

Online Quantitative Trading Strategies

An Empirical Performance Overview Through Backtesting

Glucksman Fellowship



NYU

**LEONARD N. STERN
SCHOOL OF BUSINESS**

Niko Lahanis¹, Alice Liu¹, and Zhengyuan Zhou²

¹Researchers and Contributors, niko_lahanis@stern.nyu.edu, alice.liu@stern.nyu.edu

²Professors and Mentors, zzhou@stern.nyu.edu

Contents

1. Introduction	4
2. Input Data	4
3. Backtesting Architecture	4
3.1. Data Collection and Portfolio Initialization	4
3.2. Performance Metrics and Evaluation	5
3.3. Scalability and Practicality	5
4. Algorithms Tested	6
4.1. Benchmarks	6
4.2. Follow-the-Winner	7
4.3. Follow-the-Loser	8
4.4. Pattern Matching	8
4.5. Meta-Learning	10
5. Experiments and Results	11
5.1. Training Methodology	11
5.1.1 Hyperparameter Tuning	11
5.1.2 Walk-Forward Validation	11
5.2. Output	12
5.2.1 Follow-the-Winner	12
5.2.2 Follow-the-Loser	13
5.2.3 Pattern Matching	14
5.2.4 Meta-Learning	15
5.3. Performance Summary	17
5.4. Comparison with Existing Literature	17
6. Experience	17
6.1. Challenges	17
6.2. Approach Adjustments	17
6.3. Key Strengths	17
6.4. Research Limitations	18
7. Conclusion and Future Work	18
8. Appendix	20

Abstract

*This paper systematically evaluates the empirical performance of online quantitative trading strategies through a standardized Python-based backtesting framework. We aim to answer a central question: **which online portfolio selection methods consistently deliver optimal risk-adjusted returns under realistic market conditions?** We investigate four primary families of algorithms—including momentum-based Follow-the-Winner, mean-reversion-based Follow-the-Loser, pattern-matching techniques, and ensemble meta-learning strategies—benchmarking them against traditional strategies such as Buy-and-Hold, Best Stock, and Constantly Rebalanced Portfolios (CRP). Rigorous hyperparameter tuning via grid search optimizes each method to maximize key performance metrics, specifically cumulative wealth, exponential growth rate, maximum drawdown, and the Sharpe ratio. Our findings highlight that adaptive strategies with regularization (e.g., Follow-the-Regularized-Leader) and confidence-based mean-reversion methods (e.g., CWMR, PAMR) provide superior risk-adjusted returns. However, these and other more advanced strategies often incur substantial computational costs, which may limit their practicality for intraday trading. We conclude by discussing practical implications and outlining directions for future work, including enhanced computational efficiency, cloud deployment, and broader asset class applications.*

1. Introduction

Online portfolio selection (OPS) represents a foundational paradigm in quantitative finance, enabling dynamic asset allocation through online learning mechanisms that adapt to streaming market data. By continuously updating portfolio allocations based on new market data, this adaptive method captures trends and exploits market inefficiencies in real time. Inspired by the original work [20] titled “Online Portfolio Selection – A Survey,” this paper implements and compares a wide array of online portfolio selection algorithms discussed in that survey. Our objective is to evaluate not only their risk-adjusted performance—measured through metrics such as cumulative wealth, exponential growth, maximum drawdown and the Sharpe ratio—but also their computational efficiency in realistic trading scenarios.

To achieve this, we developed a unified Python-based backtesting infrastructure that standardizes data preprocessing, portfolio management, and performance evaluation across multiple algorithmic families. These include momentum-based methods (Follow-the-Winner), mean-reversion strategies (Follow-the-Loser), pattern-matching techniques, and meta-learning ensemble approaches. By systematically tuning hyperparameters and benchmarking these strategies against traditional methods like Buy-and-Hold, Best Stock, and Constant Rebalancing, our study bridges theoretical insights with practical applications in quantitative trading.

More concisely, this paper addresses a critical research gap: *determining which algorithms reliably offer superior risk-adjusted returns in realistic trading environments*. For access to the complete GitHub repository and installation instructions, please visit the following link: [Public Repo](#) [22].

2. Input Data

For this research, the input data is derived from QuantQuote Minute Market Data [25], a comprehensive intra-day stock database. This dataset includes minute-by-minute trading data for securities listed on NASDAQ, NYSE, and AMEX from 1998 to the present, along with select ETFs traded on NYSE ARCA. The data spans regular trading hours from 9:30 AM to 4:00 PM, with optional coverage for pre-market and after-market trading sessions. Key features of the dataset include:

- **Survivorship Bias-Free Indexes:** The dataset maintains survivorship bias-free lists of major indexes.
- **Data Adjustments:** Prices and volumes are automatically adjusted for stock splits and dividends, ensuring continuity in the data.
- **Earnings Data:** Earnings announcements are recorded with their timing, including indicators for pre-market and post-market releases, allowing differentiation of trading behavior during these periods.
- **Dividend Adjustments:** Dividends are factored into historical prices to reflect price continuity, with corresponding values provided in the dataset.
- **Symbol Tracking:** Changes in stock symbols due to mergers, acquisitions, or other corporate actions are tracked automatically. Symbol reuses are managed to avoid ambiguity.

The data for this study was sourced from minute-level historical stock data organized into subdirectories containing individual CSV files for each stock. Each file, named in the format `table_TICKER.csv`, includes key columns such as Date, Time, Open, High, Low, Close, Volume, and adjustments for splits, dividends, and earnings. In our experiments, the breadth of securities was limited to the NASDAQ 100, and the time frame was set from 1998 to 2010.

3. Backtesting Architecture

To support the exploration of online portfolio selection techniques, we developed a robust Python-based trading infrastructure. This framework integrates data preprocessing, portfolio management, and performance evaluation, enabling efficient implementation and testing of various online portfolio selection strategies. In the following sections, we detail its core components and functionalities.

3.1. Data Collection and Portfolio Initialization

The infrastructure begins with retrieving stock price data stored in structured CSV files. Using the `get_all_tickers()` function, the system scans subdirectories for files named `table-<TICKER>.csv`, extracting unique ticker symbols. Each file contains historical data for individual stocks, which are standardized with uniform column names. Data is further processed to compute price-relative vectors (x_t), where $x_{t,i}$ represents the ratio of the closing price of stock i at time t to its closing price at $t - 1$. This step ensures compatibility across datasets and prepares the data for subsequent portfolio calculations.

The portfolio is initialized as a uniform distribution across the available assets using the `initialize_portfolio()` function, where the allocation to each stock is $b_{t,i} = \frac{1}{m}$. Price-relative vectors for multiple assets are aligned side-by-side in a consolidated DataFrame using the `calculate_price_relative_vectors()` function, allowing for the computation of multi-asset portfolio returns.

For each ticker, the closing prices are extracted, and price-relative vectors are calculated by dividing each day’s closing price by the previous day’s. This normalization removes the influence of the absolute price level of a stock, enabling a direct comparison of

	AAPL	ADBE	ADI	ADP	ADSK	AEP	ALGN	AMAT	AMD
Date									
2020-03-20	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000
2020-03-23	1.001272	1.064626	1.013781	0.990994	0.990994	0.998042	0.955933	1.087877	1.070791
2020-03-24	1.077576	0.989776	1.111111	1.131817	1.038761	1.010376	1.200000	1.093676	1.087204
2020-03-25	1.006179	0.987088	0.996277	1.002895	0.993306	1.017210	1.143315	1.010011	0.993243
2020-03-26	1.045977	1.043492	0.989002	1.115085	1.076266	1.114477	1.010635	1.089866	1.043450

Figure 1. Price Relative Vectors

performance across different securities and ensuring compatibility for portfolio optimization and algorithmic trading strategies. The data is cleaned by handling missing values and adjusting non-positive values to small positive numbers to avoid computational issues. The final dataset is structured as a single Python DataFrame, with rows representing daily price relatives and columns for each ticker, as shown in Figure 1.

3.2. Performance Metrics and Evaluation

The infrastructure includes utilities for calculating key performance metrics such as:

- **Cumulative Wealth [6]:** The `calculate_cumulative_wealth()` function computes the wealth trajectory over a sequence of trading periods based on the portfolio allocation and price-relative vectors. The initial cumulative wealth value is set to 1.0.
- **Exponential Growth Rate [16]:** Defined as $W_n = \frac{1}{n} \log \left(\frac{S_n}{S_0} \right)$, this metric evaluates the long-term growth potential of a strategy.
- **Sharpe Ratio [26]:** The `compute_sharpe_ratio()` function calculates risk-adjusted performance by annualizing the excess returns over a risk-free rate (assumed to be a constant 5%).
- **Maximum Drawdown:** The `calculate_maximum_drawdown()` function quantifies the largest peak-to-trough decline in cumulative wealth, expressed as a percentage. This metric reflects the worst historical loss during the investment horizon and is critical for evaluating downside risk.

These metrics ensure a balanced assessment of both profitability and risk, while also directly aligning with industry standards. The infrastructure is designed to evaluate various online portfolio selection strategies. Parameter tuning for each algorithm is automated via a grid search over predefined hyperparameter values (e.g., β and δ), with performance primarily ranked by the Sharpe ratio.

3.3. Scalability and Practicality

To ensure scalability, the framework leverages efficient vectorized operations provided by NumPy and pandas. Its modular design allows seamless integration of additional strategies, datasets, and performance metrics. This makes the infrastructure a powerful tool for exploring state-of-the-art methods in online portfolio selection.

ALGORITHM 1: Online portfolio selection framework.

Input: $\mathbf{x}_1^n$: Historical market sequence

Output: S_n : Final cumulative wealth

Initialize $S_0 = 1$, $\mathbf{b}_1 = (\frac{1}{m}, \dots, \frac{1}{m})$

for $t = 1, 2, \dots, n$ **do**

 Portfolio manager computes a portfolio $\mathbf{b}_t$;

 Market reveals the market price relative $\mathbf{x}_t$;

 Portfolio incurs period return $\mathbf{b}_t^\top \mathbf{x}_t$ and updates cumulative return $S_t = S_{t-1} \times (\mathbf{b}_t^\top \mathbf{x}_t)$;

 Portfolio manager updates his/her online portfolio selection rules ;

end

Figure 2. Environment Architecture and Algorithm Flow

Figure 2 summarizes the entire trading environment developed in Python, illustrating the process across each trading strategy. This framework was inspired by the survey assessing trading strategies in [20].

4. Algorithms Tested

Below is an overview of the trading strategies systematically compared in this study, as discussed in [20].

Classifications	Algorithms	Representative References
Benchmarks	Buy And Hold Best Stock Constant Rebalanced Portfolios	Kelly [1956]; Cover [1991]
Follow-the-Winner	Universal Portfolios Exponential Gradient Follow the Leader Follow the Regularized Leader Aggregating-type Algorithms	Cover [1991]; Cover and Ordentlich [1996] Helmhold et al. [1996] [1998] Gaivoronski and Stella [2000] Agarwal et al. [2006] Vovk and Watkins [1998]
Follow-the-Loser	Anti Correlation Passive Aggressive Mean Reversion Confidence Weighted Mean Reversion Online Moving Average Reversion Robust Median Reversion	Borodin et al. [2003] [2004] Li et al. [2012] Li et al. [2011b] [2013] Li and Hoi [2012] Huang et al. [2013]
Pattern-Matching Approaches	Nonparametric Histogram Log-optimal Strategy Nonparametric Kernel-based Log-optimal Strategy Nonparametric Nearest Neighbor Log-optimal Strategy Correlation-driven Nonparametric Learning Strategy Nonparametric Kernel-based Semi-log-optimal Strategy Nonparametric Kernel-based Markowitz-type Strategy Nonparametric Kernel-based GV-type Strategy	Györfi et al. [2006] Györfi et al. [2006] Györfi et al. [2008] Li et al. [2011a] Györfi et al. [2007] Oftucsák and Vajda [2007] Györfi and Vajda [2008]
Meta-Learning Algorithms	Aggregating Algorithm Fast Universalization Algorithm Online Gradient Updates Online Newton Updates Follow the Leading History	Vovk [1990] [1998] Akçoglu et al. [2002] [2004] Das and Banerjee [2011] Das and Banerjee [2011] Hazan and Seshadri [2009]

Figure 3. Algorithms Tested and Their Original References

4.1. Benchmarks

To evaluate the performance of online portfolio selection algorithms, we compare them against three benchmark strategies commonly used in financial literature. These benchmarks serve as simple, interpretable reference points that help contextualize the efficacy of more sophisticated portfolio selection methods.

1. Buy-and-Hold (BAH)

The Buy-and-Hold strategy is one of the simplest investment approaches. Under this strategy, an initial allocation of wealth is distributed across assets at the beginning of the trading period, and no rebalancing occurs thereafter. The portfolio holdings evolve passively based on the relative price changes of the assets. Mathematically, the portfolio weight at time t is given by:

$$b_t = \frac{b_{t-1} \odot x_{t-1}}{\sum_{i=1}^N (b_{t-1,i} x_{t-1,i})}. \quad (1)$$

2. Best Stock Strategy

The Best Stock strategy assumes perfect hindsight by investing all wealth in the single best-performing stock over the period. This strategy provides an upper bound on the performance that any portfolio selection method can achieve based solely on past price movements. Formally, at each time step:

$$b_t = e_{\arg \max x_{t-1}}, \quad (2)$$

where e_i is the one-hot vector representing full allocation in the best-performing asset.

3. Constant Rebalanced Portfolio (CRP)

The CRP strategy maintains a fixed proportion of wealth in each asset by rebalancing at every time step. Given an initial portfolio weight vector b , the strategy enforces:

$$b_t = b, \quad \forall t. \quad (3)$$

This approach ensures stability and leverages volatility pumping, where the portfolio benefits from price fluctuations. The optimal CRP in hindsight, referred to as the Best Constant Rebalanced Portfolio (BCRP), is computed as:

$$b^* = \arg \max_{b \in \Delta_N} \sum_{t=1}^T \log(b \cdot x_t), \quad (4)$$

where Δ_N denotes the probability simplex over N assets.

4.2. Follow-the-Winner

The Follow-the-Winner (FTW) approaches allocate more wealth to assets that have historically performed well. These strategies often attempt to track or improve upon the BCRP by dynamically adjusting portfolio allocations.

1. Universal Portfolios (Approximation)

Cover's Universal Portfolio [6] is a theoretically optimal FTW strategy that continually updates portfolio weights by averaging across a continuum of constant rebalanced portfolios, weighted by their past performance. Due to computational constraints, we approximate this strategy using a Monte Carlo method by sampling M random portfolios and updating their weights based on historical returns:

$$b_t = \frac{\sum_{i=1}^M w_{t-1}^{(i)} p_{t-1}^{(i)}}{\sum_{i=1}^M w_{t-1}^{(i)}}, \quad (5)$$

where $p_t^{(i)}$ represents the i^{th} sampled portfolio and $w_t^{(i)}$ is its corresponding wealth.

2. Exponential Gradient (EG)

The Exponential Gradient (EG) algorithm, proposed by Helmbold et al. [14], updates portfolio weights using a multiplicative adjustment based on past returns:

$$b_{t+1,i} = b_{t,i} \cdot \exp \left(\eta \left(\frac{x_{t,i}}{b_t \cdot x_t} - 1 \right) \right), \quad (6)$$

where η is the learning rate. The weights are normalized to maintain a valid portfolio.

3. Follow-The-Leader (FTL)

The FTL strategy selects the portfolio that would have been optimal in hindsight, based on cumulative past log returns:

$$b_t = \arg \max_b \sum_{\tau=1}^{t-1} \log(b \cdot x_\tau). \quad (7)$$

While this method reallocates capital to the historically best-performing assets, it can be vulnerable to market instability due to overfitting on short-term trends [9].

4. Follow-the-Regularized-Leader (FTRL)

A variation of FTL, FTRL introduces a regularization term to mitigate overreaction to short-term fluctuations:

$$b_t = \arg \max_b \left\{ \sum_{\tau=1}^{t-1} \log(b \cdot x_\tau) - \frac{\beta}{2} \|b\|^2 \right\}, \quad (8)$$

where β controls the strength of regularization. This formulation, inspired by the Online Newton Step (ONS) [1], improves the stability of portfolio updates.

5. Aggregation-Based Strategy

The Aggregation-Based method [30] combines multiple base portfolios by weighting them according to their historical performance:

$$b_t = \sum_{i=1}^N w_t^{(i)} p^{(i)}, \quad (9)$$

with the weights updated using an exponential rule:

$$w_t^{(i)} = \frac{w_{t-1}^{(i)} \cdot (p^{(i)} \cdot x_{t-1})^\eta}{\sum_{j=1}^N w_{t-1}^{(j)} \cdot (p^{(j)} \cdot x_{t-1})^\eta}. \quad (10)$$

This approach provides adaptive allocation while maintaining diversification.

4.3. Follow-the-Loser

In contrast to Follow-the-Winner approaches, Follow-the-Loser strategies are based on the principle of mean reversion—that is, assets that underperform tend to revert to their historical average performance. These methods aim to exploit market inefficiencies caused by overreaction to short-term trends [5]. The main strategies implemented in this research are:

1. Anti-Correlation (Anticor)

Introduced by Borodin et al. [5], the Anticor strategy exploits statistical mean reversion by transferring wealth from assets with high recent returns to those with lower returns. It does so by computing cross-correlation matrices over a rolling window of past price relatives:

$$Mcov(i, j) = \frac{1}{w-1} \sum_{t=1}^w (y_{t,i} - \bar{y}_i)(y_{t,j} - \bar{y}_j), \quad (11)$$

where $y_{t,i}$ represents the log price relative of asset i , and w is the look-back window size.

2. Passive Aggressive Mean Reversion (PAMR)

Proposed by Li et al. [19], PAMR formalizes mean reversion in an online learning framework. It penalizes portfolios when the predicted return exceeds a predefined threshold ϵ :

$$\ell_\epsilon(b_t; x_t) = \max\{0, b_t \cdot x_t - \epsilon\}. \quad (12)$$

The next portfolio is computed by solving a constrained optimization problem that minimizes deviation from the previous allocation while enforcing the mean reversion principle.

3. Confidence Weighted Mean Reversion (CWMR)

CWMR [18] extends PAMR by incorporating confidence-weighted learning. It models the portfolio as a Gaussian distribution with mean μ_t and covariance Σ_t , allowing the strategy to adjust the confidence of weight estimates dynamically:

$$(\mu_{t+1}, \Sigma_{t+1}) = \arg \min_{\mu, \Sigma} D_{KL}(N(\mu, \Sigma) \| N(\mu_t, \Sigma_t)), \quad (13)$$

where D_{KL} denotes the Kullback-Leibler divergence.

4. Online Moving Average Reversion (OLMAR)

OLMAR [19] improves upon single-period mean reversion strategies by using a moving average of past price relatives as a forecast:

$$\hat{x}_{t+1} = \frac{1}{w} \sum_{i=1}^w x_{t-i}. \quad (14)$$

This method helps smooth out short-term noise, resulting in a more robust rebalancing strategy.

5. Robust Median Reversion (RMR)

The RMR strategy [15] further enhances mean reversion by replacing the simple moving average with an ℓ_1 median estimator:

$$\hat{x}_{t+1} = \arg \min_{\mu} \sum_{i=1}^w \|x_{t-i} - \mu\|. \quad (15)$$

This approach reduces the influence of outliers, leading to a more stable allocation.

4.4. Pattern Matching

Pattern-matching approaches leverage historical market behavior to identify and apply trading strategies that were effective under similar conditions in the past. These nonparametric techniques select historical data samples based on similarity measures and then optimize portfolio allocation accordingly. Unlike the Follow-the-Winner and Follow-the-Loser strategies, pattern matching makes no explicit assumptions about market structure and relies solely on data-driven selection criteria.

Following the classification in [20], we implement four sample selection methods and three portfolio optimization techniques:

1. Sample Selection Methods

- (a) **Histogram-Based Selection** [11]: Historical price-relative vectors are segmented into predefined bins, and past periods with similar mean behavior (falling into the same bin as the most recent market window) are selected into the candidate set C_t .
- (b) **Kernel-Based Selection** [11]: A smooth similarity function based on Euclidean distance is applied between the most recent price-relative window and all past windows. Historical windows that fall within a predefined threshold are selected.
- (c) **Nearest Neighbor Selection** [17]: Historical windows are ranked by their distance from the latest price-relative window, and the k nearest neighbors are selected.
- (d) **Correlation-Based Selection** [4]: The Pearson correlation coefficient between the most recent window and each historical window is computed, and those with correlation above a threshold ρ are included in C_t .

2. Portfolio Optimization Methods

- (a) **Log-Optimal Portfolio** [11]: The optimal portfolio b^* is determined by maximizing the expected logarithmic return over the selected sample set C_t :

$$b^* = \arg \max_b \sum_{i \in C_t} \log(b \cdot x_i), \quad (16)$$

subject to $b \in \Delta_N$.

- (b) **Semi-Log-Optimal Portfolio** [12]: This method uses a smoothed objective function $f_z(z) = z - 0.5(z - 1)^2$ in place of $\log(z)$ to reduce sensitivity to extreme values:

$$b^* = \arg \max_b \sum_{i \in C_t} f_z(b \cdot x_i). \quad (17)$$

- (c) **Markowitz Portfolio** [23]: Inspired by classical mean-variance optimization, this method balances expected return and risk using:

$$b^* = \arg \max_b [\mu_C \cdot b - \lambda b^T \Sigma_C b], \quad (18)$$

where μ_C is the mean of the selected price-relative vectors, Σ_C is their covariance matrix, and λ controls the risk-return tradeoff.

3. Pattern-Matching Portfolio Framework

- (a) Select a candidate set C_t from historical data using one of the sample selection methods.
- (b) Compute the optimal portfolio weights using one of the portfolio optimization methods.
- (c) Rebalance the portfolio to b_t .

The table below summarizes the combinations of sample selection and portfolio optimization strategies that were tested:

Sample Selection Method	Portfolio Optimization Method
Histogram-Based Selection	Log-Optimal Portfolio
Histogram-Based Selection	Semi-Log-Optimal Portfolio
Histogram-Based Selection	Markowitz Portfolio
Kernel-Based Selection	Log-Optimal Portfolio
Kernel-Based Selection	Semi-Log-Optimal Portfolio
Kernel-Based Selection	Markowitz Portfolio
Nearest Neighbor Selection	Log-Optimal Portfolio
Nearest Neighbor Selection	Semi-Log-Optimal Portfolio
Nearest Neighbor Selection	Markowitz Portfolio
Correlation-Based Selection	Log-Optimal Portfolio
Correlation-Based Selection	Semi-Log-Optimal Portfolio
Correlation-Based Selection	Markowitz Portfolio

Table 1. Pattern-Matching Strategy Combinations

By leveraging past market patterns, these nonparametric strategies provide an adaptive approach to portfolio selection. Our empirical results compare the performance of these methods against other online trading algorithms.

4.5. Meta-Learning

Meta-learning algorithms in online portfolio selection enhance performance by dynamically combining multiple base strategies. Rather than relying on a single predefined approach, these methods form an ensemble of algorithms, adjusting allocations based on past performance to optimize returns. The goal is to achieve adaptability in non-stationary market conditions while maintaining theoretical performance guarantees.

The primary meta-learning strategies implemented in this research include Aggregation Algorithms, Fast Universalization, Online Gradient Updates, Online Newton Updates, and Follow-the-Leading-History. These methods build upon earlier work on ensemble learning and regret minimization in online convex optimization [29, 2, 7, 13].

1. Aggregation Algorithm (AA)

The Aggregation Algorithm weights individual base experts using an exponentially weighted scheme, as introduced by Vovk [29, 30]. At each time step, the expert weights are updated according to:

$$w_t^{(i)} = \frac{w_{t-1}^{(i)} \exp(-\eta \ell_t^{(i)})}{\sum_j w_{t-1}^{(j)} \exp(-\eta \ell_t^{(j)})}, \quad (19)$$

where $\ell_t^{(i)}$ is the loss for expert i at time t and η is the learning rate. The overall portfolio allocation is then computed as:

$$b_t = \sum_i w_t^{(i)} b_t^{(i)}. \quad (20)$$

2. Fast Universalization (FU)

Fast Universalization accelerates convergence by efficiently combining a subset of base strategies [2]. The method initializes with uniform weights over a set of pre-selected algorithms (e.g., CWMR, FTRL, and PAMR) and dynamically reweights them based on their cumulative performance:

$$b_t = \sum_i w_t^{(i)} b_t^{(i)}, \quad \text{where } w_t^{(i)} = \frac{w_{t-1}^{(i)} S_{t-1}^{(i)}}{\sum_j w_{t-1}^{(j)} S_{t-1}^{(j)}}. \quad (21)$$

3. Online Gradient Updates (OGU)

Inspired by exponential gradient methods in online convex optimization [7], Online Gradient Updates adjust the portfolio allocation via gradient descent updates over expert predictions:

$$b_{t+1,i} = b_{t,i} \cdot \exp \left(\eta \left(\frac{x_{t,i}}{b_t \cdot x_t} - 1 \right) \right). \quad (22)$$

The weights are normalized to ensure a valid portfolio.

4. Online Newton Updates (ONU)

ONU incorporates second-order gradient information to adaptively adjust the learning rate [7]. The portfolio update follows the Newton step:

$$b_t = b_{t-1} - \eta A^{-1} \nabla \ell_t, \quad (23)$$

where A is the Hessian matrix of historical losses.

5. Follow-the-Leading-History (FTLH)

The FTLH algorithm [13] maintains a set of historical experts, updating their portfolios using the Online Newton Step (ONS). At each time step, a new expert is introduced, and the Weighted Majority Algorithm (WMA) is applied to combine past expert performances:

$$w_t^{(i)} = \frac{w_{t-1}^{(i)} e^{-\eta \ell_t^{(i)}}}{\sum_j w_{t-1}^{(j)} e^{-\eta \ell_t^{(j)}}}. \quad (24)$$

Experts with poor historical performance (below a threshold) are pruned, ensuring the system focuses on the more effective strategies.

Meta-learning approaches enable dynamic portfolio adaptation by combining multiple experts into a robust framework. In our experiments, ensemble methods combining high-performing base strategies (CWMR, FTRL, PAMR) using gradient or Newton-based updates demonstrate enhanced risk-adjusted performance.

5. Experiments and Results

5.1. Training Methodology

5.1.1 Hyperparameter Tuning

Before officially comparing performances, we conducted hyperparameter tuning using a grid search. During initial experiments, we discovered that the performance of some algorithms was highly sensitive to their parameter settings. Although default parameters are convenient, they may not be optimal for our historical market data, potentially leading to suboptimal performance. By systematically exploring a range of candidate values with a grid search, we ensured that each algorithm was finely tuned to maximize key performance metrics such as the Sharpe ratio, enabling a fair and robust comparison across all methods.

Using the Anti-Correlation (Anticor) algorithm as an example, we defined a grid of candidate values for its key hyperparameters, including the rolling window size for computing correlations, the α parameter governing the degree of adjustment, and the correlation threshold that determines when to trigger wealth transfers between assets. As shown in Figure 4, this grid search resulted in eight unique hyperparameter configurations.

```
anticor_grid = {  
    'window_size': [3, 5],  
    'alpha': [1.5, 2.0],  
    'corr_threshold': [0.4, 0.5]  
}
```

Figure 4. Hyperparameter Grid Search Example – 8 Combinations

For each combination, the algorithm was executed on historical price-relative vectors, and its performance was evaluated using metrics such as cumulative wealth, exponential growth rate, and the Sharpe ratio (see Section 3.2). The hyperparameter set yielding the highest Sharpe ratio was selected as optimal.

Window Size	Alpha	Corr. Threshold	Final Wealth	Exp. Growth	Sharpe Ratio
3	2.0	0.5	36.88	0.00195042	1.01752
3	1.5	0.5	36.88	0.00115903	1.01751
3	2.0	0.4	35.57	0.00118301	1.00748
3	1.5	0.4	35.59	0.00116523	1.00777
5	2.0	0.5	28.65	0.00111968	0.93937
5	1.5	0.5	28.65	0.00110729	0.93727
5	2.0	0.4	26.50	0.00099422	0.89344
5	1.5	0.4	28.50	0.00110062	0.93440

Table 2. Anticor – Hyperparameter Tuning Results

This systematic tuning procedure was applied across all algorithms in the study, ensuring that each method was optimized under comparable conditions. To expedite the grid search, we employed the joblib library [28] for parallelization, significantly reducing computation time. For a comprehensive list of the optimal hyperparameters for each algorithm, please refer to Exhibit A in the Appendix.

5.1.2 Walk-Forward Validation

To enhance the robustness and practical relevance of our empirical results, we implemented a rolling-window walk-forward validation framework—commonly advocated in quantitative finance—to reflect non-stationary conditions and reduce look-ahead bias during hyperparameter tuning. This approach simulates real-world trading constraints by partitioning the historical dataset into multiple overlapping windows, each consisting of a distinct training and validation segment that advances sequentially through time. Specifically, we divide the data into up to five rolling validation windows, where each window holds out the final 20% of the segment for validation while using a preceding 30–40% segment for training. The windows roll forward chronologically, shifting both the training and validation periods to capture multiple temporal slices of market behavior.

The primary motivation for employing walk-forward validation—particularly in a rolling configuration—is to mitigate the risk of overfitting and to ensure that hyperparameter optimization generalizes across varying market regimes. Unlike conventional cross-validation—which may violate the temporal dependencies intrinsic to financial time series—walk-forward validation preserves the chronological structure of the data, training only on past information and validating on future, unseen data. This setup more accurately reflects the conditions under which trading strategies operate in practice and has been advocated in recent quantitative finance literature as a safeguard against backtest overfitting and spurious strategy performance [3, 10].

For each rolling window, the strategy is trained and then evaluated independently. Key performance metrics such as Sharpe ratio, maximum drawdown, and cumulative wealth are computed for the validation period. The final performance of a given parameter configuration is then averaged across all validation windows to provide a robust, time-consistent estimate of generalization ability.




Figure 5. Rolling-Window Walk-Forward Validation Visual [24]

By incorporating rolling-window walk-forward validation, we are able to evaluate not only the in-sample effectiveness of each algorithm but also its adaptability to shifting market dynamics. This procedure strengthens the external validity of our findings and enables more informed conclusions regarding the stability and robustness of each strategy under realistic deployment conditions.

5.2. Output

5.2.1 Follow-the-Winner

Figure 6 presents the cumulative wealth (on a logarithmic scale) for the Follow-the-Winner strategies over the sample period, compared against benchmark strategies (BAH, Best Stock, and CRP). The main observations are as follows:

- **Follow-the-Regularized Leader (FTRL)** stands out with a final wealth of 15.249 and the highest Sharpe ratio of 1.0395, indicating superior risk-adjusted performance. Although it lags in the early years, its robust final performance suggests that incorporating regularization into the allocation process significantly improves stability and returns.
- **Aggregation-Based Simple, Constant Rebalancing (CRP), and Exponential Gradient (EG)** exhibit nearly identical performance, with final wealth around 12.1584 and Sharpe ratios of approximately 0.7553. **Universal Portfolios** deliver solid results with a final wealth of 11.7161 and a Sharpe ratio of 0.7515, albeit at a longer runtime (0.0768 seconds).
- **Follow the Leader (FTL)** underperforms relative to the other dynamic strategies, achieving a final wealth of 8.43 and a Sharpe ratio of 0.5256, which suggests that its simplistic momentum capturing mechanism may not fully capture market dynamics.

Table 3 summarizes the performance metrics for each strategy, ordered by their Sharpe ratio. The results underscore the effectiveness of regularization in online portfolio selection, as evidenced by the performance of Follow-the-Regularized Leader.

Strategy	Final Wealth	Exponential Growth	Sharpe Ratio	Max Drawdown	Time (sec)
Follow the Regularized Leader	15.2490	0.0009	1.0395	-42.02%	0.6271
Aggregation-Based Simple	12.9385	0.0008	0.7655	-45.70%	0.0263
Constant Rebalancing	12.1584	0.0008	0.7552	-45.19%	0.0070
Exponential Gradient	12.1584	0.0008	0.7552	-45.19%	0.0655
Universal Portfolios	12.4189	0.0008	0.7531	-45.19%	0.1240
Buy and Hold	9.1742	0.0007	0.5796	-68.14%	0.0254
Follow the Leader	8.4369	0.0007	0.5256	-79.23%	3.5386
Best Stock	0.0153	-0.0014	0.0905	-99.95%	0.0147

Table 3. Performance Summary – Follow-the-Winner Strategies (Ordered by Sharpe Ratio)

In summary, meta-strategies that adaptively adjust weights based on recent performance—particularly those employing regularization (e.g., FTRL)—tend to deliver significantly improved risk-adjusted returns compared to both traditional benchmarks and less adaptive methods.




Figure 6. Cumulative Wealth of Follow-the-Winner Algorithms vs. Benchmarks (Log Scale)

5.2.2 Follow-the-Loser

Figure 7 displays the cumulative wealth (log scale) for the Follow-the-Loser strategies over the sample period, compared against benchmark strategies (BAH, Best Stock, and CRP). Key observations include:

- **Confidence Weighted Mean Reversion (CWMR)** achieves the strongest performance among all strategies, with the highest final wealth (1590.3818) and Sharpe ratio (1.7481). Its use of confidence-weighted updates enables effective exploitation of mean-reverting signals.
- **Passive Aggressive Mean Reversion (PAMR)** follows closely with a Sharpe ratio of 1.6292 and final wealth of 788.2571. By penalizing large deviations from expected returns, PAMR adapts quickly to reversals while maintaining control over volatility.
- **Anti-Correlation** achieves a Sharpe ratio above 1.0, highlighting the effectiveness of reallocating capital away from recent winners toward negatively correlated assets in capturing short-term reversals.
- **Robust Median Reversion (RMR)** and **Online Moving Average Reversion (OLMAR)** provide moderate risk-adjusted returns, indicating that simpler or more rigid formulations of reversion signals may underperform more adaptive approaches like CWMR and PAMR.

Table 4 summarizes the performance metrics for each Follow-the-Loser strategy, ordered by Sharpe ratio. The results indicate that advanced mean-reversion methods—especially those incorporating adaptive or probabilistic elements (CWMR, PAMR)—can substantially outperform traditional benchmarks.




Figure 7. Cumulative Wealth of Follow-the-Loser Strategies vs. Benchmarks (Log Scale)

Strategy	Final Wealth	Exp. Growth	Sharpe Ratio	Max Drawdown	Time (sec)
CWMR	1590.3818	0.0024	1.7481	-51.34%	1.2274
PAMR	788.2571	0.0022	1.6292	-47.04%	0.0688
Anti-Correlation	36.8868	0.0012	1.0174	-40.32%	0.8193
RMR	18.9278	0.0010	0.9611	-38.83%	0.7387
OLMAR	12.1584	0.0008	0.7552	-45.19%	0.0400
Constant Rebalancing	12.1584	0.0008	0.7552	-45.19%	0.0081
Buy and Hold	9.1742	0.0007	0.5796	-68.14%	0.0254
Best Stock	0.0153	-0.0014	0.0905	-99.95%	0.0148

Table 4. Performance Summary – Follow-the-Loser Strategies (Ordered by Sharpe Ratio)

5.2.3 Pattern Matching

Figure 8 illustrates the cumulative wealth (log scale) of the pattern-matching strategies over the sample period, compared against benchmark strategies (BAH, Best Stock, and CRP). The key findings are:

- **Histogram-Based Selection** combined with any of the three portfolio optimization methods (Log-Optimal, Semi-Log-Optimal, Markowitz) significantly outperforms other pattern-matching strategies, achieving final wealth between 578 and 610 and Sharpe ratios above 1.14. This validates the effectiveness of discretizing historical price relatives and targeting similar market states, though at the cost of extended runtimes (350–490 seconds).
- **Kernel-Based Selection** yields modest returns, with final wealth near 4.2–4.5 and Sharpe ratios just below 0.41. Despite smoother generalization compared to Nearest Neighbor methods, it suffers from large drawdowns (over 84%), which limit its risk-adjusted performance.
- **Nearest Neighbor Selection** results in low final wealth (1.4–2.2) and Sharpe ratios in the 0.39–0.43 range. The method's reliance on exact past similarities appears overly sensitive and vulnerable to noise or non-stationarity in the data.
- **Correlation-Based Selection** remains the weakest pattern-matching approach, with extremely low final wealth (0.12–0.15), minimal growth, and Sharpe ratios between 0.09 and 0.11. These results suggest that using correlation as the sole similarity metric is insufficient for reliable strategy selection.




Figure 8. Cumulative Wealth of Pattern-Matching Strategies vs. Benchmarks (Log Scale)

Table 5 summarizes these results, with strategies ordered by Sharpe ratio. While Histogram-based methods deliver the strongest performance overall, they also come with substantial computational overhead—highlighting the trade-off between precision in historical matching and runtime feasibility.

Strategy	Final Wealth	Exp. Growth	Sharpe Ratio	Max Drawdown	Time (sec)
Histogram-Based Selection + Semi-Log-Optimal Portfolio	610.4945	0.0021	1.1583	-75.49%	468.3940
Histogram-Based Selection + Markowitz Portfolio	578.2083	0.0021	1.1493	-76.11%	351.6113
Histogram-Based Selection + Log-Optimal Portfolio	598.7398	0.0021	1.1490	-77.52%	489.6378
Constant Rebalancing	12.1584	0.0008	0.7552	-45.19%	0.0057
Nearest Neighbor Selection + Markowitz Portfolio	2.2253	0.0003	0.4332	-96.78%	223.1461
Nearest Neighbor Selection + Semi-Log-Optimal Portfolio	1.6012	0.0002	0.4053	-97.09%	433.6794
Kernel-Based Selection + Markowitz Portfolio	4.5210	0.0005	0.4055	-84.10%	153.1853
Nearest Neighbor Selection + Log-Optimal Portfolio	1.4205	0.0001	0.3948	-97.25%	213.1959
Kernel-Based Selection + Semi-Log-Optimal Portfolio	4.1948	0.0005	0.3940	-84.02%	167.0283
Kernel-Based Selection + Log-Optimal Portfolio	4.2759	0.0005	0.3973	-84.18%	162.3842
Correlation-Based Selection + Markowitz Portfolio	0.1473	-0.0006	0.1091	-99.32%	280.6663
Correlation-Based Selection + Semi-Log-Optimal Portfolio	0.1282	-0.0007	0.0954	-99.40%	297.4214
Correlation-Based Selection + Log-Optimal Portfolio	0.1225	-0.0007	0.0913	-99.40%	297.7024
Best Stock	0.0153	-0.0014	0.0905	-99.95%	0.0104
Buy and Hold	9.1742	0.0007	0.5796	-68.14%	0.0223

Table 5. Performance Summary – Pattern-Matching Strategies (Ordered by Sharpe Ratio)

5.2.4 Meta-Learning

Figure 9 shows the cumulative wealth (log scale) of the meta-learning strategies over the sample period, compared with benchmarks (BAH, Best Stock, and CRP). In our experiments, we tested various combinations of base experts, eventually finding that an ensemble combining **CWMR**, **FTRL**, and **PAMR** yielded consistently superior performance. The key findings are:

- **Online Gradient Update (OGU)** and **Fast Universalization** emerge as the top-performing meta-learning strategies, with Sharpe ratios of 1.2631 and 1.2580 and final wealth values around 66–68. These results highlight the effectiveness of adaptive portfolio weighting based on recent expert performance, even when relying solely on first-order updates or efficient averaging techniques.

- **Online Newton Update (ONU)** also delivers strong risk-adjusted returns (Sharpe ratio 1.1379) with a slightly lower final wealth of 56.8144. While incorporating second-order information can improve convergence, the additional computational cost may not always translate into superior outcomes compared to OGU or Fast Universalization.
- **Aggregation-Based Generalized** and **Constant Rebalancing (CRP)** exhibit nearly identical results (final wealth = 12.16, Sharpe = 0.755), indicating that while ensemble-based diversification offers some benefits, the lack of adaptive updating limits performance under dynamic market conditions.
- **Follow the Leading History (FTLH)** underperforms relative to the other meta-learning approaches, ending with a Sharpe ratio of 0.6685 and a final wealth of 16.5305. Its tendency to follow recent winners may cause overfitting to short-term market noise rather than capturing durable trends.




Figure 9. Cumulative Wealth of Meta-Learning Strategies vs. Benchmarks (Log Scale)

Table 6 summarizes the performance metrics for the meta-learning strategies, sorted by Sharpe ratio. Notably, the best-performing approaches (ONU, OGU, and Fast Universalization) all rely on adaptively shifting capital among high-performing base portfolios (CWMR, FTRL, PAMR), demonstrating the effectiveness of ensemble methods in capturing diverse market behaviors. It is worth mentioning that although these ensemble methods exhibit attractive Sharpe ratios, they are significantly more computationally intensive compared to some of the more individual algorithms.

Strategy	Final Wealth	Exp. Growth	Sharpe Ratio	Max Drawdown	Time (sec)
Online Gradient Update (OGU)	65.7971	0.0014	1.2631	-44.79%	1685.0398
Fast Universalization	67.9772	0.0014	1.2580	-45.11%	2233.5067
Online Newton Update (ONU)	56.8144	0.0013	1.1379	-43.09%	1597.8124
Constant Rebalancing	12.1584	0.0008	0.7552	-45.19%	0.0075
Aggregation-Based Generalized	12.1514	0.0008	0.7550	-45.19%	0.0688
Follow the Leading History (FTLH)	16.5305	0.0009	0.6685	-64.37%	339.4211
Buy and Hold	9.1742	0.0007	0.5796	-68.14%	0.0251
Best Stock	0.0153	-0.0014	0.0905	-99.95%	0.0128

Table 6. Performance Summary – Meta-Learning Strategies (Ordered by Sharpe Ratio)

5.3. Performance Summary

Across the algorithm groups, our results reveal clear trends regarding performance and computational efficiency:

- **Follow-the-Winner:** The adaptive method *Follow-the-Regularized Leader (FTRL)* consistently delivers higher risk-adjusted returns, underscoring the benefits of incorporating regularization into the rebalancing process.
- **Follow-the-Loser:** Mean-reversion strategies that incorporate adaptive or confidence-based adjustments—especially *CWMR* and *PAMR*—achieve outstanding final wealth and Sharpe ratios by effectively exploiting mean-reversion opportunities.
- **Pattern Matching:** Histogram-Based Selection methods significantly outperform other pattern-matching variants in terms of returns and risk-adjusted performance, albeit at a substantially higher computational cost.
- **Meta-Learning:** Ensemble methods that combine high-performing base portfolios (*CWMR*, *FTRL*, *PAMR*) using gradient-based updates or Fast Universalization deliver strong risk-adjusted performance. However, these approaches are considerably more computationally intensive than many individual algorithms.

5.4. Comparison with Existing Literature

Our empirical results generally align with the established literature on online portfolio selection and quantitative trading. In particular, mean-reversion strategies—such as Confidence Weighted Mean Reversion (*CWMR*) and Passive-Aggressive Mean Reversion (*PAMR*)—consistently achieve high Sharpe ratios and cumulative wealth in mean-reverting market regimes, confirming earlier findings [8]. Likewise, momentum-based approaches like Follow-the-Regularized-Leader (*FTRL*) demonstrate robust performance when appropriate regularization is applied, in agreement with both academic studies and industry insights [21]. Although pattern-matching methods yield exceptional risk-adjusted returns, their high computational cost has been documented in previous work [27].

Our investigation into meta-learning ensembles—where base strategies such as *CWMR*, *PAMR*, and *FTRL* are combined using online gradient or Newton updates—indicates that adaptive weighting can enhance performance by capturing both momentum and mean-reversion effects. While these ensemble methods require considerable computational resources, their ability to dynamically adjust to changing market conditions confirms and extends prior work on expert aggregation techniques [30, 13]. Notably, this work integrates these diverse approaches into a unified framework that bridges theoretical performance with real-world deployment contexts in quantitative trading.

6. Experience

This section outlines the key challenges encountered during the research process, adjustments made to improve technical performance, and important considerations regarding the scope and applicability of the study.

6.1. Challenges

- **Limited Computing Resources:** Reliance on local processing made it difficult to efficiently handle large datasets and execute compute-intensive tasks, particularly during hyperparameter tuning and validation phases.
- **Data Storage Constraints:** Managing and accessing the historical data—structured across hundreds of CSV files—introduced bottlenecks that impacted preprocessing speed and overall runtime.

6.2. Approach Adjustments

- **Parallelization:** To accelerate hyperparameter searches, the `joblib` library [28] was used to implement parallel execution across CPU cores.
- **Modularization:** The codebase was restructured into modular components—data ingestion, algorithm logic, and evaluation—enabling cleaner experimentation and easier debugging.
- **Vectorization:** Computational performance was improved by refactoring core operations using `NumPy` and `pandas`, significantly reducing runtime during batch calculations.

6.3. Key Strengths

- **Repeatable Infrastructure:** The framework supports consistent evaluation of diverse trading strategies, enabling reproducible experiments and comparisons.
- **Robust Hyperparameter Tuning:** A comprehensive grid search, integrated with walk-forward validation, ensured that all algorithms were tested under realistic and comparable conditions.
- **Standardized Evaluation:** By applying uniform metrics and validation protocols, the study enables fair benchmarking of methods across algorithm families.

6.4. Research Limitations

- **Limited Market Scope:** The analysis was restricted to NASDAQ 100 stocks from 1998–2010. While this choice maintained data quality and computational feasibility, it limits generalizability to broader market conditions or more recent regimes.
- **Absence of Execution Constraints:** The study assumes frictionless trading and does not incorporate transaction costs, slippage, or live execution dynamics such as latency and position sizing. As a result, strategies with high turnover may appear more viable than they would be in a real-world setting.

7. Conclusion and Future Work

This study introduced a comprehensive Python-based backtesting infrastructure designed to evaluate a diverse range of online quantitative trading strategies. Through systematic implementation and extensive hyperparameter optimization, we benchmarked four canonical families of online portfolio selection algorithms—including momentum-driven Follow-the-Winner, mean-reversion-based Follow-the-Loser, pattern-matching techniques, and meta-learning ensembles—against conventional benchmarks like Buy-and-Hold, Best Stock, and Constant Rebalancing.

Our experiments, conducted on daily historical stock data, demonstrated that adaptive approaches (particularly those incorporating regularization, such as Follow-the-Regularized Leader, or confidence-based adjustments, such as CWMR and PAMR) consistently achieve superior risk-adjusted performance. Additionally, ensemble methods that combine multiple strong base strategies via gradient-based updates or Fast Universalization emerged as robust solutions for dynamically capturing diverse market behaviors.

However, the study also uncovered a practical trade-off: the substantial computational overhead of pattern-matching and meta-learning strategies may limit their application in high-frequency or intraday trading. The modular infrastructure developed in this work offers a repeatable approach to online portfolio selection and effectively bridges theoretical insights with practical financial applications.

Future Work – The challenges and results presented in this study provide several promising directions for further research:

- **Computational Efficiency:** Leverage GPU acceleration, distributed computing, and more efficient algorithmic techniques to significantly reduce runtime, especially for computationally intensive strategies.
- **Cloud Deployment:** Transition the framework to a cloud-native environment using containerization, CI/CD pipelines, and restful API frameworks to enable scalable, resilient, and collaborative system development.
- **Additional Algorithms and Metrics:** Expand the framework to include advanced techniques—such as deep reinforcement learning and evolutionary algorithms—and incorporate additional performance metrics (e.g. Conditional Value at Risk).
- **Broader Market Data and Asset Classes:** Evaluate strategies across alternative asset classes (cryptocurrencies, commodities, ESG-focused portfolios) and different time scales (intraday, weekly, monthly) to improve generalization and robustness.
- **Real-Time Integration:** Integrate the framework with real-time data feeds and simulate live trading conditions by accounting for transaction costs, order-book dynamics, latency, and slippage.
- **Explainability and Interpretability:** Enhance transparency and user trust by incorporating interpretability tools (such as SHAP, LIME, or attention mechanisms) to clarify the factors influencing algorithmic decisions.

Overall, our findings contribute to a deeper understanding of online portfolio selection while emphasizing the promise of adaptive, ensemble-based approaches in quantitative trading. The proposed future work also provides a road map for further refining these strategies and bridging the gap between theoretical performance and real-world applicability.

References

- [1] Amit Agarwal, Elad Hazan, Satyen Kale, and Robert E. Schapire. Algorithms for portfolio management based on the newton method. *ICML*, 2006. 7
- [2] M. Akcoglu et al. Fast universalization algorithm. *Journal of Financial Optimization*, 2002. 10
- [3] David H. Bailey, Jonathan M. Borwein, Marcos López de Prado, and Qiji Jim Zhu. The probability of backtest overfitting. *Journal of Computational Finance*, 20(4):39–69, 2014. 11
- [4] Steven C. H. Hoi Bin Li and Peilin Zhao. Confidence weighted mean reversion strategy for online portfolio selection. *ACM Transactions on Intelligent Systems and Technology*, 3(2):1–28, 2011. 9
- [5] Allan Borodin, Ran El-Yaniv, and Vincent Gogan. Can we learn to beat the best stock? *Journal of Artificial Intelligence Research*, 18:221–287, 2003. 8
- [6] Thomas M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991. 5, 7
- [7] P. Das and A. Banerjee. Online gradient and newton updates for portfolio selection. *Journal of Machine Learning Research*, 2011. 10
- [8] Ali Fereydooni and Others. Pattern-matching approaches in portfolio optimization: A comparative study. *Journal of Financial Data Science*, 5(1):34–56, 2023. 17
- [9] Alexei Gaivoronski and Francesco Stella. Adaptive portfolio optimization. *Annals of Operations Research*, 99(1-4):165–187, 2000. 7
- [10] Cédric Geurts, Jan De Spiegeleer, Dilip B. Madan, and Wim Schoutens. Backtesting systematic trading strategies: A quantitative approach. *Quantitative Finance*, 22(2):311–332, 2022. 11
- [11] László Györfi, Gábor Ottucsák, and Harro Walk. *Machine Learning for Financial Engineering*. World Scientific, 2006. 9
- [12] László Györfi and Harro Walk. Empirical portfolio selection based on log-optimal strategy. *Mathematical Finance*, 17(3):493–514, 2007. 9
- [13] E. Hazan and C. Seshadhri. Follow the leading history for online portfolio selection. *Mathematical Finance*, 2009. 10, 17
- [14] 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. 7
- [15] Steven C. H. Hoi Jiangang Huang and Bin Li. Robust median reversion strategy for online portfolio selection. *ACM Transactions on Intelligent Systems and Technology*, 5(1), 2013. 8
- [16] John L. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35(4):917–926, 1956. 5
- [17] György Ottucsák László Györfi and Harro Walk. Nearest neighbor based portfolio selection strategies. *Statistics and Decisions*, 26:145–157, 2008. 9
- [18] Bin Li and Steven C.H. Hoi. Confidence weighted mean reversion strategy for online portfolio selection. *ACM Transactions on Intelligent Systems and Technology*, 4(4), 2013. 8
- [19] Bin Li, Steven C.H. Hoi, and Doyen Sahoo. Passive aggressive mean reversion strategy for portfolio selection. *Machine Learning and Knowledge Discovery in Databases (ECML PKDD)*, pages 77–92, 2012. 8
- [20] Bin Li and Steven C. H. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys*, 46(3):1–36, 2013. Available at: <http://arxiv.org/abs/1212.2129v2>. 4, 5, 6, 8
- [21] Tobias J. Moskowitz, Yongheng Ooi, and Lasse H. Pedersen. Time series momentum. *Journal of Financial Economics*, 104(2):228–250, 2012. 17
- [22] nglahani. Online quantitative trading strategies. <https://github.com/nglahani/Online-Quantitative-Trading-Strategies>, 2020. GitHub repository (accessed: 2025-02-06). 4
- [23] György Ottucsák and István Vajda. An efficient implementation of markowitz portfolio selection. *Annals of Operations Research*, 151(1):181–192, 2007. 9
- [24] José Manuel Pino, Daniel Morales, David Salinas, Julián Agudelo, Álvaro Orozco, and David Restrepo. Forecasting financial time series through causal and dilated convolutional neural networks. *Entropy*, 22(10):1094, 2020. Accessed: 2025-05-05. 12
- [25] Inc. QuantQuote. Quantquote historical market data, 2024. Accessed: January 28, 2025. 4
- [26] William F. Sharpe. The sharpe ratio. *The Journal of Portfolio Management*, 21(1):49–58, 1994. 5
- [27] Tim van Erven, Wouter Koolen, et al. *Online Learning: Theory, Algorithms, and Applications*. Cambridge University Press, 2020. 17
- [28] Gaël Varoquaux. Joblib: a set of tools to provide lightweight pipelining in python. <https://joblib.readthedocs.io/>, 2013. 11, 17
- [29] V. Vovk. Aggregating strategies. *Proceedings of COLT*, 1990. 10
- [30] Vladimir Vovk and Chris Watkins. Universal portfolio selection. *Theoretical Computer Science*, 244(1-2):37–58, 1998. 7, 10, 17

8. Appendix

Table Exhibit A. Optimal Hyperparameters for All Algorithms

Category	Algorithm	Optimal Parameters
Follow-the-Loser		
	Anticorrelation (Anticor)	window_size = 3, $\alpha = 2.5$, corr_threshold = 0.5
	PAMR	$\epsilon = 0.9$, C = 10.0
	CWMR	$\epsilon = 0.89$, $\theta = 0.92$, $\eta = 0.93$
	OLMAR	window_size = 2, $\epsilon = .8$, $\eta = 20$
	RMR	window_size = 8, $\epsilon = 1.1$, $\eta = 20$
Follow-the-Winner		
	Universal Portfolios	num_portfolios = 3, $\tau = 0.3$
	Exponential Gradient (EG)	learning_rate = 0.05, smoothing = 0.0
	Follow-the-Leader (FTL)	$\gamma = .8$, $\alpha = 1.5$
	Follow-the-Regularized-Leader (FTRL)	$\beta = 0.13$, $\delta = 0.925$, ridge_const = 0.015
	Aggregation-Based Simple	learning_rate = 0.4, num_base_portfolios = 3
Meta-Learning		
	Aggregation Algorithm (Generalized)	learning_rate = 0.005, $\gamma = 0.3$
	Fast Universalization	learning_rate = 0.1, base_experts = {CWMR, FTRL, PAMR}
	Online Gradient Update Meta	learning_rate = 0.01, base_experts = {CWMR, FTRL, PAMR}
	Online Newton Update Meta	learning_rate = 0.01, base_experts = {CWMR, PAMR, RMR, OLMAR}
	Follow-the-Leading-History (FTLH)	$\eta = 0.35$, learning_rate = 0.07, drop_threshold = 0.55
Pattern Matching		
	Histogram-Based Selection	w = 4, bins = {(0,0.5), (0.5,1), (1,1.5)}
	Kernel-Based Selection	w = 5, threshold = 0.1
	Nearest Neighbor Selection	w = 3, num_neighbors = 3
	Correlation-Based Selection	w = 3, $\rho = 0.6$
	Log-Optimal Portfolio	(No additional parameters)
	Semi-Log-Optimal Portfolio	(No additional parameters)
	Markowitz Portfolio	$\lambda = 0.7$
	Pattern-Matching Portfolio Master	w = 4, methods as defined