

2025 S.T. Yau High School Science Award (Asia)

Research Report

Authorship

Registration number: Econ-044

Name: Junqin Yan

School: Melbourne High School

Location: Melbourne, Victoria, Australia

Supervisor: Bill Theodoropoulos (bill.theodoropoulos@mhs.vic.edu.au),
Careers and Pathways Coordinator, Melbourne High School, Melbourne, Australia.

Title of Research Report

Momentum Investing: A Comparison of Machine Learning and Markowitz Theory

Declaration

This research has been submitted previously to another competition. However, this version has been substantially revised, with improved clarity, conciseness and relevance of information.

Momentum Investing: A Comparison of Machine Learning and Markowitz Theory

Junqin Yan

Econ-044

Melbourne High School

Melbourne, Australia

junqinyan1@gmail.com

Abstract

This paper presents a comparative analysis of two momentum-based trading strategies: a traditional rule-based approach using mean-variance optimization and a machine learning approach utilizing the XGBoost algorithm. The rule-based strategy selects stocks based on recent short-term returns and applies mean variance optimization to enhance portfolio construction. Although this method achieved lower overall returns (17.55%), it demonstrated a more stable risk-adjusted performance, with a maximum drawdown of -68.35% and a return-to-drawdown ratio of 0.26. However, its effectiveness is constrained by idealized assumptions, such as normally distributed returns and static covariances, that often fail under real-world market dynamics.

In contrast, the XGBoost-based strategy, which learns from historical price and volume data to forecast short-term returns, achieved stronger predictive performance. It generated a cumulative return of 40.43% and a higher return-to-drawdown ratio of 0.31, although this came with a significantly greater downside risk, reflected in a maximum drawdown of -84.45%. These results underscore the fragility of machine learning models in volatile environments, despite their ability to uncover complex return signals.

Overall, this study highlights the trade-offs between traditional financial models and machine learning approaches in momentum investing. While each offers distinct advantages, future research should explore hybrid strategies that combine statistical discipline with machine learning flexibility, as well as alternative models such as neural networks to improve robustness in dynamic market conditions.

Keywords

Momentum

Machine-Learning

Optimization

Commitments on Academic Honesty and Integrity

Research Report

2025 S.T. Yau High School Science Award (Asia)

Commitments on Academic Honesty and Integrity


We hereby declare that we

1. are fully committed to the principle of honesty, integrity and fair play throughout the competition.
2. actually perform the research work ourselves and thus truly understand the content of the work.
3. observe the common standard of academic integrity adopted by most journals and degree theses.
4. have declared all the assistance and contribution we have received from any personnel, agency, institution, etc. for the research work.
5. undertake to avoid getting in touch with assessment panel members in a way that may lead to direct or indirect conflict of interest.
6. undertake to avoid any interaction with assessment panel members that would undermine the neutrality of the panel member and fairness of the assessment process.
7. observe the safety regulations of the laboratory(ies) where the we conduct the experiment(s), if applicable.
8. observe all rules and regulations of the competition.
9. agree that the decision of YHSA(Asia) is final in all matters related to the competition.

We understand and agree that failure to honour the above commitments may lead to disqualification from the competition and/or removal of reward, if applicable; that any unethical deeds, if found, will be disclosed to the school principal of team member(s) and relevant parties if deemed necessary; and that the decision of YHSA(Asia) is final and no appeal will be accepted.

(Signatures of full team below)

X


Name of team member: Jungin Yan

X

Name of team member:

X

Name of team member:

X


Name of supervising teacher:
Bill Theodoropoulos

Noted and endorsed by



(signature)

Name of school principal: PAUL DREW

Contents

1	Introduction	4
1.1	Overview	4
2	Approach	5
2.1	Markowitz Based Rule Approach	5
2.2	Machine Learning Based Approach: XGBoost	6
2.3	Evaluation Metrics	8
3	Results	8
3.1	Rule-Based Strategy	8
3.1.1	Mean-Variance Weighting	10
3.1.2	Comparison	11
3.2	Machine Learning Based Strategy	12
3.2.1	Feature Importance	12
4	Conclusion	14

1 Introduction

Momentum investing, anchored in the observation that assets with strong recent performance tend to continue outperforming in the short term, remains a cornerstone of quantitative finance (1). Momentum strategies exploit behavioral inefficiencies such as investor herding and delayed price reactions. While empirical studies have validated the profitability of these strategies across time periods and markets (2), their effectiveness often diminishes in the presence of sudden reversals, high volatility, or regime shifts (3).

Traditional implementations typically rely on simple rules, such as ranking stocks by recent returns and forming equally weighted portfolios from the top or bottom performers. While transparent and easy to execute, these rule-based approaches may overlook important risk factors and fail to capture the complexity of real-world market behavior.

Recent research has shown that machine learning models, particularly decision trees and neural networks, significantly outperform traditional linear return forecasting methods in empirical asset pricing (4). However, there is still a lack of direct comparison between these models and classical portfolio construction frameworks, such as Markowitz's mean-variance optimization, particularly within the context of momentum-based investing.

To address these limitations, this study compares two enhanced approaches to momentum investing: one grounded in classical financial theory and the other driven by modern machine learning.

The first enhancement applies Markowitz's mean-variance optimization to a momentum selected stock pool, assigning risk adjusted weights in an effort to improve portfolio efficiency. This method aims to maximize expected return for a given level of volatility. However, its practical effectiveness is often constrained by strong assumptions, such as normally distributed returns, static covariances and frictionless markets, which rarely hold in dynamic financial environments.

The second approach employs XGBoost, a powerful ensemble learning algorithm increasingly used in finance for return prediction, risk classification and portfolio construction. Capable of modeling complex, non-linear relationships in historical price and volume data, XGBoost provides strong predictive performance and insights through feature importance ranking. Nevertheless, it also presents challenges: a tendency to overfit noisy financial data, limited interpretability and the absence of built-in risk controls that are essential in real-world portfolio management.

By evaluating these two enhancements within the context of short-term momentum strategies, this paper aims to assess their relative effectiveness in balancing predictive accuracy with practical risk management. Ultimately, these findings can contribute to the broader discussion on how traditional financial models and machine learning can be integrated in the evolution of quantitative investing.

1.1 Overview

This paper investigates the comparative effectiveness of two enhancements to momentum-based investing: a risk-optimized approach using classical Markowitz's Modern Portfolio Theory and a predictive, data-driven method using XGBoost. While both techniques are applied within a short-term momentum framework, they reflect contrasting philosophies: one grounded in statistical portfolio construction, the other in machine learning-driven return forecasting.

The objective is to assess how each approach balances return potential with risk control when applied to momentum investing. Through empirical evaluation, the study aims to clarify their respective

advantages and limitations in order to contribute to the evolving dialogue on blending traditional financial models with modern data science in quantitative asset management.

2 Approach

The baseline momentum strategy in this study involves systematically buying stocks with the highest or lowest returns over the past five trading days, assuming that recent trends may continue. Portfolios are rebalanced every five days, selecting ten stocks based on their short-term performance and allocating equal weights. This rule-based approach serves as a control to assess the value added by more advanced techniques like mean-variance optimization and machine learning.

The data used for all strategies in this study is drawn from China's A-share market, covering a 17 years time period from 2007 to 2024. This dataset was chosen for its long historical span and the A-share market's relative structural stability compared to more volatile indices such as the NASDAQ, making it a suitable environment for testing momentum-based strategies.

This section outlines the two primary methodologies evaluated in this study: a rule-based mean-variance optimized approach and a machine learning approach utilizing XGBoost. The rule-based method selects stocks based on historical return rankings, offering simplicity and transparency, while the XGBoost model leverages data-driven insights to forecast short-term returns with greater flexibility and predictive accuracy. Each strategy is used to identify and re-balance a momentum driven portfolio and is assessed for its ability to capture return patterns and manage risk.

2.1 Markowitz Based Rule Approach

This strategy enhances a traditional momentum screen by combining quantitative analysis to identify optimal return percentiles with Markowitz's Modern Portfolio Theory, aiming to improve portfolio efficiency through both targeted stock selection and risk-adjusted weight allocation. Specifically, statistical evaluations are used to determine which decile ranges historically offer the best trade-off between return and risk, while mean-variance optimization further refines the portfolio by adjusting weights to optimize the balance between expected return and risk.

- **Data collection and return calculation:** Every five trading days, closing prices are collected for all stocks. Returns are computed using a 5-day look-back window:

$$R_i = \frac{P_{i,t} - P_{i,t-5}}{P_{i,t-5}}$$

Where R_i is the return of stock i and $P_{i,t}$, $P_{i,t-5}$ denote its price at time t and $t - 5$, respectively.

- **Stock ranking and selection:** Stocks are ranked by 5-day returns and grouped into deciles. A quantitative screen evaluates each group's historical performance, identifying the decile with the best return-risk profile. A simple backtest of cumulative return and maximum drawdown for each decile group was used to evaluate risk-return trade-offs and select the most favorable percentile range. The 10 stocks from this decile are then selected for the portfolio.
- **Portfolio optimization (mean-variance):** For the selected stocks, an additional 30 days of historical data is collected. Expected returns and the covariance matrix are computed to capture both average performance and risk relationships among assets.

Using this information, weights are optimized via the Markowitz mean-variance framework to maximize expected return for a given level of risk:

$$\min_w \frac{1}{2} w^T \Sigma w - \lambda w^T \mu \quad \text{subject to} \quad \sum w_i = 1, \quad w_i \geq 0$$

Where μ is the vector of expected returns, Σ is the covariance matrix and λ denotes risk aversion.

- **Portfolio construction:** Funds are allocated to the selected stocks according to the optimized weights. The portfolio is held for five trading days.
- **Rebalancing:** After five days, all positions are liquidated. The process then restarts using updated return rankings.
- **Performance evaluation:** The strategy is evaluated over time using cumulative return, maximum drawdown and Sharpe ratio:

– **Sharpe ratio:**

$$S = \frac{R_p - R_f}{\sigma_p}$$

Where R_p is portfolio return, R_f is the risk-free rate and σ_p is the standard deviation of returns.

Although default hyperparameters were used, the model's performance could potentially be improved with hyperparameter tuning and cross-validation, which were left for future work due to computational constraints. This rule-based approach retains the momentum signal but refines portfolio construction through statistical optimization, aiming to balance return and risk more effectively than equal weighting.

2.2 Machine Learning Based Approach: XGBoost

XGBoost (Extreme Gradient Boosting) is a high-performance ensemble learning algorithm known for its predictive accuracy and scalability. Unlike traditional statistical models, XGBoost constructs an additive ensemble of decision trees, where each successive tree corrects the residuals of the current model (5). In this study, XGBoost is employed as a supervised learning approach aimed at predicting short-term stock returns using historical price and volume data. This allows for the identification of the most promising stocks to hold over 5-day trading windows.

The prediction for a given input x_i is modeled as the sum of outputs from K regression trees:

$$\hat{y}_i = \sum_{k=1}^K f_k(x_i), \quad f_k \in \mathcal{F} \quad (1)$$

The objective function to be minimized includes a loss function and a regularization term:

$$\text{Obj} = \sum_{i=1}^n \ell(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (2)$$

Here, ℓ is typically the mean squared error for regression tasks. The regularization term $\Omega(f_k)$ penalizes model complexity:

$$\Omega(f_k) = \gamma T + \frac{1}{2} \lambda \sum_{j=1}^T w_j^2 \quad (3)$$

Where T is the number of leaves, w_j the weight of leaf j and γ and λ are regularization parameters. This framework balances predictive accuracy and model generalizability.

- **Data collection:** Gather 30 days of historical K-line data (open, close, high, low prices and volume) for each stock.
- **Input construction:** The XGBoost model is trained on 30 days of market data per stock, including price points and volume. Each row captures a stock's daily behavior, with the output predicting its movement over the next five days.

Table 1: Structure of the Input and Output Data Used for XGBoost Modeling

$P_{1,o}^1$	$\dots$	$P_{1,o}^{30}$	$P_{1,c}^1$	$\dots$	$P_{1,c}^{30}$	$P_{1,h}^1$	$\dots$	$P_{1,h}^{30}$	$P_{1,l}^1$	$\dots$	$P_{1,l}^{30}$	V_1^1	$\dots$	V_1^{30}	Y_{11}
$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$P_{2,o}^1$	$\dots$	$P_{2,o}^{30}$	$P_{2,c}^1$	$\dots$	$P_{2,c}^{30}$	$P_{2,h}^1$	$\dots$	$P_{2,h}^{30}$	$P_{2,l}^1$	$\dots$	$P_{2,l}^{30}$	V_2^1	$\dots$	V_2^{30}	Y_{21}
$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$P_{n,o}^1$	$\dots$	$P_{n,o}^{30}$	$P_{n,c}^1$	$\dots$	$P_{n,c}^{30}$	$P_{n,h}^1$	$\dots$	$P_{n,h}^{30}$	$P_{n,l}^1$	$\dots$	$P_{n,l}^{30}$	V_n^1	$\dots$	V_n^{30}	Y_{n1}
$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
Input															Output

Table 1 represents the structure of the input and output data used in XGBoost modeling where input variables P denote price and V denote volume. The superscript indicates the trading day. The first subscript refers to stock number while the second subscript o, c, h, l represents, open, close, high and low respectively. Y is the output, indicating the predicted price movement.

- **Labeling:** Define the output as the future 5-day return.
- **Training:** Train the XGBoost model using labeled A-shares stock data.
- **Prediction and ranking:** Predict returns for each stock and rank them by forecasted return.
- **Portfolio construction:** The top 10 stocks are selected and invested in with equal weighting for a 5-day holding period.
- **Remodeling:** After five days, all positions are liquidated and the process restarts using updated stock data.
- **Performance evaluation** The strategy is evaluated over time using cumulative return and maximum drawdown:

– **Cumulative return (machine learning):**

$$R_t = \prod_{i=1}^T (1 + r_i) - 1$$

Where r_i is the portfolio return for each 5-day window. Stocks are held equally for each cycle based on the top 10 predicted returns.

2.3 Evaluation Metrics

Cumulative Return: Measures the total return from the start of an investment to a given point, reflecting overall strategy performance (6). Calculated as:

$$R_{\text{cumulative}} = \frac{V_t - V_0}{V_0} \quad (4)$$

Where V_t is the value of the portfolio at time t and V_0 is the initial investment amount.

Maximum Drawdown: Indicates the largest drop from peak to trough, used to capture the portfolio's risk. Calculated as:

$$\text{Max Drawdown} = \max_{t \in [0, T]} \left(\frac{V_{\text{peak}} - V_t}{V_{\text{peak}}} \right) \quad (5)$$

Where V_{peak} is the highest portfolio value during the period and V_t is the portfolio value at time t .

3 Results

My analysis began with an exploration of recurring price patterns in 5-day candlestick sequences, which revealed four typical motifs: consecutive gains, gains followed by losses, losses followed by rebounds and consecutive losses. These patterns suggest a mix of short-term trend persistence and rapid reversals, enough to support the use of a 5-day momentum screen, but also volatile enough to highlight the importance of managing downside risk.

To test the viability of short-term momentum, I examined two baseline strategies: selecting the top 10 highest-returning stocks and the 10 lowest over the prior 5-day window. While the low-buy strategy appeared promising at first, both approaches ultimately performed poorly over the 17-year backtest. The high-buy strategy exhibited minimal growth from the outset and the low-buy portfolio eventually reversed its early gains.

These results demonstrate the limitations of relying solely on recent returns as a predictive signal. Short-term performance often reflects temporary cyclical expansion rather than sustained structural trends where prices frequently revert as market sentiment stabilizes. This cyclical behavior, often disconnected from underlying fundamentals, can erode momentum-based returns if left unchecked.

While momentum remains a well-established phenomenon in financial literature, these findings suggest that naive implementations, such as selecting stocks purely on 5-day past returns, may fail to capture its deeper structure and are particularly vulnerable to short-term volatility and noise: an observation consistent with Barroso and Santa-Clara's findings on momentum crash risk and fragility in volatile regimes (3).

Recognizing this, I introduced mean-variance optimization and machine learning based stock selection as enhancements to the core momentum framework. These methods aim to counter the instability of raw return signals by improving portfolio construction and sharpening prediction accuracy, providing a more robust foundation for consistent, long-term performance.

3.1 Rule-Based Strategy

This section evaluates the performance of the rule-based strategy. Stocks were first screened using a short-term momentum signal, based on 5-day return rankings, and grouped into deciles to identify favorable return-risk profiles. From the best-performing decile, a portfolio of ten stocks was constructed using equal weighting. The strategy's effectiveness was assessed through key metrics

including cumulative return, maximum drawdown and return-to-drawdown ratio, providing insight into the strengths and limitations of naive momentum-based selection when applied to dynamic market conditions.




Figure 1: Cumulative return for high-to-low and low-to-high ranking approaches.




Figure 2: Maximum drawdown for high-to-low and low-to-high ranking approaches.




Figure 3: Return-to-drawdown ratio for high-to-low and low-to-high ranking approaches.


Figure 1 analyzes cumulative return for the two ranking approaches across 1,100 stocks grouped into deciles: (a) shows the high to low ranking based on 5-day past return, while (b) presents the low-to-high ranking. Figure 2 presents maximum drawdown for the same approaches, with (a) corresponding to high-to-low and (b) to low-to-high. The return-to-drawdown ratio in Figure 3 follows the same structure, with (a) for high-to-low and (b) for low-to-high. For the high-to-low approach, all three metrics—cumulative return, maximum drawdown, and return-to-drawdown ratio—exhibit a broadly similar pattern. Each begins at a relatively low value in the first decile group, peaks in the 71st-80th interval, and then declines in the later deciles. By contrast, the low-to-high approach displays the reverse behavior: values for all three metrics start at their lowest point in the early intervals and gradually increase across the ranking spectrum, suggesting a more monotonic improvement as one moves toward higher-ranked deciles.

These patterns suggest that stocks temporarily depressed by cyclical factors, such as sector rotation or short-term sentiment, tend to rebound, while those suffering from deeper structural issues, like deteriorating fundamentals or inefficiencies, are more likely to experience lasting declines. This is reflected in (Figure 1b), where the initial low-return group likely consists of structurally impaired firms that have failed to recover. The subsequent spike in returns appears to capture stocks rebounding from cyclical lows, an effect that is often short-lived. Conversely, (Figure 1a) shows that stocks at the start of the distribution likely benefited from temporary cyclical strength, which quickly fades over time. Given the long-term stability of the A-shares index, the relative balance observed may result from structural losses in these declining firms being offset by gains from new market entrants. As a result, stocks in the mid-range growth deciles are less influenced by structural extremes and more shaped by short-term cyclical fluctuations. This reinforces the notion that extreme recent performance, whether high or low, often fails to persist in the short term, limiting the reliability of simple momentum-based selection.

3.1.1 Mean-Variance Weighting

Figure 4 depicts the efficient frontier, illustrating the trade-off between risk and return for portfolios composed of 10 stocks selected during the quantitative screening stage on June 24, 2020. Each point corresponds to a unique weight allocation, with the red star marking the portfolio that achieves the highest Sharpe ratio. This point represents the optimal balance between return and risk, offering the most efficient risk-adjusted performance among all portfolios considered.

Figure 4: Efficient frontier showing the risk-return profiles of 10 selected stocks.



This reflects the momentum-based strategy's premise that optimizing weights via mean-variance methods can improve performance. By calibrating the balance between return and risk at this optimal point, such portfolio construction may enhance robustness under varying market conditions.

3.1.2 Comparison

Figure 5: Comparison of cumulative return, weighted and unweighted (71st-80th).




Figure 6: Comparison of cumulative return, weighted and unweighted (61st-70th)




Table 2: Comparison of portfolio performance metrics for weighted and unweighted strategies (Figures 5 and 6).

Metric	Figure 5 (71st-80th)		Figure 6 (61st-70th)	
	Weighted	Unweighted	Weighted	Unweighted
Cumulative Net Value	7.33	11.18	8.52	6.85
Annualized Return	16.21%	19.99%	17.55%	15.63%
Maximum Drawdown	-59.24%	-60.33%	-68.35%	-64.76%
Return/Drawdown Ratio	0.27	0.33	0.26	0.24

Figures 5 and 6 compare cumulative returns for weighted versus unweighted portfolios across two stock intervals, 71st-80th and 61st-70th percentiles, identified as optimal ranges for growth. In both cases, the orange line represents the unweighted strategy, while the blue line depicts the mean-

variance-weighted portfolio. Spanning from 2010 to 2023, the graphs show similar upward trends, punctuated by notable spikes in 2015 and 2022.

Detailed in Table 2, in the 71st-80th percentile group (Figure 5), the unweighted portfolio outperforms the weighted one, ending with a cumulative return around 11.18% versus 7.33%. In contrast, the 61st-70th percentile group (Figure 6) shows the opposite: the weighted portfolio achieves a higher final return, 8.52%, compared to the unweighted version, 6.85%. These mixed results reinforce the notion that the Markowitz mean-variance framework is limited in capturing the complexity of dynamic market conditions (7).

3.2 Machine Learning Based Strategy

I apply Extreme Gradient Boosting (XGBoost), a scalable tree boosting framework introduced by Tianqi Chen and Carlos Guestrin (5), known for its strong performance in supervised learning. XGBoost builds shallow decision trees sequentially, offering high accuracy with reduced overfitting in both classification and regression tasks. For regression, it minimizes squared loss to predict continuous outcomes, while allowing user-defined objectives and evaluation metrics for flexibility. This research uses the open-source Python implementation with default hyperparameters, providing a standardized baseline for model training.

The model is trained to predict 5-day forward stock returns based on historical pricing and volume data. The model was retrained at each rebalancing point using a rolling window of the most recent 30 days of stock data to reflect evolving market conditions. Once trained, it generates return predictions for each stock in the testing window. Stocks are ranked by their predicted returns, and the top 10 are selected to form the investment portfolio. These stocks are held for a 5-day trading cycle, after which the portfolio is liquidated and the process repeated.

This predictive allocation process is repeated across multiple trading cycles, allowing for dynamic portfolio composition that adapts to recent market patterns. Model performance is assessed based on cumulative return, annualized return, maximum drawdown, and the return-to-drawdown ratio, offering a holistic view of both profitability and risk exposure under this machine learning based strategy.

3.2.1 Feature Importance

Table 3: Feature importance by category in XGBoost regression and classification models (in percentage)

Feature Category	Regression (%)	Classification (%)
Closing Price	24.0	21.5
High Price	23.5	19.5
Low Price	20.5	21.0
Opening Price	17.0	19.0
Volume	15.0	19.0

In the regression model, which predicts numerical return value over 5 days, the closing price is the most influential feature (24%), followed by high (23.5%) and low prices (20.5%). This prominence of the closing price likely stems from its role as a summary of daily market consensus, making

it particularly reflective of recent sentiment and momentum. High and low prices add context by capturing volatility and price extremes, helping to model upside and downside risk.

In the classification model, which predicts directional movements (rise/fall), importance is more evenly spread. The closing price remains most significant (21.5%), but low price (21%), volume (19%), and opening price (19%) contribute nearly as much. Notably, volume assumes a more prominent role here than in the regression model, indicating the classification task's greater sensitivity to market dynamics such as liquidity shifts, investor reactions, and short-term trading behavior.

These differences show how regression focuses on trend magnitude using price continuity, while classification emphasizes directionality using both price and behavioral signals. This insight informs more effective feature engineering and model selection when designing momentum-based trading strategies tailored to specific forecasting objectives.




Figure 7: Cumulative return curve of a top-10 stock strategy based on XGBoost predictions, with key performance metrics shown in the inset.

After generating return predictions, the XGBoost based strategy was tested using real market data. In each 5-day trading cycle, the top 10 stocks with the highest predicted returns were selected for investment, and performance was evaluated using cumulative return, annualized return, maximum drawdown, and return-to-drawdown ratio.

The strategy achieved a cumulative return of 40.43% and an annualized return of 26.02%, indicating strong short-term growth potential. However, a maximum drawdown of 84.45% reveals substantial downside risk. Despite this, the return-to-drawdown ratio of 0.31 suggests that the strategy still maintained a reasonable trade-off between return and risk. This aligns with existing findings that momentum-driven or machine learning-based strategies, while capable of delivering outsized returns, often incur high volatility and exposure to market reversals (1; 4).

These sharp drawdowns likely stem from the model's sensitivity to recent return patterns, which can overweight transient market noise or reversals, particularly during volatile regimes. To mitigate this, incorporating dynamic position sizing, where allocation adjusts based on prediction confidence or recent volatility, can help reduce exposure during uncertain periods. Likewise, applying volatility filters can limit trades in high-risk environments, improving stability without sacrificing upside. Together, such enhancements may strengthen the robustness of the strategy by aligning risk with conviction and market conditions (8; 9).

4 Conclusion

This study examined two momentum-based trading strategies: a traditional mean-variance optimized rule-based strategy and a machine learning approach utilizing XGBoost. Both were evaluated in terms of cumulative return, drawdown, and overall risk-adjusted performance.

The results reveal that a simple 5-day momentum strategy performs best within the mid-high return decile (61st-70th), while extreme past winners and losers tend to mean-revert, limiting their long-term profitability. While mean-variance re-weighting is theoretically expected to enhance the return to risk trade-off, its real-world performance proved mixed in this context. This outcome reflects the broader challenges noted in the literature—such as estimation error and shifting market dynamics—that often limit the practical benefits of optimized weighting over short horizons (7).

The XGBoost-based strategy demonstrated strong predictive capability, achieving a cumulative return of 40.43% and an annualized return of 26.02%. However, these gains were accompanied by substantial tail risk, with a maximum drawdown of 84.45%, underscoring the volatility and fragility often associated with machine learning models in financial markets. This reflects broader findings in the literature, which note that while machine learning algorithms can uncover complex, non-linear patterns, they are also prone to overfitting, unstable generalization, and heightened sensitivity to regime shifts (4).

Ultimately, while machine learning presents powerful tools for uncovering complex return signals, their practical success depends on effective risk management. This comparison underscores the value of integrating predictive models with classical financial principles. Future research should investigate hybrid approaches that merge the statistical rigor of frameworks like Markowitz mean-variance optimization with advanced algorithms, such as reinforcement learning or deep neural networks, that can adapt to complex and evolving market conditions. Such methods could directly address the framework's limitations identified in this study, though their implementation in this research was constrained due to computational resources.

Acknowledgment

I would like to express my heartfelt gratitude to my parents for their unwavering support and encouragement, which enabled me to undertake and complete this research. I am also grateful to my supervisor, Bill Theodoropoulos, for his invaluable guidance and expertise throughout this journey. Finally, I thank Melbourne High School for providing me with the opportunity and resources to pursue this project.

References

- [1] Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, 48(1):65–91, 1993. <https://doi.org/10.1111/j.1540-6261.1993.tb04702.x>.
- [2] Clifford S Asness, Tobias J Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The journal of finance*, 68(3):929–985, 2013. <https://doi.org/10.1111/jofi.12021>.
- [3] Pedro Barroso and Pedro Santa-Clara. Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120, 2015. <https://doi.org/10.1016/j.jfineco.2014.11.010>.

- [4] Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273, 2020. <https://doi.org/10.1093/rfs/hhaa009>.
- [5] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016. <https://doi.org/10.1145/2939672.2939785>.
- [6] Andrew W Lo. The statistics of sharpe ratios. *Financial analysts journal*, 58(4):36–52, 2002. <https://doi.org/10.2469/faj.v58.n4.2453>.
- [7] Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953, 2009. <https://doi.org/10.1093/rfs/hhm075>.
- [8] Kent Daniel and Tobias J Moskowitz. Momentum crashes. *Journal of Financial economics*, 122(2):221–247, 2016. <https://doi.org/10.1016/j.jfineco.2015.12.002>.
- [9] Bruce D Grundy and J Spencer Martin Martin. Understanding the nature of the risks and the source of the rewards to momentum investing. *The Review of Financial Studies*, 14(1):29–78, 2001. <https://doi.org/10.1093/rfs/14.1.29>.