COST-SENSITIVE ONLINE WINDOW SIZE SELECTION FOR PORTFOLIO MANAGEMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper investigates cost-sensitive online window size selection for portfolio management under changing market conditions. Specifically, we propose a two-level framework that constructs portfolios using candidate window sizes and dynamically aggregates them through online learning. By treating candidate window sizes as "experts," we dynamically update their aggregation weights using turnover-inclusive losses. Moreover, we derive finite-horizon cost-sensitive tracking-regret bounds that account for turnover of the aggregated portfolio, with static regret as a special case. Under bounded losses and cost rates, suitably tuned Fixed Share achieves asymptotically no tracking regret for sublinear switching budgets, with Hedge covering the static case.

***Keywords*—** Online Learning, Tracking Regret, Transaction Costs, Window Size Selection, Portfolio Selection, Fixed Share Algorithm

1 INTRODUCTION

A fundamental issue in financial forecasting is how much historical information to retain. Long windows can smooth short-term noise and capture persistent patterns, whereas short windows can respond more rapidly to recent information. This trade-off is particularly important in rolling portfolio management, where the choice of estimation window directly affects portfolio decisions. When portfolios constructed from different estimation windows are combined online, changing their aggregation weights can induce turnover even if the individual portfolios remain unchanged. Hence, each window expert's transaction costs do not, by themselves, account for the cost of the aggregate portfolio. This raises the question of how to track changing expert performance while controlling regret that includes the turnover of the aggregate portfolio.

In particular, financial time-series forecasting often involves complex and time-varying dynamics, making the choice of estimation window especially consequential; see, e.g., Peters (1994) Historically, window sizes in financial applications were often chosen empirically or based on forecasting performance. For example, Molodtsova & Papell (2009) used a fixed 10-year rolling window for monthly exchange-rate forecasting, and Stock & Watson (2007) used quarterly units to predict inflation. Pesaran & Timmermann (2007) proposed a forecaster that weighted averages with different window sizes to mitigate model uncertainty, a concept detailed in Pesaran & Pick (2011) and finalized with optimal weighting in Pesaran et al. (2013). Rossi & Inoue (2012) generally discussed the topic and provided a statistically robust window selection method. Inoue et al. (2017) identified the solution for window size that is asymptotically equivalent to minimizing Mean Squared Forecast Error (MSFE), which is theoretically unsolvable. However, these approaches do not formulate window-size selection as a sequential learning problem that adapts to evolving markets with performance guarantees.

Dynamic window selection has also been considered in financial prediction and portfolio problems. A common approach is to evaluate several candidate windows and select the best-performing one. For example, Jeon & McCurdy (2017) dynamically weighted different sub-window sizes via probability distributions to predict correlations between financial instruments, although the resulting computational burden limits frequent updating. Rajabi et al. (2022) used multi-layer perceptrons to determine dynamic windows for Bitcoin prediction, but focused on very short time horizons. Thus, a general framework for online window-size adaptation with rigorous performance guarantees remains lacking.

Parallel to this literature, online portfolio selection has developed extensively since the Universal Portfolio of Cover (1991), itself motivated by Kelly's growth-optimal betting framework Kelly (1956). Follow-the-Winner methods favor assets or portfolios that have performed well historically. Examples include Exponential Gradient (EG) Helmbold et al. (1998), the Aggregating Algorithm Vovk & Watkins (1998), and Variable Rebalanced Portfolios (VRP) Gaivoronski & Stella (2000). By contrast, Follow-the-Loser methods exploit mean reversion,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

including Passive Aggressive Mean Reversion (PAMR) Li et al. (2012), Confidence Weighted Mean Reversion (CWMR) Li et al. (2013), and OLMAR Li & Hoi (2014).

However, how to choose the window size, an important parameter for time series prediction, is less emphasized in online portfolio selection. Gaivoronski & Stella (2000) points out the need to consider sliding window size. Gaivoronski & Stella (2003) studies adaptive portfolio selection with transaction costs. Li & Hoi (2012) propose a window-size-sensitive model OLMAR for trading the portfolio. However, that work does not establish a tracking-regret bound that accounts for turnover of the aggregated portfolio.

We formulate window-size adaptation as the online aggregation of sliding-window portfolio experts within the prediction-with-expert-advice framework Cesa-Bianchi & Lugosi (2006). Each candidate window generates a portfolio, and the algorithm combines these portfolios using weights updated from the experts' turnover-inclusive losses. We employ Fixed Share Herbster & Warmuth (1998) to track changing expert performance, with Hedge Littlestone & Warmuth (1994) as the static-regret baseline.

Our main contribution is a finite-horizon tracking-regret bound that accounts for turnover of the deployed aggregate portfolio. The comparator is the best expert sequence under a prescribed switching budget, with each selected expert evaluated by its own turnover-inclusive loss. The resulting bound makes the dependence on the transaction-cost rate and switching budget explicit, yields a cost-dependent learning-rate choice, and recovers the static-regret bound for Hedge as a special case.

2 Problem Formulation

Consider a portfolio $\mathcal{M} := \{1, \dots, m\}$ with $m \geq 2$ assets, including a risk-free asset. Let $S_i(t) > 0$ denote the adjusted closing price of asset $i \in \mathcal{M}$ at time $t \geq 1$. For asset $i \in \mathcal{M}$, its return at time t is denoted by

$$y_i(t) := \frac{S_i(t+1) - S_i(t)}{S_i(t)}.$$

Let $\mathbf{y}(t) := [y_1(t), \dots, y_m(t)]^\top \in \mathbb{R}^m$ be the return vector for m assets at time $t \geq 1$.

2.1 Sliding-Window Expert and Cost-Sensitive Mean-Variance Model

We extend online window-size selection within the expert-advice framework Cesa-Bianchi & Lugosi (2006); Orabona (2026) to account for turnover induced by both expert portfolio updates and changes in aggregation weights. Let $\mathcal{N}$ be a finite set of $n \geq 2$ candidate window sizes, where each $j \in \mathcal{N}$ satisfies $j \geq 2$ and is treated as an *expert*. For expert $j \in \mathcal{N}$, consider Markowitz's Mean-Variance (MV) optimization problem with turnover costs to obtain the advice $\mathbf{w}_j(t) \in \mathcal{W}$; the weight vector optimizes the last j-days of data prior to time t; see Markowitz (1952); Luenberger (2013). Given an initial history of $\max \mathcal{N}$ returns, define the rolling estimators for expert $j \in \mathcal{N}$ by

$$\widehat{\boldsymbol{\mu}}_j(t) := \frac{1}{j} \sum_{s=t-j}^{t-1} \mathbf{y}(s) \in \mathbb{R}^m,$$

$$\widehat{\Sigma}_j(t) := \frac{1}{j-1} \sum_{s=t-j}^{t-1} \big(\mathbf{y}(s) - \widehat{\boldsymbol{\mu}}_j(t)\big)\big(\mathbf{y}(s) - \widehat{\boldsymbol{\mu}}_j(t)\big)^\top \in \mathbb{S}_+^m.$$

Here $\mathbb{S}_+^m$ denotes the set of real symmetric positive semidefinite $m \times m$ matrices.

Problem 2.1 (Cost-Sensitive Mean-Variance Optimization)**.** Let $\mathbf{w}_j(t-1) \in \mathcal{W}$ denote the portfolio weights held by expert j immediately before rebalancing at time t, with $\mathbf{w}_j(0) := \mathbf{w}_0$ for a prescribed initial portfolio $\mathbf{w}_0 \in \mathcal{W}$, and let $c(t) \geq 0$ be the proportional transaction cost rate. The *cost-sensitive mean-variance* model for expert j can be written as

$$\max_{\mathbf{w} \in \mathcal{W}} \ \widehat{\boldsymbol{\mu}}_j(t)^\top \mathbf{w} - \frac{1}{2} \mathbf{w}^\top \widehat{\Sigma}_j(t) \mathbf{w} - c(t) \|\mathbf{w} - \mathbf{w}_j(t-1)\|_1, \qquad j \in \mathcal{N} \tag{1}$$

where $\mathcal{W} := \left\{ \mathbf{w} \in \mathbb{R}_+^m : \mathbf{w}^\top \mathbf{1} = 1 \right\}$. Here $\mathbf{1}$ denotes the all-ones vector of the appropriate dimension.

Let $\mathbf{w}_j(t)$, the optimal advice solving the problem for j-day data at time t, be the advice for expert j. Note that Problem 2.1 is a concave program, which can be solved efficiently by a standard solver such as CVXPY; see Diamond & Boyd (2016).

After obtaining the expert advice, we assign an *aggregation weight vector* in the simplex:

$$\mathbf{q}(t) \in \Delta_n := \left\{ \mathbf{q} \in \mathbb{R}_+^n : \sum_{j \in \mathcal{N}} q_j = 1 \right\} \subseteq \mathbb{R}^n$$

and form the *aggregateportfolio*

$$\widehat{\mathbf{w}}(t) := \sum_{j \in \mathcal{N}} q_j(t)\mathbf{w}_j(t).$$

Thus, the method dynamically aggregates the window experts rather than necessarily selecting a single window. Here $\mathbf{w}_j(t)$ and $\widehat{\mathbf{w}}(t)$ are portfolio weight vectors over assets, whereas $\mathbf{q}(t)$ is the aggregation weight vector over window experts.

2.2 Tracking Regret

Consider a sequential decision-making framework over a discrete time horizon $t = 1, 2, \ldots$. Let $\ell : \mathcal{W} \times \mathcal{Y} \to \mathbb{R}$ be a loss function, assumed to be convex in its first argument. At each time t, the algorithm first selects a decision $\widehat{\mathbf{w}}(t) \in \mathcal{W}$. After the decision is made, the market outcome $\mathbf{y}(t) \in \mathcal{Y} \subseteq \mathbb{R}^m$ is revealed, and the algorithm incurs the prediction loss $\ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t))$.

Recall that each $j \in \mathcal{N}$ indexes the expert associated with a j-day sliding window; at each time t, this expert provides advice $\mathbf{w}_j(t) \in \mathcal{W}$. For a finite horizon $T \geq 1$, define the set of all expert sequences over T periods as

$$\mathcal{N}^T := \underbrace{\mathcal{N} \times \cdots \times \mathcal{N}}_{T \text{ times}}.$$

Write $\mathbf{j} := (j_1, \ldots, j_T) \in \mathcal{N}^T$ for an arbitrary comparator sequence of experts, where j_t is the expert selected by the comparator at time t. For an expert sequence $\mathbf{j} = (j_1, \ldots, j_T) \in \mathcal{N}^T$, define its number of switches by

$$S_T(\mathbf{j}) := \sum_{t=1}^{T-1} \mathbb{1}_{\{j_t \neq j_{t+1}\}},$$

where $\mathbb{1}_A$ denotes the indicator function, equal to 1 if the event A occurs and 0 otherwise. For a switching budget $K \in \{0, 1, \ldots, T-1\}$, define

$$\mathcal{J}_{T,K} := \left\{ \mathbf{j} \in \mathcal{N}^T : S_T(\mathbf{j}) \leq K \right\},$$

which represents the *comparator class*, consisting of all expert sequences that switch at most K times over the horizon T. The switching budget K is prescribed for each horizon T and may depend on T. We use K in finite-horizon statements and write $K(T)$ when considering the limit $T \to \infty$. In asymptotic statements, "fixed K" means that the budget is independent of T.

Definition 2.2. For a given horizon $T \geq 1$ and a switching budget $K \in \{0, 1, \ldots, T-1\}$, the *tracking regret* with switching budget K is defined as

$$R_{\text{track}}(T, K) := \sum_{t=1}^{T} \ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) - \min_{\mathbf{j} \in \mathcal{J}_{T,K}} \sum_{t=1}^{T} \ell(\mathbf{w}_{j_t}(t), \mathbf{y}(t)).$$

Each sequence $\mathbf{j} \in \mathcal{J}_{T,K}$ partitions $\{1, \ldots, T\}$ into at most $K+1$ consecutive blocks, within each of which the comparator follows one fixed expert. The minimization over $\mathcal{J}_{T,K}$ selects the best such expert sequence in hindsight; see Cesa-Bianchi & Lugosi (2006).

Remark 2.3 (Static Regret as a Special Case). When $K = 0$, the comparator must follow one fixed expert throughout the horizon; i.e., $\mathcal{J}_{T,0} = \{\mathbf{j} = (j, j, \ldots, j) \in \mathcal{N}^T : j \in \mathcal{N}\}$. Hence, static regret is the special case

$$R_{\text{static}}(T) := R_{\text{track}}(T, 0) = \sum_{t=1}^{T} \ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) - \min_{j \in \mathcal{N}} \sum_{t=1}^{T} \ell(\mathbf{w}_j(t), \mathbf{y}(t)).$$

2.3 Regret Minimization Problem

Our goal is to dynamically aggregate window-size experts and control the cost-sensitive tracking regret relative to the best expert sequence under a prescribed switching budget. We explicitly account for the turnover of the aggregated portfolio in the regret criterion. Assume further that $a \leq \ell(\mathbf{w}, \mathbf{y}) \leq b$ for all $(\mathbf{w}, \mathbf{y}) \in \mathcal{W} \times \mathcal{Y}$, where $a < b$ are fixed finite constants. Let $c(t) \geq 0$ denote the *transaction cost rate*. We now formalize our main problem.

Problem 2.4 (Cost-Sensitive Regret Minimization)**.** Define the *cost-sensitive loss of expert j at time t* as their prediction loss plus their incurred transaction cost:

$$\tilde{\ell}_j(t) := \ell(\mathbf{w}_j(t), \mathbf{y}(t)) + c(t)\|\mathbf{w}_j(t) - \mathbf{w}_j(t-1)\|_1.$$

Given a time horizon $T \geq 1$ and a switching budget $K \in \{0, \ldots, T-1\}$, our goal is to select $\mathbf{q}(t) \in \Delta_n$ before $\mathbf{y}(t)$ is revealed and establish a worst-case upper bound on the resulting *cost-sensitive tracking regret*:

$$R^c_{\text{track}}(T, K) := \sum_{t=1}^{T} [\ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) + c(t)\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1] - \min_{\mathbf{j} \in \mathcal{J}_{T,K}} \sum_{t=1}^{T} \tilde{\ell}_{j_t}(t),$$

where $\widehat{\mathbf{w}}(0) = \mathbf{w}_0$ and $\mathcal{J}_{T,K}$ consists of all expert sequences that switch at most K times over the horizon T.

Hannan Consistency for Cost-Sensitive Tracking Regret. Beyond finite-horizon regret bounds, we seek asymptotic no regret. We formalize this objective through a cost-sensitive tracking analog of Hannan consistency Hannan (1957).

Definition 2.5 (Cost-Sensitive Hannan Consistency)**.** Given a prescribed switching-budget sequence $K(T) \in \{0, 1, \ldots, T-1\}$, we say that an algorithm is *Hannan consistent for cost-sensitive tracking regret* with respect to $\{\mathcal{J}_{T,K(T)}\}_{T\geq 1}$ if, for every outcome sequence $\{\mathbf{y}(t)\}_{t\geq 1} \subseteq \mathcal{Y}$ and every transaction-cost sequence $\{c(t)\}_{t\geq 1} \subseteq [0, 1]$,

$$\limsup_{T\to\infty} \frac{R^c_{\mathrm{track}}(T, K(T))}{T} \leq 0.$$

Remark 2.6 (Role of the Switching Budget). The switching budget restricts the comparator sequences, not the evolution of market returns. For asymptotic analysis, we consider prescribed budgets satisfying

$$K(T) = o(T). \tag{2}$$

No-regret guarantees for these budgets depend on the loss assumptions and the algorithm's parameter choices.

3 Theoretical Results

3.1 Cost-Sensitive Regret Bounds

Using the cost-sensitive expert losses defined in Problem 2.4, we establish the following regret decomposition for any aggregation rule. For notational convenience, set $\mathbf{q}(0) := \mathbf{q}(1)$.

Theorem 3.1 (Cost-Sensitive Tracking Regret Bound)**.** *For any time horizon $T \geq 1$, let $c(t) \geq 0$ be a transaction cost rate at time t. Consider any algorithm $\mathcal{A}$ that generates aggregation weights $\mathbf{q}(t) \in \Delta_n$. For any minimizing benchmark sequence of experts in $\mathcal{J}_{T,K}$, the cost-sensitive tracking regret $R^c_{\mathrm{track}}(T, K)$ satisfies*

$$R^c_{\mathrm{track}}(T, K) \leq \underbrace{\sum_{t=1}^{T} \left[\left(\sum_{j\in\mathcal{N}} q_j(t)\tilde{\ell}_j(t) \right) - \tilde{\ell}_{j_t}(t) \right]}_{\text{Expert-loss regret bound}} + \underbrace{\sum_{t=1}^{T} c(t)\|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1}_{\text{Additional turnover cost regret bound}}. \tag{3}$$

The coefficient 1 of the additional turnover cost bound is sharp; i.e., it cannot be reduced uniformly over all admissible expert portfolios and aggregation rules.

Proof. Let $\widehat{\mathbf{w}}(t) = \sum_{j\in\mathcal{N}} q_j(t)\mathbf{w}_j(t)$ denote the aggregate portfolio. Define the combined loss of expert j at time t as their prediction loss plus their incurred transaction cost at step t:

$$\tilde{\ell}_j(t) = \ell(\mathbf{w}_j(t), \mathbf{y}(t)) + c(t)\|\mathbf{w}_j(t) - \mathbf{w}_j(t-1)\|_1.$$

Because the original loss is bounded in $[a, b]$ and the maximum l_1 distance between any two vectors on the probability simplex is 2, the combined loss at time t lies in the interval $\tilde{\ell}_j(t) \in [a, b+2c(t)]$. The cost-sensitive tracking regret against a minimizing benchmark sequence of optimal experts $j_1, \ldots, j_T$ with at most K shifts is defined as:

$$R^c_{\mathrm{track}}(T, K) = \sum_{t=1}^{T} [\ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) + c(t)\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1] - \sum_{t=1}^{T} \tilde{\ell}_{j_t}(t). \tag{4}$$

Since the loss function ℓ is convex in its first argument, Jensen's inequality guarantees that:

$$\ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) = \ell\left(\sum_{j\in\mathcal{N}} q_j(t)\mathbf{w}_j(t), \mathbf{y}(t) \right) \leq \sum_{j\in\mathcal{N}} q_j(t)\ell(\mathbf{w}_j(t), \mathbf{y}(t)). \tag{5}$$

On the other hand, observe that $c(t)\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1$ in (4) can be written as

$$c(t)\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1 = c(t) \left\| \sum_{j\in\mathcal{N}} q_j(t)\mathbf{w}_j(t) - \sum_{j\in\mathcal{N}} q_j(t-1)\mathbf{w}_j(t-1) \right\|_1.$$

By adding/subtracting $\sum q_j(t)\mathbf{w}_j(t-1)$ yields:

$$
\begin{aligned}
& c(t)\|\widehat{\mathbf{w}}(t)-\widehat{\mathbf{w}}(t-1)\|_1 \\
&= c(t)\left\|\sum_{j\in\mathcal{N}} q_j(t)\mathbf{w}_j(t)-\sum q_j(t)\mathbf{w}_j(t-1)+\sum q_j(t)\mathbf{w}_j(t-1)-\sum_{j\in\mathcal{N}} q_j(t-1)\mathbf{w}_j(t-1)\right\|_1 \\
&= c(t)\left\|\sum_{j\in\mathcal{N}} q_j(t)\left(\mathbf{w}_j(t)-\mathbf{w}_j(t-1)\right)+\sum_{j\in\mathcal{N}}\left(q_j(t)-q_j(t-1)\right)\mathbf{w}_j(t-1)\right\|_1 \\
&\le c(t)\sum_{j\in\mathcal{N}} q_j(t)\|\mathbf{w}_j(t)-\mathbf{w}_j(t-1)\|_1+c(t)\sum_{j\in\mathcal{N}}|q_j(t)-q_j(t-1)|\|\mathbf{w}_j(t-1)\|_1. \\
&= c(t)\sum_{j\in\mathcal{N}} q_j(t)\|\mathbf{w}_j(t)-\mathbf{w}_j(t-1)\|_1+c(t)\|\mathbf{q}(t)-\mathbf{q}(t-1)\|_1. \qquad (6)
\end{aligned}
$$

where the second-to-last inequality follows from the triangle inequality and $q_j(t)\in[0,1]$.

Because each expert's portfolio $\mathbf{w}_j(t-1)\in\mathcal{W}$, which lies on the probability simplex, its l_1-norm is exactly 1. Therefore, combining the (5) and the (6) allows us to reconstruct the expected combined loss $\tilde{\ell}_j(t)$:

$$
\ell(\widehat{\mathbf{w}}(t),\mathbf{y}(t))+c(t)\|\widehat{\mathbf{w}}(t)-\widehat{\mathbf{w}}(t-1)\|_1\le\sum_{j\in\mathcal{N}} q_j(t)\tilde{\ell}_j(t)+c(t)\|\mathbf{q}(t)-\mathbf{q}(t-1)\|_1. \qquad (7)
$$

After substituting (7) into (4) and rearranging, the statement is proved.

To establish sharpness, consider two assets and two experts, relabeled as 1 and 2 for this example, with $T=2$, $\ell\equiv 0$, and $c(1)=0<c(2)$. Let the two experts hold $\mathbf{e}_1$ and $\mathbf{e}_2$, respectively, at both periods; i.e., $\mathbf{w}_1(t)=\mathbf{e}_1$ and $\mathbf{w}_2(t)=\mathbf{e}_2$ for all $t=1,2$. Take $\mathbf{q}(1)=(1,0)$ and $\mathbf{q}(2)=(0,1)$ Then $\widehat{\mathbf{w}}(1)=\sum_{j=1}^2 q_j(1)\mathbf{w}_j(1)=\mathbf{e}_1$ and $\widehat{\mathbf{w}}(2)=\sum_{j=1}^2 q_j(2)\mathbf{w}_j(2)=\mathbf{e}_2$. Hence, for $j=1,2$,

$$
\tilde{\ell}_j(1)=0+c(1)\|\mathbf{e}_j-\mathbf{w}_0\|_1=0 \quad \text{and} \quad \tilde{\ell}_j(2)=0+c(2)\|\mathbf{e}_j-\mathbf{e}_j\|_1=0. \qquad (8)
$$

Hence, any admissible benchmark sequence has zero cumulative expert loss: $\min_{\mathbf{j}\in\mathcal{J}_{2,K}}\sum_{t=1}^2\tilde{\ell}_{j_t}(t)=0$. The corresponding cost-sensitive tracking regret

$$
\begin{aligned}
R^c_{\text{track}}(2,K) &= \underbrace{c(1)\|\widehat{\mathbf{w}}(1)-\widehat{\mathbf{w}}(0)\|_1}_{=0}+c(2)\|\widehat{\mathbf{w}}(2)-\widehat{\mathbf{w}}(1)\|_1-0 \\
&= 0+c(2)\|\mathbf{e}_2-\mathbf{e}_1\|_1 \\
&= 2c(2),
\end{aligned}
$$

On the other hand, applying the setting above to the derived regret bound (3) yields

$$
\begin{aligned}
&\sum_{t=1}^2\left(\sum_{j=1}^2 q_j(t)\tilde{\ell}_j(t)-\tilde{\ell}_{j_t}(t)\right)+\sum_{t=1}^2 c(t)\|\mathbf{q}(t)-\mathbf{q}(t-1)\|_1 \\
&=\left(\tilde{\ell}_1(1)-\tilde{\ell}_{j_1}(1)\right)+\left(\tilde{\ell}_2(2)-\tilde{\ell}_{j_2}(2)\right)+c(2)\|\mathbf{q}(2)-\mathbf{q}(1)\|_1 \\
&=0+0+2c(2)
\end{aligned}
$$

where the last equality holds by (8). We see hence that the cost-senstive tracking regret attains the derived bound, which completes the proof. □

3.2 Fixed Share for Window Selection

We now specialize the aggregation rule to the Fixed Share algorithm Herbster & Warmuth (1998), with aggregation weights updated from the turnover-inclusive expert losses $\tilde{\ell}_j(t)$. Initialize $q_j(1)=1/n$ for every $j\in\mathcal{N}$. After observing $\mathbf{y}(t)$, the algorithm updates the aggregation weights for stage $t+1$ in two steps:

Step 1: Exponential weighting. For expert j at time t, compute the intermediate weight $q_j^{\exp}(t)$ using the rule:

$$
q_j^{\exp}(t)=\frac{q_j(t)e^{-\eta\tilde{\ell}_j(t)}}{\sum_{k\in\mathcal{N}} q_k(t)e^{-\eta\tilde{\ell}_k(t)}}, \qquad (9)
$$

where $\eta>0$ is the learning rate.[1]

[1]The learning rate η controls the trade-off between exploiting historically strong experts and adapting to changes in expert performance. A higher η means faster adaptation. Conversely, a smaller η leads to more stable weights that change slowly.

Step 2: Sharing. For a mixing parameter $\alpha \in [0,1]$, retain a $(1-\alpha)$ fraction of the intermediate weight and redistribute the remaining α fraction across all n experts uniformly:

$$\begin{aligned} q_j(t+1) &:= (1-\alpha) q_j^{\exp}(t) + \frac{\alpha}{n} \sum_{k \in \mathcal{N}} q_k^{\exp}(t) \\ &= (1-\alpha) q_j^{\exp}(t) + \frac{\alpha}{n}, \end{aligned} \tag{10}$$

where the last equality follows from $\sum_{k \in \mathcal{N}} q_k^{\exp}(t) = 1$.

Remark 3.2. The Hedge Algorithm Littlestone & Warmuth (1994) is the special case of Fixed Share with $\alpha = 0$, for which $q_j(t+1) = q_j^{\exp}(t)$, exactly the (9). Hedge competes with the best fixed expert in hindsight. For $\alpha > 0$, the mixing step in Fixed Share guarantees $q_j(t+1) \geq \alpha/n$, ensuring a positive weight for every expert.

3.3 Fixed Share Regret Bounds

To apply Theorem 3.1 to Fixed Share, we first bound the variation of its aggregation weights.

Lemma 3.3 (One-Step Weight Variation). *Let $\mathbf{q}(t) = (q_j(t))_{j \in \mathcal{N}}$ be generated by (9) and (10) with learning rate $\eta > 0$ and parameter $\alpha \in [0,1]$. Then, for every $t \geq 2$, we have*

$$\|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1 \leq \eta \left(\max_{j \in \mathcal{N}} \tilde{\ell}_j(t-1) - \min_{j \in \mathcal{N}} \tilde{\ell}_j(t-1) \right) + 2\alpha \left(1 - \frac{1}{n}\right).$$

Proof. The proof is given in Appendix A. □

Combining Theorem 3.1 and Lemma 3.3 with the standard Fixed Share regret bound for the losses $\tilde{\ell}_j(t)$ yields the following result.

Corollary 3.4 (Cost-Sensitive Tracking Regret Bound for Fixed Share). *Assume Theorem 3.1 holds. Consider the Fixed Share algorithm with mixing parameter $\alpha \in [0,1)$, the cost-sensitive tracking regret $R^c_{\text{track}}(T,K)$ satisfies*

$$\begin{aligned} R^c_{\text{track}}(T,K) \leq \frac{\mathcal{C}(K,\alpha,n,T)}{\eta} + \frac{\eta}{8} \sum_{t=1}^{T} (b-a+2c(t))^2 \\ + \eta \sum_{t=2}^{T} c(t)(b-a+2c(t-1)) + 2\alpha \left(1 - \frac{1}{n}\right) \sum_{t=1}^{T} c(t), \end{aligned}$$

where the term $\mathcal{C}(K,\alpha,n,T)$ is given by

$$\mathcal{C}(K,\alpha,n,T) := \begin{cases} \log n & \text{if } \alpha = 0 \text{ and } K = 0 \\ \infty & \text{if } \alpha = 0 \text{ and } K > 0 \\ (K+1)\log n + K \log \frac{1}{\alpha} + (T-K-1)\log \frac{1}{1-\alpha} & \text{otherwise.} \end{cases}$$

Proof. Recall Theorem 3.1, for the cost-sensitive tracking regret we have

$$R^c_{\text{track}}(T,K) \leq \underbrace{\sum_{t=1}^{T} \left[\left(\sum_{j \in \mathcal{N}} q_j(t) \tilde{\ell}_j(t) \right) - \tilde{\ell}_{j_t}(t) \right]}_{\text{Expert-loss regret bound}} + \underbrace{\sum_{t=1}^{T} c(t) \|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1}_{\text{Additional turnover cost regret bound}}.$$

For the Regret Bound for $\mathcal{A}$, evaluated on losses in $[a, b+2c(t)]$, by Cesa-Bianchi & Lugosi (2006), standard analysis for the Fixed Share algorithm bounds the regret by:

$$\text{Expert-loss regret bound } \leq \frac{\mathcal{C}(K,\alpha,n,T)}{\eta} + \frac{\eta}{8} \sum_{t=1}^{T} (b-a+2c(t))^2.$$

For the Regret Bound for Transaction Cost, using $\mathbf{q}(0) = \mathbf{q}(1)$ and Lemma 3.3, we obtain

$$\begin{aligned}\text{Additional Turnover cost regret bound} &= \sum_{t=1}^{T} c(t)\|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1 \\ &= 0 + \sum_{t=2}^{T} c(t)\|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1 \\ &\leq \sum_{t=2}^{T} c(t)\left[\eta\left(\max_{j\in\mathcal{N}} \tilde{\ell}_j(t-1) - \min_{j\in\mathcal{N}} \tilde{\ell}_j(t-1)\right) + 2\alpha\left(1 - \frac{1}{n}\right)\right] \\ &\leq \eta \sum_{t=2}^{T} c(t)(b - a + 2c(t-1)) + 2\alpha\left(1 - \frac{1}{n}\right)\sum_{t=1}^{T} c(t).\end{aligned}$$

The last inequality follows from $\tilde{\ell}_j(t-1) \in [a, b + 2c(t-1)]$ and $c(t) \geq 0$.

Combining both components yields the final regret bound. □

3.4 Parameter Selection and Asymptotic No-Regret

We can further optimize the cost-sensitive tracking regret bounds obtained in Corollary 3.4 when the transaction costs $c(t)$ are known for the entire horizon T.

Proposition 3.5 (Optimal Regret Bound)**.** *Assume the conditions of Corollary 3.4 hold, with $T \geq 2$ and $K \in \{0, \ldots, T-2\}$. Define $A_T := \frac{1}{8}\sum_{t=1}^{T}(b - a + 2c(t))^2 + \sum_{t=2}^{T} c(t)(b - a + 2c(t-1))$, and $C_T := \sum_{t=1}^{T} c(t)$. Let $H(x) = -x\log x - (1-x)\log(1-x)$ be the binary entropy function defined for $x \in [0, 1]$ with the convention $0 \log 0 := 0$. The cost-sensitive tracking regret $R^c_{\text{track}}(T, K)$ is optimally bounded as follows:*

Case 1: For $C_T = 0$. With the optimal mixing parameter $\alpha^* = \frac{K}{T-1} \in [0, 1)$ and the optimal learning rate η is chosen as:

$$\eta^* = \frac{1}{b-a}\sqrt{\frac{8}{T}\left((K+1)\log n + (T-1)H\left(\frac{K}{T-1}\right)\right)},$$

the resulting bound is

$$R^c_{\text{track}}(T, K) \leq (b-a)\sqrt{\frac{T}{2}\left((K+1)\log n + (T-1)H\left(\frac{K}{T-1}\right)\right)}.$$

Case 2: For $C_T > 0$ and $K = 0$. With $\alpha^* = 0$ and $\eta^* = \sqrt{\frac{\log n}{A_T}}$, the bound is

$$R^c_{\text{track}}(T, 0) \leq 2\sqrt{A_T \log n}.$$

Case 3: For $C_T > 0$ and $K \geq 1$. Let $\alpha^* \in \left(0, \frac{K}{T-1}\right)$ be the unique solution to the equation

$$\frac{K - (T-1)\alpha}{\alpha(1-\alpha)} = 2\left(1 - \frac{1}{n}\right) C_T \sqrt{\frac{\mathcal{C}(K, \alpha, n, T)}{A_T}}.$$

For $\eta^* = \sqrt{\frac{\mathcal{C}(K,\alpha^*,n,T)}{A_T}}$, the bound is

$$R^c_{\text{track}}(T, K) \leq 2\sqrt{A_T \cdot \mathcal{C}(K, \alpha^*, n, T)} + 2\left(1 - \frac{1}{n}\right)\alpha^* C_T.$$

Proof. The proof is given in Appendix A. □

Remark 3.6 (Cost-Sensitive Static Regret Bound). Setting $K = 0$ and $\alpha = 0$ in Corollary 3.4 and Proposition 3.5 yields the cost-sensitive static regret for the Hedge algorithm and its optimal learning rate.

Corollary 3.7 (Asymptotic Cost-Sensitive No-Regret)**.** *Assume the conditions of Corollary 3.4, with n, a, b fixed and $c(t) \in [0, 1]$ for all t. Let $K = K(T) \in \{0, \ldots, T-1\}$ satisfy $K(T) = o(T)$. For each sufficiently large T, choose the parameters as in Proposition 3.5. Then the cost-sensitive tracking regret satisfies*

$$\limsup_{T\to\infty} \frac{R^c_{\text{track}}(T, K(T))}{T} \leq 0.$$

Proof. Since $c(t) \in [0,1]$, we have $A_T = \mathcal{O}(T)$ and $C_T \le T$. For sufficiently large T, the choice $\bar{\alpha} = K/(T-1)$ belongs to $[0,1)$. By the definition of $\mathcal{C}$, using its first case when $K = 0$ and its last case when $K > 0$, we obtain

$$B_T := \mathcal{C}(K, \bar{\alpha}, n, T) = (K+1)\log n + (T-1)H\left(\frac{K}{T-1}\right),$$

where $H(0) = 0$. Since $K = o(T)$ and $H(x) \to 0$ as $x \to 0$, we have $B_T = o(T)$. The optimized upper bound is no larger than the bound obtained at $\bar{\alpha}$ and $\bar{\eta} = \sqrt{B_T/A_T}$. Hence,

$$\frac{R^c_{\text{track}}(T,K)}{T} \le 2\sqrt{\frac{A_T}{T}\frac{B_T}{T}} + 2\left(1 - \frac{1}{n}\right)\frac{K}{T-1}\frac{C_T}{T}.$$

Both terms on the right converge to zero. □

4 Simulation Results

After making decision $\widehat{\mathbf{w}}(t) \in \mathcal{W}$ at time t, the trader's account value evolves according to

$$\begin{aligned} V(t+1) &= V(t) + V(t)\big(\widehat{\mathbf{w}}(t)^\top \mathbf{y}(t)\big) - c\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1 V(t) \\ &= V(t)\big(1 + \widehat{\mathbf{w}}(t)^\top \mathbf{y}(t) - c\,\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1\big) \end{aligned}$$

where $V(1) > 0$ and $\widehat{\mathbf{w}}(0) = \mathbf{w}_0$. For this evaluation, we assume $1 + \widehat{\mathbf{w}}(t)^\top \mathbf{y}(t) - c\|\widehat{\mathbf{w}}(t) - \widehat{\mathbf{w}}(t-1)\|_1 > 0$ at every trading period. This condition ensures positive account values.

4.1 Synthetic Data

The Setup. To illustrate our two-level framework, we first construct a synthetic environment that includes a prescribed regime shift. We simulate daily prices for an artificial market of 32 stocks (`SYN_1` to `SYN_32`) over 2,000 business days, beginning January 1, 2020. For this portfolio setting, we define our reference expert set as $\mathcal{N} = \{5, 10, 21, 63\}$. The transaction cost rate is set to $c = 0.001$ (10 bps).

Data Generating Process. Let $\boldsymbol{\mu} \in \mathbb{R}^{32}$ be the baseline daily expected return, $\alpha \in (0,1)$ control the ratio of signal, and $\epsilon(t)$ is s noise vector where $\epsilon(t) \sim \mathcal{N}(0, \sigma^2)$. Let $\mathbf{y}(t) \in \mathbb{R}^{32}$ denote the return vector of the stocks at time t, which is defined as:

$$\mathbf{y}(t) = \boldsymbol{\mu} + \alpha\left(\frac{1}{j_t}\sum_{k=1}^{j_t}\mathbf{y}(t-k) - \boldsymbol{\mu}\right) + \epsilon(t)$$

where $j_t \in \mathcal{N}$ represents the target window size. This construction induces a controlled change in the horizon that governs the conditional mean, allowing us to examine how the proposed cost-sensitive online aggregation framework responds to a known regime shift. In our setup, the market undergoes a sharp regime shift precisely halfway through the dataset: for the first 1,000 days, the returns are driven by a $j_t = 63$, and $j_t = 5$ for the remaining days.

We then execute the Hedge and Fixed Share algorithm. Setting the cost-sensitive loss function to negative PnL, Profit and Loss, with rate $c \ge 0$:

$$\ell(\mathbf{f}_j(t), \mathbf{y}(t)) = -\mathbf{f}_j(t)^\top \mathbf{y}(t) + c\|\mathbf{f}_j(t) - \mathbf{f}_j(t-1)\|_1.$$

Since the number of experts n, the time horizon T, and the transaction cost rate c are given, all the parameters needed can be optimized. Besides the two-level framework we set, we compare our work with some standard online portfolio selection benchmarks, including UP (Cover (1991)), EG (Helmbold et al. (1998)), PAMR (Li et al. (2012)), CWMR (Li et al. (2013)) and OLMAR (Li & Hoi (2014). For these comparators, we initialize their parameters to the theoretically optimal values prescribed in their respective foundational papers. Note that these standard configurations were originally derived under frictionless assumptions. We evaluate them under these standard configurations not to claim they are fully optimized for our specific setting, but to illustrate the structural necessity of incorporating cost-sensitive directly into the financial optimization process.

Performance Evaluation. Figure 1 shows the cost-sensitive regret defined in Problem 2.4 with synthetic data. Although both algorithms exhibit a temporary rising in cost-sensitive regret at the designated moments of optimal expert switching, they rapidly recover and restabilize. CHHSIEH: where is the spike? Notably, the Fixed Share algorithm reallocates its weights to the newly dominant expert faster than Hedge.

Figure 3 further shows the evolution of the aggregation weights $q_j(t)$, with the thickness of each colored band representing the weight assigned to one window-size expert. Building upon this baseline, we introduce variations to the experimental parameters: Figure 4 illustrates the effects of altering the timing of the regime shift, while Figure 2 demonstrates the impact of expanding the expert set to $\mathcal{N} = \{5, 10, 21, 30, 41, 50, 63, 126\}$.

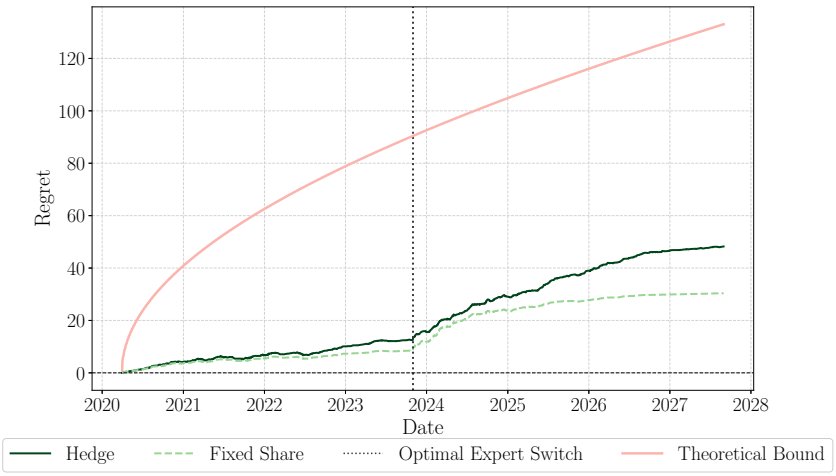

Figure 1: Cost-Sensitive Tracking Regret for Hedge and Fixed Share with Negative PnL Loss

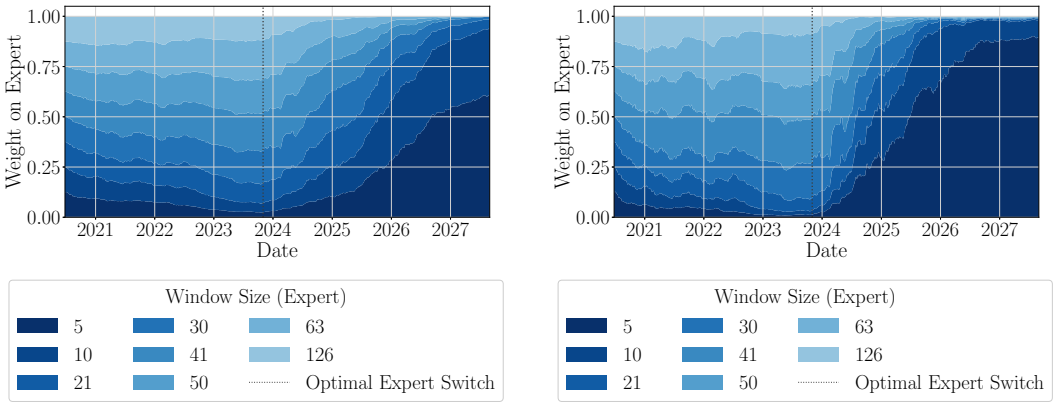

(a) Hedge

(b) Fixed Share

Figure 2: Evolution of the aggregation weights $q_j(t)$ on eight window sizes.

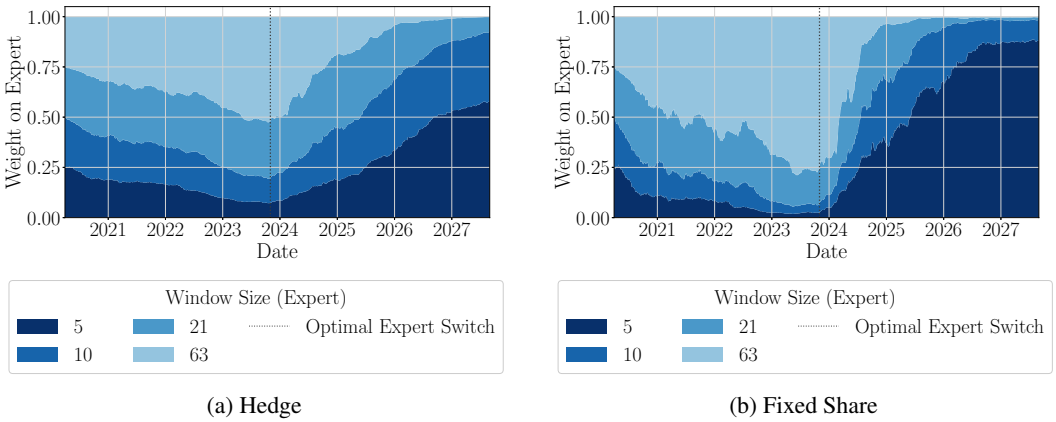

(a) Hedge

(b) Fixed Share

Figure 3: Evolution of the aggregation weights $q_j(t)$ assigned to the four window-size experts $j \in \{5, 10, 21, 63\}$ under (a) Hedge and (b) Fixed Share.

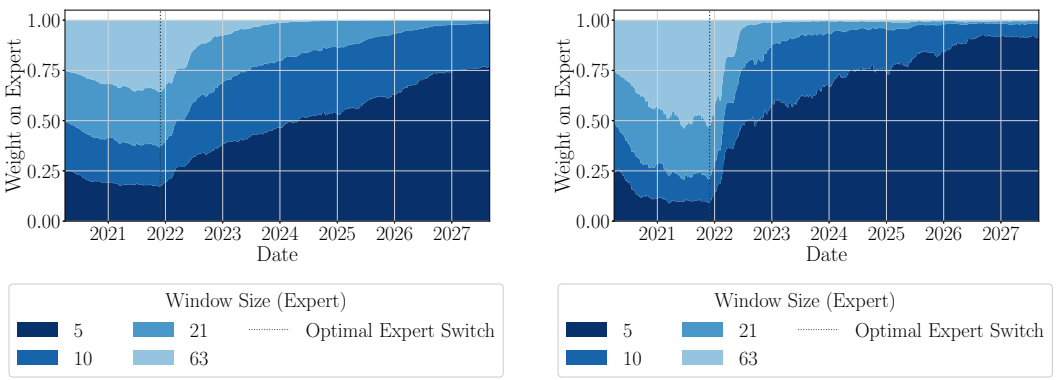

(a) Hedge

(b) Fixed Share

Figure 4: Evolution of the aggregation weights $q_j(t)$ on four Window Sizes for different switching timing.

Comparison with Online Portfolio Selection Algorithms. Additionally, Figure 5 compares our two-level framework against several benchmark algorithms commonly used in online portfolio selection. For this comparison, we define $R_{\text{stock}}(T, K)$ as the tracking regret against the best stock sequence with at most K switches.

Consider the portfolio $\mathcal{M} = \{1, 2, \ldots, m\}$ with $m \geq 2$. We treat each asset $i \in \mathcal{M}$ as an expert. Write $\mathbf{i} := (i_1, \ldots, i_T) \in \mathcal{M}^T$ for an arbitrary comparator sequence of experts, where i_t is the expert selected by the comparator at time t. For the switching budget K, define

$$\mathcal{I}_{T,K} := \Big\{ \mathbf{i} \in \mathcal{M}^T : S_T(\mathbf{i}) \leq K, \Big\}$$

where $S_T(\mathbf{i})$ counts the switches in $\mathbf{i}$. For a given horizon $T \geq 1$ and a switching budget $K \in \{0, 1, \ldots, T-1\}$, the tracking regret R_{stock} against stock sequence is defined as

$$R_{\text{stock}}(T, K) := \sum_{t=1}^{T} \ell(\widehat{\mathbf{w}}(t), \mathbf{y}(t)) - \min_{\mathbf{i} \in \mathcal{I}_{T,K}} \sum_{t=1}^{T} \ell(\mathbf{w}_{i_t}(t), \mathbf{y}(t)).$$

The minimization over $\mathcal{I}_{T,K}$ selects the best such expert sequence in hindsight.

Consistent with existing literature, the "experts' advice" is formulated as a buy-and-hold strategy for individual stocks, and the loss function is the negative logarithmic return. Under this configuration, our two-level framework achieves a lower R_{stock} than the compared algorithms in this synthetic experiment. Negative regret indicates that our framework incurs lower cumulative loss than the best stock sequence with at most K switches.

4.2 Empirical Studies: Dow Jones 30

This section presents an empirical study using historical data to validate our proposed framework.

The Setup. For the portfolio setting, we use the *Dow Jones Industrial Average* and invest directly in its constituent stocks to benchmark performance against the `DJI` index itself. We choose the reference expert set $\mathcal{N} = \{5, 10, 21, 30, 41, 50, 63, 126\}$. Because the `DJI` periodically updates its components, we isolate a period free of any asset additions or removals to simplify our historical data. From 2021/03/04 to 2024/02/25, we construct the portfolio using the 30 exact components of the `DJI` index over that period. The portfolio additionally includes `BIL`, a 1–3 Month Treasury Bill ETF. The asset tickers are listed in Appendix B. We set the transaction cost rate $c = 0.001$ (10 bps) for turnover trades. Since α^* and η^* depend on the switching times, we choose a prescribed switching budget $K(T)$ satisfying (2). For any constant $\varepsilon > 0$, if sufficiently large time horizon T satisfies $\log(T)\big(\log(\log(T))\big)^{1+\varepsilon} > 1$, we define the switching budget $K(T)$ as:

$$K(T) = \left\lfloor \frac{T}{\log(T)\big(\log(\log(T))\big)^{1+\varepsilon}} \right\rfloor.$$

Since it follows that:

$$\lim_{T\to\infty} \frac{\frac{T}{\log(T)(\log(\log(T)))^{1+\varepsilon}}}{\frac{T}{\log(T)\log(\log(T))}} = \lim_{T\to\infty} \frac{1}{\big(\log(\log(T))\big)^{\varepsilon}} = 0.$$

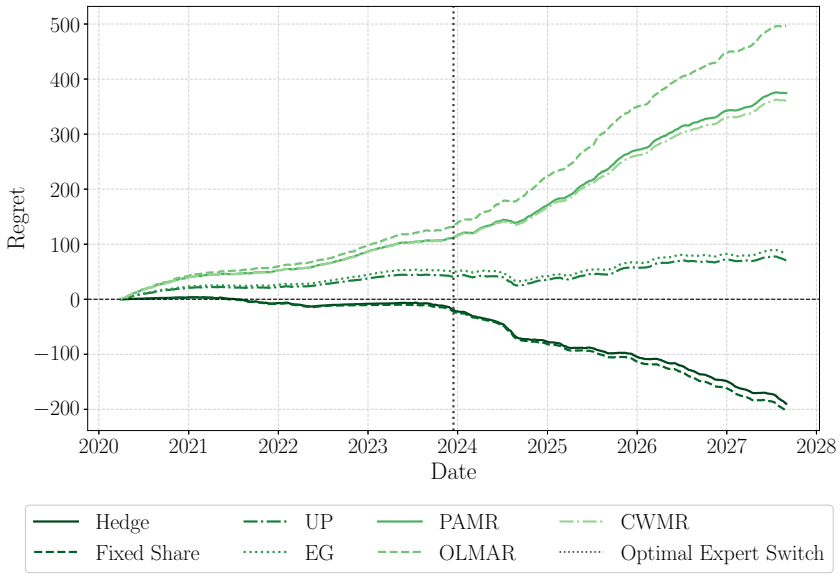

Figure 5: R_{stock} for Hedge, Fixed Share, and Comparators with Negative Log Return Loss

In the following studies, we use this $K(T)$ as our switching budget K estimator.

Performance Evaluation. Figure 6 presents the performance of Cost-Sensitive Regret evaluated on the Dow Jones Industrial Average (DJIA). As depicted in Figure 7, the weight distribution remains largely uniform, with the proportion assigned to $j = 126$ being marginally higher than that of the others. Figure 8 confirms that the R_{stock} of our proposed method remains consistently lower than that of the comparative algorithms. Furthermore, Table 1 and Figure 9 illustrate the trading performance of our method relative to the baselines. Figure 9 demonstrates that our two-level framework progressively accumulates a higher account value, even when subjected to these transaction cost penalties. Table 1 additionally highlights that the two-level framework exhibits the lowest Maximum Drawdown (MDD), a critical risk management metric in algorithmic trading.

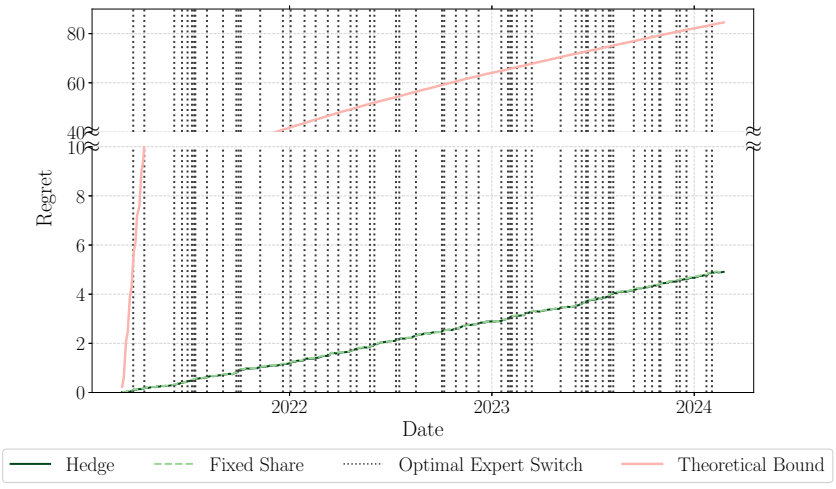

Figure 6: Cost-Sensitive Tracking Regret for Hedge and Fixed Share with Negative PnL Loss (2021/03/04-2024/02/25).

4.3 Empirical Studies: S&P 500

The Setup. Keep the expert setting $\mathcal{N}$ and the transaction cost rate as above. For the portfolio setting, we use the *S&P 500* (Ticker: `GSPC`) and invest directly in its constituent stocks to benchmark performance against

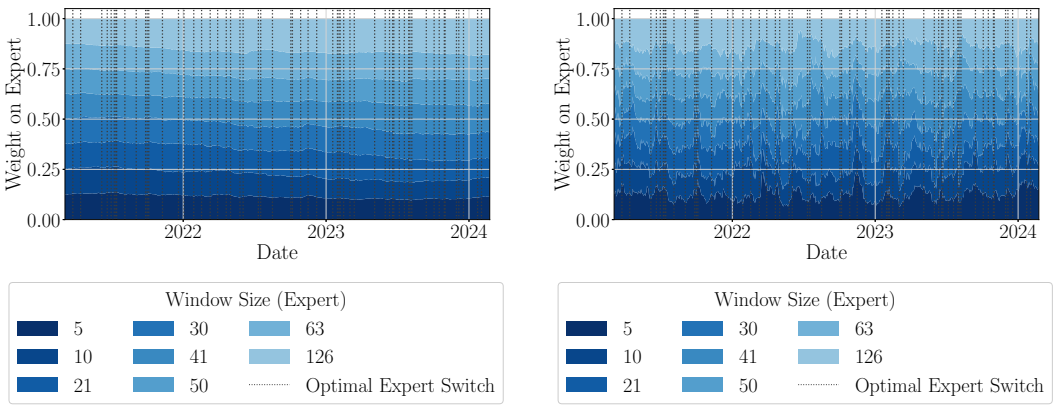

(a) Hedge (b) Fixed Share

Figure 7: Evolution of the aggregation weights $q_j(t)$ assigned to window-size experts (2021/03/04-2024/02/25).

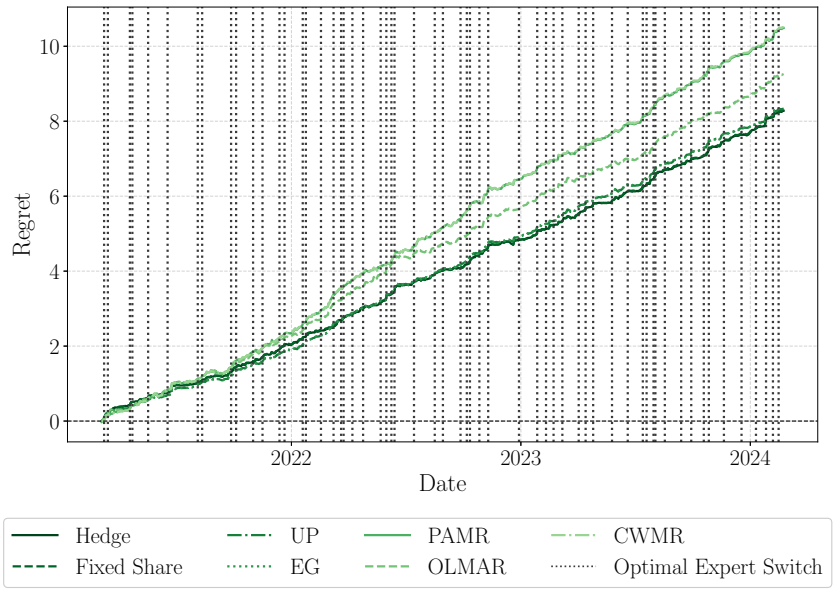

Figure 8: R_{stock} for Hedge, Fixed Share, and Comparators with Negative Log Return Loss (2021/03/04-2024/02/25).

Table 1: Trading Performance Metric for different Window Sizes (2021/03/04-2024/02/25).

	5-day	10-day	21-day	30-day	41-day
Cumulative Return(%)	109.8576	21.0539	7.6831	57.0513	85.2425
Sample Mean of Return(%)	0.1142	0.0394	0.0216	0.0740	0.0943
Sample Standard Deviation	0.0174	0.0166	0.0153	0.0165	0.0155
Annualized Sharpe Ratio	-1.0535	-1.8158	-2.1586	-1.4931	-1.3843
Maximum Drawdown(%)	-22.5798	-28.3883	-34.2497	-24.3892	-22.2521

	50-day	63-day	126-day	Hedge	Fixed Share
Cumulative Return(%)	50.2176	35.3917	157.8256	41.6977	38.1074
Sample Mean of Return(%)	0.0672	0.0526	0.1367	0.0533	0.0500
Sample Standard Deviation	0.0160	0.0155	0.0140	0.0115	0.0117
Annualized Sharpe Ratio	-1.6089	-1.8141	-1.0478	-2.4214	-2.4308
Maximum Drawdown(%)	-26.1014	-26.8669	-22.2768	-12.6326	-12.9282

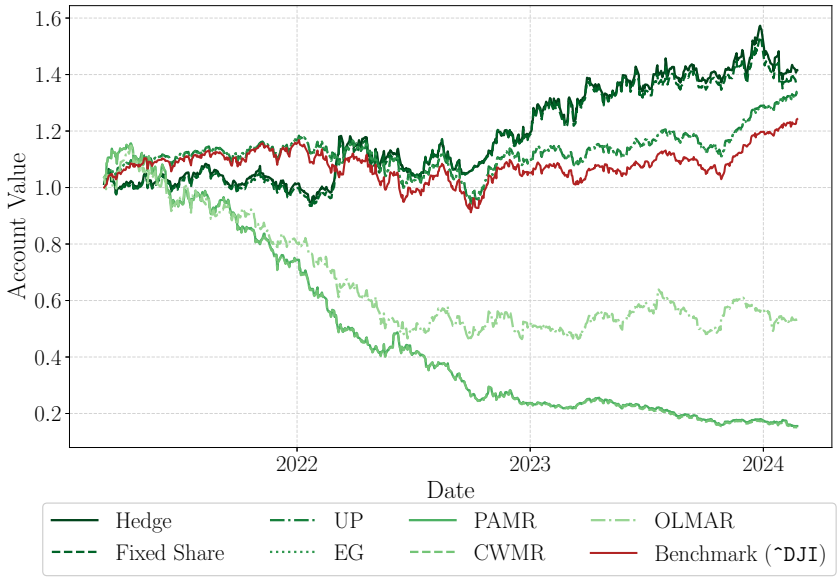

Figure 9: Account Value for Hedge, Fixed Share and benchmark DJI(2021/03/04-2024/02/25).

	UP	EG	PAMR	CWMR	OLMAR
Cumulative Return(%)	48.9870	48.5459	-88.2410	-88.4022	-38.8878
Sample Mean of Return(%)	0.0499	0.0495	-0.2262	-0.2278	-0.0365
Sample Standard Deviation	0.0093	0.0093	0.0190	0.0190	0.0200
Annualized Sharpe Ratio	-3.0639	-3.0720	-3.8025	-3.8207	-2.1119
Maximum Drawdown(%)	-20.3914	-20.3837	-88.9765	-89.1155	-61.0217

the GSPC index itself. Because the GSPC periodically updates its components, from 2020/01/02 to 2026/08/01, we remove all the changed stocks and construct our portfolio using the 481 components of the GSPC during that window (Tickers are listed in Appendix B.) and the risk-free asset BIL, which represents the 1-3 Month T-Bill ETF.

Performance Evaluation. Figure 10 displays the performance of the Cost-Sensitive Regret evaluated on the S&P 500 index. Consistent with previous observations, Figure 11 shows a generally uniform weight allocation, with the proportion for $j = 126$ remaining slightly elevated. Figure 12 indicates that our method continues to yield a marginally lower R_{stock} compared to the baseline algorithms, while Figure 13 demonstrates that our method achieves a higher overall account value, and Table 2 concludes the financial performance.

It is important to note that these empirical studies are intended solely to demonstrate the practical application of our framework over a restricted timeframe, utilizing Hedge and Fixed Share as illustrative examples to account for window sizes and transaction costs. The primary objective is to underscore the necessity of incorporating such considerations, rather than to assert any predictive capability regarding the future performance of specific financial instruments.

Table 2: Trading Performance Metric for different Window Sizes (2020/01/02-2026/08/01).

	5-day	10-day	21-day	30-day	41-day
Cumulative Return(%)	256.0875	640.7924	361.6846	261.9684	1173.9300
Sample Mean of Return(%)	0.1840	0.2313	0.2013	0.1857	0.2605
Sample Standard Deviation	0.0468	0.0476	0.0473	0.0469	0.0464
Annualized Sharpe Ratio	-0.1767	-0.0160	-0.1167	-0.1708	0.0837
Maximum Drawdown(%)	-62.5062	-63.1732	-74.2201	-76.1460	-64.1657

5 Concluding Remarks

We use Hedge and Fixed Share to dynamically aggregate window-size experts for portfolio selection while accounting for turnover costs. We derive a cost-sensitive tracking-regret bound for Fixed Share, with the static-

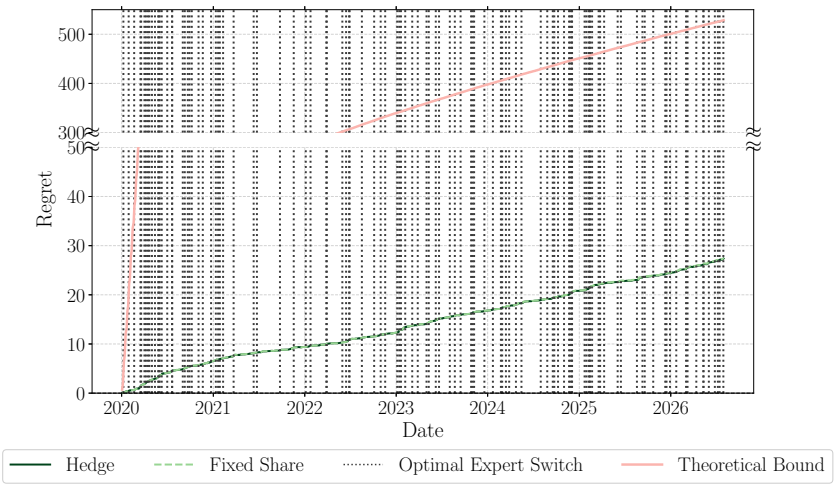

Figure 10: Cost-Sensitive Tracking Regret of Expert for Hedge and Fixed Share with Negative PnL Loss (2020/01/02-2026/08/01).

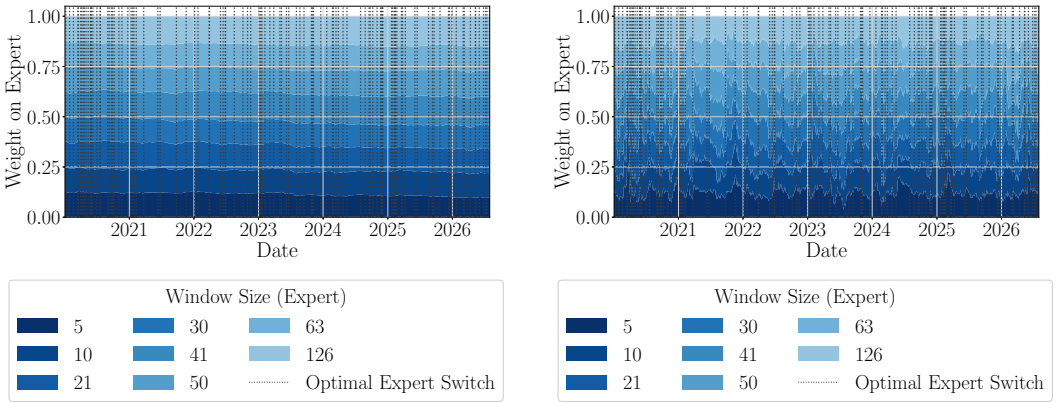

(a) Hedge (b) Fixed Share

Figure 11: Evolution of the aggregation weights $q_j(t)$ assigned to window-size experts (2020/01/02-2026/08/01).

	50-day	63-day	126-day	Hedge	Fixed Share
Cumulative Return(%)	385.5686	316.8115	1881.7486	578.6997	471.0522
Sample Mean of Return(%)	0.1995	0.1851	0.2771	0.1816	0.1737
Sample Standard Deviation	0.0460	0.0445	0.0440	0.0364	0.0372
Annualized Sharpe Ratio	-0.1261	-0.1817	0.1480	-0.2373	-0.2663
Maximum Drawdown(%)	-78.5935	-67.1939	-60.4582	-56.0818	-57.8097

	UP	EG	PAMR	CWMR	OLMAR
Cumulative Return(%)	205.2483	205.7900	-52.0775	205.5993	-44.0283
Sample Mean of Return(%)	0.0711	0.0712	0.0094	0.0827	0.0826
Sample Standard Deviation	0.0129	0.0129	0.0321	0.0199	0.0501
Annualized Sharpe Ratio	-2.0168	-2.0159	-1.1188	-1.2206	-0.4843
Maximum Drawdown(%)	-38.3269	-38.2862	-84.5912	-59.5749	-81.6301

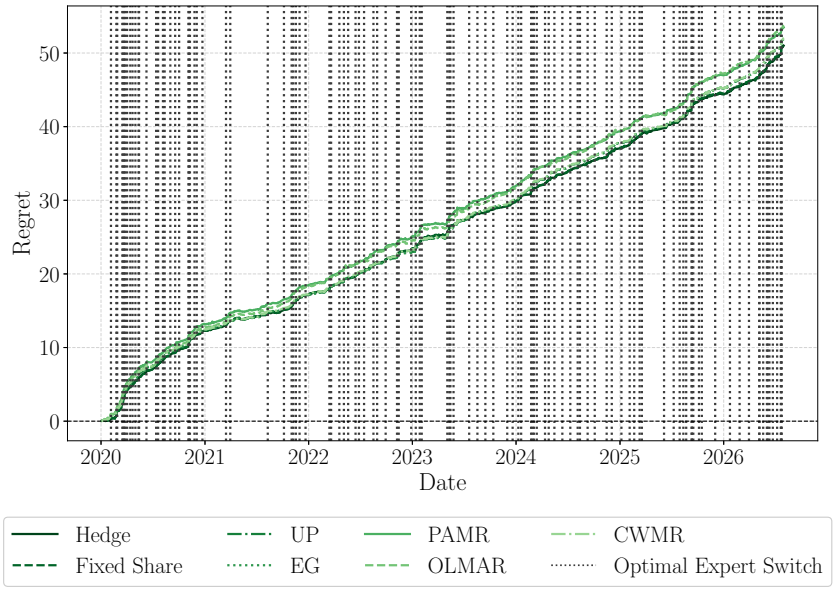

Figure 12: R_{stock} for Hedge, Fixed Share, and Comparators with Negative Log Return Loss (2020/01/02-2026/08/01).

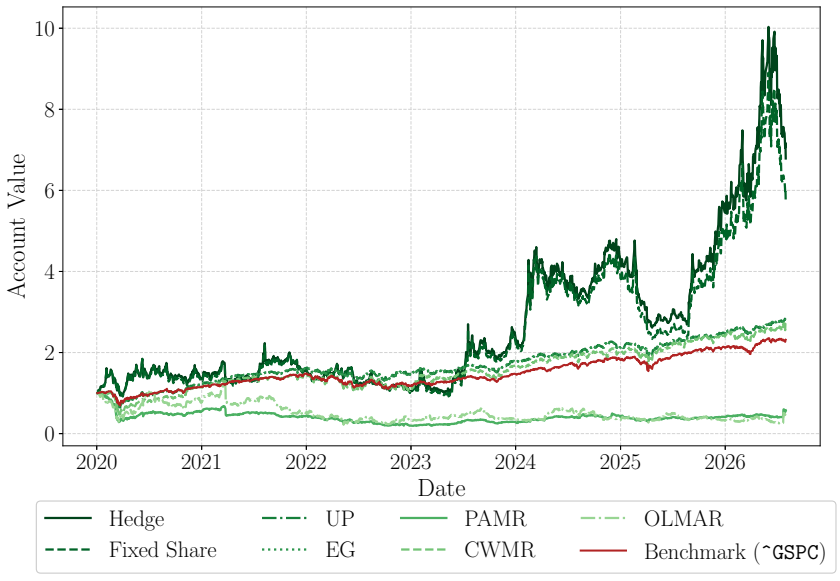

Figure 13: Account Value for Hedge, Fixed Share and benchmark `GSPC`(2020/01/02-2026/08/01).

regret bound for Hedge recovered as a special case. The bound explicitly characterizes the dependence on the transaction-cost rate and guides the choice of the learning rate.

This paper focuses on tracking regret for expert sequences relative to expert sequences under a prescribed switching budget. A complementary extension would be to embed the proposed cost-sensitive learner in a strongly adaptive meta-algorithm Daniely et al. (2015), seeking regret guarantees uniformly over all subintervals. Such an extension requires a separate treatment of interval-specific parameterization and transaction costs and is left for future research.

AI USE STATEMENT

Generative AI tools (ChatGPT 5.6 sol and Gemini 3.8 Flash) were used to assist with language editing and presentation, and to provide feedback on mathematical derivations and proofs. The authors reviewed and independently verified all AI-assisted content. The authors take full responsibility for the final content of this work.

REFERENCES

Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Thomas M Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.

Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1405–1411, Lille, France, 07–09 Jul 2015. PMLR.

Steven Diamond and Stephen Boyd. CVXPY: A Python-Embedded Modeling Language for Convex Optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

Alexei A Gaivoronski and Fabio Stella. Stochastic Nonstationary Optimization for Finding Universal Portfolios. *Annals of Operations Research*, 100(1):165–188, 2000.

Alexei A Gaivoronski and Fabio Stella. On-line Portfolio Selection Using Stochastic Programming. *Journal of Economic Dynamics and Control*, 27(6):1013–1043, 2003.

James Hannan. Approximation to Bayes Risk in Repeated Play. *Contributions to the Theory of Games*, 3(2): 97–139, 1957.

David P Helmbold, Robert E Schapire, Yoram Singer, and Manfred K Warmuth. On-Line Portfolio Selection using Multiplicative Updates. *Mathematical Finance*, 8(4):325–347, 1998.

Mark Herbster and Manfred K Warmuth. Tracking the Best Expert. *Machine Learning*, 32(2):151–178, 1998.

Atsushi Inoue, Lu Jin, and Barbara Rossi. Rolling window selection for out-of-sample forecasting with time-varying parameters. *Journal of econometrics*, 196(1):55–67, 2017.

Yoontae Jeon and Thomas H McCurdy. Time-Varying Window Length for Correlation Forecasts. *Econometrics*, 5(4):54, 2017.

John L Kelly. A New Interpretation of Information Rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.

Bin Li and Steven C. H. Hoi. On-Line Portfolio Selection with Moving Average Reversion. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012. URL https://icml.cc/2012/papers/168.pdf.

Bin Li and Steven C. H. Hoi. Online Portfolio Selection: A survey. *ACM Computing Surveys*, 46(3), January 2014. ISSN 0360-0300. doi: 10.1145/2512962.

Bin Li, Peilin Zhao, Steven CH Hoi, and Vivekanand Gopalkrishnan. PAMR: Passive Aggressive Mean Reversion Strategy for Portfolio Selection. *Machine Learning*, 87(2):221–258, 2012.

Bin Li, Steven CH Hoi, Peilin Zhao, and Vivekanand Gopalkrishnan. Confidence weighted mean reversion strategy for online portfolio selection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 7 (1):1–38, 2013.

N. Littlestone and M.K. Warmuth. The Weighted Majority Algorithm. *Information and Computation*, 108(2): 212–261, 1994. ISSN 0890-5401. doi: https://doi.org/10.1006/inco.1994.1009.

David G Luenberger. *Investment Science*. Oxford University Press, 2013.

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

Tanya Molodtsova and David H Papell. Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals. *Journal of International Economics*, 77(2):167–180, 2009.

Francesco Orabona. Online Learning: A Modern Introduction Using Convex Optimization. *arXiv preprint arXiv:1912.13213*, 2026.

M. Hashem Pesaran and Andreas Pick. Forecast Combination Across Estimation Windows. *Journal of Business & Economic Statistics*, 29(2):307–318, 2011. doi: 10.1198/jbes.2010.09018.

M Hashem Pesaran and Allan Timmermann. Selection of estimation window in the presence of breaks. *Journal of Econometrics*, 137(1):134–161, 2007.

M Hashem Pesaran, Andreas Pick, and Mikhail Pranovich. Optimal forecasts in the presence of structural breaks. *Journal of Econometrics*, 177(2):134–152, 2013.

Edgar E Peters. *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*. John Wiley & Sons, 1994.

Shahab Rajabi, Pardis Roozkhosh, and Nasser Motahari Farimani. MLP-Based Learnable Window Size for Bitcoin Price Prediction. *Applied Soft Computing*, 129:109584, 2022.

Barbara Rossi and Atsushi Inoue. Out-of-Sample Forecast Tests Robust to the Choice of Window Size. *Journal of Business & Economic Statistics*, 30(3):432–453, 2012.

James H Stock and Mark W Watson. Why has us inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39:3–33, 2007.

Volodya Vovk and Chris Watkins. Universal Portfolio Selection. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, pp. 12–23, 1998.

A Technical Proofs

Proof of Lemma 3.3. By adding and subtracting the intermediate weight vector $\mathbf{q}^{\exp}(t-1)$ and applying the triangle inequality, the total variation can be decomposed into the variation from the loss update and the variation from the sharing step:

$$\begin{aligned}\|\mathbf{q}(t)-\mathbf{q}(t-1)\|_1 &= \|\mathbf{q}(t)-\mathbf{q}^{\exp}(t-1)+\mathbf{q}^{\exp}(t-1)-\mathbf{q}(t-1)\|_1 \\ &\le \underbrace{\|\mathbf{q}^{\exp}(t-1)-\mathbf{q}(t-1)\|_1}_{\text{Loss Update Variation}} + \underbrace{\|\mathbf{q}(t)-\mathbf{q}^{\exp}(t-1)\|_1}_{\text{Sharing Variation}}.\end{aligned} \tag{11}$$

To bound the Loss Update Variation, define an auxiliary continuous function $\tilde{q}_j(s)$ parameterized by $s \in [0,\eta]$ that interpolates the weights during the exponential update:

$$\tilde{q}_j(s) := \frac{q_j(t-1)e^{-s\tilde{\ell}_j(t-1)}}{Z(s)}, \tag{12}$$

where $Z(s) = \sum_{k\in\mathcal{N}} q_k(t-1)e^{-s\tilde{\ell}_k(t-1)}$. Note that $\tilde{q}_j(0) = q_j(t-1)$, $\tilde{q}_j(\eta) = q_j^{\exp}(t-1)$ and $\sum_{j\in\mathcal{N}}\tilde{q}_j(s) = 1$.

For each expert j, by the Fundamental Theorem of Calculus, the absolute change is bounded by the integral of the absolute derivative:

$$|q_j^{\exp}(t-1)-q_j(t-1)| = \left|\int_0^\eta \frac{d}{ds}\tilde{q}_j(s)\,ds\right| \le \int_0^\eta \left|\frac{d}{ds}\tilde{q}_j(s)\right| ds. \tag{13}$$

Using the logarithmic derivative identity, $\frac{d}{ds}\log\tilde{q}_j(s) = \frac{1}{\tilde{q}_j(s)}(\frac{d}{ds}\tilde{q}_j(s))$, then:

$$\begin{aligned}\frac{1}{\tilde{q}_j(s)}\left(\frac{d}{ds}\tilde{q}_j(s)\right) &= \frac{d}{ds}\log\tilde{q}_j(s) \\ &= -\tilde{\ell}_j(t-1) - \frac{d}{ds}\log Z(s).\end{aligned} \tag{14}$$

where the last equality follows from differentiating (12).

Evaluating the derivative of the log-partition function $\log Z(s)$ yields the negative expected loss under the interpolated weights $\{\tilde{q}_j(s)\}_{j\in\mathcal{N}}$; i.e.,

$$\begin{aligned}\frac{d}{ds}\log Z(s) &= \frac{1}{Z(s)}\sum_{k\in\mathcal{N}} q_k(t-1)(-\tilde{\ell}_k(t-1))e^{-s\tilde{\ell}_k(t-1)} \\ &= -\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1). \end{aligned} \tag{15}$$

Substituting (15) back to (14) and isolating the derivative gives:

$$\frac{d}{ds}\tilde{q}_j(s) = \tilde{q}_j(s)\left(\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1) - \tilde{\ell}_j(t-1)\right). \tag{16}$$

Applying (16) to (13) yields

$$\begin{aligned}|q_j^{\exp}(t-1) - q_j(t-1)| &\le \int_0^\eta \left|\tilde{q}_j(s)\left(\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1) - \tilde{\ell}_j(t-1)\right)\right| ds \\ &\le \int_0^\eta \tilde{q}_j(s)\left|\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1) - \tilde{\ell}_j(t-1)\right| ds. \end{aligned} \tag{17}$$

Summing up over all n experts yields:

$$\begin{aligned}\|\mathbf{q}^{\exp}(t-1) - \mathbf{q}(t-1)\|_1 &\le \sum_{j\in\mathcal{N}}\int_0^\eta \tilde{q}_j(s)\left(\left|\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1) - \tilde{\ell}_j(t-1)\right|\right) ds \\ &= \int_0^\eta \sum_{j\in\mathcal{N}} \tilde{q}_j(s)\left(\left|\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1) - \tilde{\ell}_j(t-1)\right|\right) ds \\ &\le \int_0^\eta \sum_{j\in\mathcal{N}} \tilde{q}_j(s)\left(\max_{k\in\mathcal{N}}\tilde{\ell}_k(t-1) - \min_{k\in\mathcal{N}}\tilde{\ell}_k(t-1)\right) ds \\ &= \eta\left(\max_{k\in\mathcal{N}}\tilde{\ell}_k(t-1) - \min_{k\in\mathcal{N}}\tilde{\ell}_k(t-1)\right)\end{aligned}$$

where the last inequality follows because both $\sum_{k\in\mathcal{N}} \tilde{q}_k(s)\tilde{\ell}_k(t-1)$ and $\tilde{\ell}_j(t-1)$ lie in the interval $[\min_{k\in\mathcal{N}}\tilde{\ell}_k(t-1), \max_{k\in\mathcal{N}}\tilde{\ell}_k(t-1)]$, and the last equality uses $\sum_{j\in\mathcal{N}}\tilde{q}_j(s) = 1$.

Next, we bound the Sharing Variation term in (11) by substituting the update step $q_j(t) = (1-\alpha)q_j^{\exp}(t-1) + \frac{\alpha}{n}$:

$$\begin{aligned}\|\mathbf{q}(t) - \mathbf{q}^{\exp}(t-1)\|_1 &= \left\|(1-\alpha)\mathbf{q}^{\exp}(t-1) + \frac{\alpha}{n}\mathbf{1} - \mathbf{q}^{\exp}(t-1)\right\|_1 \\ &= \left\|\frac{\alpha}{n}\mathbf{1} - \alpha\mathbf{q}^{\exp}(t-1)\right\|_1 \\ &= \alpha\left\|\frac{1}{n}\mathbf{1} - \mathbf{q}^{\exp}(t-1)\right\|_1 \\ &\le 2\alpha\left(1 - \frac{1}{n}\right),\end{aligned}$$

where the last inequality follows from the convexity of the ℓ_1-norm and the fact that $\|\frac{1}{n}\mathbf{1} - \mathbf{e}_j\|_1 = 2(1-\frac{1}{n})$ for every $j \in \mathcal{N}$. Here $\mathbf{e}_j \in \Delta_n$ assigns unit mass to expert j and zero mass to every other expert. The bound is sharp over the probability simplex, with equality at its vertices.

Combining the bounds for both components yields the final result:

$$\|\mathbf{q}(t) - \mathbf{q}(t-1)\|_1 \le \eta\left(\max_{j\in\mathcal{N}}\tilde{\ell}_j(t-1) - \min_{j\in\mathcal{N}}\tilde{\ell}_j(t-1)\right) + 2\alpha\left(1 - \frac{1}{n}\right),$$

which completes the proof. □

Proof of Proposition 3.5. By Corollary 3.4, the regret boundcan be expressed as a joint function of η and α:

$$R(\eta, \alpha) = \frac{\mathcal{C}(K, \alpha, n, T)}{\eta} + \eta A_T + 2\alpha \left(1 - \frac{1}{n}\right) C_T.$$

For $C_T = 0$. Since there are no transaction costs, the setting reduces to the standard Fixed Share algorithm; see Herbster & Warmuth (1998), Cesa-Bianchi & Lugosi (2006).

For $K = 0$, if $\alpha > 0$, the derivative with respect to α is

$$\frac{\partial R}{\partial \alpha} = \frac{1}{\eta}\left(\frac{(T-1)\alpha - K}{\alpha(1-\alpha)}\right) + 2\left(1 - \frac{1}{n}\right) C_T,$$

which is strictly positive for all $\alpha \in (0, 1)$ with $T > 1$ and $\eta > 0$. The regret function R is strictly increasing in α. Therefore, it cannot have a minimum where the derivative is zero; the minimum must lie on the boundary constraint, i.e., $\alpha = 0$. With $K = 0$ and $\alpha = 0$, the complexity term simplifies to $\mathcal{C}(0, 0, n, T) = \log n$. The regret bound reduces to a function of η alone:

$$R(\eta, 0) = \frac{\log n}{\eta} + \eta A_T.$$

Taking the derivative with respect to η and setting it to zero yields $\eta^* = \sqrt{\frac{\log n}{A_T}}$. Substituting η^* back into $R(\eta, 0)$ gives the minimized bound $R^c_{\text{track}}(T, 0) \leq 2\sqrt{A_T \log n}$.

For the rest of the cases, $K > 0$, $\alpha \in (0, 1)$ and $\eta > 0$, we optimize $R(\eta, \alpha)$ sequentially. For any fixed $\alpha \in (0, 1)$, the first and second partial derivatives with respect to η are:

$$\frac{\partial R}{\partial \eta} = -\frac{\mathcal{C}(K, \alpha, n, T)}{\eta^2} + A_T, \quad \frac{\partial^2 R}{\partial \eta^2} = \frac{2\mathcal{C}(K, \alpha, n, T)}{\eta^3}.$$

Since $\mathcal{C}(K, \alpha, n, T) > 0$ and $\eta > 0$, the second partial derivative is strictly positive. Thus, $R(\eta, \alpha)$ is strictly convex in η, and setting $\frac{\partial R}{\partial \eta} = 0$ yields the unique conditional minimum:

$$\eta^*(\alpha) = \sqrt{\frac{\mathcal{C}(K, \alpha, n, T)}{A_T}}.$$

Substituting $\eta^*(\alpha)$ back into the objective function yields the profile function $g(\alpha)$ depending solely on α:

$$\begin{aligned} g(\alpha) &:= R(\eta^*(\alpha), \alpha) \\ &= \frac{\mathcal{C}(K, \alpha, n, T)}{\sqrt{\mathcal{C}(K, \alpha, n, T)/A_T}} + A_T\sqrt{\frac{\mathcal{C}(K, \alpha, n, T)}{A_T}} + 2\alpha\left(1 - \frac{1}{n}\right) C_T \\ &= 2\sqrt{A_T \cdot \mathcal{C}(K, \alpha, n, T)} + 2\alpha\left(1 - \frac{1}{n}\right) C_T. \end{aligned}$$

To find the minimum of $g(\alpha)$, we compute its first derivative:

$$g'(\alpha) = \sqrt{\frac{A_T}{\mathcal{C}(K, \alpha, n, T)}}\left(-\frac{K}{\alpha} + \frac{T - K - 1}{1 - \alpha}\right) + 2\left(1 - \frac{1}{n}\right) C_T.$$

Setting $g'(\alpha) = 0$ yields the stationarity condition:

$$\frac{K - (T-1)\alpha}{\alpha(1-\alpha)} = 2C_T\left(1 - \frac{1}{n}\right)\sqrt{\frac{\mathcal{C}(K, \alpha, n, T)}{A_T}}.$$

To establish that the root α^* is the unique global minimum, we prove that $g(\alpha)$ is strictly convex on the domain $\alpha \in \left(0, \frac{K}{T-1}\right)$. Since $g(\alpha)$ is the sum of a linear term and $2\sqrt{A_T \cdot \mathcal{C}(K, \alpha, n, T)}$, it suffices to show that $\sqrt{\mathcal{C}}$ is strictly convex. The second derivative of $\sqrt{\mathcal{C}}$ is strictly positive if and only if $2\mathcal{C}\mathcal{C}'' > (\mathcal{C}')^2$.

The domain constraint $\alpha < \frac{K}{T-1}$. By subtracting αK at the both sides of $\alpha(T-1) < K$, we algebraically rearranges to $K(1-\alpha) > (T - K - 1)\alpha$, which implies

$$\frac{K}{\alpha} > \frac{T - 1 - K}{1 - \alpha} > 0. \tag{18}$$

The derivatives of $\mathcal{C}$ with respect to α can be expressed as $\mathcal{C}' = -\frac{K}{\alpha} + \frac{T-1-K}{1-\alpha}$. Because (18), we have

$$(\mathcal{C}')^2 < (\frac{K}{\alpha})^2 = \frac{K^2}{\alpha^2}.$$

Furthermore, since $\frac{T-1-K}{(1-\alpha)^2} > 0$, we know

$$\mathcal{C}'' = \frac{K}{\alpha^2} + \frac{T-1-K}{(1-\alpha)^2} > \frac{K}{\alpha^2}. \tag{19}$$

By definition, since other term are always positive in $\alpha \in (0, \frac{K}{T-1})$, we have

$$\mathcal{C}(\alpha) > (K+1)\log n. \tag{20}$$

Combining (19) and (20) yields:

$$2\mathcal{C}\mathcal{C}'' > \frac{2K(K+1)\log n}{\alpha^2}.$$

For any $n \geq 2$, we have $\log n \geq \log 2 > 0.5$, which guarantees the coefficient $2(K+1)\log n > K+1 > K$, yields

$$2\mathcal{C}\mathcal{C}'' > \frac{K^2}{\alpha^2} > (\mathcal{C}')^2.$$

Thus, $\sqrt{\mathcal{C}}$ and g are strictly convex on $(0, \alpha_0)$, where $\alpha_0 := K/(T-1)$. Moreover,

$$\lim_{\alpha \downarrow 0} g'(\alpha) = -\infty, \qquad g'(\alpha_0) = 2\left(1 - \frac{1}{n}\right) C_T > 0.$$

By continuity and strict monotonicity of g' on $(0, \alpha_0)$, there exists a unique root $\alpha^* \in (0, \alpha_0)$. For $\alpha \in [\alpha_0, 1)$, we have $\mathcal{C}'(\alpha) \geq 0$, and hence $g'(\alpha) > 0$. Therefore, α^* is the unique global minimizer of g on $(0, 1)$. Evaluating $g(\alpha^*)$ completes the proof. □

B Assets Used in the Empirical Studies

Dow Jones 30. The tickers of the Dow Jones 30 stocks considered in Section 4.2 are listed below: `WBA, CRM, HON, AMGN, DOW, AAPL, GS, V, NKE, UNH, TRV, CSCO, CVX, VZ, MSFT, HD, INTC, JNJ, WMT, CAT, JPM, DIS, BA, KO, MCD, AXP, IBM, MRK, MMM, PG` and the risk-free asset `BIL`, which represents the 1-3 Month T-Bill ETF.

S&P 500. The tickers of the S&P 500 stocks considered in Section 4.3 are listed below:

```
A, AAPL, ABBV, ABT, ACGL, ACN, ADBE, ADI, ADM, ADP, ADSK, AEE, AEP, AES,
AFL, AIG, AIZ, AJG, AKAM, ALB, ALGN, ALL, ALLE, AMAT, AMCR, AMD, AME,
AMGN, AMP, AMT, AMZN, ANET, AON, AOS, APA, APD, APH, APO, APTV, ARE,
ARES, ATO, AVGO, AVY, AWK, AXON, AXP, AZO, BA, BAC, BALL, BAX, BBY, BDX,
BEN, BF-B, BG, BIIB, BKNG, BKR, BLDR, BLK, BMY, BNY, BR, BRK-B, BRO, BSX,
BWA, BX, BXP, C, CAG, CAH, CARR, CAT, CB, CBOE, CBRE, CCI, CCL, CDNS,
CDW, CE, CEG, CF, CFG, CHD, CHRW, CHTR, CI, CINF, CL, CLX, CMA, CMCSA,
CME, CMG, CMI, CMS, CNC, CNP, COF, COO, COP, COR, COST, CPB, CPRT, CPT,
CRWD, CSCO, CSGP, CSX, CTAS, CTC, CTRA, CTSH, CTVA, CVS, CVX, CZR, D,
DAL, DD, DE, DFS, DG, DGX, DHI, DHR, DIS, DOC, DOV, DOW, DPZ, DRI, DTE,
DUK, DVA, DVN, DXCM, EA, EBAY, ECL, ED, EFX, EG, EIX, EL, ELV, EMN, EMR,
ENPH, EOG, EPAM, EQT, ERIE, ES, ESS, ETN, ETR, EVRG, EW, EXC, EXPD, EXPE,
EXR, F, FANG, FAST, FCX, FDS, FDX, FE, FFIV, FI, FICO, FIS, FITB, FLT,
FMC, FOX, FOXA, FRT, FSLR, FTNT, GE, GEHC, GEV, GILD, GIS, GL, GLW, GM,
GNRC, GOOG, GOOGL, GPC, GPN, GRMN, GS, GWW, HAL, HAS, HBAN, HCA, HD, HES,
HIG, HII, HLT, HOLX, HON, HPE, HPQ, HRL, HSIC, HST, HSY, HUBB, HUM, HWM,
IBM, ICE, IDXX, IEX, IFF, ILMN, INCY, INTC, INTU, INVU, IP, IPG, IQV,
IR, IRM, ISRG, IT, ITW, IVZ, J, JBHT, JCI, JKHY, JNJ, JNPR, JPM, K, KDP,
KEY, KEYS, KHC, KIM, KLAC, KMB, KMI, KMX, KO, KR, KKV, L, LDOS, LEN, LH,
LHX, LIN, LKQ, LLY, LMT, LNT, LOW, LRCX, LUV, LVS, LYB, LYV, MA, MAA,
MAR, MAS, MCD, MCHP, MCK, MCO, MDLZ, MDT, MET, META, MGM, MHK, MKC, MKTX,
MLM, MMC, MMM, MNST, MO, MOH, MOS, MPC, MPWR, MRK, MRO, MS, MSCI, MSFT,
MSI, MTB, MTCH, MTD, MU, NCLH, NDAQ, NDSN, NEE, NEM, NFLX, NI, NKE, NOC,
NOW, NRG, NSC, NTAP, NTRS, NUE, NVDA, NVR, NWS, NWSA, NXPI, O, ODFL, OKE,
OMC, ON, OKE, ORCL, ORLY, OTIS, OXY, PANW, PARA, PAYC, PAYX, PCAR, PCG,
PEAK, PEG, PEP, PFE, PFG, PG, PGR, PH, PHM, PKG, PLD, PM, PNC, PNR, PNW,
PODD, POOL, PPG, PPL, PRU, PSA, PTC, PTEN, PWR, PXD, PYPL, QCOM, QRVO,
RCL, REG, REGN, RF, RHI, RJF, RL, RMD, ROK, ROL, ROP, ROST, RSG, RTX,
RVTY, SBUX, SCHW, SHW, SJM, SLB, SMCI, SNA, SNPS, SO, SPG, SPGI, SRE,
STE, STLD, STT, STX, STZ, SWK, SWKS, SYF, SYK, SYY, T, TAP, TARE, TDG,
```

TDY, TECH, TEL, TER, TFC, TFX, TGT, TJX, TMO, TMUS, TPR, TRMB, TROW, TRV, TSCO, TSLA, TSN, TT, TTWO, TXN, TXT, TYL, UAL, UBER, UDR, UHS, ULTA, UNH, UNP, UPS, URI, USB, V, VLO, VLTO, VMC, VRSK, VRSN, VRTX, VTR, VZ, WAB, WAT, WBA, WBD, WDC, WEC, WELL, WFC, WHR, WM, WMB, WMT, WRB, WST, WTW, WY, WYNN, XEL, XOM, XYL, XYZ, YUM, ZBH, ZBRA, ZTS) and the risk-free asset BIL, which represents the 1-3 Month T-Bill ETF.