  $\bar{\mathbf{r}}_t$ , gradient functions  $\bar{\mathbf{g}}_t(\mathbf{w})$ ,  $\bar{\mathbf{g}}_t(z)$ ,  $\mathbf{g}_{t-1}(\mathbf{w})$  and  $\mathbf{g}_{t-1}(z)$ .  $M, \delta, \alpha$  are defined the same as Algorithm 2.

**Output:** Portfolio  $\mathbf{w}_t^h$  and auxiliary parameter  $z_t^h$

- 1 Estimate  $\bar{\mathbf{w}}_t^h = \Pi_{\mathcal{W}}[\mathbf{w}_{t-1}^h - \eta_h \mathbf{g}_{t-1}(\mathbf{w}_{t-1}^h)]$  and  $\bar{z}_t^h = \Pi_{\mathcal{Z}}[z_{t-1}^h - \eta_h \mathbf{g}_{t-1}(z_{t-1}^h)]$ .
- 2 Amend  $\mathbf{w}_t^h = \Pi_{\mathcal{W}}[\bar{\mathbf{w}}_t^h - \eta_h \bar{\mathbf{g}}(\bar{\mathbf{w}}_t^h)]$  and  $z_t^h = \Pi_{\mathcal{Z}}[\bar{z}_t^h - \eta_h \bar{\mathbf{g}}(\bar{z}_t^h)]$ .
- 3 Estimate  $\bar{\mathbf{g}}_t(\bar{\mathbf{w}}_t^h)$  and  $\bar{\mathbf{g}}_t(\bar{z}_t^h)$ , and estimate  $\epsilon_t^h$  by Eq. (5).
- 4 **if**  $\|(\mathbf{w}_t^h, z_t^h) - (\bar{\mathbf{w}}_t^h, \bar{z}_t^h)\|_2 \leq \frac{\epsilon_t^h \eta_h}{2} + \sqrt{\frac{(\epsilon_t^h)^2 \eta_h^2}{2} + \eta_h \delta}$   
and  $\sqrt{|\bar{\mathbf{g}}_t(\bar{\mathbf{w}}_t^h)|^2 + |\bar{\mathbf{g}}_t(\bar{z}_t^h)|^2} \geq \epsilon_t^h$  **then return**  
 $\mathbf{w}_t^h, z_t^h$ .
- 5 **else return**  $\bar{\mathbf{w}}_t^h, \bar{z}_t^h$ .

adjustment tendency may be too optimistic because some news may be noisy. Therefore, we devise a mechanism to filter the news-driven adjustment tendency for each instance.

Specifically, each instance  $h$  first updates its portfolio to be  $\bar{\mathbf{w}}_t^h = \Pi_{\mathcal{W}}[\mathbf{w}_{t-1}^h - \eta_h \mathbf{g}_{t-1}(\mathbf{w}_{t-1}^h)]$  and the auxiliary parameter to be  $\bar{z}_t^h = \Pi_{\mathcal{Z}}[z_{t-1}^h - \eta_h \mathbf{g}_{t-1}(z_{t-1}^h)]$ . Note that  $\Pi_{\mathcal{W}}$  and  $\Pi_{\mathcal{Z}}$  refer to the Euclidean projection into  $\mathcal{W}$  and into  $\mathcal{Z}$  respectively. After that, CAND selectively determine whether to accept the news-driven amendment by  $\mathbf{w}_t^h = \Pi_{\mathcal{W}}[\bar{\mathbf{w}}_t^h - \eta_h \bar{\mathbf{g}}(\bar{\mathbf{w}}_t^h)]$  and  $z_t^h = \Pi_{\mathcal{Z}}[\bar{z}_t^h - \eta_h \bar{\mathbf{g}}(\bar{z}_t^h)]$  based on a novel mechanism. This mechanism consists of three steps: (1) quantifies the estimated price growth rates' error  $\epsilon_t$  caused by the noisy news; (2) estimates the error  $\epsilon_t^h$  of each instance  $h$  such that  $\epsilon_t^h \geq |\bar{\mathbf{g}}_t(\bar{\mathbf{w}}_t^h) - \mathbf{g}_t(\bar{\mathbf{w}}_t^h) + \bar{\mathbf{g}}_t(\bar{z}_t^h) - \mathbf{g}_t(\bar{z}_t^h)|$ ; and (3) filters the adjustment tendency for each instance.

First, we estimate  $\epsilon_t$  such that for all asset  $i$ ,  $|\bar{r}_{t,i} - r_{t,i}| \leq \epsilon_t$ . Because  $\bar{\mathbf{r}}_t$  is estimated by an online multivariate ridge regression, with probability  $1 - \beta$ ,

$$\epsilon_t = \|\mathbf{n}_t\|_{\mathbf{A}_t^{-1}} (\sqrt{2 \log(\sqrt{\det(\mathbf{A}_t)}/\sqrt{d}\beta)} + \sqrt{d}), \quad (4)$$

where  $\mathbf{A}_t = \sum_{\tau=1}^t \mathbf{n}_\tau \mathbf{n}_\tau^\top$  and  $d$  is the dimension of  $\mathbf{n}_t$ .

Second,  $\epsilon_t^h$  of each instance  $h$  can be estimated as

$$\epsilon_t^h = \begin{cases} 0 & \text{if } s_1 \leq \bar{z}_t^h \text{ and } s_2 \leq \bar{z}_t^h \\ \sqrt{\frac{M \epsilon_t^2}{\alpha^2}} & \text{if } s_1 \text{ and } s_3 > \bar{z}_t^h \\ \sqrt{\frac{|\bar{\mathbf{r}}|^2 + 1}{\alpha^2}} & \text{if } s_3 < \bar{z}_t^h < s_1 \\ \sqrt{\frac{|\bar{\mathbf{r}}|^2 + M \epsilon_t^2 + 2\sqrt{M \epsilon_t^2} |\bar{\mathbf{r}}| + 1}{\alpha^2}} & \text{otherwise,} \end{cases} \quad (5)$$

where  $s_1 = -\bar{\mathbf{w}}_t^h \bar{\mathbf{r}}$ ,  $s_2 = -\bar{\mathbf{w}}_t^h \bar{\mathbf{r}} - \epsilon_t \mathbf{1}$ , and  $s_3 = -\bar{\mathbf{w}}_t^h \bar{\mathbf{r}} + \epsilon_t \mathbf{1}$ .

Third, the adjustment tendency for each based model instance can be filtered by considering the integrated decisions  $\bar{\mathbf{x}}_t^h = (\bar{\mathbf{w}}_t^h, \bar{z}_t^h)$  based on Lemma 3.1. The detailed

**Algorithm 2: CAND**

**Input:**  $T$  trading rounds,  $M$  assets, news feature dimension  $d$ , filtering threshold  $\delta$ , confidence level of new-driven prediction  $\beta$ , CVaR's probability level  $\alpha$ .

**Output:** Portfolio  $\mathbf{w}_t$  in each trading round.

- 1 **Initialize:**  $G = \sqrt{\max(|C_1|, |C_2|)^2 + (1 - \alpha)^2} / \alpha$ ,  
 $D = \sqrt{1 + (C_2 - C_1)^2}$ ,  $\gamma = \sqrt{2/(TG^2 D^2)}$ ,  
 $N = \lceil \frac{1}{2} \log_2(1 + (4T)/7) \rceil + 1$ ,  $\mathcal{H} = \{1, 2, \dots, N\}$   
 $\nabla_{\mathbf{w}} \ell_0(\mathbf{w}, z) = \mathbf{0}_M$ ,  $\nabla_z \ell_0(\mathbf{w}, z) = 0$ ,  
 $\forall h \in \mathcal{H}$ : adjustment scale  $\eta_h = 2^{h-1} \sqrt{7} D / (G \sqrt{2T})$ ,  
 $\forall h \in \mathcal{H}$ ,  $b_0^h = \frac{(|\mathcal{H}|+1)/|\mathcal{H}|}{h(h+1)}$ ,  $z_0^h = 0$ ,  $\mathbf{w}_0^h = \{\frac{1}{M}, \dots, \frac{1}{M}\}$ ,  
 $\mathbf{A}_1 = \lambda \mathbf{I}_d$ ,  $\mathbf{B}_1 = \mathbf{0}_{d \times M}$ , and  $\mathbf{C}_1 = \mathbf{A}_1^{-1} \mathbf{B}_1$ .
- 2 **for**  $t = 1, 2, \dots, T$  **do**
  - # Adjustment tendency module (Section 3.3)
  - 3 Observe news and its feature  $\mathbf{n}_t$ .
  - 4  $\bar{\mathbf{r}}_t = \mathbf{n}_t^\top \mathbf{C}_t$  and  $\ell_t(\mathbf{w}, z) = \frac{1}{\alpha} [-\mathbf{w}^\top \bar{\mathbf{r}}_t - z]_+ + z$ .
  - 5  $\bar{\mathbf{g}}_t(\mathbf{w}) = \nabla_{\mathbf{w}} \ell_t(\mathbf{w}, z)$  and  $\bar{\mathbf{g}}_t(z) = \nabla_z \ell_t(\mathbf{w}, z)$ .
  - # Adjustment scale module (Section 3.4)
  - 6  $\mathbf{g}_{t-1}(\mathbf{w}) = \nabla_{\mathbf{w}} \ell_{t-1}(\mathbf{w}, z)$ ,  
 $\mathbf{g}_{t-1}(z) = \nabla_z \ell_{t-1}(\mathbf{w}, z)$ .
  - 7  $\forall h \in \mathcal{H}$ : estimate  $\hat{\ell}_{t-1}^h$  according to Eq. (3).
  - 8  $\forall h \in \mathcal{H}$ :  $b_t^h = \frac{b_{t-1}^h e^{-\gamma \hat{\ell}_{t-1}^h}}{(\sum_{i \in \mathcal{H}} b_{t-1}^i e^{-\gamma \hat{\ell}_{t-1}^i})}$ .
  - # Filtering module (Section 3.5)
  - 9 Estimate  $\epsilon_t$  according to Eq. (4)
  - 10  $\forall h \in \mathcal{H}$ :  $\mathbf{x}_t^h = (\mathbf{w}_t^h, z_t^h) =$   
UF( $\mathbf{w}_{t-1}^h, z_{t-1}^h, \eta_h, \epsilon_t, \bar{\mathbf{r}}_t, \bar{\mathbf{g}}_t(\mathbf{w}), \bar{\mathbf{g}}_t(z)$ ,  
 $\mathbf{g}_{t-1}(\mathbf{w}), \mathbf{g}_{t-1}(z), M, \delta, \alpha$ ).
  - 11  $\mathbf{w}_t = \sum_{h \in \mathcal{H}} b_t^h \mathbf{w}_t^h$  and  $z_t = \sum_{h \in \mathcal{H}} b_t^h z_t^h$ .
  - 12 Output portfolio  $\mathbf{w}_t$ , observe  $\mathbf{r}_t$  and  $\ell_t(\mathbf{w}_t, z_t)$ .
  - 13 Update news predicting parameters by  
 $\mathbf{A}_{t+1} = \mathbf{A}_t + \mathbf{n}_t \mathbf{n}_t^\top$ ,  $\mathbf{B}_{t+1} = \mathbf{B}_t + \mathbf{n}_t \mathbf{r}_t$ , and  
 $\mathbf{C}_{t+1} = \mathbf{A}_{t+1}^{-1} \mathbf{B}_{t+1}$ .

proof of Lemma 3.1 can be found in a longer version of this paper. The filtering mechanism based on Lemma 3.1 is to accept the adjustment tendency and the amendment  $\mathbf{x}_t^h = (\mathbf{w}_t^h, z_t^h)$  for each base model instance  $h$  only if  $\sqrt{|\bar{\mathbf{g}}_t(\bar{\mathbf{w}}_t^h)|^2 + |\bar{\mathbf{g}}_t(\bar{z}_t^h)|^2} \geq \epsilon_t^h$  and  $\|(\bar{\mathbf{w}}_t^h, \bar{z}_t^h) - \mathbf{x}_t^h\| \leq f(\epsilon_t^h)$ . The online updating and filtering process of each instance is presented in Algorithm 1.

**Lemma 3.1.** *In round  $t$ , if  $\sqrt{|\bar{\mathbf{g}}_t(\bar{\mathbf{w}}_t^h)|^2 + |\bar{\mathbf{g}}_t(\bar{z}_t^h)|^2} \geq \epsilon_t^h$  and  $\|(\bar{\mathbf{w}}_t^h, \bar{z}_t^h) - \mathbf{x}_t^h\| \leq \epsilon_t^h \eta_h / 2 + \sqrt{(\epsilon_t^h)^2 \eta_h^2 / 2 + \eta_h \delta} = f(\epsilon_t^h)$  for an instance  $h$ , where  $\delta > 0$  is the filtering threshold, then the amended update reduces the original loss by  $\delta$  with probability  $1 - \beta$ .*

Finally, CAND constructs portfolio  $\mathbf{w}_t = \sum_{h \in \mathcal{H}} b_t^h \mathbf{w}_t^h$ , and updates  $z_t = \sum_{h \in \mathcal{H}} b_t^h z_t^h$  in round  $t$ .

## 4 Theoretical Analysis

The performance of CAND in addressing the CS-OLPS problem in non-stationary markets is quantified by pseudo-CVaR dynamic regret. The definition and the result are as follows.

**Theorem 4.1.** *Let  $\mathbf{w}_t^* = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \operatorname{CVaR}_t(\mathbf{w})$  be the optimal portfolio and  $\mathbf{w}_t$  be a model’s portfolio in trading round  $t$ . Then, the pseudo-CVaR dynamic regret after  $T$  rounds is defined as  $R_T = \sum_{t=1}^T \operatorname{CVaR}_t(\mathbf{w}_t) - \sum_{t=1}^T \operatorname{CVaR}_t(\mathbf{w}_t^*)$ .*

*Let  $P_T = \sum_{t=2}^T |\mathbf{w}_t^* - \mathbf{w}_{t-1}^*|$  denote the changing scale of the optimal portfolio over  $T$  rounds. Let  $|\mathcal{H}| = N = \lceil \frac{1}{2} \log_2(1 + (4T)/7) \rceil + 1$ . Then, with probability  $1 - \beta$ , CAND’s pseudo-CVaR dynamic regret is upper bounded by*

$$R_T \leq \frac{GD\sqrt{2T}}{2} [1 + 2 \ln(k + 1)] + \frac{3G}{4} \sqrt{2T(7D^2 + 4D(P_T + Q_T))} - v\delta, \quad (6)$$

where  $G = \frac{\sqrt{\max(-C_1, C_2)^2 + (1-\alpha)^2}}{\alpha}$ ,  $k = \lceil \frac{1}{2} \log_2(1 + \frac{4(P_T + Q_T)}{7D}) \rceil + 1$ ,  $D = \sqrt{1 + (C_2 - C_1)^2}$ ,  $Q_T = \sum_{t=2}^T |z_t^* - z_{t-1}^*|$  with  $z_t^*$  be the optimal auxiliary parameter that minimize Eq.(2) in round  $t$ ,  $\delta$  is the filtering threshold, and  $v$  is the total number of rounds where all instances accept the news-driven adjustment tendency.

The detailed proof is presented in a longer version of this paper.

From Theorem 4.1, we can make three conclusions.

First, CAND’s pseudo-CVaR dynamic regret scales sub-linearly in  $T$  with order  $O(\sqrt{T})$ , which demonstrates that the number of sub-optimal portfolio from CAND decreases to zero when  $T$  goes to infinity. In other words, CAND can effectively update its portfolio construction over time for good performance in non-stationary financial markets.

Second, CAND’s pseudo-CVaR dynamic regret scales sub-linearly in  $P_T$  with order  $O(\sqrt{P_T})$ , which demonstrates that CAND can effectively track the optimal portfolio in non-stationary financial markets as long as the degree of market changes is sub-linear in  $T$ . Note that this result is achieved even though CAND does not know the  $P_T$  in advance, which indicates its adaptiveness to the market changes.

Third, the news-driven adjustment tendencies can further reduce CAND’s CVaR dynamic regret according to the last term in Eq. (6). According to the proofs in a longer version of this paper, the mechanism of the filtering module is more likely to accelerate the convergence of the dynamic regret at the beginning when the news-driven adjustment tendencies are more useful to base model instances.

## 5 Experiments

In this section, we present the comprehensive experiments conducted on three real-world datasets to answer the following questions: **Q1:** How does CAND perform on accumulating wealth (see Result 1)? **Q2:** How does CAND perform on reducing risk (see Result 2)? **Q3:** How does the three key modules contribute to CAND’s performance (see Result 3)?

Dataset	Time Period	# Rounds	# News	# Assets
SP500	10/20/2006 - 11/20/2013	1,784	49,725	500
DJIA	08/08/2008 - 11/01/2016	1,889	106,521	30
COVID	11/01/2019 - 12/31/2020	294	94,078	30

Table 1: Datasets descriptions

### 5.1 Experimental Settings

**Datasets Descriptions.** We provide a summary of the three real-world datasets of our experiments in Table 1. In particular, SP500 [Ding *et al.*, 2014] and DJIA<sup>1</sup> are widely used datasets for news-driven financial modeling. COVID<sup>2</sup> is a newly collected dataset, which contains news and historical prices of Dow Jones 30 constituent stocks. Among the three datasets, the SP500 and COVID datasets cover the 2007 - 2008 financial crisis and the COVID-19 recession respectively, which are the well-known periods of violent non-stationary markets, and the DJIA dataset covers some mini-recessions only and thus is not much non-stationary. In each dataset, we use the data in the earliest 30% days to train the news feature generation model for CAND and conduct the model training required by the comparison methods.

**News Feature Generation.** Similar to the settings in [Ye *et al.*, 2020], the news feature  $\mathbf{n}_t$  of trading round  $t$  is generated by (1) utilizing the unsupervised pre-trained Word2Vec model [Mikolov *et al.*, 2013] to embed the news articles during the previous  $t-12$  days, (2) feeding the embedding into a hybrid attention network (HAN) [Hu *et al.*, 2018], which is a trained classifier to predict the price movement of dataset’s index, and (3) extracting the last layer, i.e., a 100-dimensional vector, in HAN’s discriminative network as news feature  $\mathbf{n}_t$  in round  $t$ . In the above process, a dataset’s index refers to the equally weighted average of the price growth rates of the dataset’s constituent assets, and price movement refers to rise, down, or preserve of the index as defined in HAN. Notably, the main difference between our news feature generation and that of [Ye *et al.*, 2020] is that we substitute the classifier in [Ye *et al.*, 2020] by HAN because HAN is specifically designed for leveraging financial news, which can model sequential news dependency and diverse news influence in financial markets.

**Evaluation Metrics.** We use five standard metrics [Li *et al.*, 2012] to measure portfolios’ performance. Cumulative wealth (CW) measures the portfolios’ returns. Volatility (VO) and Maximum drawdown (MD) measure the portfolios’ risks while Sharpe ratio (SR) and Calmar ratio (CR) measure the risk-adjusted returns. Overall, higher CW, SR and CR indicate better performance while lower MD and VO values indicate better performance.

**Comparison Methods.** We compare CAND with methods in two categories.

<sup>1</sup>www.kaggle.com/aaron7sun/stocknews.

<sup>2</sup>News are collected from aylien.com and covid19-archive.com.




Figure 3: Results on cumulative wealth.

Dataset	Metrics	News-free methods			News-driven methods			Our proposed methods			Improvement <sup>1</sup>	
		MCVaR	SOCO	CANN	HAN	SARL	SE	CAND	Ablation methods			
								w/o-Tend w/o-Scale w/o-Filter				
SP500	VO	0.051	0.058	0.047*	0.051	0.057	0.048	<b>0.045</b>	0.051	0.074	0.062	3.94%
	MD	0.283	0.300	0.282*	0.283	0.283	0.452	<b>0.218</b>	0.248	0.294	0.337	22.69%
	SR	1.364*	1.312	1.085	1.230	1.262	0.855	<b>1.548</b>	1.345	0.843	0.875	13.49%
	CR	1.124	1.165	0.832	1.019	1.168*	1.005	<b>1.573</b>	1.280	0.969	0.739	34.64%
DJIA	VO	0.018	0.012*	0.032	0.017	0.060	0.012	<b>0.012</b>	0.012	0.013	0.012	1.73%
	MD	0.159	0.057*	0.266	0.122	0.486	0.058	<b>0.055</b>	0.056	0.059	0.058	4.25%
	SR	2.232	3.648	0.466	1.299	0.312	3.745*	<b>4.073</b>	3.695	3.539	3.971	8.76%
	CR	1.146	3.459*	0.255	0.834	0.177	3.440	<b>4.129</b>	3.704	3.659	3.752	19.35%
COVID	VO	0.027	0.028	0.026*	0.031	0.075	0.086	<b>0.025</b>	0.029	0.028	0.028	3.82%
	MD	0.026	0.026	0.024*	0.053	0.158	0.193	<b>0.020</b>	0.026	0.026	0.024	17.77%
	SR	6.076	6.083	7.062*	2.831	3.056	2.122	<b>10.060</b>	6.834	6.430	9.623	42.43%
	CR	28.937	29.378	37.150*	7.612	6.605	4.338	<b>64.781</b>	34.868	31.607	53.099	74.37%

\* Results of the best-performing comparison methods. <sup>1</sup> Improvement of CAND over the best-performing comparison methods.

Table 2: Results on risk, i.e., volatility (VO), maximum drawdown (MD), sharpe ratio (SR), and calmar ratio (CR).

(1) *News-free portfolio selection methods*: **MCVaR** [Iyengar and Ma, 2013] is a conventional offline CVaR-sensitive portfolio method. **SOCO** [Cardoso and Xu, 2019] is a state-of-the-art CVaR-sensitive portfolio method based on convex optimization. **CANN** [Uziel and El-Yaniv, 2018] is a state-of-the-art CS-OLPS method.

(2) *News-driven portfolio selection methods*: **HAN** [Hu *et al.*, 2018] is a conventional method. **SE** [Du and Tanaka-Ishii, 2020] is a state-of-the-art method that incorporates news into the mean-variance minimization model. **SARL** [Ye *et al.*, 2020] is a state-of-the-art method based on deep reinforcement learning.

**Parameter Settings.** Following the settings in [Uziel and El-Yaniv, 2018; Abbasi-Yadkori *et al.*, 2011], we set  $\alpha$  to be 0.05,  $\beta$  to be 0.0001. The news-filtering threshold  $\delta$  is set to be 0.06, which we found to give reasonable performance. The rest of CAND’s parameters are set to be their theoretically optimal values (see Line 1, Algorithm 2). The parameters of the comparison methods are set to be the optimal values reported in their publications. To model real-world financial markets, we follow the proportional transaction cost model and set the transaction cost rate to be 0.25% [Ye *et al.*, 2020].

## 5.2 Performance Comparison and Analysis

### Result 1: Performance on Cumulative Wealth (for Q1).

To answer question **Q1**, we present the CW values achieved by CAND and the six comparison methods in Figure 3. CAND achieves the highest CW values on SP500, DJIA, and COVID with improvements of 0.18%, 12.07%, and 10.64% respectively over the best-performing comparison methods. CAND outperforms all news-free comparison methods because the news-free comparison methods fail to adapt to non-stationary financial markets where CAND leverages news to adapt to non-stationary financial markets. Besides, CAND outperforms all news-driven methods because of its effective filtering mechanism to identify useful news.

### Result 2: Performance on Risk (for Q2).

To answer question **Q2**, we present the risk and risk-adjusted return metrics achieved by CAND and the six comparison methods in Table 2. In Table 2, CAND achieves the lowest VO and MD values on all three real-world datasets with an average improvement of 9.03% (ranging from 1.72% to 22.67%) over the best-performing comparison methods, which demonstrates that CAND performs well on reducing risks. In addition, CAND achieves the highest SR and CR values on all three real-world datasets with an average improvement of 32.18% (ranging

from 8.76% to 74.38%) over the best-performing comparison models, which demonstrates that CAND performs well on balancing returns and risks.

**Result 3: Ablation Experiments (for Q3).** To answer question Q3, we conduct ablation experiments with three simplified versions of CAND: (1) **w/o-Tend**: CAND without the adjustment tendency module, (2) **w/o-Scale**: CAND without the adjustment scale module, i.e., fixing the adjustment scale to be  $\frac{\sqrt{1+C_2-C_1}}{G} \sqrt{\frac{7}{2T+2}}$ , which corresponds to assuming the most fluctuated market according to the proof in Theorem 4.1, and (3) **w/o-Filter**: CAND without the filtering module, i.e., without using Lemma 3.1 to filter out noisy news-driven adjustment tendency for each base model instance. Results are shown in Figure 3 and Table 2. Comparing CAND with w/o-Tend, we find that CAND outperforms w/o-Tend on all five metrics with average improvements of 15.12%, 7.91%, and 28.67% on SP500, DJIA, and COVID respectively. Without the adjustment tendency estimating module, w/o-Tend fails to react to the market dynamics timely. Comparing CAND with w/o-Scale, we find that CAND outperforms w/o-Scale on all five metrics with average improvements of 53.61%, 11.22%, and 44.03% on SP500, DJIA, and COVID respectively. Without leveraging the adjustment scale selecting module, w/o-Scale can only fix a large adjustment-scale to prevent the most drastic market changes, which sacrifices its performance when the market changes slowly. Comparing CAND with w/o-Filter, we find that CAND outperforms w/o-Filter on all five metrics with average improvements of 70.31%, 5.47%, and 14.03% on SP500, DJIA, and COVID respectively, which demonstrate the effectiveness of CAND’s filtering module. Overall, the results in this section demonstrate that CAND’s three modules are indispensable and contribute to CAND’s promising performance.

## 6 Conclusions and Future Work

In this paper, we propose a novel news-driven model CAND, which adaptively leverages news to address the CS-OLPS in non-stationary financial markets. Rigorous theoretical analysis and empirical studies demonstrate the effectiveness of CAND. In the future, an interesting extension is to incorporate more sophisticated mechanism to further exploit some indirect relations between news and the CS-OLPS in non-stationary markets. Another extending direction is to explore more advanced news feature generation models such as sentiment-based [Xing *et al.*, 2018; Malandri *et al.*, 2018] and Bert-like models [Devlin *et al.*, 2019] to better exploit the information in financial news for improving the performance.

## Acknowledgments

This work was supported in part by the National Key R&D Program of China (No. 2018YFB1403001).

## References

[Abbasi-Yadkori *et al.*, 2011] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *The 25th Annual Conference*

*on Neural Information Processing Systems*, pages 2312–2320, 2011.

[Cardoso and Xu, 2019] Adrian Rivera Cardoso and Huan Xu. Risk-averse stochastic convex bandit. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 39–47, 2019.

[Cesa-Bianchi and Lugosi, 2006] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, 2019.

[Ding *et al.*, 2014] Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. Using structured events to predict stock price movement: an empirical investigation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1415–1425, 2014.

[Du and Tanaka-Ishii, 2020] Xin Du and Kumiko Tanaka-Ishii. Stock embeddings acquired from news articles and price history, and an application to portfolio optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3353–3363, 2020.

[Hu *et al.*, 2018] Ziniu Hu, Weiqing Liu, Jiang Bian, Xuanzhe Liu, and Tie-Yan Liu. Listening to chaotic whispers: a deep learning framework for news-oriented stock trend prediction. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 261–269, 2018.

[Iyengar and Ma, 2013] Garud Iyengar and Alfred Ka Chun Ma. Fast gradient descent method for mean-cvar optimization. *Annals of Operations Research*, 205(1):203–212, 2013.

[Li and Hoi, 2014] Bin Li and Steven C. H. Hoi. Online portfolio selection: a survey. *ACM Computing Surveys*, 46(3):1–36, 2014.

[Li *et al.*, 2012] Bin Li, Peilin Zhao, Steven C. H. Hoi, and Vivekanand Gopalkrishnan. PAMR: passive aggressive mean reversion strategy for portfolio selection. *Machine Learning*, 87(2):221–258, 2012.

[Malandri *et al.*, 2018] Lorenzo Malandri, Frank Z Xing, Carlotta Orsenigo, Carlo Vercellis, and Erik Cambria. Public mood-driven asset allocation: The importance of financial sentiment in portfolio management. *Cognitive Computation*, 10(6):1167–1176, 2018.

[Markowitz, 1952] Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):39–68, 1952.

[Mikolov *et al.*, 2013] Tomáš Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed

- representations of words and phrases and their compositionality. In *27th Annual Conference on Neural Information Processing Systems*, pages 3111–3119, 2013.
- [Nakagawa *et al.*, 2020] Kei Nakagawa, Shuhei Noma, and Masaya Abe. Rm-cvar: Regularized multiple  $\beta$ -cvar portfolio. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 4562–4568, 2020.
- [Rockafellar *et al.*, 2000] R. Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- [Sharma *et al.*, 2017] Amita Sharma, Sebastian Utz, and Aparna Mehra. Omega-cvar portfolio optimization and its worst case analysis. *OR spectrum*, 39(2):505–539, 2017.
- [Uziel and El-Yaniv, 2018] Guy Uziel and Ran El-Yaniv. Growth-optimal portfolio selection under cvar constraints. In *International Conference on Artificial Intelligence and Statistics*, pages 48–57, 2018.
- [Xing *et al.*, 2018] Frank Z Xing, Erik Cambria, Lorenzo Malandri, and Carlo Vercellis. Discovering bayesian market views for intelligent asset allocation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 120–135. Springer, 2018.
- [Yang *et al.*, 2016] Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: optimal dynamic regret of online learning with true and noisy gradient. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 449–457, 2016.
- [Ye *et al.*, 2020] Yunan Ye, Hengzhi Pei, Boxin Wang, Pin-Yu Chen, Yada Zhu, Ju Xiao, and Bo Li. Reinforcement-learning based portfolio management with augmented asset movement prediction states. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 1112–1119, 2020.
- [Zhang *et al.*, 2018] Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *The 31 Annual Conference on Neural Information Processing Systems*, pages 1330–1340, 2018.