UNIVERSIDADE FEDERAL DE PERNAMBUCO
CENTRO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIAS DA COMPUTAÇÃO

Matheus Dias Maciel

Evaluating Technical Analysis Indicators as Machine-Learning Features in Cryptocurrency Trading Strategies

Recife

2025

Matheus Dias Maciel

Evaluating Technical Analysis Indicators as Machine-Learning Features in Cryptocurrency Trading Strategies

Trabalho apresentado ao Programa de Pós-graduação em Ciência da Computação do Centro de Informática da Universidade Federal de Pernambuco, como requisito parcial para obtenção do grau de Mestre em Ciência da Computação.

Área de Concentração: Inteligência Computacional

Orientador (a): Daniel Carvalho da Cunha

Recife

2025

.Catalogação de Publicação na Fonte. UFPE - Biblioteca Central

Maciel, Matheus Dias.
Evaluating technical analysis indicators as machine-learning features in cryptocurrency trading strategies / Matheus Dias Maciel. - Recife, 2025.
89f.: il.

Dissertação (Mestrado)- Universidade Federal de Pernambuco, Centro de Informática, Pós-Graduação em Ciência da Computação, 2025.
Orientação: Daniel Carvalho da Cunha.

1. Machine learning; 2. Algorithmic trading; 3. Technical analysis; 4. Cryptocurrencies. I. Cunha, Daniel Carvalho da. II. Título.

UFPE-Biblioteca Central

Matheus Dias Maciel

“Evaluating Technical Analysis Indicators as Machine-Learning Features in Cryptocurrency Trading Strategies”

Dissertação de mestrado apresentada ao Programa de Pós-Graduação em Ciência da Computação da Universidade Federal de Pernambuco, como requisito parcial para a obtenção do título de Mestre em Ciência da Computação. Área de Concentração: Inteligência Computacional.

Aprovado em: 24/07/2025.

BANCA EXAMINADORA

Prof. Dr. Paulo Salgado Gomes de Mattos Neto
Centro de Informática / UFPE

Prof. Dr. Eraylson Galdino da Silva
Escola Politécnica de Pernambuco / UPE

Prof. Dr. Daniel Carvalho da Cunha
Centro de Informática / UFPE
(**orientador**)

Dedico esta dissertação à minha esposa Maria Júlia pelo apoio, compreensão e amor direcionados a mim durante o desenvolvimento do trabalho. Dedico também à minha família pela minha criação e pelo imenso suporte ao longo dos anos.

AGRADECIMENTOS

O autor gostaria de agradecer ao *Centro de Informática* da *Universidade Federal de Pernambuco* e em especial ao professor Daniel Carvalho da Cunha pela infraestrutura, flexibilidade e apoio concedidos para o desenvolvimento deste trabalho. O autor também agradece aos contribuidores de código aberto da linguagem de programação Python e dos pacotes utilizados neste estudo pela ajuda no desenvolvimento dos programas empregados nesta pesquisa. Esta pesquisa foi realizada como parte da dissertação de mestrado na *Universidade Federal de Pernambuco* sem bolsas de agências governamentais ou de empresas privadas. De forma alguma, o autor possui qualquer associação com a empresa Binance, exceto pelo uso da API pública de dados para a obtenção do dataset bruto.

"An investment in knowledge pays the best interest."(FRANKLIN, 1969).

RESUMO

Esta dissertação investiga a eficácia dos indicadores de análise técnica (AT) como variáveis engenheiradas em pipelines de machine learning financeiro para negociação de criptomoedas. Um modelo de regressão XGBoost é empregado para prever os retornos do próximo período no horizonte de candles de um minuto, contrastando estratégias construídas com variáveis de AT em relação àquelas baseadas exclusivamente em dados brutos de mercado (candles e informações agregadas de order book). Para mitigar formas perigosas de viés, adota-se um critério de informação mútua que remove variáveis suspeitas de look-ahead bias, garantindo que os sinais preditivos reflitam dependências genuínas em vez de correlações espúrias com informações futuras.

O conjunto de dados abrange criptomoedas de maior capitalização de mercado, proporcionando uma avaliação robusta em condições voláteis. Os resultados indicam que os indicadores de AT aumentam a acurácia preditiva e melhoram a separação de regimes em certas configurações quando comparados às variáveis brutas, embora os desfechos variem entre os ativos. Os backtests mostram ainda que estratégias com AT podem gerar alfas positivos antes de custos; contudo, ao se incorporarem taxas de transação e slippage realistas, esses retornos excedentes desaparecem, em consonância com a Hipótese de Mercado Eficiente em sua forma fraca, no horizonte analisado. Conclui-se que, embora os indicadores de AT possam aprimorar o aprendizado do modelo, eles não resultam em lucros persistentes após custos, oferecendo subsídios relevantes para o desenho de estratégias de engenharia de variáveis em negociação de criptoativos. Possibilidades de aprimoramento em estudos futuros também são discutidas.

Palavras-chaves: Aprendizagem de Máquina; Análise Técnica; Criptomoedas; Trading algoritmico; Hipótese do mercado eficiente;

ABSTRACT

This dissertation investigates the effectiveness of technical analysis (TA) indicators as engineered features in financial machine learning pipelines for cryptocurrency trading. An XG-Boost regression model is employed to forecast next-period returns at the one-minute candle horizon, contrasting strategies built on TA-augmented features with those relying exclusively on raw market data (candlestick and aggregated order-book information). To mitigate dangerous forms of bias, a mutual information criterion is introduced to filter out variables suspected of look-ahead bias, ensuring that predictive signals capture genuine dependencies rather than spurious foresight.

The dataset for this study comprises high-capitalization cryptocurrencies, providing robust evaluation in volatile conditions. Results indicate that TA indicators enhance predictive accuracy and improve regime separation in certain configurations when compared with raw features alone, though outcomes vary across assets. Backtesting further shows that TA-based strategies can generate positive alphas before costs; however, once realistic transaction fees and slippage are incorporated, these excess returns vanish, consistent with the weak-form Efficient Market Hypothesis in the examined time frame. The findings suggest that while TA features may improve model learning, they do not yield persistent profits after costs, offering relevant insights for feature-engineering design in crypto trading. Directions for future research are also discussed.

Keywords: Machine Learning; Technical Analysis; Cryptocurrencies; Algorithmic Trading; Efficient Market Hypothesis;

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

ARIMA Autoregressive integrated moving average

CNN Convolutional Neural Network

EMH Efficient Market Hypothesis

GARCH Generalized autoregressive conditional heteroskedasticity

GBM Gradient Boosting Machine

LSTM Long-short term memory - A type of RNN

MAE Mean Absolute Error

ML Machine Learning

MSE Mean Squared Error

PAC Probably approximately correct

R2 R-squared - The coefficient of determination

RL Reinforcement Learning

RNN Recurrent Neural Network

RSI Relative strength index - One TA indicator

SNR Signal-to-noise ratio

SVM Support Vector Machine

TA Technical Analysis

XGBoost eXtreme Gradient Boosting

LIST OF SYMBOLS

λ	Greek letter *lambda* - Used to represent exponent power transform.
σ	Greek letter *sigma* - Used to represent standard deviation.
μ	Greek letter *mu* - Used to represent mean values.
$\mathbb{E}$	Blackboard E letter - Used to represent expected value operators.
$\sum$	Greek letter capital *sigma* - Used to represent sum operators.
$\prod$	Greek letter capital *pi* - Used to represent product operators.

TABLE OF CONTENTS

1 INTRODUCTION

The rapid development of machine learning (ML) has reshaped financial research and practice, offering advanced techniques to analyze, forecast, and optimize market behavior. In particular, cryptocurrency markets have drawn increasing academic and practitioner interest because of their decentralized structure, continuous trading hours, and pronounced volatility. These features make them fertile ground for ML-based forecasting, while simultaneously introducing significant challenges in terms of noise, non-stationarity, and market microstructure effects (GU; KELLY; XIU, 2020; MALLQUI; FERNANDES, 2019). The efficiency of financial markets remains a central topic of debate. The weak form of the Efficient Market Hypothesis (EMH) asserts that asset prices already incorporate all available historical information (FAMA, 1970). If this holds, neither raw historical data nor features engineered from technical analysis (TA) should provide systematic predictive value for future returns. Despite this theoretical expectation, many studies in the literature integrate TA indicators into ML pipelines without rigorously testing whether they offer incremental value over raw market features such as candlestick or order book data (DROZDZ et al., 2020; FISCHER; KRAUSS, 2018).

This dissertation investigates whether TA features improve the predictive and economic performance of ML-based trading strategies in cryptocurrency markets. The study focuses on forecasting next-period returns at the one-minute horizon using an XGBoost regression model. Two datasets are constructed: one composed solely of raw market information (candlestick and aggregated order book data) and another augmented with TA indicators generated via the *pandas_ta* library. To mitigate bias, a mutual information criterion is employed to filter out features suspected of look-ahead behavior, ensuring that retained variables capture genuine predictive dependencies rather than spurious correlations with future outcomes (BROWN et al., 2012; VERGARA; ESTÉVEZ, 2014). The empirical analysis uses data from Binance (Binance, 2024b) covering the largest cryptocurrencies by market capitalization, including Bitcoin, Ethereum, Solana, and Binance Coin, in both spot and perpetual futures markets. Predictions are translated into systematic trading strategies that are evaluated under backtesting conditions incorporating realistic frictions such as slippage and transaction costs. In doing so, the study directly tests whether TA-based feature engineering contributes to persistent improvements in predictive accuracy and trading profitability, or whether results remain consistent with the weak-form EMH in a high-frequency setting.

1.1 HYPOTHESIS

The weak form of the EMH states that current prices fully reflect historical information, including past prices and volumes (FAMA, 1970). This implies that strategies based only on historical data should not achieve consistent predictive or economic gains. The central hypothesis of this dissertation is that the blind inclusion of TA indicators as features in ML models does not systematically improve the performance of trading strategies. While certain configurations may display temporary improvements in predictive accuracy or profitability, these effects are expected to be inconsistent across assets and to disappear once realistic market frictions are considered. This hypothesis is tested through a direct comparison of raw-feature and TA-augmented models, combined with rigorous feature selection to avoid bias, with a special target for look-ahead bias.

1.2 OBJECTIVES

The primary objective of this thesis is to assess whether TA indicators, when used as engineered features, materially enhance ML-based trading strategies in cryptocurrency markets. Three specific goals are pursued:

1. **Model development**: Construct and train XGBoost regression models to forecast next-period returns using both raw and TA-augmented datasets at the one-minute horizon.
2. **Bias mitigation**: Apply a mutual information criterion to identify and exclude features suspected of look-ahead bias, ensuring that predictive signals capture genuine relationships.
3. **Strategy evaluation**: Translate model outputs into systematic trading strategies and assess their profitability under realistic market conditions, including transaction costs and slippage.

Through these steps, the dissertation aims to establish whether TA-based feature engineering provides genuine incremental value, or whether the results remain consistent with the weak-form EMH in volatile cryptocurrency markets.

2 THEORETICAL FOUNDATION AND RESEARCH LITERATURE

2.1 GENERAL BACKGROUND

2.1.1 Financial Markets Overview

Financial markets provide the infrastructure for trading financial instruments such as equities, bonds, derivatives, and commodities, thereby enabling price discovery, liquidity provision, and risk transfer. The foundational Efficient Market Hypothesis (EMH), formulated by Fama (FAMA, 1970), posits that asset prices incorporate all available information, rendering consistent outperformance unattainable. While EMH remains influential, subsequent work in behavioral finance (e.g., (SHILLER, 2003)) has emphasized systematic departures from rational expectations, highlighting the role of investor sentiment, biases, and bounded rationality.

Derivatives represent contracts whose value depends on underlying assets, ranging from commodities and equities to interest rates. A landmark contribution was the Black–Scholes model (BLACK; SCHOLES, 1973), which introduced a closed-form solution for option pricing and remains central to modern risk management. Extensions such as stochastic volatility models (HESTON, 1993) and jump-diffusion models (MERTON, 1976) sought to reconcile theoretical assumptions with empirical regularities, including volatility clustering and heavy-tailed return distributions.

The digitization of markets has transformed execution, with electronic trading gradually replacing traditional floor-based systems. Within this environment, algorithmic trading—defined as the automation of order execution via pre-programmed strategies—has become dominant (ALDRIDGE, 2013; CARTEA; JAIMUNGAL; PENALVA, 2015). These strategies range from rules based on technical indicators to predictive models leveraging machine learning (CHAN, 2013), underscoring the increasingly data-driven nature of modern financial markets.

2.1.2 Cryptocurrencies

Cryptocurrencies constitute a novel asset class built on decentralized, cryptographic protocols. Nakamoto's white paper (NAKAMOTO, 2008) introduced Bitcoin, pioneering a peer-to-peer payment system secured by blockchain consensus. Since then, digital assets have proliferated, with market capitalization exceeding $3 trillion at its peak in 2021 (CoinMarketCap, 2021). Compared to traditional assets, cryptocurrencies are characterized by 24/7 trading, fragmented exchanges, elevated volatility, and evolving regulation (CORBET et al., 2020).

The unique market microstructure of crypto—absence of a central authority, heterogeneous participants, and frequent speculative behavior—poses both challenges and opportunities for empirical finance. Research highlights their role as speculative assets rather than hedging instruments (BORRI, 2019), while others explore spillovers between crypto and traditional markets (DYHRBERG, 2016). For algorithmic trading, the high-frequency, continuously open nature of crypto markets presents fertile ground for applying machine learning models to short-horizon prediction and anomaly detection (ALESSANDRETTI et al., 2018).

2.1.3 Technical Analysis

Technical analysis (TA) studies historical price and volume data to extract patterns believed to have predictive value. Popular tools include moving averages, Relative Strength Index (RSI), and Bollinger Bands. Despite its wide adoption among practitioners, TA has been contested within the EMH framework (FAMA, 1970), which questions the ability of past prices to predict future movements. Academic assessments remain mixed: early surveys find little statistical evidence of predictive power (MALKIEL, 2003), while later studies suggest certain indicators may capture behavioral or structural inefficiencies in specific markets (LO; MAMAYSKY; WANG, 2000; PARK; IRWIN, 2007).

Recent machine learning applications revisit TA not as stand-alone strategies but as engineered features within predictive models (KRAUSS; DO; HUCK, 2017; BALLINGS et al., 2015). In this context, TA indicators can enrich the input space by encoding domain-specific transformations of raw market data. This thesis adopts the latter perspective: TA is used to generate structured features, and its incremental value is systematically evaluated against models trained solely on raw microstructure data.

2.2 STATISTICAL INFERENCE

2.2.1 Feature Importance and Information Theory

A central task in statistical inference and machine learning is to evaluate the relationship between candidate explanatory variables and the target variable of interest. In the simplest case, when both variables are continuous and jointly normally distributed, the Pearson correlation coefficient is often used as a measure of linear dependence. However, correlation captures only linear relationships and may underestimate the true statistical dependence when the association is nonlinear.

To address this limitation, information-theoretic measures such as *Mutual Information* (MI) are employed. Given two random variables X and Y with joint probability distribution $p(x, y)$ and marginal distributions $p(x)$ and $p(y)$, the mutual information is defined as

$$I(X;Y) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right), \tag{2.1}$$

which quantifies the reduction in uncertainty about one variable given knowledge of the other. Mutual information is always non-negative and equals zero if and only if X and Y are statistically independent.

For practical applications such as feature selection, it is often desirable to normalize MI to obtain a bounded measure. One common normalization is the *Normalized Mutual Information* (NMI), defined as

$$\mathrm{NMI}(X;Y) = \frac{I(X;Y)}{\sqrt{H(X)H(Y)}}, \tag{2.2}$$

where $H(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x)$ denotes the Shannon entropy of X. This normalization yields values in the interval $[0, 1]$, where 0 indicates independence and 1 indicates perfect association. Good examples of formal properties and applications of NMI can be found on (COVER; THOMAS, 2006; VINH; EPPS; BAILEY, 2009).

In the context of machine learning, NMI provides a flexible tool to assess feature importance. Unlike correlation-based measures, which are limited to linear dependence, NMI captures arbitrary nonlinear relationships between features and the target variable. Features with high NMI values are considered informative, whereas features with low NMI values provide little predictive signal. This makes NMI a useful filter criterion in preprocessing pipelines where potentially redundant or biased features (e.g., technical indicators with implausibly strong associations) are to be excluded before model training.

2.2.2 Machine Learning

Machine learning is a field of study withing artificial intelligence that provides a set of computational techniques and statistical models that can learn patterns from data, generalize to unseen data, and thus make predictions or decisions without being explicitly programmed. Early groundwork for machine learning can be traced back to pioneers in statistics and computer science, who developed methods to enable computers to learn from examples or discover structure in unlabeled data. Within computational learning theory, it is analyzed mathematically with the framework of probably approximately correct (PAC) learning (VALIANT, 1984). It is broadly divided in three categories: supervised learning, unsupervised learning, and reinforcement learning. Our problem of predicting the next period return is a type of problem concerned to the supervised learning theory. Supervised learning involves training models on labeled datasets where the desired output (or "label") is known. Once trained, the models can generalize from their training experiences to predict or classify new, unseen data. In financial applications, supervised learning is commonly employed for classification and regression tasks. Classification models predict discrete or categorical outcomes, such as whether a cryptocurrency's price will go up or down, or which volatility regime the market currently occupies. Regression models predict continuous or numeric outcomes, such as the level of asset prices, the magnitude of price changes, or estimates of market volatility. Some popular models are explained shortly below.

- **Linear Regression:** Rooted in the least squares method from the 18th and 19th centuries, linear regression remains a fundamental approach for modeling the relationship between one or more explanatory variables and a continuous outcome.

- **Logistic Regression:** One of the earliest and most widely used classification techniques, logistic regression models the probability of a binary outcome using the logistic (sigmoid) function. Its theoretical underpinnings were developed in the mid-20th century for analyzing binary data (COX, 1958).

- **Decision Trees:** Decision trees partition the feature space into regions, making predictions by traversing a series of yes/no questions. Early influential work on decision-tree induction was introduced by Quinlan (QUINLAN, 1986).

- **Random Forests:** Introduced by Breiman (BREIMAN, 2001), random forests are an ensemble method that constructs a large number of decision trees on bootstrapped samples of the data and aggregates their predictions. By averaging across many decorrelated trees, random forests reduce variance and typically outperform single decision trees in predictive performance.

- **Gradient Boosting Machines (GBMs):** Proposed by Friedman (FRIEDMAN, 2001), gradient boosting is an ensemble method that builds a strong learner by iteratively adding weak learners (often decision trees). Each new tree is fit on the residual errors of the ensemble, leading to improved predictive accuracy. Modern implementations such as XGBoost (CHEN; GUESTRIN, 2016; XGBoost Developers, 2025) introduced regularization, efficient handling of sparse data, and parallelization, making GBMs one of the most widely adopted machine learning methods in practice.

2.3 FINANCIAL STATISTICS AND TIME SERIES ANALYSIS

Financial market data typically arrive in the form of time series, meaning that observations are recorded sequentially over time (e.g., daily stock prices, minute-by-minute cryptocurrency trades). A distinctive characteristic of such data is its non-i.i.d. (independent and identically distributed) nature, often resulting from temporal correlations, volatility clustering, and other market microstructure effects (TSAY, 2010). As a result, traditional statistical models require adaptation to handle both the dynamics and the stochastic properties of financial time series. Financial time series also often exhibit a famous *low signal-to-noise ratio*, meaning that genuine predictable patterns can be easily overwhelmed by random fluctuations. Market movements can be influenced by macroeconomic news, company-specific news, general sentiment, different market actors, and market microstructure noise (TSAY, 2010), making it challenging to discern meaningful signals. This is particularly true in high-frequency trading environments, where rapidly generated ticks may reflect spurious noise rather than actionable patterns.

2.3.1 Stationarity and volatility clustering

One of the key concepts in time series analysis is *stationarity*. Intuitively, a stationary process is one whose statistical properties like mean and variance do not change over time. In financial contexts, raw price data rarely exhibit strict stationarity due to evolving market conditions, structural breaks, and trends. Consequently, practitioners often transform prices into *relative returns* or *log-returns*, both represented below in equation form, where P_t is the asset price at time t.

$$r_t = \left(\frac{P_t}{P_{t-1}}\right) - 1 \tag{2.3}$$

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \tag{2.4}$$

This target transformation is commonly employed to achieve (or approximate) stationarity (BOX et al., 1976), facilitating the application of numerous statistical and econometric models, such as ARIMA, GARCH, or machine learning approaches that assume stationary inputs. However, this transformation does not make the data stationary still, as the market conditions such as average return and volatility are non-static. This property of many financial time series where the volatility changes is called *volatility clustering*, it is characterized for the fact that periods of high volatility tend to be followed by further high volatility, while periods of relative calm tend to persist, illustrating a form of temporal dependence in second moments (ENGLE, 1982; BOLLERSLEV, 1986). Given this, a model trained on past returns data may generalize out-of-sample only for so long. Therefore, the models must eventually be retrained once its out-of-sample performance starts to degrade. While this may happen in other domains, financial market conditions will change much faster than object shapes used in object recognition models.

2.3.2 Skew and Fat Tails

Additionally, financial returns often exhibit heavier tails (i.e. positive kurtosis) compared to a Gaussian distribution, that means that the financial returns have higher order moments that are not negligible. This implies a higher probability of extreme price moves such as market shocks, which are important factors to consider when building robust machine learning or statistical forecasting models for algorithmic trading. The figure 1 showcases what are the effects of skew and kurtosis in probability distributions. Usually returns data is slightly positive skewed and with positive kurtosis.

With this in mind, researchers often make use of power transformations to create a dataset that is closer to normality, which is often a requirement for certain types of machine learning algorithms to generalize out-of-sample. The *Yeo–Johnson* transform (YEO; JOHNSON, 2000) is often the power transform of choice for returns data because it allows non-positive values for

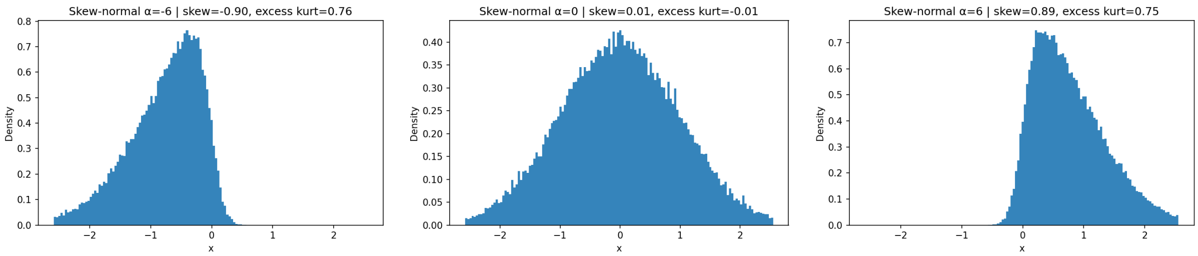

(a) Distributions with varying skewness (skew-normal with $\alpha = -6, 0, 6$).

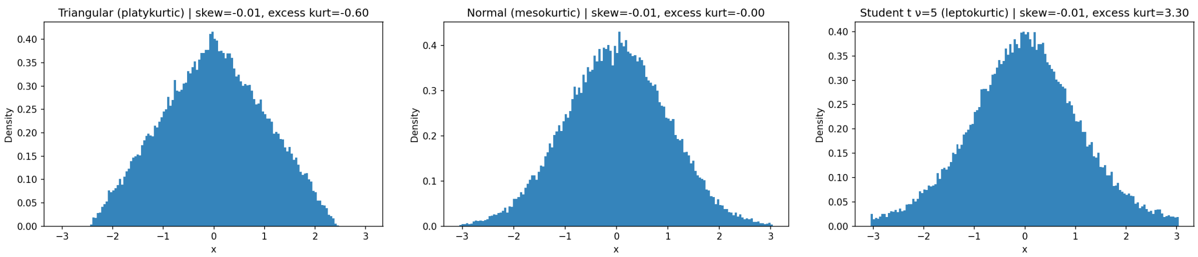

(b) Distributions with varying kurtosis (triangular, normal, Student-t).

Figure 1 – Illustration of how skewness (top) and kurtosis (bottom) affect the shape of probability distributions.

the transformation input. This transform is defined by the following formula:

$$y_i^{(\lambda)} = \begin{cases} \dfrac{(y_i+1)^{\lambda}-1}{\lambda}, & \lambda \neq 0,\ y_i \geq 0, \\[2ex] \ln(y_i+1), & \lambda = 0,\ y_i \geq 0, \\[2ex] -\dfrac{(-y_i+1)^{2-\lambda}-1}{2-\lambda}, & \lambda \neq 2,\ y_i < 0, \\[2ex] -\ln(-y_i+1), & \lambda = 2,\ y_i < 0. \end{cases} \tag{2.5}$$

2.3.3 Statistical Tests and Robustness Checks

Because financial returns are noisy, non-i.i.d. and prone to data-snooping bias, evaluating a predictive model requires more than a simple in-sample accuracy metric, or even a simple out-of-sample metric. The most widely-used techniques fall into three complementary groups:

- **Predictive-accuracy tests:**
  - Diebold–Mariano (DM) test (DIEBOLD; MARIANO, 2002) for pairwise forecast-error comparison.
  - Clark–West test (CLARK; WEST, 2007) for nested models (e.g. raw vs. TA-enriched).
  - Mean-Squared-Prediction-Error (MSPE) bootstraps for small samples.

- **Multiple-model reality checks:**
  - White's Reality Check (WHITE, 2000). Designed to control data-snooping when many models/strategies are compared.
  - Hansen's Superior Predictive Ability (SPA) test (HANSEN, 2005).
  - Model Confidence Set (MCS) by Hansen, Lunde, and Nason (HANSEN; LUNDE; NASON, 2011).
- **Economic-significance tests:**
  - Jobson–Korkie (JOBSON; KORKIE, 1981) and Ledoit–Wolf (LEDOIT; WOLF, 2008) corrections for Sharpe-ratio differences.
  - Deflated Sharpe Ratio (DSR) of Bailey & López de Prado (BAILEY; PRADO, 2014) to adjust for back-test overfitting and selection bias.

Robustness is further enhanced by *rolling-origin* or *purged/embargoed* cross-validation schemes (PRADO, 2018). Block and stationary bootstraps preserve time dependence when resampling, while Monte-Carlo permutation tests (randomly shuffling or phase-scrambling returns) help gauge how likely a reported performance could arise by chance.

2.3.4 Market Impact and Liquidity Constraints

Even if a trading signal were perfectly predictive, turning that signal into profit requires executing orders in a market that is only *imperfectly* liquid. Submitting sizable buy (sell) orders pushes the mid-price upward (downward), a phenomenon known as *market impact*. Kyle's seminal model (KYLE, 1985) formalizes this feedback loop: an informed trader optimally balances trade size against the price concession demanded by liquidity providers. Empirical work shows that impact has both a *temporary* component—prices partially revert after execution—and a *permanent* component that reflects information revelation (ALMGREN; CHRISS, 2001; GATHERAL, 2010). For high-turnover strategies, ignoring impact leads to systematically overstated back-test Sharpe ratios. In practice, traders often cap their *participation rate* (e.g. $<10\%$ of market volume) or use volume-weighted execution tactics to mitigate costs. Because major cryptocurrencies exhibit thinner order books and intermittent liquidity compared with large-cap equities, impact can be especially severe in this domain, further weakening the real-world viability of strategies that look attractive in a friction-free simulation.

2.4 RECENT LITERATURE

A comprehensive review of recent research on stock-market forecasting, as presented in (VERMA; SAHU; SAHU, 2023), provides an essential background for this study. In that review, 50 research articles published between 2018 and 2021 were analyzed to understand trends and shortcomings in the field. One of the key findings was that while historical price data is widely used (approximately 78% of reviewed studies), only about 34% incorporate technical-analysis (TA) indicators as additional features. Moreover, merely 30% of the articles employed statistical tests to rigorously validate their models.

These findings suggest that although TA features remain popular among researchers, their true efficacy is often inadequately assessed. Many studies integrate TA indicators without robust statistical validation, raising concerns about the reliability of the developed models. Additionally, methods based on machine- and deep-learning, such as Long Short-Term Memory networks (LSTMs) (GHOSH; NEUFELD; SAHOO, 2022; LIU, 2019) and Convolutional Neural Networks (CNNs) (CHANDAR, 2022), are increasingly prevalent, yet the advantages of incorporating TA indicators remain unclear. For instance, Chandar (2022) employed CNNs alongside TA indicators to predict directional stock movement, demonstrating some predictive capability but without conclusive statistical confirmation. Gang et al. (2022) (JI et al., 2022) further explored feature-selection techniques, applying adaptive methodologies to identify optimal TA features for Random-Forest models. Their results indicated potential advantages in adaptive feature selection, though statistical validation of overall model improvements was limited.

Several empirical studies focused on cryptocurrencies echo the pattern of limited—or nonexistent—incremental value from TA indicators:

- **Sentiment vs. TA.** Ortu et al. (2022) fuse TA indicators with social-media sentiment in deep-learning classifiers; Wilcoxon signed-rank tests show that almost all incremental predictive power comes from the sentiment variables, not the TA indicators (ORTU et al., 2022).
- **Marginal gains from TA.** Goutte et al. (2023) combine deep-learning models with a standard TA feature set and find only marginal, statistically insignificant improvements once bootstrapped t-tests are applied (GOUTTE et al., 2023).
- **Transformer–based hybrids.** Labbaf Khaniki and Manthouri (2024) augment raw

OHLCV data with 15 classical TA indicators in a Performer-BiLSTM hybrid; Diebold–Mariano tests show no significant MAE reduction from the extra indicators (KHANIKI; MANTHOURI, 2024).

- **Data-type sensitivity analyses.** Tanrikulu and Pabuccu (2025) compare seven ML algorithms on datasets with (i) raw OHLCV, (ii) 12 canonical TA indicators, and (iii) macro–sentiment variables. Paired bootstrapped *t*-tests reveal that TA features help only in very short Bitcoin histories, with no benefit over seven-year windows (TANRIKULU; PABUCCU, 2025).

- **Direct ML–vs.–TA strategy benchmarks.** Islas Anguiano and García-Medina (2025) benchmark an LSTM strategy against two rule-based TA strategies on Bitcoin. While the LSTM beats both baselines, adding the same TA indicators does *not* improve the out-of-sample Sharpe ratio at the 5% level (ANGUIANO; GARCÍA-MEDINA, 2025).

- **Feature-selection focus.** El Youssefi et al. (2025) test three feature-selection schemes on 130 TA indicators across BTC, ETH, and BNB. Momentum and volatility measures matter most, but pruning 80–85% of indicators usually *maintains or improves* accuracy—implying many TA variables add noise, not signal (YOUSSEFI et al., 2025).

- **Cross-asset robustness checks.** Lee (2025) trains a Temporal Fusion Transformer on fifteen crypto assets, integrating on-chain metrics with seven TA indicators. Only 4/15 assets see validation-loss improvements, and those disappear after Bonferroni-corrected Wilcoxon tests (LEE, 2025).

Collectively, the evidence from the studies above corroborates earlier suspicions: once rigorous out-of-sample testing and appropriate multiple-comparison corrections are applied, TA indicators rarely supply universal, reproducible gains. Instead, hybrid feature sets that include order-book microstructure variables or exogenous sentiment signals appear more promising. The literature now spans classical statistical learning, deep learning, and transformer-based approaches, yet consensus on TA's incremental value remains elusive.

The recurring methodological gap—insufficient statistical testing and the absence of realistic cost assumptions—motivates the experimental design of this thesis, which emphasizes:

1. Constructing parallel raw-data and TA-enriched datasets;

2. Implementing rigorous cross-validation and Bayesian hyper-parameter optimization;

3. Applying robust tests (paired Wilcoxon, McNemar, Diebold–Mariano) to compare models;

4. Implementing strategies with common industry techniques such as using volatility based position sizing or stop-losses.

5. Incorporating realistic frictions (slippage, commissions, funding fees) when evaluating derived trading strategies.

6. If necessary evaluate the strategies with robustness checks.

By drawing on the mixed yet increasingly rigorous body of work summarized above, the study aims to deliver a decisive assessment of whether TA features offer statistically and economically significant benefits in algorithmic cryptocurrency trading. Notably, even among the few recent papers that claim sizable gains, most stop short of applying the full suite of robustness tests outlined in section 2.3.3, underscoring the need a more stringent evaluation protocol.

3 METHODOLOGY

The overarching goal of this study is to evaluate whether the sole usage of technical analysis (TA) indicators provides meaningful performance improvements in predictive performance for algorithmic trading strategies based on machine learning models that predict next period returns. This work adopts an experimental design that constructs two datasets: (1) a baseline *raw* dataset, and (2) a *TA-enriched* dataset derived from one-minute frequency data obtained via the Binance exchange public historical data (Binance, 2024b).

The choice of 1-minute candle data in this study was motivated by a trade-off between market accessibility and automation feasibility. On one hand, the 1-minute timeframe is too fast for manual trading by retail investors, inherently requiring algorithmic execution, which aligns with the study's focus on automated strategies. On the other hand, this timeframe remains within the reach of retail traders' resources, as it does not demand low-level programming languages or specialized low-latency infrastructure typical of high-frequency trading. This makes the timeframe practical for individual traders looking to deploy systematic strategies using standard retail brokerage APIs and widely available technologies. The steps followed in the investigation and experiments are described in the upcoming sections of this chapter.

All data was collected from January 2020 to mid-2024, capturing a period characterized by substantial market volatility in cryptocurrency markets. Data for the top four non-stable cryptocurrency crosses by market capitalization as of January 2024 (**BTC**, **ETH**, **SOL**, and **BNB**, each against the staple stable-coin **USDT**) were collected for both spot and perpetual markets via the Binance public API (Binance, 2024a). The selection was based on the CoinMarketCap snapshot of January 1, 2024, where these four assets ranked top 4 in market cap excluding stable-coins such as USDT and USDC from consideration (CoinMarketCap, 2024). While Ripple's XRP has at times surpassed BNB and SOL in capitalization, its lower relative ranking at the beginning of 2024, coupled with ongoing regulatory uncertainty surrounding the asset, justified its exclusion. Furthermore, SOL and BNB provide diversification by representing distinct blockchain ecosystems (high-throughput Layer-1 smart contracts in Solana and the exchange-native utility token of Binance Chain), thereby complementing BTC and ETH and offering broader coverage of heterogeneous market dynamics beyond simple capitalization.

3.1 EXPLORATORY DATA ANALYSIS

The close prices for the collected cryptocurrencies which are part of the candle data from the raw dataset and used to generate most of the TA indicators are displayed on the figure 2 for the BTC spot. Between the currencies the plots share similarities to the crypto market being very affected by trends or market regimes, however, all of the currencies have different paths due to their own particularities. The histogram of the returns graph in the figure 3 for the BTC spot shows that even though the distribution resembles a bell shape it has skew and fat tails. The volatility of returns of all the currencies, presented in the figures 4 for the BTC spot, showcases how the volatility changes in the dataset as expected to financial data. For the graphs representing the other currencies one can look at the appendix B.

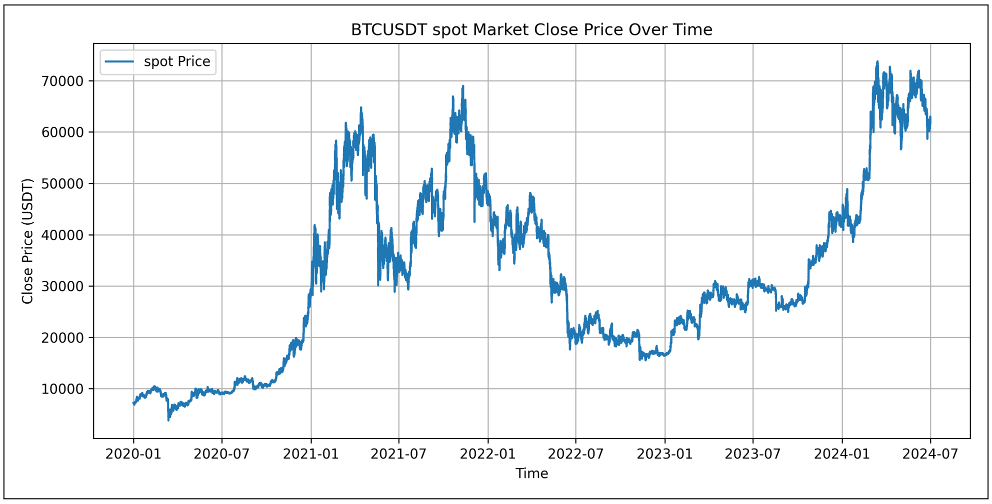

Figure 2 – BTCUSDT Spot

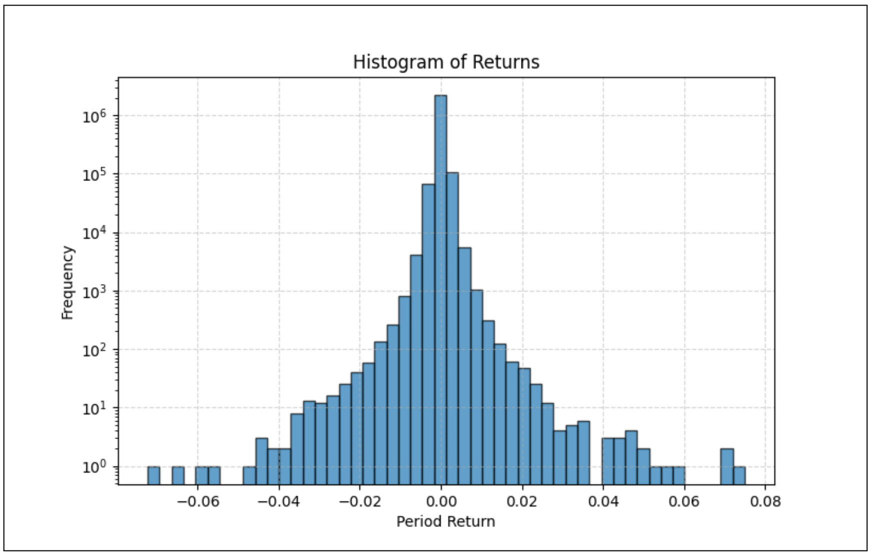

Figure 3 – BTCUSDT returns histogram

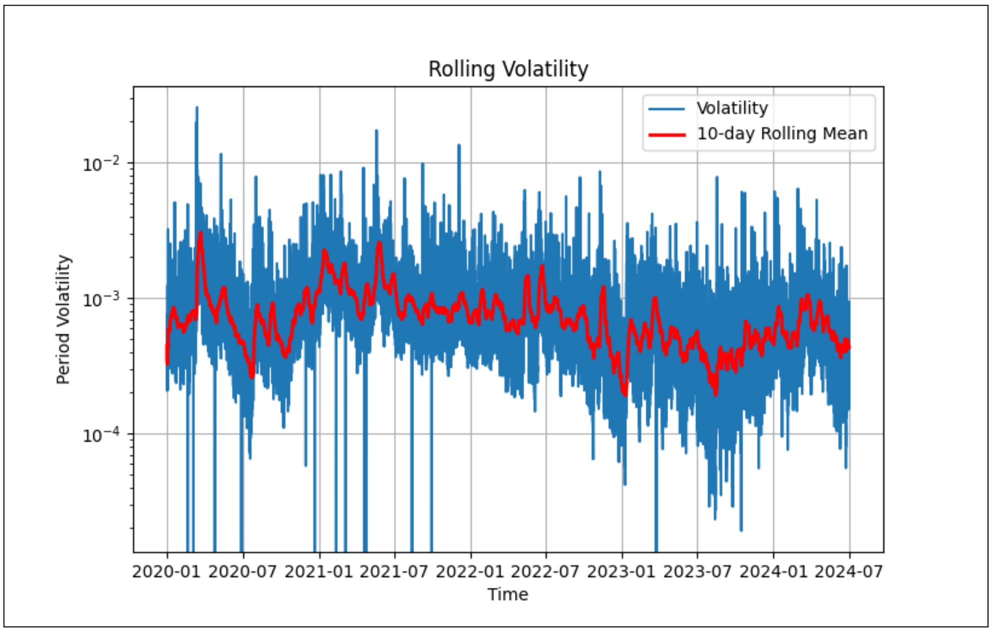

Figure 4 – BTCUSDT average volatility

3.2 FEATURE ENGINEERING AND SELECTION

The primary data source is Binance historical market data. We retrieve OHLCV candles (Table 1) and aggregated trade or order-book statistics at the 1-minute frequency, which are used to construct additional microstructure features. The *raw* feature set therefore includes the candle fields together with simple, well-defined statistics (for example `price_acceleration_mean` and `vwap_deviation`) derived from the aggregates, as summarized in Table 2.

Technical-analysis (TA) features are built from the candle series using TA-Lib (FORTIER; CONTRIBUTORS, 2024) and related rolling statistics. We audited implementations to avoid look-ahead effects and aligned all rolling windows with explicit one-step lags. Categories include momentum, volatility, trend, volume, and selected candlestick patterns; an overview is given in Table 3 and a bit more detailed description can be found over the Appendix A. It is worth highlight that not all indicators from the library are shown in the table, since many were filtered by criteria to be described further in this section.

To limit biases, we focus on four large and continuously listed instruments, which reduces survivorship concerns. The main threat in this setting is look-ahead bias. We therefore applied a quantitative screen based on **normalized mutual information** (NMI) computed between each candidate feature and the target, using only the training folds in time order. *NMI was the exclusion criterion.* Any indicator whose NMI exceeded the highest NMI observed among the raw OHLCV plus microstructure features by a wide margin was flagged as suspicious and removed. The rationale is that derived indicators should not exhibit substantially greater predictive information than the underlying price and volume dynamics unless they inadvertently leak future information.

We did not pursue further feature selection beyond this NMI-based leakage screen. The predictive model employs regularization (for example L1 and L2 penalties and column subsampling in gradient-boosted trees), which mitigates redundancy at training time. More importantly, the goal of this study is to assess whether statistically reliable *edge* exists when adding TA features, not to maximize that edge through aggressive feature pruning. A deeper feature-selection study is left for future work.

In summary, for each cryptocurrency we construct two datasets:

1. **Raw dataset:** OHLCV candles and aggregated-trade or order-book statistics, plus simple derived microstructure features (Tables 1 and 2).

2. **TA-enriched dataset:** The raw dataset (Tables 1 and 2) extended with vetted TA indicators from TA-Lib (Table 3), after excluding indicators by the NMI criterion described above.

Table 1 – Raw Dataset Features and Their Descriptions

Data field	Explanation
open	Opening price of the asset in the given time interval.
high	Highest price reached during the time interval.
low	Lowest price reached during the time interval.
close	Closing price of the asset at the end of the time interval.
volume	Total trading volume of the asset.
quote_volume	Total quote volume traded in the quote currency.
count	Number of trades executed during the time interval.
taker_buy_volume	Volume of trades executed by buyers acting as takers.
taker_buy_quote_volume	Quote volume corresponding to taker buy orders.

Table 2 – Raw Dataset order-book aggregate features and their descriptions

Feature	Explanation
price_mean	Average price during the interval.
price_std	Standard deviation of prices during the interval.
quantity_sum	Total quantity traded over the interval.
quantity_mean	Average quantity traded per transaction.
quantity_max	Maximum quantity traded in a single transaction.
price_change_mean	Mean of the price changes within the interval.
price_change_std	Standard deviation of the price changes.
vwap_last	Last value of the Volume Weighted Average Price (VWAP).
trade_count_sum	Sum of trade counts over the interval.
buyer_maker_ratio_mean	Average ratio of buyer to maker trades.
cum_volume_last	The last cumulative volume value recorded.
volume_weighted_price_mean	Average price weighted by volume over the interval.
trade_size_mean	Average size of trades.
trade_size_max	Maximum trade size recorded.
price_direction_mean	Mean of the price directional changes.
price_direction_last	Direction of the price change in the last trade of the interval.
price_acceleration_mean	Average rate of change in price differences.
price_acceleration_std	Standard deviation of price acceleration values.
volume_imbalance_sum	Sum of volume imbalances, indicating the disparity between buying and selling pressure.
is_buyer_maker_first	Indicator (boolean) if the first trade was executed by a buyer-maker.
is_buyer_maker_last	Indicator (boolean) if the last trade was executed by a buyer-maker.
vwap_deviation	Deviation of the current price from the VWAP.
volume_intensity	Measure of trading intensity based on volume and frequency.
trade_size_variation	Variation in the sizes of trades over the interval.

Table 3 – TA indicators by category

Category	TA Indicators
candles	cdl_2crows, cdl_3blackcrows, cdl_3inside, cdl_3linestrike, cdl_3outside, cdl_3starsinsouth, cdl_3whitesoldiers, cdl_abandonedbaby, cdl_advanceblock, cdl_belthold, cdl_breakaway, cdl_closingmarubozu, cdl_concealbabyswall, cdl_counterattack, cdl_darkcloudcover, cdl_doji_10_0.1, cdl_dojistar, cdl_dragonflydoji, cdl_engulfing, cdl_eveningdojistar, cdl_eveningstar, cdl_gapsidesidewhite, cdl_gravestonedoji, cdl_hammer, cdl_hangingman, cdl_harami, cdl_haramicross, cdl_highwave, cdl_hikkake, cdl_hikkakemod, cdl_homingpigeon, cdl_identical3crows, cdl_inneck, cdl_inside, cdl_invertedhammer, cdl_kicking, cdl_kickingbylength, cdl_ladderbottom, cdl_longleggeddoji, cdl_longline, cdl_marubozu, cdl_matchinglow, cdl_mathold, cdl_morningdojistar, cdl_morningstar, cdl_onneck, cdl_piercing, cdl_rickshawman, cdl_risefall3methods, cdl_separatinglines, cdl_shootingstar, cdl_shortline, cdl_spinningtop, cdl_stalledpattern, cdl_sticksandwich, cdl_takuri, cdl_tasukigap, cdl_thrusting, cdl_tristar, cdl_unique3river, cdl_upsidegap2crows, cdl_xsidegap3methods
momentum	ao, apo, bias, bop, brar, cci, cfo, cg, cmo, coppock, cti, er, eri, fisher, inertia, kdj, kst, macd, mom, pgo, ppo, psl, pvo, qqe, roc, rsi, rsx, rvgi, slope, smi, stc, stoch, stochrsi, trix, tsi, uo, willr
overlap	alma, dema, ema, fwma, hilo, hma, jma, kama, linreg, midpoint, midprice, pwma, rma, sinwma, sma, ssf, supertrend, swma, t3, tema, trima, vwap, vwma, wcp, wma, zlma
performance	log_return, percent_return
statistics	entropy, kurtosis, mad, median, quantile, skew, stdev, tos_stdevall, variance, zscore
trend	adx, amat, aroon, chop, cksp, decay, decreasing, increasing, psar, qstick, ttm_trend, vhf, vortex
volatility	aberration, accbands, atr, bbands, donchian, hwc, kc, massi, natr, pdist, rvi, thermo, true_range, ui
volume	ad, adosc, cmf, efi, kvo, mfi

3.3 DATASET SPLITTING

We use a three-stage walk-forward design that separates model estimation, strategy fitting, and final testing. For each ticker, the series is partitioned chronologically into *Train 1* (70%), *Validation 1* (10%), *Validation 2* (10%), and *Test* (10%), with no shuffling. All preprocessing is fitted only on the training side of each split and applied forward to avoid leakage; targets are never transformed.

- **Stage 1 - model training**: On *Train 1* we run five-fold time-series cross-validation with expanding windows to select XGBoost hyperparameters by minimizing the median mean-squared error across folds. The selected model is then refit on the full *Train 1* and evaluated on *Validation 1* to record out-of-sample regression metrics.
- **Stage 2 - strategy optimization**: Using the tuned model, we refit on *Train 1* ∪ *Validation 1* and generate predictions on *Validation 2*. Trading rules are parameterized by a small set of interpretable variables. We select parameters under several transaction-cost tiers by maximizing a robustness-oriented objective (median Sharpe across contiguous validation blocks) subject to a minimum trade-count constraint. Costs follow Binance schedules (Binance, 2024b) and simple slippage assumptions (small constant slippage).
- **Stage 3 - final evaluation**: We refit the predictive model on *Train 1* ∪ *Validation 1* ∪ *Validation 2*, freeze the strategy parameters chosen in Stage 2, and evaluate on the unseen *Test* segment. Reported headline results refer to this hold-out test and include return and risk measures (cumulative and annualized return, Sharpe, maximum drawdown, Calmar), trade statistics, and order-level summaries.

Split boundaries are never crossed during scaler fitting or model training. Rolling statistics used inside the backtest are computed with one-step lags and seeded with a short trailing window from the training side to provide context without peeking. The optimization loop in Stage 2 does not access the *Test* data at any point. The full procedure is applied uniformly across four liquid tickers (BTCUSDT, ETHUSDT, SOLUSDT, BNBUSDT), two feature sets (raw and technical-analysis), and three trading modalities (spot, perpetual futures, and a spot–futures spread). Results are stored in a structured table for downstream analysis and reproducibility.

A diagram that was crafted to better illustrate the whole splitting process can be found in the Figure 5, it makes use of blocks and colors (green for fit, blue for evaluation, grey for no usage) to showcase what each subset of the original data is used for what purpose in each step.

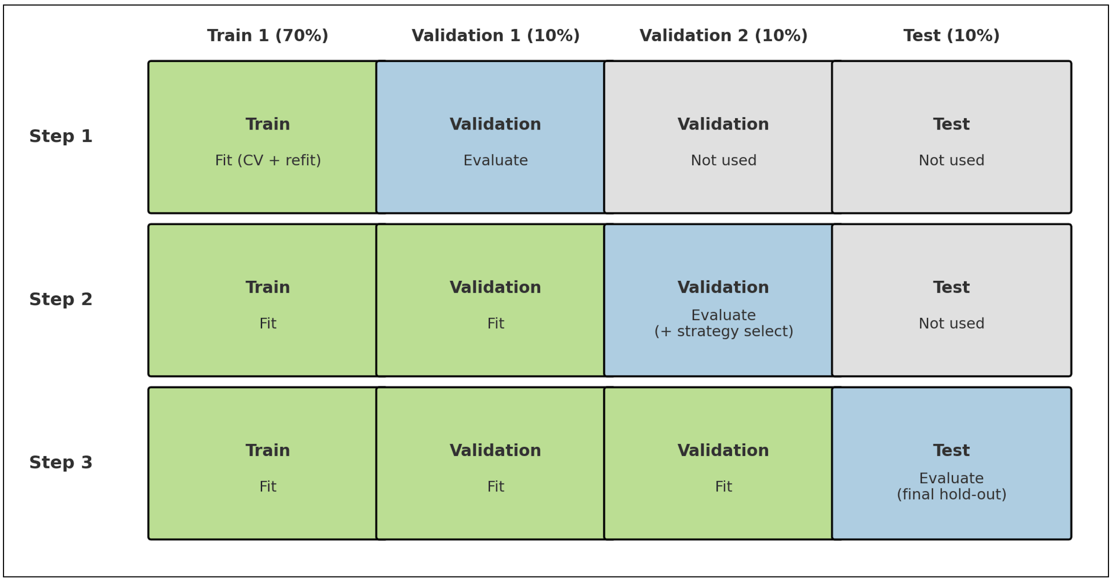

Figure 5 – Usage of each dataset slice by step

3.4 DATA PREPROCESSING

Supervised learning algorithms (especially those reliant on gradient-based methods or Euclidean distances) often benefit from normally distributed or near-normally distributed input features. In particular, attributes such as volume, volatility, and certain TA indicators can exhibit skewed or heavy-tailed distributions. Given that, it was opted to use power transforms such as the *Yeo–Johnson* transform (YEO; JOHNSON, 2000) since it can handle zero and negative values, making it suitable for returns-based data. his transform is very effective in reducing skew (makes the distribution symmetric) and can reduce kurtosis by a bit (makes the distribution less *fat-tailed*). It is also worth noting that when using a parametric transform such as the power transform the parameters are inferred on the training set and carried to the test set, this way no look-ahead bias is introduced in the test set transformations. Also to avoid algorithms being dominated by very large numbers, clipping and substitution of invalid values was done.

3.5 MODEL DEVELOPMENT AND SELECTION

This study asks a narrow question: do TA indicators provide any incremental edge when added to a compact raw feature set under a disciplined walk-forward design. The objective is not to squeeze the last basis point by heavy tuning, but to test whether an edge exists at all. To avoid overfitting by search, we limit model families, keep optimization strictly inside the training and validation partitions defined in Section 3.3, and favor methods that are fast to train and easy to interpret.

- **Baseline and model family choice:** A plain linear regression on both raw and TA-enriched datasets produced out-of-sample performance close to a coin toss, which indicated the relationship is non-linear. We therefore considered tree models that naturally capture interactions and non-linearities and can be explained with TreeSHAP (LUNDBERG; LEE, 2017). Two families were compared: Random Forests and gradient-boosted trees (XGBoost (CHEN; GUESTRIN, 2016)). Boosted trees achieved higher out-of-sample R^2 with lower training time, so XGBoost was chosen for this proof-of-concept check.
- **Training protocol:** Aligned with the three stage split described in the section 3.3.
- **Regularization and hyperparameter search.** To contain variance and discourage complexity, the optimization grid included regularization and capacity controls: L_2 penalty (`lambda`), L_1 penalty (`alpha`), `max_depth`, `min_child_weight`, learning rate, `subsample`, `colsample_bytree`, and `gamma`. These were tuned by an automated search procedure within Stage 1 only. No parameters were hand-picked from Test.
- **Feature selection stance.** Beyond the normalized mutual information leakage screen described in Section 3.2, no additional feature selection was applied. The intent was to test if TA features yield a detectable edge, not to maximize that edge. Redundancy is addressed implicitly through model regularization. A more aggressive selection program is left for future work.
- **Interpretability and diagnostics.** TreeSHAP explanations for the chosen tree model provide local and global attributions without retraining, enabling sanity checks that learned signals are consistent with market intuition (LUNDBERG; LEE, 2017).

3.6 STRATEGY DEVELOPMENT AND EVALUATION

Once the predictive model was developed and evaluated, it was integrated into a simulated trading framework designed to reflect realistic execution conditions (KISSELL, 2013; ALMGREN; CHRISS, 2001; PRADO, 2018). The workflow for strategy evaluation included:

- Translating model outputs into trading signals via systematic, rule-based logic (thresholds, direction, and sizing rules).
- Incorporating realistic trading assumptions for the chosen timeframe. Fees follow Binance schedules, and a simple market-friction baseline is applied so that evaluation reflects deployable conditions (Binance, 2024b). A fuller microstructure model (order-execution lag, slippage curves, liquidity constraints) is left to future work (KISSELL, 2013; BOUCHAUD et al., 2018).
- Implementing position sizing and risk overlays (stop-loss, take-profit, leverage caps) to control downside risk and emulate practical operation.
- Running walk-forward backtests across the four cryptocurrency pairs and two asset types with a true out-of-sample test segment and multiple cost tiers (Section 3.3).
- Comparing strategies trained on raw data versus TA-enriched data to assess the practical value-add of technical indicators.

Common evaluation metrics include Sharpe and Calmar ratios, which account for both return and risk and are more informative than raw cumulative returns (SHARPE, 1998; YOUNG, 1991). We report:

- **Cumulative return** and **net profitability** after fees and slippage:

$$\text{Cumulative Return} = \prod_{t=1}^{T}(1 + r_t) - 1. \tag{3.1}$$

- **Sharpe ratio** (risk-adjusted return) (SHARPE, 1998), **maximum drawdown** (worst peak-to-valley loss) (MAGDON-ISMAIL; ATIYA, 2004), and **Calmar ratio** (drawdown-adjusted return) (YOUNG, 1991):

$$\text{Sharpe Ratio} = \frac{\mathbb{E}[\, r_t - r_f \,]}{\sigma_{r_t}}, \quad r_f \approx 0 \text{ at intraday horizons.} \tag{3.2}$$

$$\text{Maximum Drawdown} = \max_{t\in[1,T]} \left(\frac{\max_{\tau\in[1,t]} P_\tau - P_t}{\max_{\tau\in[1,t]} P_\tau} \right). \tag{3.3}$$

$$\text{Calmar Ratio} = \frac{\prod_{t=1}^{T}(1 + r_t) - 1}{\text{Maximum Drawdown}}. \tag{3.4}$$

- **Order hit rate** (fraction of profitable trades) and other portfolio risk diagnostics:

$$\text{Hit Rate} = \frac{\#\{\text{ profitable trades }\}}{\#\{\text{ total trades }\}}. \tag{3.5}$$

- **Turnover** (GRINOLD; KAHN, 2000) and cost sensitivity for scalability and robustness:

$$\text{Turnover} = \frac{1}{T} \sum_{t=1}^{T} \frac{|Q_t - Q_{t-1}|}{AUM}. \tag{3.6}$$

All metrics are reported for in-sample (training/validation) and out-of-sample (test) segments per Section 3.3, so results reflect genuine predictive and trading performance without parameter leakage (PRADO, 2018).

4 MODEL DEVELOPMENT

4.1 SCOPE AND EXPERIMENTAL GRID

Leveraging the experimental pipeline in Chapter 3, We benchmark a gradient-boosted tree model with regularization on raw and TA-enriched one-minute datasets to assess whether TA indicators add predictive value beyond candle and order-book microstructure features. Models are implemented in *XGBoost* (v2.0.3), with hyperparameters tuned via Optuna's Tree-structured Parzen Estimator (TPE). Only one supervised task is considered: regression on the next-minute relative return r_{t+1}, with directional-accuracy classification derived by taking the sign of the predicted value. Metrics are Mean Squared Error (MSE), R^2, and accuracy. Paired comparison between raw and TA datasets uses McNemar's test on classification accuracy and the Wilcoxon signed-rank test on absolute regression errors.

$$\mathcal{L}_{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{4.1}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \tag{4.2}$$

$$\text{Accuracy} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}!\left(\mathrm{sign}(\hat{y}_i) = \mathrm{sign}(y_i)\right) \tag{4.3}$$

The hyper-parameters grid for the XGBoost model is described below and its cardinality is calculated in Equation 4.4:

- **objective:** `reg:squarederror` (regression task).
- **booster:** `gbtree` (tree-based boosting).
- **lambda:** L2 regularization term on weights, searched as 10^{-i} where $i \in [1, 8]$.
- **alpha:** L1 regularization term on weights, searched as 10^{-i} where $i \in [1, 8]$.
- **subsample:** Subsampling ratio of the training instance, $\{0.4, 0.6, 0.8, 1.0\}$.
- **colsample_bytree:** Subsampling ratio of columns when constructing each tree, $\{0.4, 0.6, 0.8, 1.0\}$.
- **max_depth:** Maximum tree depth, integer in $[3, 10]$.
- **n_estimators:** Number of boosting rounds, integer in $[100, 1000]$ (step of 100).

- **learning_rate:** Step size shrinkage, $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$.
- **min_child_weight:** Minimum sum of instance weight needed in a child, integer in $[1, 7]$.
- **gamma:** Minimum loss reduction required for further partitioning, $\{0, 10^{-9}, 10^{-6}, 10^{-3}, 1.0\}$.

$$
\begin{aligned}
|\mathcal{H}| &= (8_\lambda) \cdot (8_\alpha) \cdot (4_{\text{subsample}}) \cdot (4_{\text{colsample\_bytree}}) \\
&\quad \cdot (8_{\text{max\_depth}}) \cdot (10_{\text{n\_estimators}}) \cdot (4_{\text{learning\_rate}}) \\
&\quad \cdot (7_{\text{min\_child\_weight}}) \cdot (5_\gamma) \\
&= 11{,}468{,}800.
\end{aligned}
\tag{4.4}
$$

4.2 EMPIRICAL RESULTS

All models are trained on the 70% segment and evaluated on the first 10% evaluation slice, with the final two evaluation and test slices reserved for strategy evaluation. Tables 4 and 5 reports regression metrics (MSE, R^2) and classification accuracy from the sign of predictions for both datasets respectively. Table 6 shows the paired tests between raw and TA datasets, using McNemar's test for classification accuracy and Wilcoxon signed-rank test for regression errors. The table shows that the Wilcoxon test results indicate statiscally significant differences for all currencies and McNemar test results indicate significant differneces for SOL and BNB.

Table 4 – Training and test results by cryptocurrency

Raw dataset (Candles and order book)			
Ticker	**Train MSE**	**Train R^2-score**	**Train Accuracy**
BTC	**1.215e-06**	-1.19e-08	0.501
ETH	1.906e-06	1.23e-02	0.513
SOL	4.869e-06	1.90e-02	**0.551**
BNB	2.063e-06	**2.56e-02**	0.547
Ticker	**Test MSE**	**Test R^2-score**	**Test Accuracy**
BTC	3.617e-07	-1.24e-07	0.523
ETH	3.923e-07	-3.43e-05	0.523
SOL	**1.348e-06**	1.55e-02	0.581
BNB	4.209e-07	**4.26e-02**	**0.603**

Table 5 – Training and test results by cryptocurrency

TA indicators dataset			
Ticker	**Train MSE**	**Train R^2-score**	**Train Accuracy**
BTC	**1.212e-06**	2.94e-03	0.501
ETH	1.917e-06	6.95e-03	0.506
SOL	4.896e-06	**1.36e-02**	0.544
BNB	2.088e-06	1.38e-02	**0.545**
Ticker	**Test MSE**	**Test R^2-score**	**Test Accuracy**
BTC	**3.617e-07**	4.97e-05	0.523
ETH	3.924e-07	-2.56e-04	0.510
SOL	1.357e-06	8.40e-03	0.585
BNB	4.266e-07	**2.94e-02**	**0.604**

Table 6 – Statistical tests (raw vs. TA) (McNemar's for accuracy, Wilcoxon for errors)

Ticker	**Train**		**Test**	
McNemar	**Test stat**	**p-value**	**Test stat**	**p-value**
BTC	2.117	0.0214	1.000	0.500
ETH	-13.248	1.000	-8.764	1.000
SOL	-40.155	1.000	8.127	2.29e-16
BNB	-17.310	1.000	4.548	3.13e-06
Wilcoxon	**Test stat**	**p-value**	**Test stat**	**p-value**
BTC	6.827e11	1.09e-06	1.378e10	7.48e-11
ETH	6.826e11	5.12e-07	1.390e10	3.45e-03
SOL	4.833e11	0.0000	9.224e09	0.0000
BNB	6.433e11	0.0000	1.199e10	0.0000

While the results presented showcase that TA overlays do not improve predictive accuracy metrics (both MSE and R2-score) across the set of currencies, it may not show the full picture. The statistical tests produces statistically significant paired-test differences between raw and TA datasets for the regression errors in Table 6 but the sign of this difference is not always decreasing in favor of TA. In the next chapter we set to test is whether model signals translate into any risk-adjusted *trading* gains once fees and simple frictions are applied.

5 STRATEGY DEVELOPMENT AND EVALUATION

This chapter showcases the usage of the `XGBoost` *regressor* model explained in Chapter 4 into an executable trading strategy, optimizes the strategy parameters on a walk-forward validation window, and evaluates the resulting rules on an unseen hold-out test set both with metrics and with statistical tests. It is worth noting that the regressor was preferred over a classifier for directional accuracy because the *regression* task on the next one-minute log-return enables the strategy to exploit both direction and expected magnitude.

5.1 TRANSLATING PREDICTIONS INTO TRADES

Predicted next-minute returns $\hat{r}_{t+1}$ are mapped to trades via the pipeline described below. All design choices are derived from, or directly implemented in Python code. The strategy developed was a long-short trading strategy with simple entry and exit rules but with position sizing based on risk measurements. The parameters of the strategy are all tuned in the second validation set as described before, and the sets of possible values are defined below. The entry and exit conditions are based on thresholds and are described also below.

- **Entry threshold:** A rolling 6000-sample window (equivalent to 100 hours) computes the empirical q quantile of historical predictions. Long positions are opened when $\hat{r}_{t+1} \geq \theta_q$; short positions symmetrically when $\hat{r}_{t+1} \leq -\theta_q$. The values of q in the set $[0.9, 0.99, 0.99]$, and θ_q is the quantile function.
- **Exit threshold:** Positions are exited when the signal decays to $\theta_q \times \phi$ with $\phi \in [0.25, 0.75]$, when a *stop loss* SL in the set $[0.005, 0.01, 0.015, 0.02, 0.025]$, or *take-profit* $\mathrm{TP} = \kappa \times \mathrm{SL}$ with κ in the set $[1, 5]$ is hit.

Leverage L_t is computed based on the equations below where μ_{r_t} represents the mean returns from the rolling 6000-sample window, σ_{r_t} represents the volatility of returns from the same window, η_t represents signal strength (probability content of a standard normal CDF), $\hat{\sigma}_{r_t}$ represents the current volatility estimation from a rolling 60-sample window, ξ_t inverse volatility cap, ensuring portfolio volatility stabilizes around 15 % p.a, and $L_{\max}$ represents the maximum leverage allowed tuned from the set $[1, 5]$.

$$L_t = L_{\max} \times \min(1, \eta_t \times \xi_t) \tag{5.1}$$

$$\eta_t = \Phi\left(\frac{\hat{r}_{t+1} - \mu_{r_t}}{\sigma_{r_t}}\right) \tag{5.2}$$

$$\xi_t = \frac{0.15}{\max(10^{-3}, \hat{\sigma}_{r_t})} \tag{5.3}$$

5.2 BACK-TESTING FRAMEWORK

Total return and Sharpe ratio are notoriously sensitive to both model predictions and strategy hyper-parameters. To keep reported performance credible and not specific to a specific set of market conditions, the following techniques to imrpove robustness were implemented:

1. **Rolling/expanding time-series cross-validation:** Hyper-parameter search for the model uses uses a five-fold expanding window from the time-series cross validation. The objective of the optimization is to minimize the *median* mean-squared error (MSE) across folds. The median was chosen over the arithmetic mean because it is less variance-sensitive.

2. **Three-way walk-forward split:** Data is partitioned according to the strategy described in the section 3.3 with the first two sets being used to train the trading model and the last two sets are used for strategy optimization and final out-of-sample validation.

The three-stage walk-forward procedure ensures that the strategy is tuned, validated, and finally tested on data that were never used for model estimation or parameter selection. This framework guarantees that the strategy is always evaluated on genuinely out-of-sample data and under realistic cost assumptions, while keeping the optimization loop entirely separate from the final performance measurement.

1. **Part 1 – Model tuning:** The model's hyper-parameters are selected on *Train 1* (70 % of the series) via a five-fold expanding-window cross-validation that minimizes the *median* mean-squared error, thereby dampening the influence of outliers. It is then validated in the *Validation 1* set as described in the Chapter 4.

2. **Part 2 – Strategy optimization:** Using the best models from Part 1, we re-fit them on *Train 1* $\cup$ *Validation 1* and generate predictions on *Validation 2*. A brute-force grid of 1875 parameter tuples $\left(q, \phi, \mathrm{SL}, \kappa, L_{\max}\right)$ is evaluated according to the description below and its grid size cardinality is showcased in Equation 5.4

   - Entry quantile $q \in \{0.999, 0.99, 0.9\}$;
   - Exit fraction $\phi \in [0.25, 0.75]$ (five steps);
   - Stop-loss $\mathrm{SL} \in [0.5\%, 2.5\%]$ (five steps);
   - Take-profit multiplier $\kappa \in \{1, 2, 3, 4, 5\}$
   - Maximum leverage $L_{\max} \in \{1, 2, 3, 4, 5\}$.

$$|\mathcal{S}| = (3_q) \cdot (5_\phi) \cdot (5_{\mathrm{SL}}) \cdot (5_\kappa) \cdot (5_{L_{\max}}) = 1{,}875. \tag{5.4}$$

   Each tuple is scored on the *median* Sharpe ratio across four quarterly blocks of *Validation 2*; parameter sets yielding fewer than 50 trades are discarded. The best tuple is selected separately for each of the five fee tiers.

3. **Part 3 – Hold-out test:** The final model is re-trained on *Train 1* $\cup$ *Validation 1* $\cup$ *Validation 2* and the pre-selected strategy parameters are applied to the unseen *Test* set (10% of the data). All headline results reported later in this chapter refer to this hold-out test.

As expressed from previous chapter, this work evaluated the strategies in the presence of market friction. Spot and perpetual fee schedules are taken directly from the Binance web site (snapshot: March 2025) (Binance, 2024b). Five tiers are modeled: retail taker, VIP-5, VIP-9, "slippage only" (1 bp each side), and an idealized "friction-free" tier. Slippage itself is represented simply by a flat 1 bp haircut on entry and exit for the "slippage only" tier. Given that preliminary results were already unpromising for raw and TA-based strategies, a more elaborate market-impact model was deferred to future work.

5.3 MODEL AND STRATEGY RESULTS

The results of the trading experiments are listed in the following tables and figures. Tables 7, 11, 13, 15, 17, 19, 21, and 9 report the three successive `XGBoost` regressors trained on *Train 1*, *Train 1* $\cup$ *Val 1*, and *Train 1* $\cup$ *Val 1* $\cup$ *Val 2* for the all the four currencies and the two datasets. Tables 8, 12, 14, 16, 18, 20, 22, and 10 report the strategy metrics for the 5 friction tiers optimized on the *Val 2* set and tested in *Test* set for all the four currencies and the two datasets. Figures 7, 20, 22, 24, 26, 28, 30, 9 represent the cummulative returns of the strategies when no friction is present. Figures 6, 19, 21, 23, 25, 27, 29, 8 represent the cummulative returns of the strategies in the scenarios with friction that were tested.

Table 7 – BTCUSDT / raw / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.215e-06	1.109e-06	1.027e-06
Rscore Train	-5.773e-14	-1.998e-14	2.384e-07
Mse Test	3.617e-07	3.733e-07	5.433e-07
Rscore Test	-3.043e-07	-5.375e-07	-7.153e-07

Table 8 – BTCUSDT / raw / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	4.072e-01	4.495e-01	4.772e-01	4.893e-01	1.086e+00
Annualized Return	1.136e+00	1.281e+00	1.379e+00	1.423e+00	4.121e+00
Sharpe Ratio	1.687e+00	1.810e+00	1.889e+00	1.923e+00	1.949e+00
Max Drawdown	-2.541e-01	-2.442e-01	-2.378e-01	-2.351e-01	-8.604e-01
Calmar Ratio	4.470e+00	5.246e+00	5.798e+00	6.053e+00	4.789e+00
Win Rate	2.742e-01	2.742e-01	2.742e-01	2.742e-01	5.088e-01
Num Trades	63	63	63	63	41619

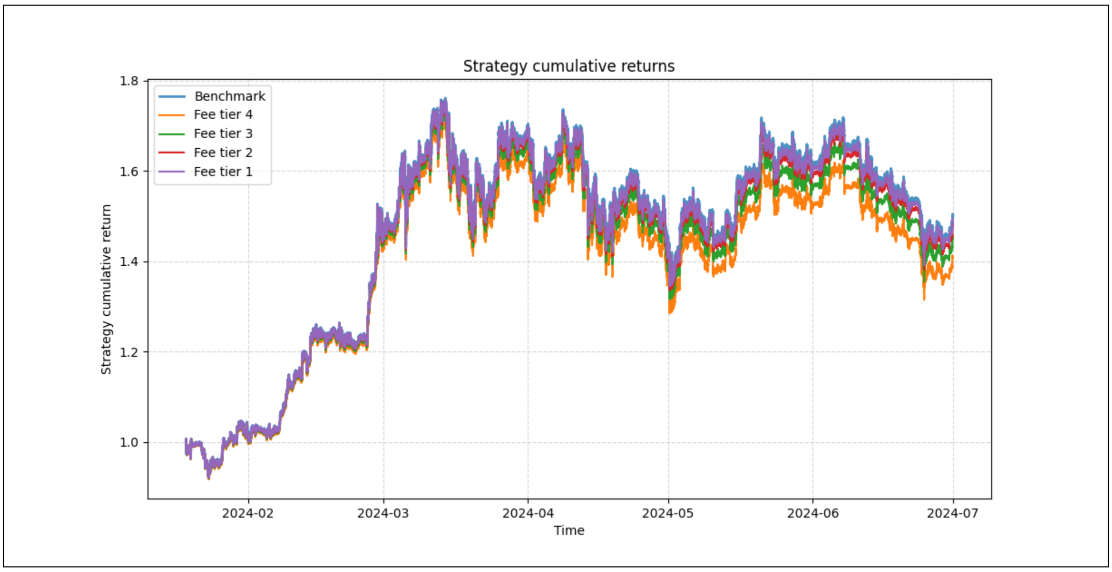

Figure 6 – BTCUSDT Spot Raw Strats with fees

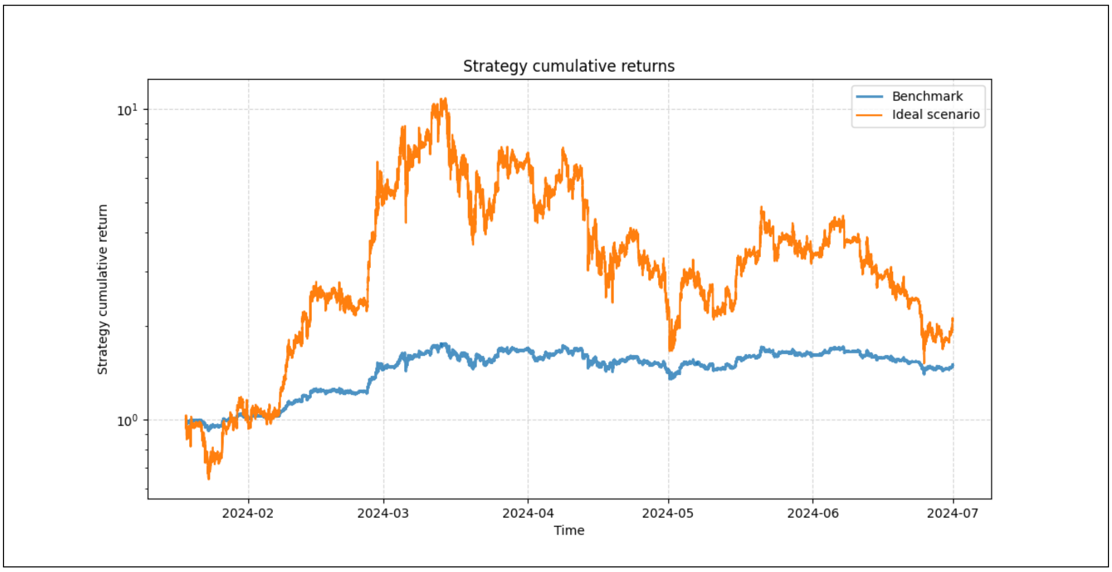

Figure 7 – BTCUSDT Spot Raw Strats in Ideal Scenario

Table 9 – BNBUSDT / ta / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.131e-06	1.917e-06	1.764e-06
Rscore Train	3.579e-03	4.531e-03	1.068e-02
Mse Test	4.252e-07	6.220e-07	7.698e-07
Rscore Test	2.878e-02	2.651e-02	7.629e-03

Table 10 – BNBUSDT / ta / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	-7.292e-01	-4.621e-01	-1.666e-01	-1.000e+00	1.233e+10
Annualized Return	-9.490e-01	-7.564e-01	-3.396e-01	-1.000e+00	9.596e+22
Sharpe Ratio	-9.669e+00	-4.569e+00	-1.242e+00	-1.597e+01	2.497e+01
Max Drawdown	-7.734e-01	-5.857e-01	-3.933e-01	-1.000e+00	-3.699e-01
Calmar Ratio	1.227e+00	1.291e+00	8.634e-01	1.000e+00	2.594e+23
Win Rate	1.255e-01	1.929e-01	3.408e-01	4.275e-01	4.991e-01
Num Trades	757	757	757	47646	47646

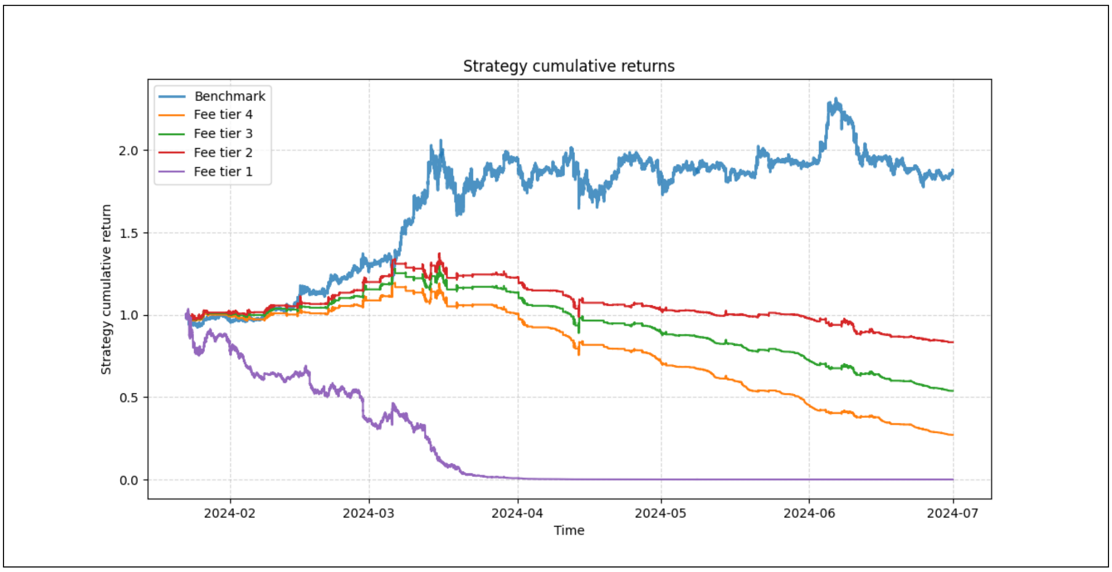

Figure 8 – BNBUSDT Spot TA Strats with fees

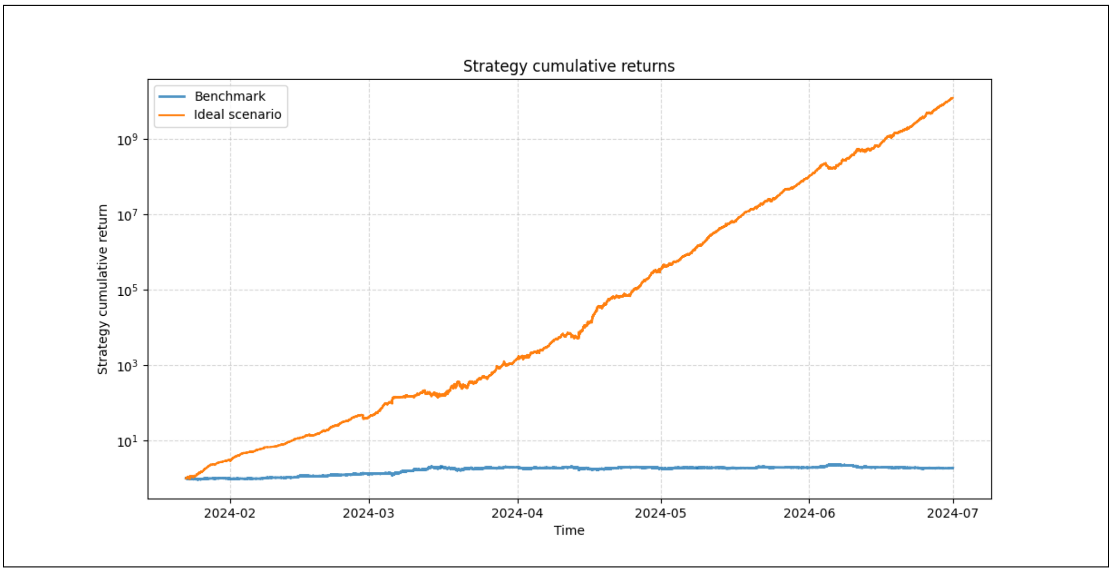

Figure 9 – BNBUSDT Spot TA Strats in Ideal Scenario

5.4 RESULTS EVALUATION

The empirical evidence presented in the preceding tables and figures is consistent with the working hypothesis articulated in Chapter 1: *canonical TA overlays provide, at best, sporadic and fee-fragile gains at the one-minute horizon*. In particular the points below were observed:

- **Ideal, friction-free setting:** Eight strategy configurations were tested (raw vs. TA for four tickers). In the absence of costs, only **three** configurations produced Sharpe ratios that exceeded a buy-and-hold benchmark by very little edge, *BTC (raw)* 7, *BTC (TA)* 20, and *SOL (TA)* 28. ETH models never surpassed buy-and-hold 22 24, and SOL (raw) 26 25 fell only slightly short. BNB strategies 30 9 generated out-sized benchmark performance, but did so via tens of thousands of micro-trades that would be impossible to replicate under real liquidity constraints.

- **Introducing realistic friction:** Once Binance taker fees and a 1 bp slippage haircut were applied, virtually all alphas evaporated. BTC 6 19 strategies morphed into fee-encumbered buy-and-hold replicas; ETH (raw) 21 and SOL (TA) 27 turned strictly loss-making; ETH (TA) 23 achieved benchmark-like returns only in the "slippage-only" tier; BNB 29 8 Sharpe ratios collapsed across every paid tier due to micro-trades edge being eliminated.

- **TA overlays offer no cost-robust edge:** Situations in which TA marginally improved returns (SOL, ETH) did so only in the ideal tier and only for one of the four assets. Nowhere did TA maintain an advantage once any fee was charged.

To explore whether lower taker fees in perpetual contracts could rescue the strategy, the identical pipeline was re-run on Binance USDT-quoted perpetuals (UM) and its full results appear in Appendix D. It is worth noting that In the friction-free tier, only **BTC (TA)** 34 posted a positive Sharpe above the buy-and-hold baseline; all other futures strategies mirrored, or under-performed, their spot counterparts. The moment even discounted VIP-9 fees were imposed, BTC (TA) 33 reverted to near-zero alpha, indicating that the apparent edge was fee-fragile. In short, neither the smaller fees of perpetuals compared to spot products nor the addition of TA overlays produced a cost-resilient advantage in any of the scenarios.

6 CONCLUSION

The conclusions of this study are found below. The limitations and suggestions for future works are also highlighted.

6.1 KEY FINDINGS

1. **Predictive signals exists but are razor thin:** One-minute returns exhibit near-zero R^2, yet a carefully tuned gradient-boosted regressor can extract a modest edge, sufficient to generate a decent sharpe ratio for a few scenarios using TA indicators in a ideal no-friction environment.

2. **TA indicators rarely help:** Across 40 model–asset-friction combinations (4 tickers $\times$ 2 feature sets $\times$ 5 friction tiers), raw micro-structure features equaled TA-enriched sets in most of configurations in the ideal case and never displayed fee-robust inferiority.

3. **Market friction is the great equalizer:** Virtually all apparent alpha vanished when realistic taker fees (2–30 bp round-trip) were applied. Strategies that relied on thousands of micro-trades (e.g. BNB Spot) proved especially vulnerable.

6.2 LIMITATIONS OF THE STUDY

- **Simple slippage model:** A flat 1 bp haircut cannot capture the non-linear impact of consuming large volumes in a thin order book.

- **Single-exchange scope:** Only Binance spot and perpetual data were examined; cross-venue liquidity and latency effects remain untested.

- **Fixed execution template:** The strategy employs a TWAP with capped participation. Adaptive order types (e.g. iceberg, POV) may alter realised costs.

- **Only one trading timeframe explored:** All the experiments were realized in the *1m* timeframe, shorter and longer timeframes are not analyzed.

6.3 SUGGESTIONS FOR FUTURE RESEARCH

1. **Longer horizons:** Extend the framework to longer horizions such as 15-minute or 1-week bars where microstructure noise is lower and fee drag is proportionally smaller.
2. **Feature learning:** Replace hand-crafted order-book aggregations with representation learning (e.g. temporal CNNs on the raw L2 grid) to investigate whether richer micro-patterns exist.
3. **Sequence learning:** Replace TA indicators, which are the source of information past the last tick in the current setup, with sequence models such as LSTM or Transformers.
4. **Reinforcement learning:** Replace fixed trading logic with a reinforcement learning model that learns from the trading environment and acts on it.
5. **Order-book simulation:** Integrate a depth-of-book simulator that replays historical liquidity to price in non-linear market impact.
6. **Cross-exchange routing:** Test whether smart-order routing across venues can recover alpha eroded by taker fees on a single venue.

6.4 FINAL REMARKS

In conclusion, this thesis finds no general, fee-robust evidence that adding canonical TA indicators enhance one-minute crypto-trading strategies based only on market data. It is also found that even the small edge the models can extract from the market data is very fragile once realistic frictions are applied, at least on the 1 minute trading timeframe.

REFERENCES

ALDRIDGE, I. *High-frequency trading: a practical guide to algorithmic strategies and trading systems*. [S.l.]: John Wiley & Sons, 2013.

ALESSANDRETTI, L.; ELBAHRAWY, A.; AIELLO, L. M.; BARONCHELLI, A. Anticipating cryptocurrency prices using machine learning. *Complexity*, Wiley Online Library, v. 2018, n. 1, p. 8983590, 2018.

ALMGREN, R.; CHRISS, N. Optimal execution of portfolio transactions. *Journal of Risk*, v. 3, n. 2, p. 5–39, 2001.

ANGUIANO, J. Á. I.; GARCÍA-MEDINA, A. *Technical Analysis Meets Machine Learning: Bitcoin Evidence*. 2025. Preprint. Disponível em: <https://www.preprints.org/manuscript/202506.1866/v1>.

BAILEY, D. H.; PRADO, M. López de. The deflated sharpe ratio: Correcting for selection bias, backtest overfitting and non-normality. *Journal of Portfolio Management*, v. 40, n. 5, p. 94–107, 2014.

BALLINGS, M.; POEL, D. Van den; HESPEELS, N.; GRYP, R. Evaluating multiple classifiers for stock price direction prediction. *Expert systems with Applications*, Elsevier, v. 42, n. 20, p. 7046–7056, 2015.

Binance. *Binance API Documentation*. 2024. <https://binance-docs.github.io/apidocs/>. Accessed: 2024-01-05.

Binance. *Binance Official Website*. 2024. <https://www.binance.com/>. Accessed: 01-05-2024.

BLACK, F.; SCHOLES, M. The pricing of options and corporate liabilities. *Journal of Political Economy*, v. 81, n. 3, p. 637–654, 1973. Disponível em: <https://www.journals.uchicago.edu/doi/10.1086/260062>.

BOLLERSLEV, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, v. 31, n. 3, p. 307–327, 1986. Disponível em: <https://www.sciencedirect.com/science/article/pii/0304407686900631>.

BORRI, N. Conditional tail-risk in cryptocurrency markets. *Journal of Empirical Finance*, Elsevier, v. 50, p. 1–19, 2019.

BOUCHAUD, J.-P.; BONART, J.; DONIER, J.; GOULD, M. *Trades, Quotes and Prices: Financial Markets Under the Microscope*. [S.l.]: Cambridge University Press, 2018.

BOX, G. E. P.; JENKINS, G. M.; REINSEL, G. C.; LJUNG, G. M. *Time Series Analysis: Forecasting and Control*. [S.l.]: Holden-Day, 1976.

BREIMAN, L. Random forests. *Machine Learning*, v. 45, n. 1, p. 5–32, 2001. Disponível em: <https://link.springer.com/article/10.1023/A:1010933404324>.

BROWN, G.; POCOCK, A.; ZHAO, M.-J.; LUJÁN, M. Conditional likelihood maximisation: a unifying framework for information theoretic feature selection. *The journal of machine learning research*, JMLR. org, v. 13, n. 1, p. 27–66, 2012.

CARTEA, Á.; JAIMUNGAL, S.; PENALVA, J. *Algorithmic and high-frequency trading*. [S.l.]: Cambridge University Press, 2015.

CHAN, E. *Algorithmic trading: winning strategies and their rationale*. [S.l.]: John Wiley & Sons, 2013.

CHANDAR, S. K. Convolutional neural network for stock trading using technical indicators. *Automated Software Engineering*, v. 29, p. 16, 2022. Disponível em: <https://link.springer.com/article/10.1007/s10515-021-00303-z>.

CHEN, T.; GUESTRIN, C. Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. [s.n.], 2016. p. 785–794. Disponível em: <https://dl.acm.org/doi/10.1145/2939672.2939785>.

CLARK, T. E.; WEST, K. D. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, v. 138, n. 1, p. 291–311, 2007.

CoinMarketCap. *Global Cryptocurrency Market Capitalization*. 2021. <https://coinmarketcap.com>. Accessed: 2025-09-19.

CoinMarketCap. *Historical Snapshot - January 01, 2024*. 2024. <https://coinmarketcap.com/historical/20240101/>. Accessed: 2024-01-05.

CORBET, S.; LARKIN, C.; LUCEY, B. M.; MEEGAN, A.; YAROVAYA, L. The impact of macroeconomic news on bitcoin returns. *The European Journal of Finance*, Taylor & Francis, v. 26, n. 14, p. 1396–1416, 2020.

COVER, T. M.; THOMAS, J. A. *Elements of Information Theory*. 2nd. ed. Wiley-Interscience, 2006. Disponível em: <https://onlinelibrary.wiley.com/doi/book/10.1002/047174882X>.

COX, D. R. The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B*, v. 20, n. 2, p. 215–242, 1958.

DIEBOLD, F. X.; MARIANO, R. S. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, ASA Website, v. 20, n. 1, p. 134–144, 2002.

DROZDZ, S.; MINATI, L.; OSWIECIMKA, P.; STANUSZEK, M.; WATOREK, M. Competition of noise and collectivity in global cryptocurrency trading: Route to a self-contained market. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, AIP Publishing, v. 30, n. 2, 2020.

DYHRBERG, A. H. Bitcoin, gold and the dollar–a garch volatility analysis. *Finance research letters*, Elsevier, v. 16, p. 85–92, 2016.

ENGLE, R. F. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, v. 50, n. 4, p. 987–1007, 1982. Disponível em: <https://www.jstor.org/stable/1912773>.

FAMA, E. F. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, v. 25, n. 2, p. 383–417, 1970. Disponível em: <https://www.jstor.org/stable/2325486>.

FISCHER, T.; KRAUSS, C. Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, Elsevier, v. 270, n. 2, p. 654–669, 2018.

FORTIER, M.; CONTRIBUTORS, T.-L. *TA-Lib: Technical Analysis Library*. 2024. <https://ta-lib.org/>. Accessed: 2024-01-05.

FRANKLIN, B. *The Way to Wealth: Ben Franklin on Money and Success*. [S.l.]: Dover Publications, 1969. Originally published in 1758.

FRIEDMAN, J. H. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, v. 29, n. 5, p. 1189–1232, 2001.

GATHERAL, J. No-dynamic-arbitrage and market impact. *Quantitative Finance*, v. 10, n. 7, p. 749–759, 2010.

GHOSH, P.; NEUFELD, A.; SAHOO, J. K. Forecasting directional movements of stock prices for intraday trading using lstm and random forests. *Finance Research Letters*, v. 46, n. Part A, p. 102280, 2022. Disponível em: <https://www.sciencedirect.com/science/article/pii/S1544612321003557>.

GOUTTE, S.; LE, H.-V.; LIU, F.; METTENHEIM, H.-J. von. Deep learning and technical analysis in cryptocurrency market. *Finance Research Letters*, v. 54, p. 103809, 2023. Disponível em: <https://www.sciencedirect.com/science/article/pii/S1544612323003894>.

GRINOLD, R. C.; KAHN, R. N. *Active Portfolio Management*. 2nd. ed. [S.l.]: McGraw–Hill, 2000.

GU, S.; KELLY, B.; XIU, D. Empirical asset pricing via machine learning. *The Review of Financial Studies*, Oxford University Press, v. 33, n. 5, p. 2223–2273, 2020.

HANSEN, P. R. A test for superior predictive ability. *Journal of Business & Economic Statistics*, v. 23, n. 4, p. 365–380, 2005.

HANSEN, P. R.; LUNDE, A.; NASON, J. M. The model confidence set. *Econometrica*, v. 79, n. 2, p. 453–497, 2011. Disponível em: <https://onlinelibrary.wiley.com/doi/10.3982/ECTA5771>.

HESTON, S. L. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, Oxford University Press, v. 6, n. 2, p. 327–343, 1993.

JI, G.; YU, J.; HU, K.; XIE, J.; JI, X. An adaptive feature selection schema using improved technical indicators for predicting stock price movements. *Expert Systems with Applications*, v. 200, p. 116941, 2022. Disponível em: <https://www.sciencedirect.com/science/article/pii/S0957417422013154>.

JOBSON, J. D.; KORKIE, B. M. Performance hypothesis testing with the sharpe and treynor measures. *Journal of Finance*, v. 36, n. 4, p. 889–908, 1981. Disponível em: <https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1981.tb04891.x>.

KHANIKI, M. A. L.; MANTHOURI, M. Enhancing price prediction in cryptocurrency using transformer neural network and technical indicators. *arXiv preprint arXiv:2403.03606*, 2024. Disponível em: <https://arxiv.org/abs/2403.03606>.

KISSELL, R. *The Science of Algorithmic Trading and Portfolio Management*. [S.l.]: Academic Press, 2013.

KRAUSS, C.; DO, X. A.; HUCK, N. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the s&p 500. *European Journal of Operational Research*, Elsevier, v. 259, n. 2, p. 689–702, 2017.

KYLE, A. S. Continuous auctions and insider trading. *Econometrica*, v. 53, n. 6, p. 1315–1335, 1985. Disponível em: <https://www.jstor.org/stable/1913210>.

LEDOIT, O.; WOLF, M. Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance*, v. 15, n. 5, p. 850–859, 2008.

LEE, M. C. Temporal fusion transformer-based trading strategy for multi-crypto assets using on-chain and technical indicators. *Systems*, v. 13, n. 6, p. 474, 2025. Disponível em: <https://www.mdpi.com/2079-8954/13/6/474>.

LIU, Y. Novel volatility forecasting using deep learning–long short term memory recurrent neural networks. *Expert Systems with Applications*, v. 132, p. 99–109, 2019. Disponível em: <https://www.sciencedirect.com/science/article/pii/S0957417419302147>.

LO, A. W.; MAMAYSKY, H.; WANG, J. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The journal of finance*, Wiley Online Library, v. 55, n. 4, p. 1705–1765, 2000.

LUNDBERG, S. M.; LEE, S.-I. A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems*. [s.n.], 2017. v. 30. Disponível em: <https://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.

MAGDON-ISMAIL, M.; ATIYA, A. F. Maximum drawdown. *Risk Magazine*, v. 17, n. 10, p. 99–102, 2004.

MALKIEL, B. G. The efficient market hypothesis and its critics. *Journal of economic perspectives*, American Economic Association, v. 17, n. 1, p. 59–82, 2003.

MALLQUI, D. C.; FERNANDES, R. A. Predicting the direction, maximum, minimum and closing prices of daily bitcoin exchange rate using machine learning techniques. *Applied Soft Computing*, Elsevier, v. 75, p. 596–606, 2019.

MERTON, R. C. Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, Elsevier, v. 3, n. 1-2, p. 125–144, 1976.

NAKAMOTO, S. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. White paper. Disponível em: <https://bitcoin.org/bitcoin.pdf>.

ORTU, M.; URAS, N.; CONVERSANO, C.; BARTOLUCCI, S.; DESTEFANIS, G. On technical trading and social media indicators for cryptocurrency price classification through deep learning. *Expert Systems with Applications*, v. 198, p. 116804, 2022. Disponível em: <https://www.sciencedirect.com/science/article/pii/S0957417422008985>.

PARK, C.-H.; IRWIN, S. H. What do we know about the profitability of technical analysis? *Journal of Economic surveys*, Wiley Online Library, v. 21, n. 4, p. 786–826, 2007.

PRADO, M. López de. *Advances in Financial Machine Learning*. Wiley, 2018. Disponível em: <https://www.wiley.com/en-us/Advances+in+Financial+Machine+Learning-p-9781119482086>.

QUINLAN, J. R. Induction of decision trees. *Machine Learning*, v. 1, n. 1, p. 81–106, 1986. Disponível em: <https://link.springer.com/article/10.1007/BF00116251>.

SHARPE, W. F. The sharpe ratio. *Streetwise–the Best of the Journal of Portfolio Management*, Princeton University Press NJ, v. 3, n. 3, p. 169–85, 1998.

SHILLER, R. J. From efficient markets theory to behavioral finance. *Journal of economic perspectives*, American Economic Association, v. 17, n. 1, p. 83–104, 2003.

TANRIKULU, H. M.; PABUCCU, H. The effect of data types on the performance of machine learning algorithms for cryptocurrency prediction. *Computational Economics*, 2025. Disponível em: <https://link.springer.com/article/10.1007/s10614-025-10919-y>.

TSAY, R. S. *Analysis of Financial Time Series*. John Wiley & Sons, 2010. Disponível em: <https://onlinelibrary.wiley.com/doi/book/10.1002/9780470644560>.

VALIANT, L. G. A theory of the learnable. *Communications of the ACM*, v. 27, n. 11, p. 1134–1142, 1984. Disponível em: <https://dl.acm.org/doi/10.1145/1968.1972>.

VERGARA, J. R.; ESTÉVEZ, P. A. A review of feature selection methods based on mutual information. *Neural computing and applications*, Springer, v. 24, n. 1, p. 175–186, 2014.

VERMA, S.; SAHU, S. P.; SAHU, T. P. Stock market forecasting with different input indicators using machine learning and deep learning techniques: A review. *Engineering Letters*, v. 31, n. 1, 2023.

VINH, N. X.; EPPS, J.; BAILEY, J. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In: *Proceedings of the 26th Annual International Conference on Machine Learning*. [s.n.], 2009. (ICML '09), p. 1073–1080. ISBN 9781605585161. Disponível em: <https://doi.org/10.1145/1553374.1553511>.

WHITE, H. A reality check for data snooping. *Econometrica*, v. 68, n. 5, p. 1097–1126, 2000. Disponível em: <https://onlinelibrary.wiley.com/doi/10.1111/1468-0262.00152>.

XGBoost Developers. *XGBoost Documentation*. 2025. <https://xgboost.readthedocs.io/en/stable/index.html>. Accessed: 2025-09-19.

YEO, I. K.; JOHNSON, R. A. A new family of power transformations to improve normality or symmetry. *Biometrika*, v. 87, n. 4, p. 954–959, 2000. Disponível em: <https://academic.oup.com/biomet/article-abstract/87/4/954/235233>.

YOUNG, T. W. Calmar ratio: A smoother tool. *Futures*, v. 20, n. 1, p. 40, 1991.

YOUSSEFI, A. E.; HESSANE, A.; ZEROUAL, I.; FARHAOUI, Y. Optimizing forecast accuracy in cryptocurrency markets: Evaluating feature selection techniques for technical indicators. *Computers, Materials & Continua*, v. 83, n. 2, p. 3411–3433, 2025.

Appendices

Appendix A – DESCRIPTION OF TA INDICATORS

- **Candlestick Patterns ("candles")**:
  - *cdl_2crows*: Detects the Two Crows bearish reversal pattern, involving three candles with a gap up and bearish continuation.
  - *cdl_3blackcrows*: Identifies the Three Black Crows bearish reversal pattern, signaled by three consecutive long-bodied bearish candles.
  - *cdl_3inside*: Finds the Three Inside Up/Down pattern, a trend reversal signal using three specific candlesticks.
  - *cdl_3linestrike*: Locates the Three-Line Strike pattern, which may signal trend continuation or reversal depending on the context.
  - *cdl_3outside*: Detects the Three Outside Up/Down pattern, a strong reversal formation.
  - *cdl_3starsinsouth*: Flags the Three Stars In The South bullish reversal pattern, usually after a downtrend.
  - *cdl_3whitesoldiers*: Identifies the Three White Soldiers bullish reversal pattern, three strong consecutive bullish candles.
  - *cdl_abandonedbaby*: Detects the Abandoned Baby pattern, a rare and strong reversal signal with a doji gapped between two candles.
  - *cdl_advanceblock*: Finds the Advance Block bearish reversal pattern, featuring three advancing candles with weakening momentum.
  - *cdl_belthold*: Locates the Belt Hold bullish/bearish reversal pattern, a single long candle that gaps in the direction of the trend.
  - *cdl_breakaway*: Detects the Breakaway bullish or bearish reversal, a five-candle pattern indicating a possible trend reversal.
  - *cdl_closingmarubozu*: Flags the Closing Marubozu, a candlestick with no shadow on the close side, indicating strong conviction.
  - *cdl_concealbabyswall*: Finds the Concealing Baby Swallow, a rare bullish reversal in a downtrend.

- *cdl_counterattack*: Detects the Counterattack pattern, two opposite candles with similar closes, signaling a potential reversal.
- *cdl_darkcloudcover*: Identifies the Dark Cloud Cover, a bearish reversal formed by a bullish candle followed by a bearish one opening above and closing below the midpoint.
- *cdl_doji_10_0.1*: Detects doji patterns, characterized by nearly equal open and close prices, indicating indecision.
- *cdl_dojistar*: Flags the Doji Star, a doji that forms after a significant trend, suggesting a potential reversal.
- *cdl_dragonflydoji*: Identifies the Dragonfly Doji, a doji with a long lower shadow and no upper shadow, hinting at bullish reversal.
- *cdl_engulfing*: Finds the Engulfing pattern, where a candle completely engulfs the previous one, signaling reversal.
- *cdl_eveningdojistar*: Detects the Evening Doji Star, a three-candle bearish reversal with a doji in the middle.
- *cdl_eveningstar*: Locates the Evening Star, a three-candle bearish reversal pattern.
- *cdl_gapsidesidewhite*: Flags the Gap Side-by-Side White Lines, a continuation pattern with two candles gapped in the direction of the trend.
- *cdl_gravestonedoji*: Identifies the Gravestone Doji, a doji with a long upper shadow and no lower shadow, suggesting bearish reversal.
- *cdl_hammer*: Detects the Hammer, a bullish reversal with a small body and long lower shadow.
- *cdl_hangingman*: Finds the Hanging Man, similar to the hammer but at the top of an uptrend, indicating bearish reversal.
- *cdl_harami*: Flags the Harami, a two-candle reversal pattern with the second candle inside the first.
- *cdl_haramicross*: Detects the Harami Cross, a variant of harami with a doji as the second candle.
- *cdl_highwave*: Finds the High-Wave pattern, candles with long upper and lower shadows, indicating indecision.

- *cdl_hikkake*: Identifies the Hikkake, a pattern suggesting continuation after a failed reversal.
- *cdl_hikkakemod*: Modified Hikkake, a variation with altered rules for pattern detection.
- *cdl_homingpigeon*: Finds the Homing Pigeon, a bullish reversal with two nested bearish candles.
- *cdl_identical3crows*: Flags the Identical Three Crows, three nearly identical bearish candles in a downtrend.
- *cdl_inneck*: Detects the In-Neck Line, a bearish continuation pattern with two candles.
- *cdl_inside*: Finds the Inside Bar, where the current candle's high and low are within the previous candle.
- *cdl_invertedhammer*: Identifies the Inverted Hammer, a bullish reversal with a small body and long upper shadow.
- *cdl_kicking*: Detects the Kicking pattern, two Marubozu candles with a gap between them.
- *cdl_kickingbylength*: Kicking by Length, a variant that considers candle length for signal strength.
- *cdl_ladderbottom*: Flags the Ladder Bottom, a rare five-candle bullish reversal.
- *cdl_longleggeddoji*: Identifies the Long-Legged Doji, a doji with long shadows, indicating strong indecision.
- *cdl_longline*: Finds the Long Line candles, long-bodied candles that may signal strength in a trend.
- *cdl_marubozu*: Detects the Marubozu, a candle with no shadows, showing strong momentum.
- *cdl_matchinglow*: Identifies the Matching Low, two candles with identical lows, hinting at support.
- *cdl_mathold*: Flags the Mat Hold, a bullish continuation pattern.
- *cdl_morningdojistar*: Detects the Morning Doji Star, a bullish reversal with a doji in the middle.

- *cdl_morningstar*: Finds the Morning Star, a three-candle bullish reversal pattern.
- *cdl_onneck*: On-Neck Line, a bearish continuation pattern.
- *cdl_piercing*: Identifies the Piercing Line, a two-candle bullish reversal where the second opens below and closes above the midpoint.
- *cdl_rickshawman*: Flags the Rickshaw Man, a long-legged doji near the midpoint of the range.
- *cdl_risefall3methods*: Detects the Rising/Falling Three Methods, continuation patterns with a pause before trend resumes.
- *cdl_separatinglines*: Finds the Separating Lines, two-candle continuation patterns.
- *cdl_shootingstar*: Identifies the Shooting Star, a bearish reversal with a small body and long upper shadow.
- *cdl_shortline*: Detects the Short Line candles, short-bodied candles indicating market hesitation.
- *cdl_spinningtop*: Flags the Spinning Top, candles with small bodies and long shadows, signaling indecision.
- *cdl_stalledpattern*: Identifies the Stalled Pattern, a bearish reversal signal.
- *cdl_sticksandwich*: Finds the Stick Sandwich, a three-candle reversal pattern.
- *cdl_takuri*: Flags the Takuri (Dragonfly Doji variant), a bullish reversal.
- *cdl_tasukigap*: Detects the Tasuki Gap, a continuation pattern with a gap followed by a pullback.
- *cdl_thrusting*: Finds the Thrusting Line, a bearish continuation with the second candle closing within the prior body.
- *cdl_tristar*: Identifies the Tri-Star, three consecutive doji patterns signaling reversal.
- *cdl_unique3river*: Detects the Unique Three River, a rare bullish reversal pattern.
- *cdl_upsidegap2crows*: Finds the Upside Gap Two Crows, a bearish reversal after a gap up.
- *cdl_xsidegap3methods*: Flags the Side-by-Side White/Black Lines, continuation patterns involving gaps.

- **Momentum Indicators ("momentum")**:
  - *ao*: Awesome Oscillator, measures market momentum using the difference between two moving averages.
  - *apo*: Absolute Price Oscillator, shows the absolute difference between two EMAs.
  - *bias*: BIAS indicator, the deviation of price from a moving average.
  - *bop*: Balance of Power, shows the strength of buyers vs. sellers.
  - *brar*: BRAR, assesses buying and selling pressure based on highs and lows.
  - *cci*: Commodity Channel Index, measures price deviation from average for identifying cyclical trends.
  - *cfo*: Chande Forecast Oscillator, the difference between close and linear regression forecast.
  - *cg*: Center of Gravity oscillator, detects turning points based on weighted moving averages.
  - *cmo*: Chande Momentum Oscillator, evaluates momentum as the difference between sums of recent gains and losses.
  - *coppock*: Coppock Curve, a long-term momentum indicator for spotting major bottoms.
  - *cti*: Correlation Trend Indicator, measures the trend strength using correlation coefficients.
  - *er*: Efficiency Ratio, quantifies price efficiency in trending vs. ranging periods.
  - *eri*: Elder Ray Index, combines Bull Power and Bear Power to identify trend strength.
  - *fisher*: Fisher Transform, converts prices into a Gaussian normal distribution for clearer signals.
  - *inertia*: Inertia Indicator, combines trend and momentum using linear regression.
  - *kdj*: KDJ indicator, a variation of Stochastic Oscillator with added J line.
  - *kst*: Know Sure Thing, a momentum oscillator using four different ROC calculations.

- *macd*: Moving Average Convergence Divergence, tracks trend and momentum via EMAs and their difference.
- *mom*: Momentum, the rate of change in price over a set period.
- *pgo*: Pretty Good Oscillator, difference between price and a moving average, normalized by mean deviation.
- *ppo*: Percentage Price Oscillator, similar to MACD but uses percentage difference between EMAs.
- *psl*: Price Slope, measures the linear slope of price over a period.
- *pvo*: Percentage Volume Oscillator, similar to PPO but for volume.
- *qqe*: Quantitative Qualitative Estimation, a smoothed RSI-based momentum indicator.
- *roc*: Rate of Change, the percentage change in price over a period.
- *rsi*: Relative Strength Index, measures the speed and change of price movements to identify overbought or oversold conditions.
- *rsx*: Smoothed RSI, an advanced version of RSI for smoother signals.
- *rvgi*: Relative Vigor Index, compares closing and opening prices to gauge strength.
- *slope*: Slope indicator, the angle of a regression line over a window.
- *smi*: Stochastic Momentum Index, a refined version of Stochastic Oscillator.
- *stc*: Schaff Trend Cycle, a fast and responsive MACD-based oscillator.
- *stoch*: Stochastic Oscillator, measures the position of close relative to the high-low range.
- *stochrsi*: Stochastic RSI, applies Stochastic calculation to RSI values for more sensitive signals.
- *trix*: TRIX, a triple-smoothed EMA that filters out market noise.
- *tsi*: True Strength Index, measures double-smoothed momentum.
- *uo*: Ultimate Oscillator, combines multiple timeframes to reduce false signals.
- *willr*: Williams R, indicates overbought or oversold conditions based on recent highs and lows.

- **Overlap Indicators ("overlap")**:
  - *alma*: Arnaud Legoux Moving Average, reduces lag and detects trends smoothly.
  - *dema*: Double Exponential Moving Average, offers reduced lag over standard EMAs.
  - *ema*: Exponential Moving Average, assigns greater weight to recent prices.
  - *fwma*: Fibonacci Weighted Moving Average, uses Fibonacci sequence for weighting.
  - *hilo*: High-Low Indicator, tracks the highest high and lowest low over a period.
  - *hma*: Hull Moving Average, reduces lag and smooths price data effectively.
  - *jma*: Jurik Moving Average, adaptive and smooth moving average.
  - *kama*: Kaufman Adaptive Moving Average, adapts speed based on volatility.
  - *linreg*: Linear Regression, predicts price trends using regression analysis.
  - *midpoint*: Midpoint, average of the highest high and lowest low.
  - *midprice*: Midprice, midpoint between current high and low.
  - *pwma*: Power Weighted Moving Average, applies power-based weights to prices.
  - *rma*: Rolling Moving Average, a variant of EMA.
  - *sinwma*: Sine Weighted Moving Average, uses sine function for weighting.
  - *sma*: Simple Moving Average, average price over a defined period.
  - *ssf*: Super Smoother Filter, a low-pass filter for trend detection.
  - *supertrend*: Supertrend Indicator, highlights trend direction with dynamic bands.
  - *swma*: Symmetrically Weighted Moving Average, weights values symmetrically.
  - *t3*: T3 Moving Average, uses triple EMA smoothing for signal clarity.
  - *tema*: Triple Exponential Moving Average, reduces lag by using triple EMA.
  - *trima*: Triangular Moving Average, double-smooths a simple moving average.
  - *vwap*: Volume Weighted Average Price, price average weighted by volume.
  - *vwma*: Volume Weighted Moving Average, moving average weighted by volume.
  - *wcp*: Weighted Close Price, average using high, low, and close.
  - *wma*: Weighted Moving Average, assigns more weight to recent prices.
  - *zlma*: Zero-Lag Moving Average, designed to reduce lag to a minimum.

- **Performance and Statistics ("performance"and "statistics")**:
  - *log_return*: Computes the logarithmic return of price series.
  - *percent_return*: Calculates the percentage return over a period.
  - *entropy*: Shannon entropy, quantifies the unpredictability in price data.
  - *kurtosis*: Kurtosis, measures the "tailedness"of the distribution of returns.
  - *mad*: Median Absolute Deviation, a robust measure of statistical dispersion.
  - *median*: Median, the middle value in a sorted data set.
  - *quantile*: Quantile, value below which a given percentage of observations fall.
  - *skew*: Skewness, measures the asymmetry of the return distribution.
  - *stdev*: Standard deviation, measures the amount of variation in price data.
  - *tos_stdevall*: TOS Standard Deviation All, a variant standard deviation calculation.
  - *variance*: Variance, the average squared deviation from the mean.
  - *zscore*: Z-score, quantifies how many standard deviations a value is from the mean.
- **Trend Indicators ("trend")**:
  - *adx*: Average Directional Index, quantifies the strength of a trend.
  - *amat*: Arnaud Legoux Moving Average Trend, a trend filter using ALMA.
  - *aroon*: Aroon Indicator, identifies trend changes and consolidations.
  - *chop*: Choppiness Index, measures whether the market is trending or ranging.
  - *cksp*: Kaufman's Efficiency Ratio-based Stop, identifies trend reversals.
  - *decay*: Decay indicator, measures the rate at which price momentum decays.
  - *decreasing*: Flags if price is decreasing over a period.
  - *increasing*: Flags if price is increasing over a period.
  - *psar*: Parabolic SAR, spots trend direction and potential reversals.
  - *qstick*: Qstick, a trend indicator based on candlestick body averages.
  - *ttm_trend*: TTM Trend, a trend coloring indicator.
  - *vhf*: Vertical Horizontal Filter, distinguishes between trending and ranging periods.
  - *vortex*: Vortex Indicator, spots the start of a new trend.

- **Volatility Indicators ("volatility")**:
  - *aberration*: Aberration, a volatility and trend-following indicator.
  - *accbands*: Acceleration Bands, envelopes that expand and contract with volatility.
  - *atr*: Average True Range, measures market volatility using price ranges.
  - *bbands*: Bollinger Bands, volatility bands placed above and below a moving average.
  - *donchian*: Donchian Channel, plots the highest high and lowest low over a period.
  - *hwc*: Historical Volatility (Close), standard deviation of log returns.
  - *kc*: Keltner Channels, envelopes based on ATR above and below an EMA.
  - *massi*: Mass Index, identifies potential trend reversals based on range expansions.
  - *natr*: Normalized ATR, ATR divided by closing price for scale invariance.
  - *pdist*: Price Distance, difference between price and moving average.
  - *rvi*: Relative Volatility Index, an RSI applied to standard deviation.
  - *thermo*: Thermometer, an indicator based on price ranges.
  - *true_range*: True Range, measures the greatest of current high-low, high-previous close, or low-previous close.
  - *ui*: Ulcer Index, measures price drawdowns and volatility.
- **Volume Indicators ("volume")**:
  - *ad*: Accumulation/Distribution Line, measures cumulative flow of money.
  - *adosc*: Accumulation/Distribution Oscillator, difference between fast and slow AD lines.
  - *cmf*: Chaikin Money Flow, combines price and volume to show buying/selling pressure.
  - *efi*: Elder Force Index, combines price change and volume to assess buying/selling power.
  - *kvo*: Klinger Volume Oscillator, identifies long-term money flow trends.
  - *mfi*: Money Flow Index, combines price and volume for overbought/oversold signals.

Appendix B – PRICE AND RETURNS FOR ETH, SOL, AND BNB

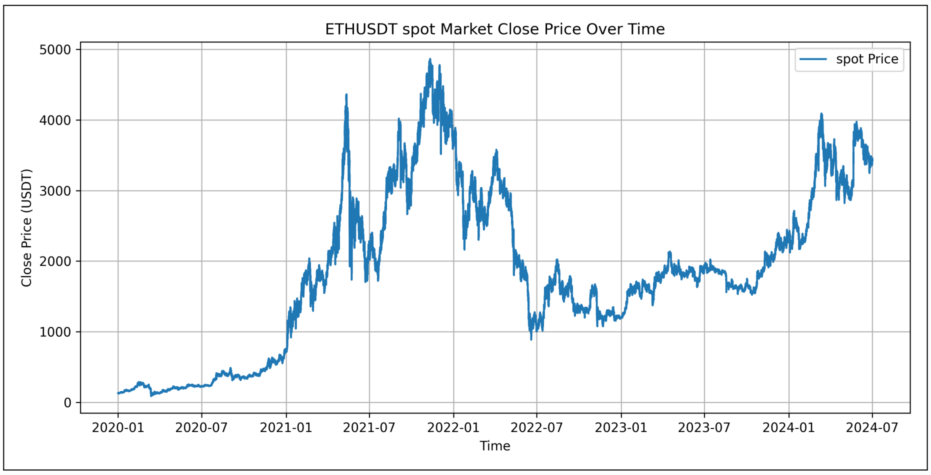

Figure 10 – ETHUSDT Spot

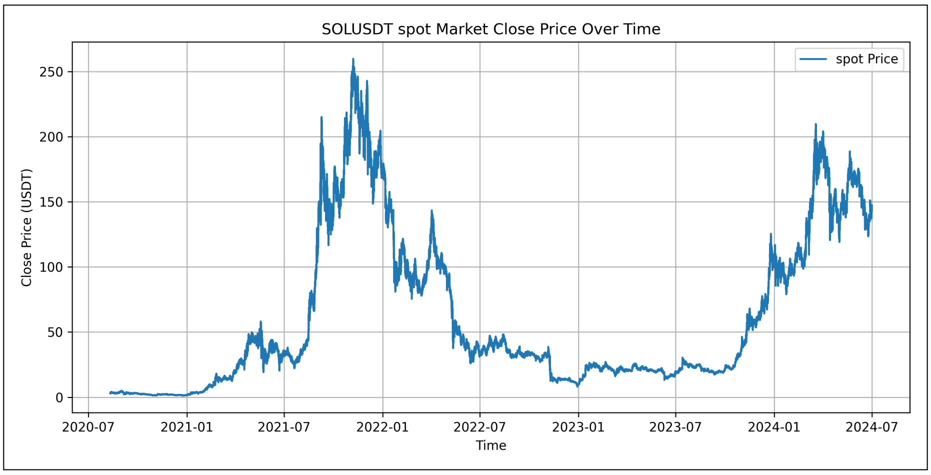

Figure 11 – SOLUSDT Spot

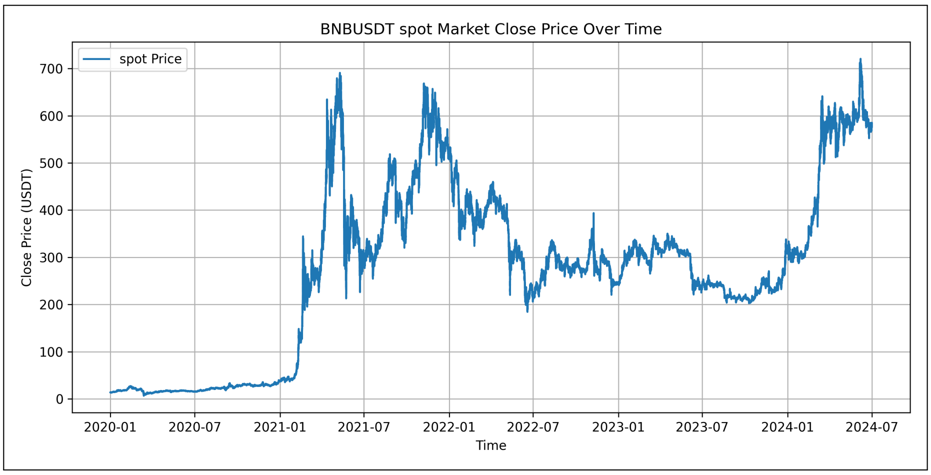

Figure 12 – BNBUSDT Spot

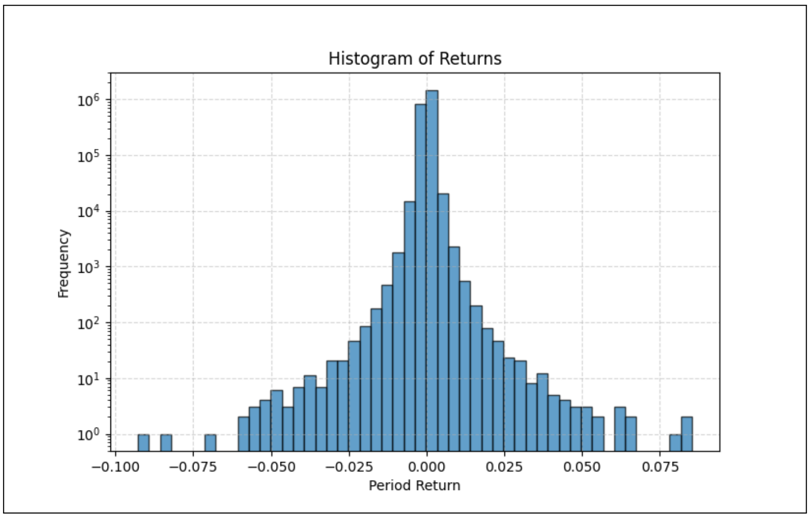

Figure 13 – ETHUSDT returns histogram

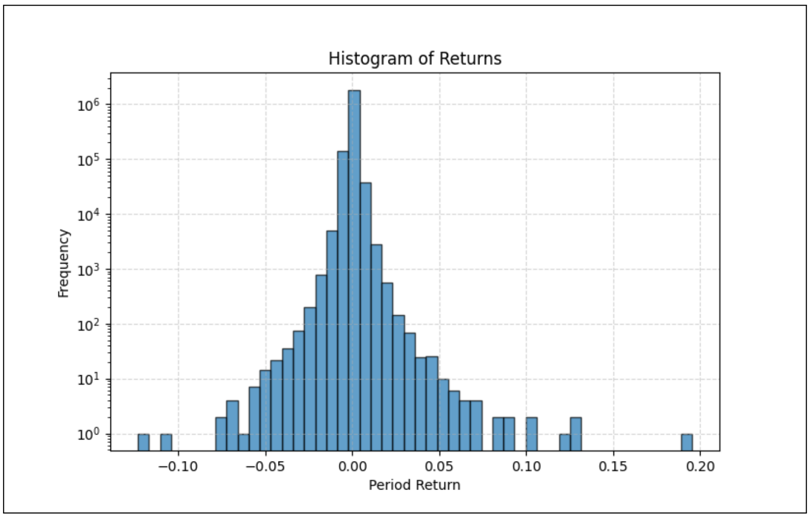

Figure 14 – SOLUSDT returns histogram

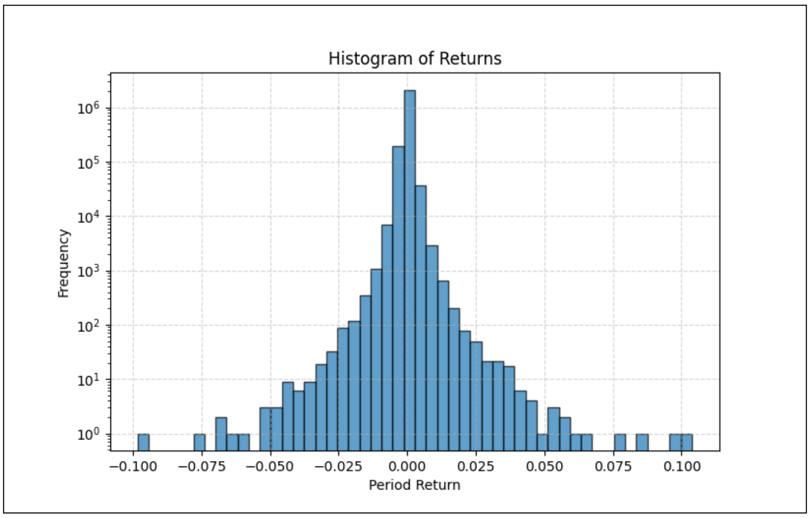

Figure 15 – BNBUSDT returns histogram

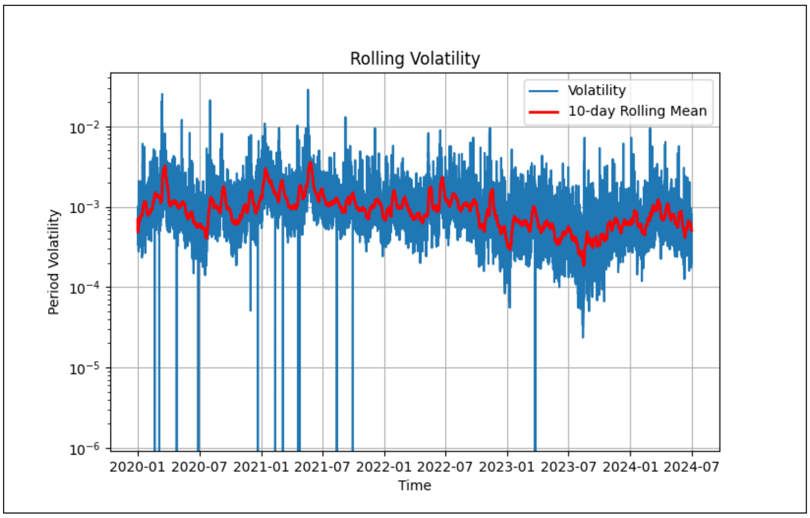

Figure 16 – ETHUSDT average volatility

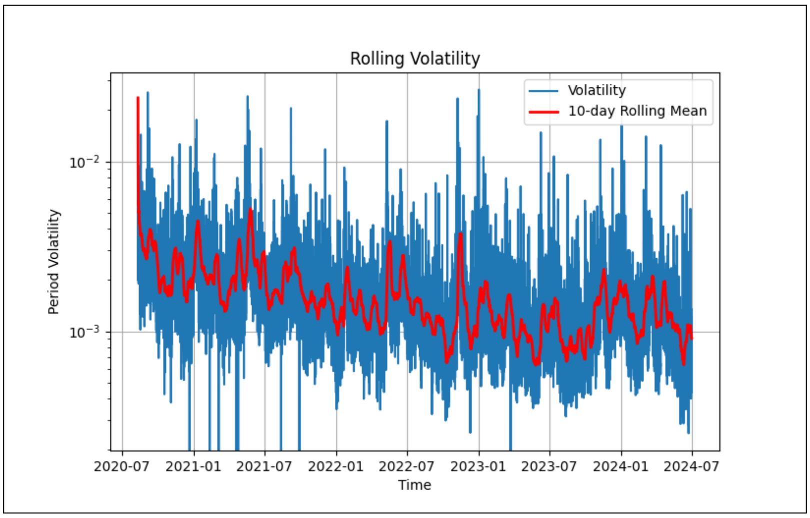

Figure 17 – SOLUSDT average volatility

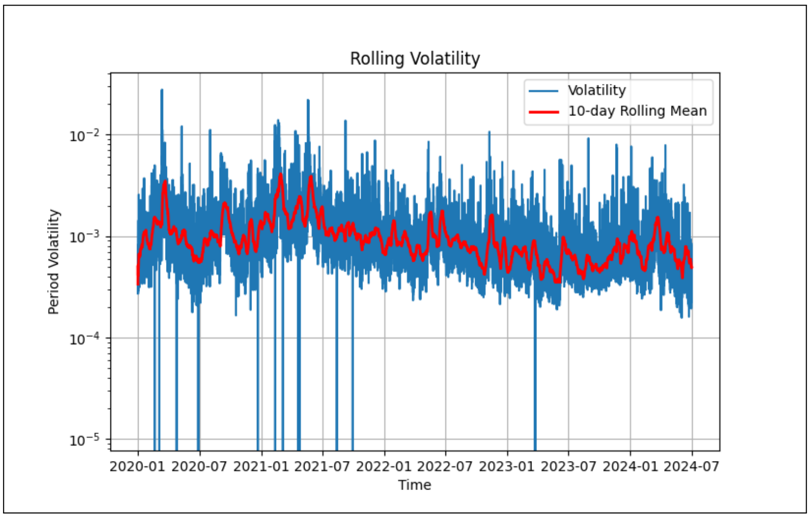

Figure 18 – BNBUSDT average volatility

Appendix C – STRATEGY RESULTS FOR SPOT

Table 11 – BTCUSDT / ta / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.215e-06	1.109e-06	1.027e-06
Rscore Train	-5.773e-14	-1.998e-14	2.384e-07
Mse Test	3.617e-07	3.733e-07	5.433e-07
Rscore Test	-3.043e-07	-5.375e-07	-7.153e-07

Table 12 – BTCUSDT / ta / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	4.072e-01	4.495e-01	4.772e-01	4.893e-01	1.086e+00
Annualized Return	1.136e+00	1.281e+00	1.379e+00	1.423e+00	4.121e+00
Sharpe Ratio	1.687e+00	1.810e+00	1.889e+00	1.923e+00	1.949e+00
Max Drawdown	-2.541e-01	-2.442e-01	-2.378e-01	-2.351e-01	-8.604e-01
Calmar Ratio	4.470e+00	5.246e+00	5.798e+00	6.053e+00	4.789e+00
Win Rate	2.742e-01	2.742e-01	2.742e-01	2.742e-01	5.088e-01
Num Trades	63	63	63	63	41619

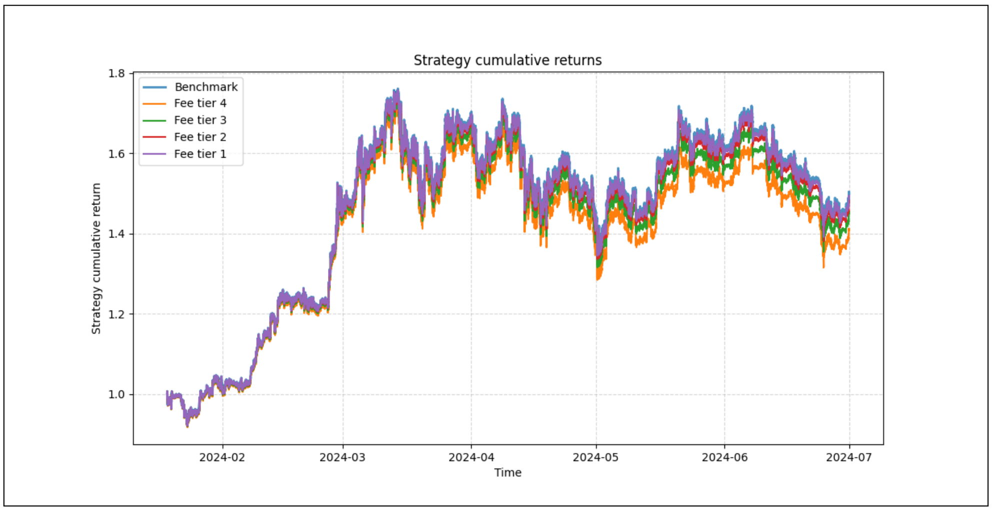

Figure 19 – BTCUSDT Spot TA Strats with fees

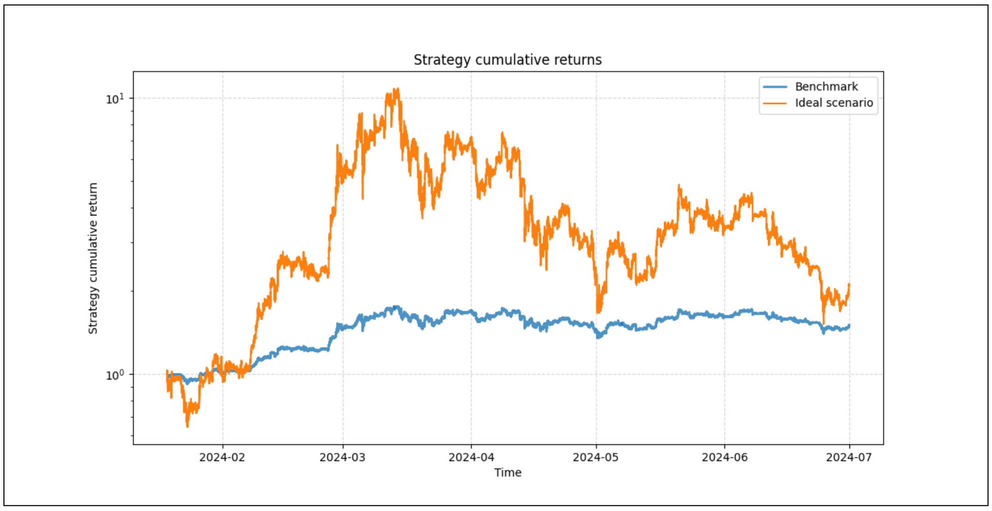

Figure 20 – BTCUSDT Spot TA Strats in Ideal Scenario

Table 13 – ETHUSDT / raw / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.930e-06	1.738e-06	1.574e-06
Rscore Train	1.317e-04	1.312e-04	1.181e-02
Mse Test	3.923e-07	4.275e-07	7.476e-07
Rscore Test	1.562e-05	6.894e-05	3.165e-04

Table 14 – ETHUSDT / raw / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	-4.923e-01	-2.914e-01	-1.759e-01	-1.203e-01	-7.487e-02
Annualized Return	-7.782e-01	-5.348e-01	-3.494e-01	-2.477e-01	-1.588e-01
Sharpe Ratio	-3.077e+00	-2.132e+00	-1.126e+00	-6.891e-01	-3.531e-01
Max Drawdown	-5.639e-01	-3.768e-01	-3.118e-01	-2.845e-01	-2.628e-01
Calmar Ratio	1.380e+00	1.419e+00	1.121e+00	8.707e-01	6.040e-01
Win Rate	2.952e-01	3.587e-01	3.877e-01	4.167e-01	4.275e-01
Num Trades	210	276	276	276	276

Table 15 – ETHUSDT / ta / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.928e-06	1.736e-06	1.582e-06
Rscore Train	1.309e-03	1.254e-03	6.503e-03
Mse Test	3.923e-07	4.275e-07	7.468e-07
Rscore Test	4.241e-05	1.905e-04	1.402e-03

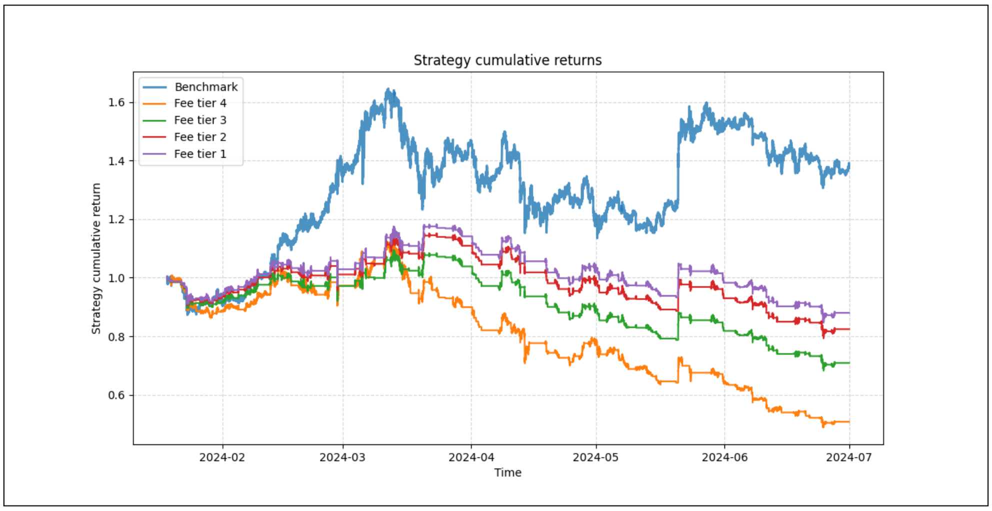

Figure 21 – ETHUSDT Spot Raw Strats with fees

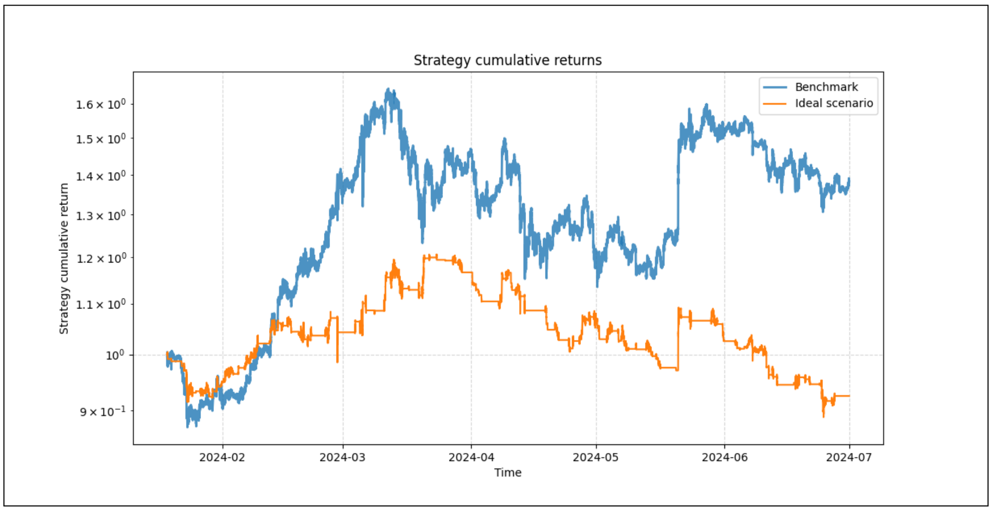

Figure 22 – ETHUSDT Spot Raw Strats in Ideal Scenario

Table 16 – ETHUSDT / ta / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	3.580e-02	2.018e-01	3.213e-01	3.768e-01	4.802e-02
Annualized Return	8.127e-02	5.043e-01	8.571e-01	1.035e+00	1.098e-01
Sharpe Ratio	4.008e-01	1.112e+00	1.566e+00	1.762e+00	5.205e-01
Max Drawdown	-3.198e-01	-2.894e-01	-2.693e-01	-2.604e-01	-3.403e-01
Calmar Ratio	2.541e-01	1.743e+00	3.183e+00	3.973e+00	3.227e-01
Win Rate	3.568e-01	3.668e-01	3.970e-01	4.121e-01	5.136e-01
Num Trades	200	200	200	200	403

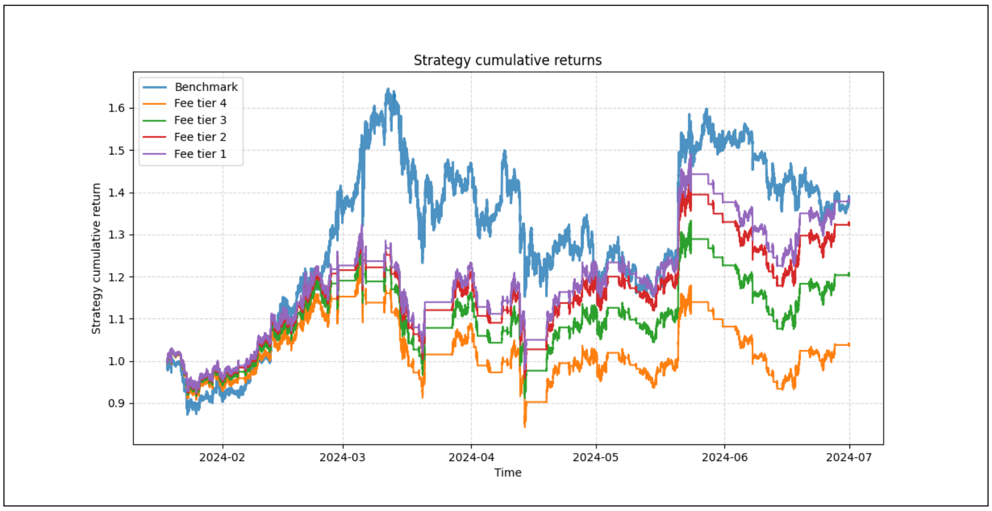

Figure 23 – ETHUSDT Spot TA Strats with fees

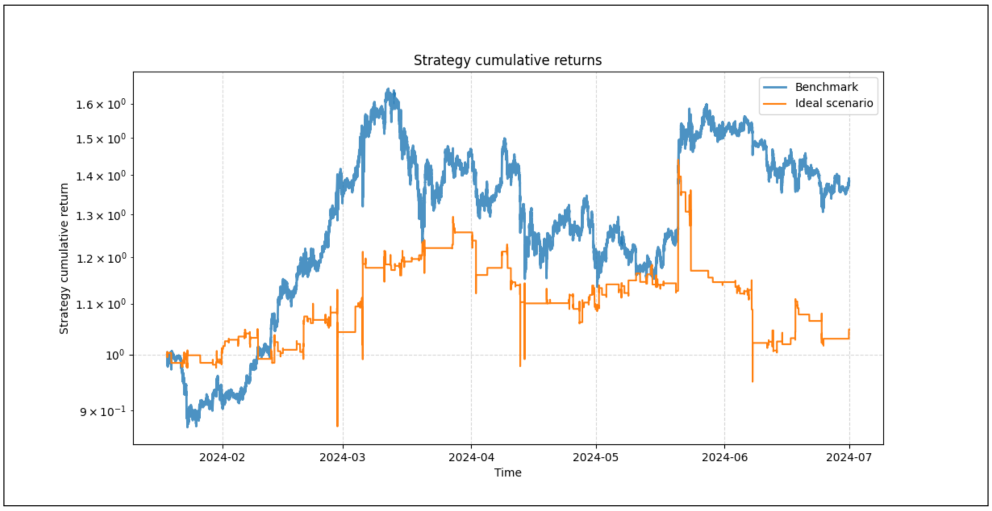

Figure 24 – ETHUSDT Spot TA Strats in Ideal Scenario

Table 17 – SOLUSDT / raw / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	4.537e-06	4.136e-06	3.930e-06
Rscore Train	8.033e-03	8.323e-03	1.488e-02
Mse Test	1.328e-06	2.533e-06	2.065e-06
Rscore Test	1.548e-02	1.445e-03	1.501e-03

Table 18 – SOLUSDT / raw / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	8.826e-02	1.733e-01	2.310e-01	1.474e-01	-6.912e-01
Annualized Return	2.496e-01	5.235e-01	7.289e-01	4.366e-01	-9.547e-01
Sharpe Ratio	6.789e-01	9.267e-01	1.085e+00	8.541e-01	4.035e-01
Max Drawdown	-3.907e-01	-3.611e-01	-3.414e-01	-3.277e-01	-9.026e-01
Calmar Ratio	6.387e-01	1.450e+00	2.135e+00	1.332e+00	1.058e+00
Win Rate	2.959e-01	2.959e-01	3.061e-01	5.068e-01	4.755e-01
Num Trades	98	98	98	221	25201

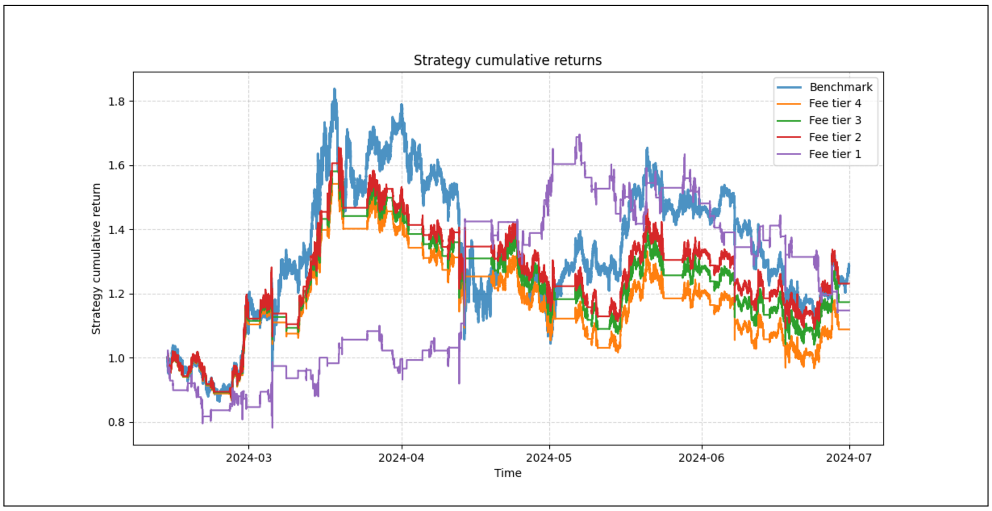

Figure 25 – SOLUSDT Spot Raw Strats with fees

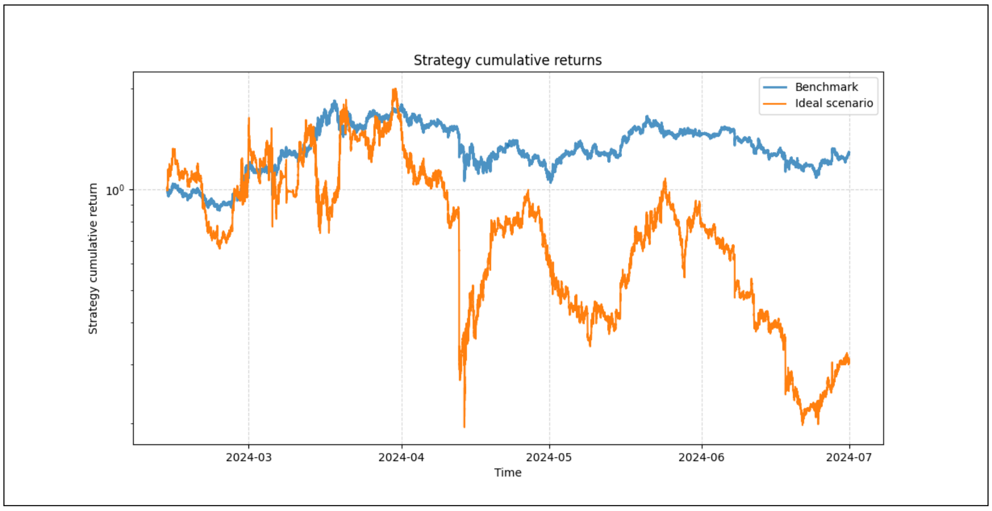

Figure 26 – SOLUSDT Spot Raw Strats in Ideal Scenario

Table 19 – SOLUSDT / ta / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	4.553e-06	4.151e-06	3.945e-06
Rscore Train	4.584e-03	4.760e-03	1.117e-02
Mse Test	1.331e-06	2.533e-06	2.064e-06
Rscore Test	1.335e-02	1.075e-03	1.763e-03

Table 20 – SOLUSDT / ta / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	-8.119e-01	-5.114e-01	-2.682e-01	-2.274e-01	2.649e+01
Annualized Return	-9.877e-01	-8.484e-01	-5.607e-01	-4.931e-01	6.178e+03
Sharpe Ratio	-5.844e+00	-4.311e+00	-1.764e+00	-2.706e-01	4.195e+00
Max Drawdown	-8.574e-01	-5.624e-01	-3.707e-01	-5.442e-01	-9.320e-01
Calmar Ratio	1.152e+00	1.508e+00	1.512e+00	9.061e-01	6.630e+03
Win Rate	1.006e-01	2.514e-01	3.851e-01	4.194e-01	4.834e-01
Num Trades	636	696	696	682	32295

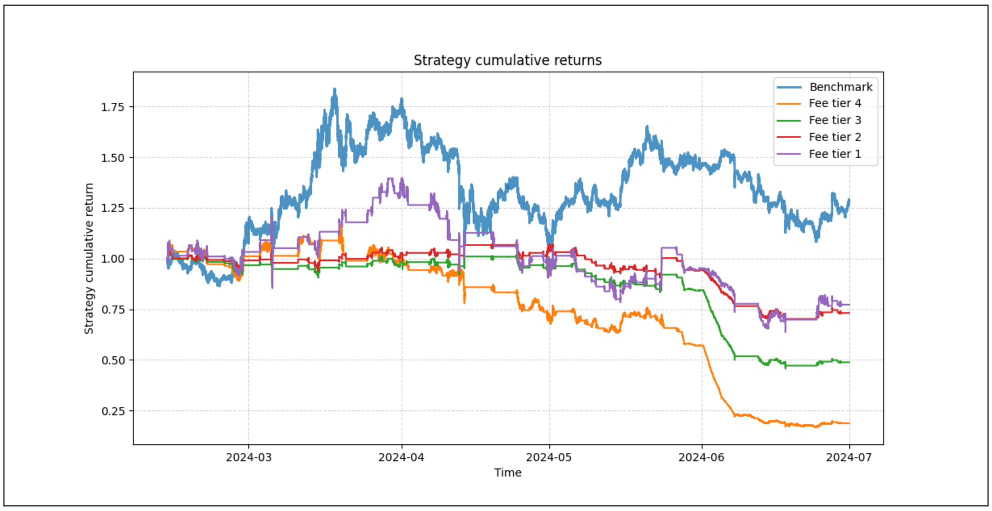

Figure 27 – SOLUSDT Spot TA Strats with fees

Table 21 – BNBUSDT / raw / spot (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.122e-06	1.909e-06	1.737e-06
Rscore Train	7.487e-03	8.489e-03	2.567e-02
Mse Test	4.182e-07	6.117e-07	7.694e-07
Rscore Test	4.473e-02	4.265e-02	8.122e-03

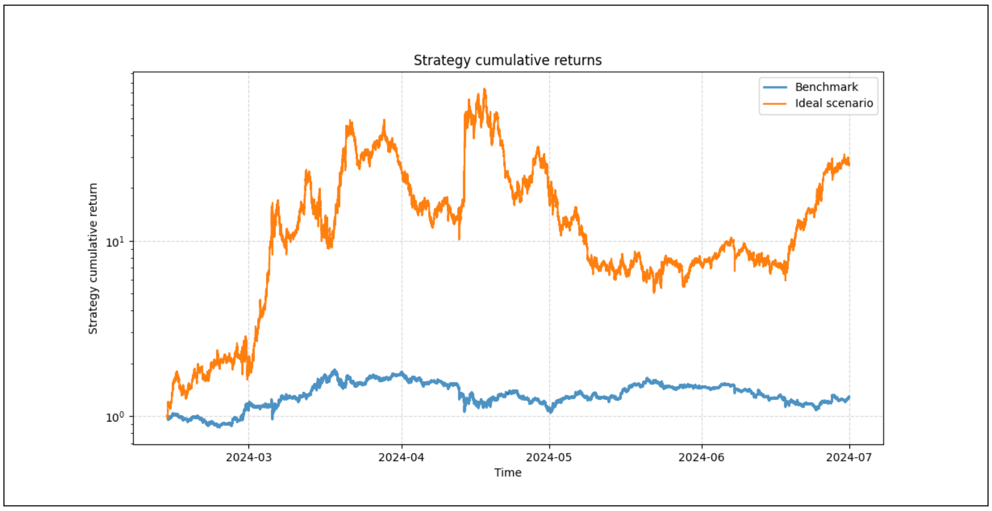

Figure 28 – SOLUSDT Spot TA Strats in Ideal Scenario

Table 22 – BNBUSDT / raw / spot (Test Strategies)

Metric	Strat tier 4	Strat tier 3	Strat tier 2	Strat tier 1	Strat Ideal
Total Return	-1.765e-01	3.746e-02	-2.241e-02	-9.949e-01	1.208e+07
Annualized Return	-3.574e-01	8.737e-02	-5.031e-02	-1.000e+00	1.347e+16
Sharpe Ratio	-1.039e+00	4.125e-01	1.471e-02	-1.066e+01	2.228e+01
Max Drawdown	-3.566e-01	-2.757e-01	-2.854e-01	-9.965e-01	-4.962e-01
Calmar Ratio	1.002e+00	3.169e-01	1.763e-01	1.003e+00	2.715e+16
Win Rate	2.662e-01	3.156e-01	3.885e-01	4.408e-01	5.118e-01
Num Trades	263	263	278	27911	28251

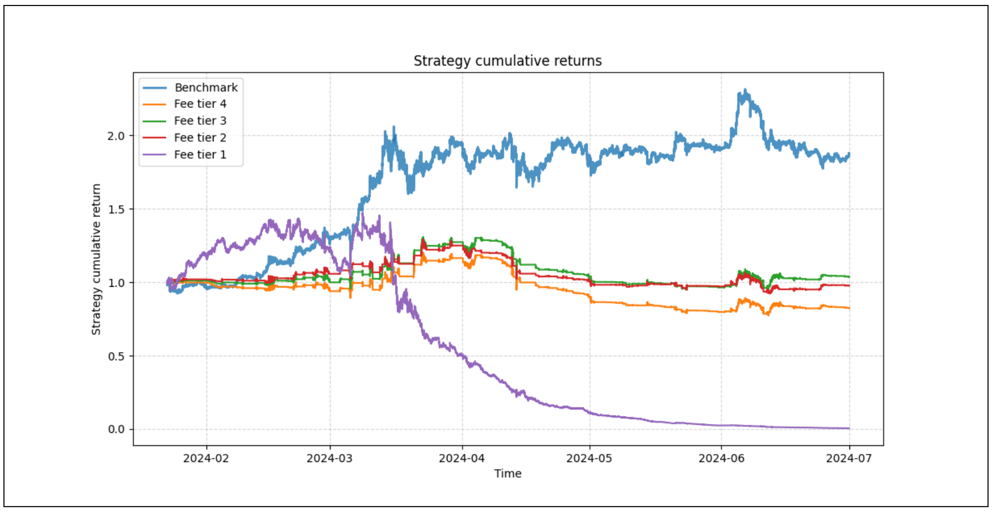

Figure 29 – BNBUSDT Spot Raw Strats with fees

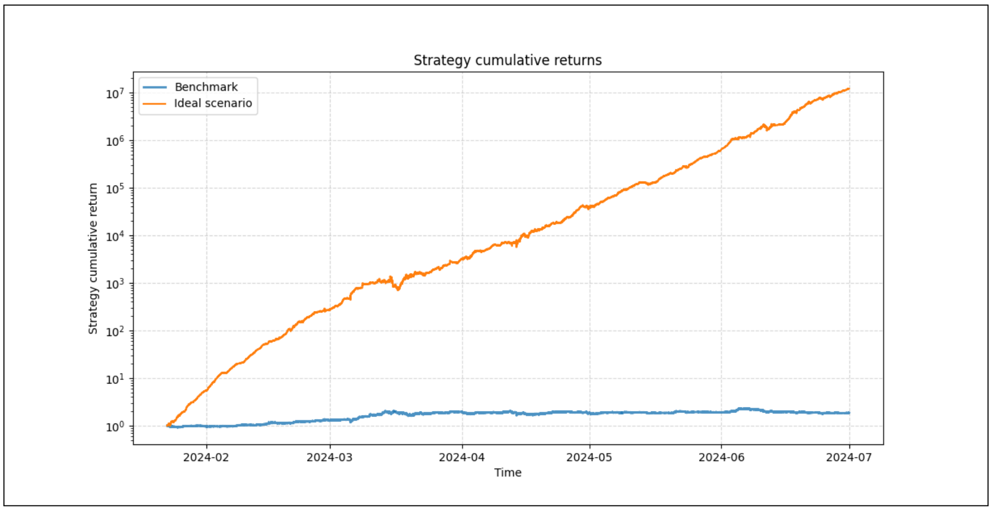

Figure 30 – BNBUSDT Spot Raw Strats in Ideal Scenario

Appendix D – STRATEGY RESULTS FOR FUTURES

Table 23 – BTCUSDT / raw / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.265e-06	1.155e-06	1.074e-06
Rscore Train	0.000e+00	-6.661e-16	-1.192e-07
Mse Test	3.885e-07	4.210e-07	5.710e-07
Rscore Test	-3.053e-07	-4.765e-07	-8.345e-07

Table 24 – BTCUSDT / raw / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	4.562e-01	4.757e-01	4.843e-01	4.904e-01	4.990e-01
Annualized Return	1.305e+00	1.374e+00	1.405e+00	1.427e+00	1.458e+00
Sharpe Ratio	1.798e+00	1.852e+00	1.876e+00	1.892e+00	1.916e+00
Max Drawdown	-2.442e-01	-2.395e-01	-2.374e-01	-2.360e-01	-2.339e-01
Calmar Ratio	5.342e+00	5.737e+00	5.916e+00	6.045e+00	6.233e+00
Win Rate	2.807e-01	2.807e-01	2.807e-01	2.807e-01	5.174e-01
Num Trades	58	58	58	58	3559

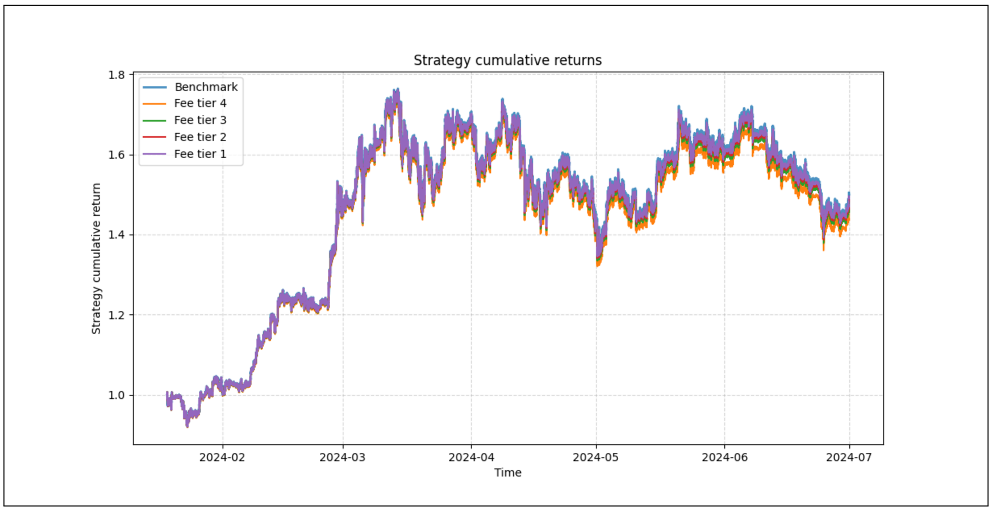

Figure 31 – BTCUSDT Future Raw Strats with fees

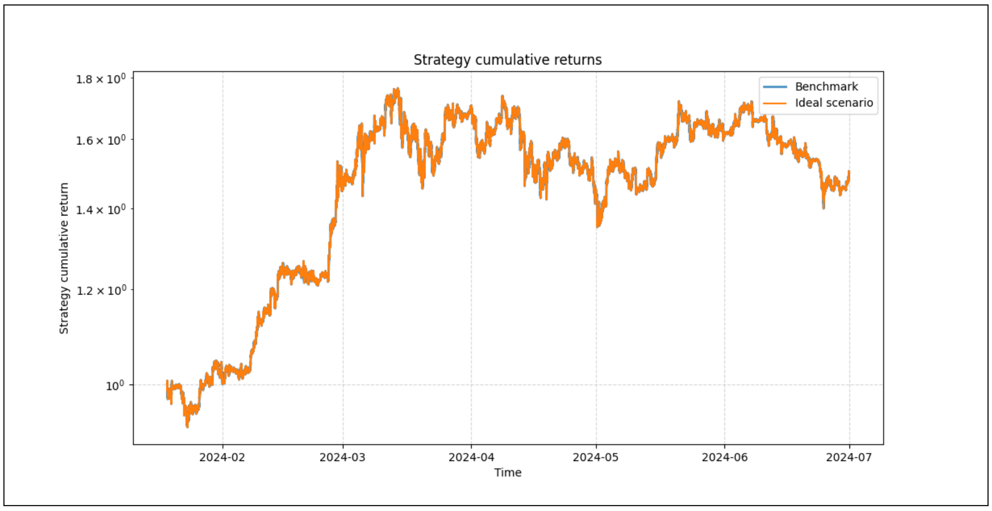

Figure 32 – BTCUSDT Future Raw Strats in Ideal Scenario

Table 25 – BTCUSDT / ta / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	1.264e-06	1.155e-06	1.057e-06
Rscore Train	4.601e-04	4.536e-04	1.562e-02
Mse Test	3.884e-07	4.209e-07	5.704e-07
Rscore Test	2.190e-04	2.409e-04	1.106e-03

Table 26 – BTCUSDT / ta / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	-5.779e-01	-4.584e-01	-3.964e-01	-3.489e-01	1.298e+01
Annualized Return	-8.528e-01	-7.439e-01	-6.742e-01	-6.145e-01	3.499e+02
Sharpe Ratio	-1.524e+00	-9.458e-01	-6.943e-01	-5.182e-01	3.688e+00
Max Drawdown	-7.047e-01	-6.502e-01	-6.356e-01	-6.258e-01	-6.719e-01
Calmar Ratio	1.210e+00	1.144e+00	1.061e+00	9.819e-01	5.207e+02
Win Rate	2.907e-01	2.996e-01	2.996e-01	3.040e-01	5.332e-01
Num Trades	227	227	227	227	26993

Table 27 – ETHUSDT / raw / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.114e-06	1.903e-06	1.747e-06
Rscore Train	0.000e+00	-2.820e-14	-1.192e-07
Mse Test	4.251e-07	4.963e-07	7.881e-07
Rscore Test	-8.801e-06	-1.330e-06	-1.192e-07

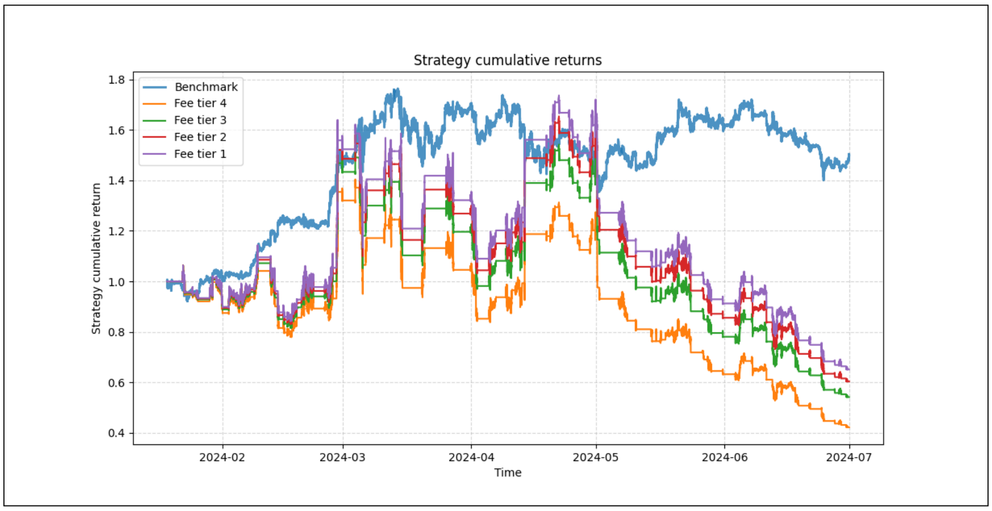

Figure 33 – BTCUSDT Future TA Strats with fees

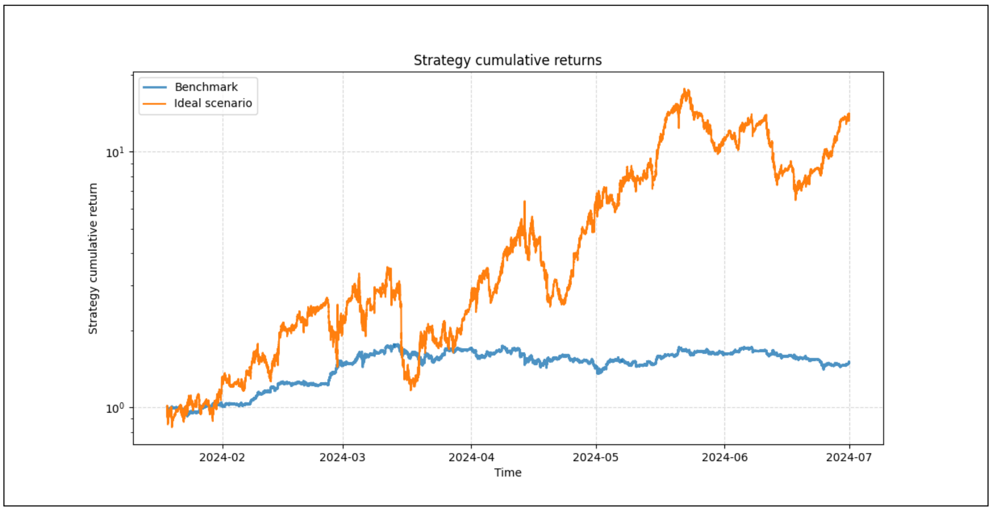

Figure 34 – BTCUSDT Future TA Strats in Ideal Scenario

Table 28 – ETHUSDT / raw / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	3.272e-01	3.522e-01	3.632e-01	3.709e-01	3.821e-01
Annualized Return	8.755e-01	9.548e-01	9.903e-01	1.016e+00	1.052e+00
Sharpe Ratio	1.299e+00	1.364e+00	1.392e+00	1.411e+00	1.439e+00
Max Drawdown	-3.246e-01	-3.188e-01	-3.163e-01	-3.145e-01	-3.120e-01
Calmar Ratio	2.697e+00	2.995e+00	3.131e+00	3.229e+00	3.373e+00
Win Rate	3.125e-01	3.125e-01	3.125e-01	3.125e-01	5.157e-01
Num Trades	81	81	81	81	4645

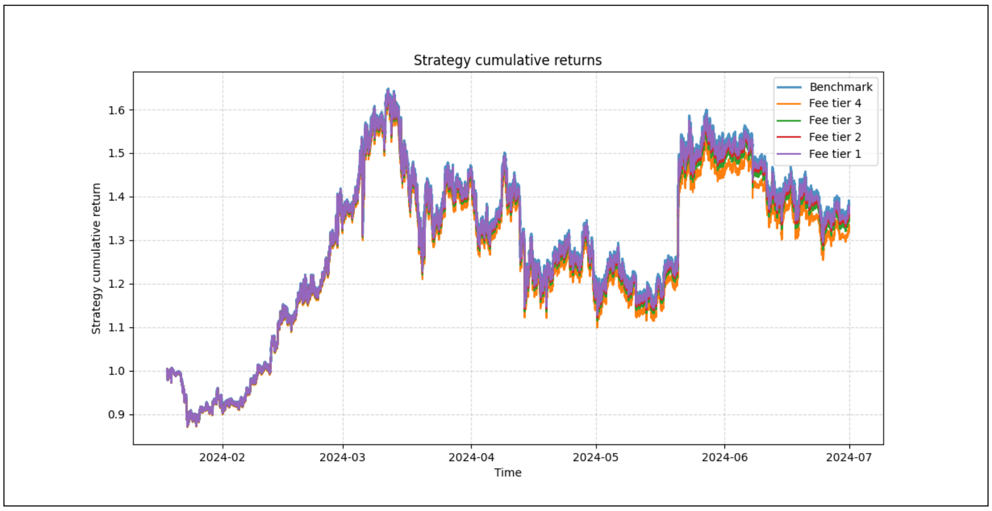

Figure 35 – ETHUSDT Future Raw Strats with fees

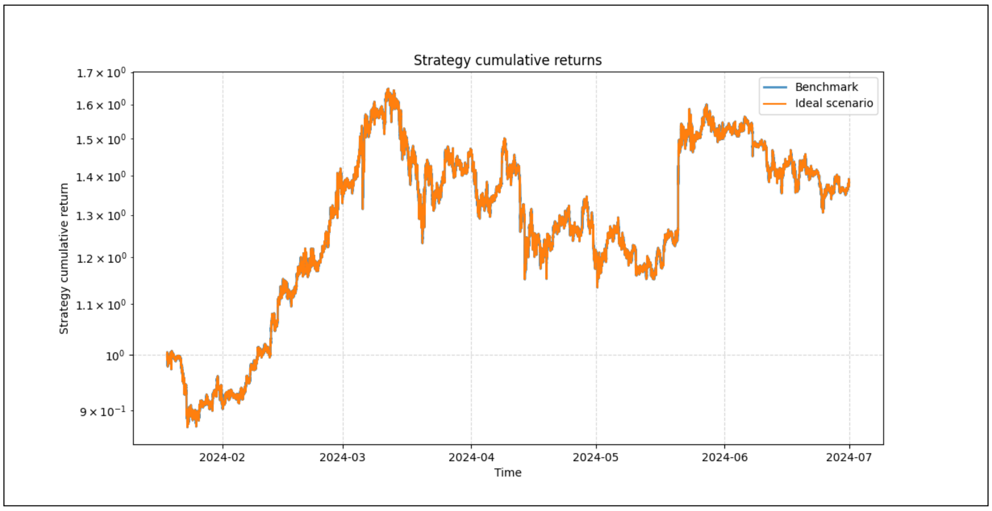

Figure 36 – ETHUSDT Future Raw Strats in Ideal Scenario

Table 29 – ETHUSDT / ta / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.114e-06	1.903e-06	1.746e-06
Rscore Train	0.000e+00	-2.820e-14	5.546e-04
Mse Test	4.251e-07	4.963e-07	7.881e-07
Rscore Test	-8.801e-06	-1.330e-06	6.592e-05

Table 30 – ETHUSDT / ta / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	2.975e-01	3.225e-01	3.336e-01	3.413e-01	3.525e-01
Annualized Return	7.836e-01	8.609e-01	8.955e-01	9.201e-01	9.559e-01
Sharpe Ratio	1.221e+00	1.287e+00	1.316e+00	1.336e+00	1.364e+00
Max Drawdown	-3.369e-01	-3.312e-01	-3.288e-01	-3.270e-01	-3.245e-01
Calmar Ratio	2.326e+00	2.599e+00	2.724e+00	2.814e+00	2.946e+00
Win Rate	3.171e-01	3.171e-01	3.171e-01	3.171e-01	5.153e-01
Num Trades	83	83	83	83	4645

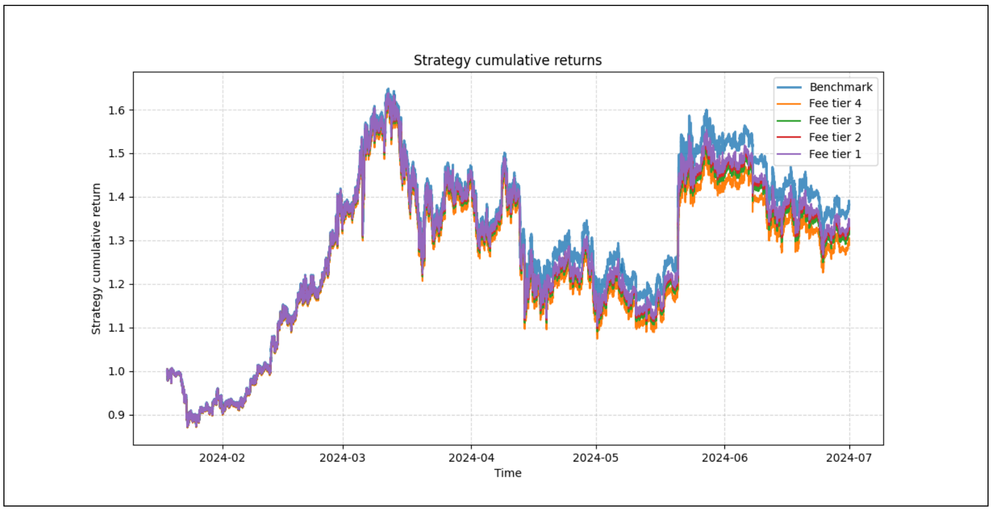

Figure 37 – ETHUSDT Future TA Strats with fees

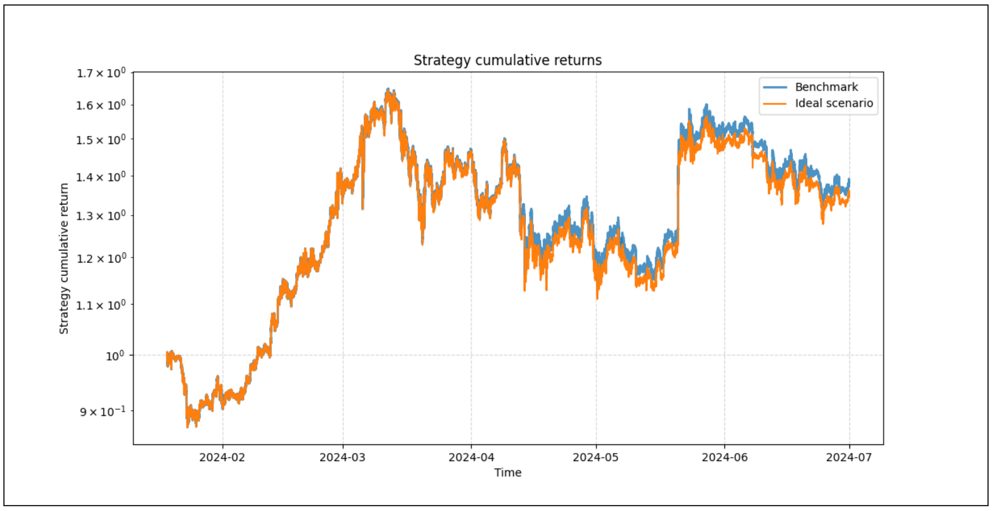

Figure 38 – ETHUSDT Future TA Strats in Ideal Scenario

Table 31 – SOLUSDT / raw / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	4.665e-06	4.251e-06	4.058e-06
Rscore Train	1.011e-03	9.480e-04	4.266e-03
Mse Test	1.354e-06	2.637e-06	2.156e-06
Rscore Test	-5.095e-05	-9.838e-06	2.942e-04

Table 32 – SOLUSDT / raw / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	1.649e-02	1.007e-01	1.395e-01	1.674e-01	2.085e-01
Annualized Return	4.403e-02	2.876e-01	4.105e-01	5.034e-01	6.469e-01
Sharpe Ratio	3.229e-01	7.839e-01	9.842e-01	1.124e+00	1.324e+00
Max Drawdown	-1.879e-01	-1.832e-01	-1.812e-01	-1.797e-01	-1.777e-01
Calmar Ratio	2.343e-01	1.570e+00	2.266e+00	2.801e+00	3.641e+00
Win Rate	5.543e-01	5.543e-01	5.543e-01	5.543e-01	5.543e-01
Num Trades	184	184	184	184	184

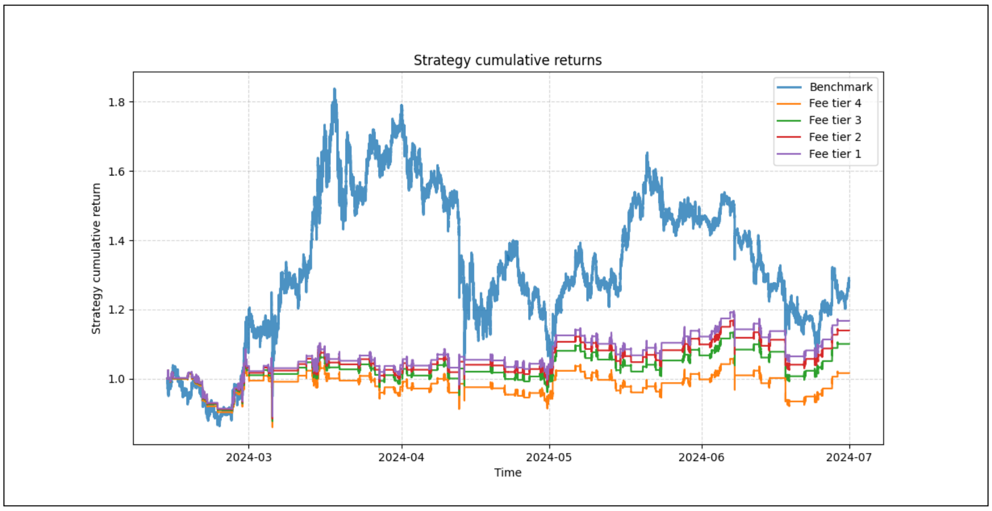

Figure 39 – SOLUSDT Future Raw Strats with fees

Table 33 – SOLUSDT / ta / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	4.668e-06	4.253e-06	4.058e-06
Rscore Train	4.569e-04	4.496e-04	4.357e-03
Mse Test	1.354e-06	2.637e-06	2.156e-06
Rscore Test	-1.635e-05	-2.489e-05	2.680e-04

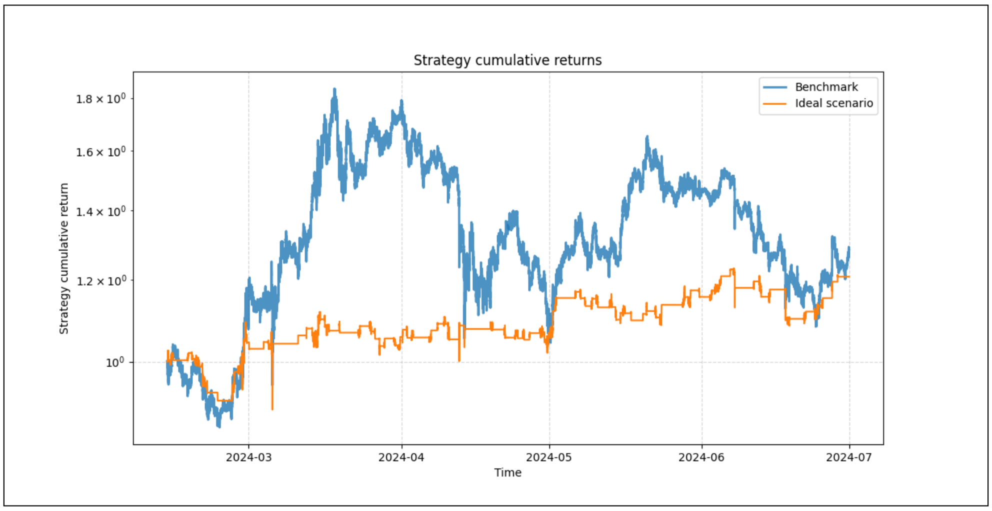

Figure 40 – SOLUSDT Future Raw Strats in Ideal Scenario

Table 34 – SOLUSDT / ta / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	-3.053e-01	-2.679e-01	-2.386e-01	-2.002e-01	-8.311e-01
Annualized Return	-6.170e-01	-5.602e-01	-5.123e-01	-4.449e-01	-9.908e-01
Sharpe Ratio	-1.104e+00	-8.968e-01	-3.458e-01	-2.016e-01	-2.206e+00
Max Drawdown	-4.820e-01	-4.640e-01	-4.268e-01	-4.173e-01	-9.022e-01
Calmar Ratio	1.280e+00	1.207e+00	1.200e+00	1.066e+00	1.098e+00
Win Rate	2.441e-01	2.520e-01	3.263e-01	3.316e-01	2.736e-01
Num Trades	128	128	191	191	837

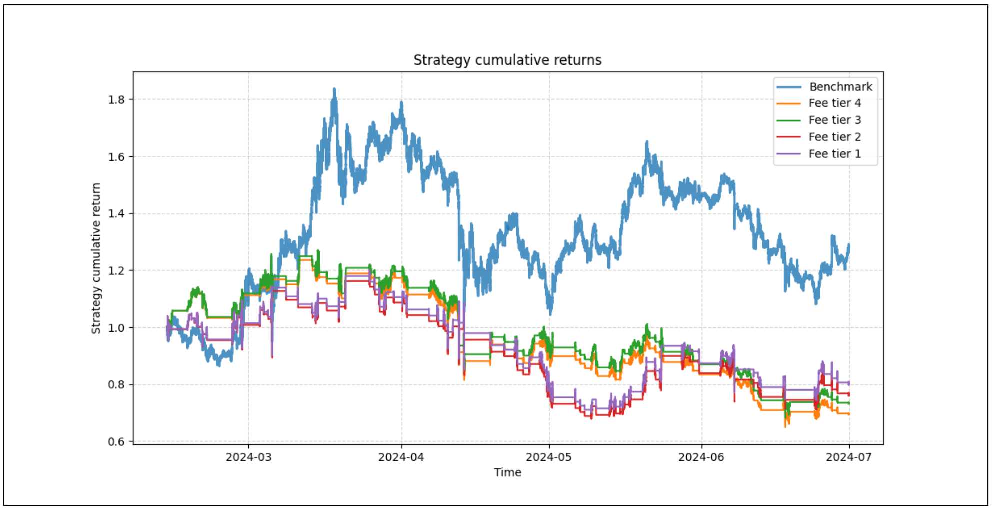

Figure 41 – SOLUSDT Future TA Strats with fees

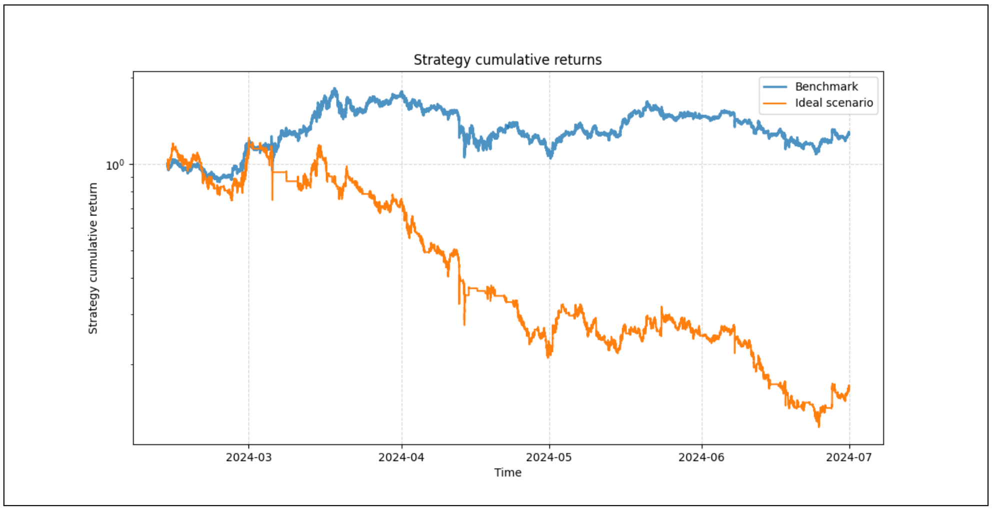

Figure 42 – SOLUSDT Future TA Strats in Ideal Scenario

Table 35 – BNBUSDT / raw / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.315e-06	2.080e-06	1.912e-06
Rscore Train	2.799e-05	2.700e-05	2.030e-03
Mse Test	4.342e-07	6.028e-07	7.886e-07
Rscore Test	2.774e-06	1.578e-05	4.554e-05

Table 36 – BNBUSDT / raw / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	2.353e-01	3.744e-01	4.397e-01	-2.296e-01	2.056e-01
Annualized Return	6.180e-01	1.063e+00	1.293e+00	-4.480e-01	5.309e-01
Sharpe Ratio	1.409e+00	2.020e+00	2.286e+00	-1.985e-01	9.284e-01
Max Drawdown	-2.177e-01	-1.911e-01	-1.828e-01	-4.834e-01	-4.397e-01
Calmar Ratio	2.838e+00	5.564e+00	7.075e+00	9.266e-01	1.208e+00
Win Rate	3.676e-01	4.228e-01	4.596e-01	4.652e-01	4.877e-01
Num Trades	273	273	273	488	488

Table 37 – BNBUSDT / ta / um (Model Metrics)

Metric	Model 1	Model 2	Model 3
Mse Train	2.315e-06	2.080e-06	1.914e-06
Rscore Train	6.029e-05	5.810e-05	8.977e-04
Mse Test	4.342e-07	6.028e-07	7.887e-07
Rscore Test	-8.560e-07	6.227e-06	-1.098e-04

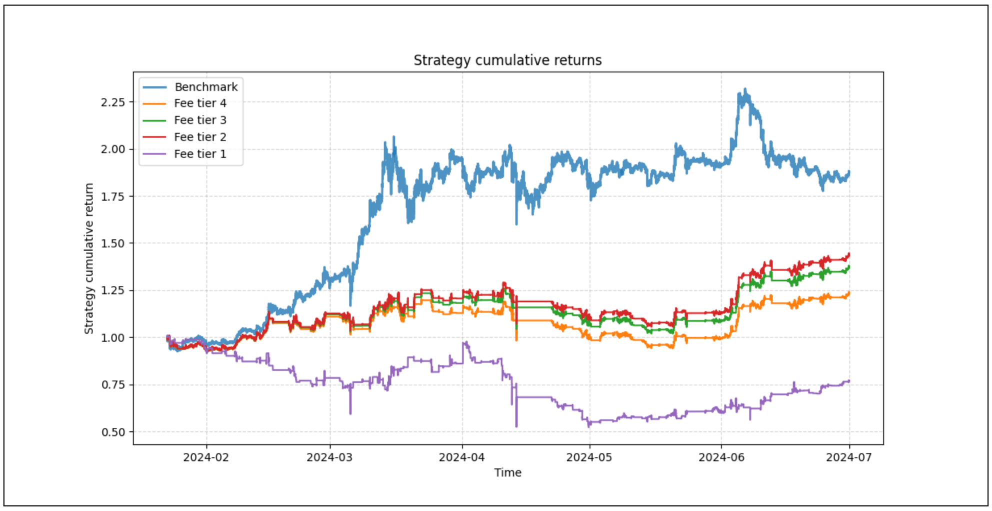

Figure 43 – BNBUSDT Future Raw Strats with fees

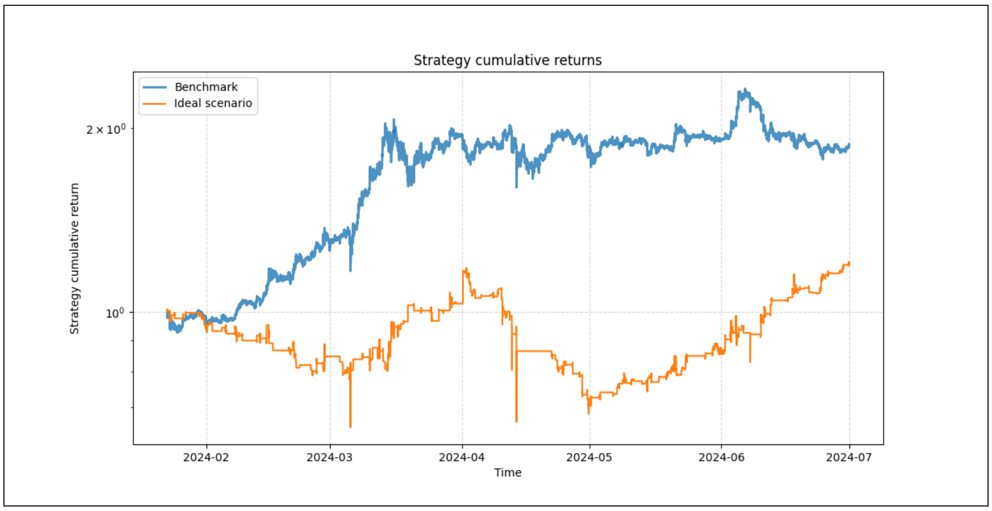

Figure 44 – BNBUSDT Future Raw Strats in Ideal Scenario

Table 38 – BNBUSDT / ta / um (Test Strategies)

Metric	Strategy 0	Strategy 1	Strategy 2	Strategy 3	Strategy 4
Total Return	4.740e-02	7.104e-02	4.536e-01	1.542e-01	2.530e-01
Annualized Return	1.112e-01	1.692e-01	1.344e+00	3.861e-01	6.714e-01
Sharpe Ratio	4.612e-01	5.625e-01	1.470e+00	8.176e-01	1.096e+00
Max Drawdown	-3.097e-01	-3.022e-01	-3.412e-01	-2.856e-01	-2.686e-01
Calmar Ratio	3.591e-01	5.598e-01	3.939e+00	1.352e+00	2.500e+00
Win Rate	2.321e-01	2.321e-01	4.245e-01	3.778e-01	4.966e-01
Num Trades	57	57	212	270	294

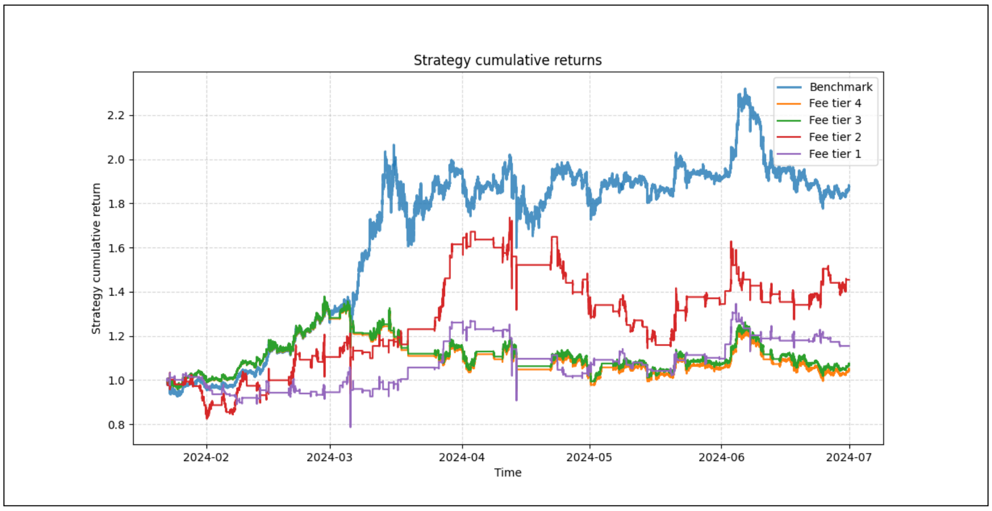

Figure 45 – BNBUSDT Future TA Strats with fees

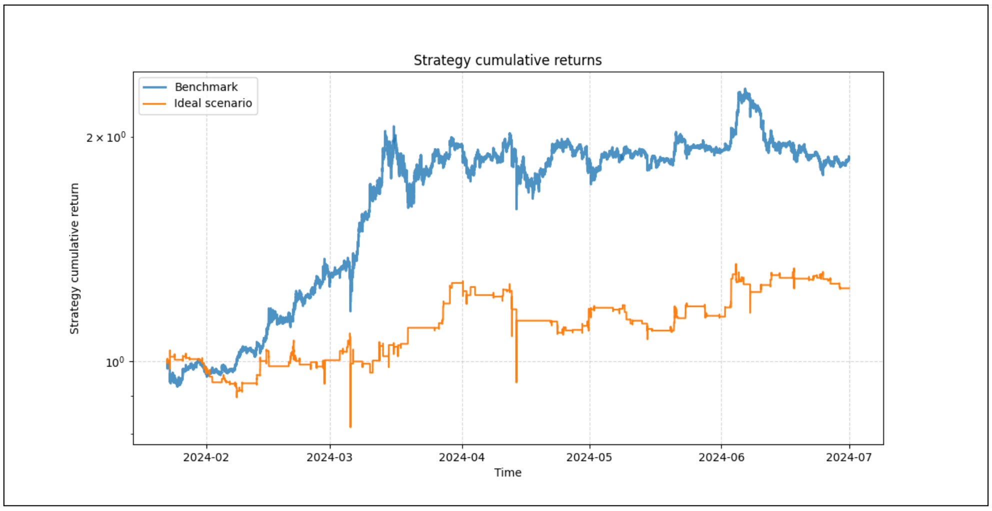

Figure 46 – BNBUSDT Future TA Strats in Ideal Scenario