Kaunas University of Technology
Faculty of Mathematics and Natural Sciences

Explainable Artificial Intelligence for the Analysis of Cryptocurrency Market Dynamics

Master's Final Project

Project prepared by
Augustė Žibėnaitė

Project supervisor
Prof. Dr. Robertas Alzbutas

Doc. Dr. Lina Sinevičienė

Kaunas, 2026

Kaunas University of Technology

Faculty of Mathematics and Natural Sciences

Explainable Artificial Intelligence for the Analysis of Cryptocurrency Market Dynamics

Master's final project

Business Big Data Analytics (6213AX001)

Project prepared by

Augustė Žibėnaitė

Project supervisor

Prof. Dr. Robertas Alzbutas

Project supervisor

Doc. Dr. Lina Sinevičienė

Project reviewed by

Doc. Dr. Šviesa Leitonienė

Project reviewed by

Doc. Dr. Mantas Landauskas

Kaunas, 2026

Kaunas University of Technology

Faculty of Mathematics ant Natural Sciences

Augustė Žibėnaitė

Explainable Artificial Intelligence for the Analysis of Cryptocurrency Market Dynamics

Declaration of academic integrity

I confirm that:

1. I have prepared this final project independently and honestly, without violating the copyright or other rights of other persons, in compliance with the Law on Copyright and Related Rights of the Republic of Lithuania, the provisions of Kaunas University of Technology (hereinafter referred to as the University) on the management and transfer of intellectual property, and the ethical requirements set out in the University's Code of Academic Ethics;

2. all data and research results presented in the final project are accurate and obtained legally, no part of the project has been plagiarised from printed or electronic sources, and all quotations and references presented in the text of the final project are indicated in the list of references;

3. I have complied with personal data protection requirements in the final project, have not used undisclosed or confidential data without legal grounds, and if I have used such data, it has been properly anonymised;

4. If I used artificial intelligence (hereinafter referred to as AI) or other automated tools in preparing the final project, I applied them in accordance with the procedures established by the University, without violating the principles of academic integrity.

5. I have not paid and am not obliged to pay any sums of money not provided for by law for the final project or any part thereof to any natural or legal person;

6. I understand that if academic dishonesty or violation of other persons' rights is discovered, I will be held accountable in accordance with the procedures established by the University

and may be expelled from the University; cases of academic dishonesty may be investigated even after graduation, initiating the procedure for revoking the degree.

Summary

Žibėnaitė, Augustė. Explainable Artificial Intelligence for the Analysis of Cryptocurrency Market Dynamics. Master's final project / project supervised by prof. dr. Robertas Alzbutas, doc. dr. Lina Sinevičienė; Kaunas University of Technology, Faculty of Mathematics and Natural Sciences.

Field of study and group of fields of study: Applied Mathematics, Mathematical Sciences.

Keywords: cryptocurrency markets, explainable artificial intelligence, Random Forest, SHAP, Hidden Markov Model, backtesting, market regimes.

Kaunas, 2026. 85 p.

Cryptocurrency markets are difficult to analyse because they are highly volatile, continuously traded, sensitive to news and affected by changing market conditions. Although machine learning models are increasingly used for cryptocurrency prediction, their value is limited when model decisions cannot be interpreted. Therefore, this thesis focuses not only on predicting market direction, but also on explaining which variables influence model decisions and whether predictions have economic meaning.

The object of the thesis is the market dynamics of Bitcoin, Ethereum and Solana. The aim is to evaluate machine learning models for analysing cryptocurrency market dynamics using explainable artificial intelligence methods. The literature review examines cryptocurrency market behaviour, price drivers, volatility, regime changes, news sentiment, machine learning-based prediction, explainability methods and economic evaluation of trading signals.

The methodology develops an empirical framework based on daily OHLCV market data, derived technical indicators, crypto-specific and macroeconomic news variables, macro-financial indicators and HMM-based market-state features. Random Forest classifiers are used to predict UP/DOWN market direction, while feature importance, SHAP analysis, PDP/ICE plots and backtesting are applied to interpret the models from technical and economic perspectives.

The results show that the Random Forest models achieved above-random directional prediction performance for all three cryptocurrencies. Bitcoin produced a weaker but meaningful result, with accuracy and macro F1 of about 0.63, while Ethereum and Solana achieved stronger results of about 0.75-0.76. Explainability analysis showed that recent asset-specific return variables were the most important predictors. Positive recent returns generally increased the predicted probability of an UP movement, while negative returns reduced it, indicating short-term momentum-like behaviour.

External variables, including news sentiment, macroeconomic indicators and HMM-based states, had weaker direct influence on model predictions than return-based variables, although they remained useful for market-context interpretation. Backtesting showed that model probabilities could be transformed into economically interpretable trading signals. The strategies did not consistently outperform buy-and-hold, but sometimes helped reduce exposure during weaker market periods. Therefore, the practical value of the thesis lies in applying explainable models to support interpretation of price-change signals, risk conditions and position-management decisions rather than in creating a universally profitable trading system.

Santrauka

Žibėnaitė, Augustė. Paaiškinamasis dirbtinis intelektas kriptovaliutų rinkos dinamikai analizuoti. Magistro studijų baigiamasis projektas / vadovas prof. dr. Robertas Alzbutas, vadovė doc. dr. Lina Sinevičienė ; Kauno technologijos universitetas, Matematikos ir gamtos mokslų fakultetas.

Studijų kryptis ir sritis (studijų krypčių grupė): Taikomoji matematika (Matematikos mokslai).

Reikšminiai žodžiai: kriptovaliutų rinkos, paaiškinamasis dirbtinis intelektas, atsitiktinis miškas, SHAP, paslėptasis Markovo modelis, grįžtamasis testavimas, rinkos režimai.

Kaunas, 2026. 85 p.

Kriptovaliutų rinkas sudėtinga analizuoti dėl didelio kintamumo, nepertraukiamos prekybos, jautrumo naujienoms ir besikeičiančių rinkos sąlygų. Nors mašininio mokymosi modeliai vis dažniau taikomi kriptovaliutų prognozavimui, jų vertė yra ribota, jei modelių sprendimai nėra interpretuojami. Todėl šiame darbe dėmesys skiriamas ne tik rinkos krypties prognozavimui, bet ir veiksnių, darančių įtaką modelių sprendimams, paaiškinimui bei ekonominės prognozių reikšmės vertinimui.

Darbo objektas yra Bitcoin, Ethereum ir Solana kriptovaliutų rinkos dinamika. Darbo tikslas - įvertinti mašininio mokymosi modelius, skirtus kriptovaliutų rinkos dinamikai analizuoti, taikant paaiškinamojo dirbtinio intelekto metodus. Literatūros analizėje nagrinėjama kriptovaliutų rinkos elgsena, kainos veiksniai, kintamumas, režimų kaita, naujienų sentimentas, mašininiu mokymusi grįstas prognozavimas, paaiškinamumo metodai ir prekybos signalų ekonominis vertinimas.

Metodologinėje dalyje parengiama empirinė tyrimo sistema, jungianti dieninius OHLCV rinkos duomenis, techninius rodiklius, naujienų kintamuosius, makrofinansinius rodiklius ir HMM pagrindu nustatytas rinkos būsenas. UP/DOWN rinkos krypčiai prognozuoti taikomi atsitiktinio miško klasifikatoriai, o modelių interpretavimui naudojama požymių svarba, SHAP analizė, PDP/ICE grafikai ir grįžtamasis testavimas.

Rezultatai parodė, kad visi trys modeliai pasiekė geresnį nei atsitiktinis krypties prognozavimo rezultatą. Bitcoin modelio tikslumas ir macro F1 siekė apie 0,63, o Ethereum ir Solana - apie 0,75-0,76. Paaiškinamumo analizė parodė, kad svarbiausi prognozavimo veiksniai buvo naujausi konkretaus aktyvo grąžos kintamieji. Teigiamos grąžos dažniausiai didino UP judėjimo tikimybę, o neigiamos ją mažino, todėl modeliai daugiausia fiksavo trumpalaikį momentum tipo elgesį.

Išoriniai kintamieji, įskaitant naujienų sentimentą, makroekonominius rodiklius ir HMM būsenas, turėjo silpnesnę tiesioginę įtaką prognozėms nei grąžos kintamieji, tačiau buvo naudingi rinkos konteksto interpretavimui. Grįžtamasis testavimas parodė, kad modelių tikimybės gali būti paverčiamos ekonomiškai interpretuojamais prekybos signalais. Strategijos ne visada pranoko pirkimo ir laikymo metodą, tačiau kai kuriais atvejais padėjo sumažinti poziciją silpnesniais rinkos laikotarpiais. Todėl praktinė darbo vertė yra ne universaliai pelningos prekybos sistemos sukūrimas, o paaiškinamų modelių taikymas kainos pokyčių signalams, rizikai ir pozicijos valdymo sprendimams interpretuoti.

Contents

List of tables

List of figures

Introduction

Investors, researchers, regulators, and tech-focused companies have all shown interest in cryptocurrency markets, which have grown to represent a significant component of the contemporary financial landscape. Cryptocurrencies, in contrast to conventional financial assets, are constantly traded, highly sensitive to market emotion, and frequently undergo abrupt shifts in price direction and volatility. Because of these features, it is challenging to assess the evolution of cryptocurrency market using solely conventional financial indicators. Simultaneously, there are chances to use machine learning techniques to find patterns in cryptocurrency behaviour due to the growing availability of market data, blockchain-related information, macroeconomic indicators, and news emotion.

However, another issue arises with using machine learning in financial market analysis. A model's judgements may still be challenging to understand even if it achieves respectable predicted performance. Knowing whether the model forecasts an upward or downward trend is insufficient in cryptocurrency markets, since price changes are extremely unpredictable and influenced by several interacting variables. Understanding which factors affect the choice, whether the model depends on significant market signals, and how its behaviour varies under various market regimes are also crucial. Explainable artificial intelligence techniques are therefore pertinent to cryptocurrency analysis as they facilitate the integration of predictive modelling outcomes with market and economic understanding.

Previous studies on cryptocurrency market prediction often focus on improving forecasting accuracy, while less attention is given to explaining how machine learning models form their decisions, whether these decisions remain consistent under different market regimes, and whether model-generated signals have practical economic meaning. This creates a scientific gap between predictive machine learning performance and the interpretability of model decisions in cryptocurrency market analysis.

The research problem of this thesis is the limited interpretability and practical evaluation of machine learning-based cryptocurrency market prediction models. Although such models may be used to predict market direction, prediction accuracy alone does not show whether the model’s decisions are stable under different market conditions, economically meaningful, or useful for decision-making. Therefore, this thesis combines directional prediction with explainable artificial intelligence methods, market-state analysis and backtesting.

The object of the thesis is the explanation of machine learning model decisions in the analysis of cryptocurrency price dynamics. The empirical analysis focuses on Bitcoin, Ethereum, and Solana, which are examined using market indicators, derived technical factors, news-related variables, macroeconomic data, and market regime characteristics. Each of these assets is analysed independently using a comparable modelling and interpretation workflow.

The objective of the thesis is to evaluate machine learning models for analysing cryptocurrency market dynamics using explainable artificial intelligence methods.

The following goals are developed in order to accomplish this aim:

1. discuss the challenges of cryptocurrency market dynamics assessment and the relevance of explainable artificial intelligence for interpreting market direction predictions;
2. review theoretical aspects of cryptocurrency market behaviour, regime shifts, machine learning-based forecasting and model explainability in financial research;
3. prepare a research methodology for evaluating cryptocurrency market dynamics using market data, news indicators, macro-financial variables, regime identification, classification models and interpretability techniques;
4. conduct an empirical study of Bitcoin, Ethereum and Solana by evaluating identified market regimes, prediction performance, explainability results and the economic interpretation of model-based trading signals.

The research applies machine learning and explainable artificial intelligence methods to analyse cryptocurrency market dynamics. The empirical analysis is based on Random Forest classification models, market regime identification using a Hidden Markov Model, SHAP-based model interpretation, and backtesting of model-generated trading signals. These methods allow the research to evaluate not only predictive accuracy, but also the interpretability and economic relevance of the obtained results.

Generative AI was used as an auxiliary tool for language editing, translation support, structural drafting and consistency checking. The scientific idea, project content, factual information, references and final conclusion were reviewed and remain the responsibility of the author of the final project. No generative AI tool is credited as an author.

1. Literature Review on Cryptocurrency Market Dynamics and Explainable Artificial Intelligence

1.1. Cryptocurrency Markets

Cryptocurrency markets differ from traditional financial markets because they combine technological decentralisation, continuous trading, high volatility, sentiment sensitivity and a still-developing regulatory environment. Bitcoin’s original proposal introduced a peer-to-peer electronic cash system in which transactions are verified without relying on a central financial intermediary [1]. However, decentralisation is not only a technological advantage. It also creates practical constraints, including scalability issues, network congestion, high resource consumption and difficulties in maintaining efficient transaction processing [2]. Therefore, cryptocurrency markets should be understood not only as technological systems, but also as markets with specific operational and economic limitations.

A further important characteristic is the trading structure of cryptocurrencies. Unlike most equity or bond markets, cryptocurrencies are traded globally and almost continuously. As a result, new information, investor expectations and external shocks can be reflected in prices at any time. Continuous trading may strengthen short-term reactions to news, sentiment and speculative behaviour, while also creating additional challenges for modelling returns, volatility and trading signals. These challenges are intensified by the fact that many cryptocurrencies do not have conventional valuation anchors such as cash flows, dividends or standard fundamental valuation models. Their prices are therefore more strongly affected by supply and demand, liquidity, speculative expectations, investor sentiment, technological developments and regulatory news [3].

These characteristics make cryptocurrency market analysis economically relevant. Crypto-assets are used in trading, portfolio management and risk-taking, but their potential returns are connected with substantial volatility and uncertainty. Levantesi, Piscopo and Roviello [3] show that cryptocurrency research places strong emphasis on volatility modelling, risk assessment and portfolio management, indicating that the economic value of cryptocurrency investment cannot be separated from risk. This is important for empirical modelling because a model that predicts price direction may still have limited practical value if its predictions are not interpreted in relation to risk, volatility and market conditions.

Cryptocurrency markets are also closely linked with speculative behaviour and investor attention. Since many crypto-assets lack traditional valuation fundamentals, their prices may be more sensitive to expectations, narratives and short-term sentiment than many conventional financial assets [5]. This supports the use of sentiment indicators, news variables, trading volume, volatility and momentum-based measures in cryptocurrency prediction. These variables are relevant not only for prediction, but also for understanding which signals machine learning models use when forming their decisions.

Another important aspect is that cryptocurrencies cannot always be treated as isolated assets. Iyer [15] shows that return and volatility spillovers between crypto-assets and broader financial markets have increased over time, especially during turbulent periods and negative economic or crypto-specific events. This means that cryptocurrency risk may affect not only individual assets, but also broader market exposure. The Bank for International

Settlements also notes that crypto-asset markets may create risks related to financial stability, market integrity, investor protection and cross-border supervision [2].

The question of market efficiency is central to cryptocurrency prediction research. In an efficient market, asset prices should reflect available information, making systematic prediction difficult [4]. However, empirical research suggests that cryptocurrency market efficiency depends on the analysed asset, time period, liquidity level and market conditions. Wei [7] finds that return predictability decreases as market liquidity increases, while Bitcoin shows stronger signs of efficiency than many smaller cryptocurrencies. Tran and Leirvik [9] also show that cryptocurrency market efficiency changes over time and tends to increase as markets mature. These findings suggest that predictability may exist, but it is unstable and dependent on market structure.

The adaptive market hypothesis helps explain this instability. Rather than assuming that markets are always efficient or inefficient, it suggests that efficiency changes with competition, investor behaviour, market shocks and institutional development. Chu, Zhang and Chan [10] show that efficiency in Bitcoin and Ethereum markets varies over time and can be related to important market events and news. This is especially relevant for cryptocurrencies, where speculative cycles, regulatory announcements, technological developments and macroeconomic shocks can quickly change the market environment.

Overall, the literature suggests that cryptocurrency markets are partially and time-varyingly efficient. Short-term predictability may arise from historical returns, technical indicators, liquidity conditions, volatility clustering, sentiment and news reactions. However, such signals may be weak, unstable, regime-dependent or economically insignificant after transaction costs and trading constraints [3,5,7,8,9,10]. Therefore, cryptocurrency prediction should not be assessed only through statistical performance metrics. This supports the need to combine predictive modelling with explainability, regime-based evaluation and backtesting, because such an approach makes it possible to examine not only whether a model predicts market direction, but also which variables drive its decisions, under which market conditions it performs better, and whether its outputs have practical economic relevance [3,6,7,9].

1.2. Cryptocurrency Market Dynamics and Price Formation

Cryptocurrency price dynamics cannot be explained by one group of variables. Previous research shows that prices are shaped by the interaction between internal market factors, investor behaviour, technological characteristics, macroeconomic conditions and regulatory information [11,12,13,15,16,18]. This distinguishes cryptocurrencies from traditional financial assets, where price formation is often linked to cash flows, interest rates, firm fundamentals or macroeconomic performance. In cryptocurrency markets, price behaviour is more strongly affected by demand, speculative activity, liquidity, attention, technological trust and expectations about future adoption.

Ciaian, Rajcaniova and Kancs [12] show that Bitcoin prices are influenced not only by supply and demand, but also by factors related to Bitcoin attractiveness for investors and users. Kristoufek [13] similarly finds that Bitcoin prices are affected by fundamental, speculative and technical factors, and that these relationships differ across time horizons. These findings suggest that cryptocurrency modelling should not assume that one fixed group of variables explains market behaviour equally well across all periods.

Internal market variables such as returns, momentum, trading volume, volatility and liquidity are widely discussed as important drivers of cryptocurrency price movements. Momentum may reflect the tendency of market participants to follow recent trends, while trading volume reflects market activity and investor participation [11,13]. Liquidity is also related to predictability. Wei [7] finds that cryptocurrency return predictability decreases as liquidity increases, suggesting that more liquid assets incorporate information more efficiently. This is relevant because Bitcoin, Ethereum and Solana differ in maturity and liquidity, which may partly explain differences in model performance across assets.

Volatility is one of the central characteristics of cryptocurrency markets. High volatility reflects uncertainty, speculative pressure and rapid changes in expectations, but it also affects the practical meaning of model predictions. Katsiampa [20] shows that Bitcoin volatility is persistent and can be better captured when both short-run and long-run components of conditional variance are considered. Baur and Dimpfl [14] also show that Bitcoin volatility is much higher than the volatility of major fiat currencies. These findings justify the inclusion of volatility-related variables in cryptocurrency modelling, especially when the aim is to interpret both prediction and risk.

External drivers are also important. Macroeconomic uncertainty, interest rates, financial market stress, inflation expectations and global risk sentiment may affect cryptocurrency prices by changing investors' willingness to hold risky assets. Iyer [15] shows that return and volatility spillovers between crypto-assets and traditional financial markets have increased over time, particularly during periods of market turbulence and negative economic or crypto-specific events. This indicates that cryptocurrencies cannot always be analysed as isolated markets.

Regulatory, technological and security-related information can also quickly affect cryptocurrency prices. Lyócsa et al. [16] find that Bitcoin volatility reacts strongly to regulation-related news, regulatory sentiment and hacking attacks on cryptocurrency exchanges. Zhang, Xu and Qi [17] further show that regulatory announcements can significantly increase cryptocurrency price, liquidity and return volatility, especially during periods of broader uncertainty such as the COVID-19 pandemic. These studies show that external information affects cryptocurrency markets through both long-term institutional expectations and short-term investor reactions.

Investor sentiment and attention form another important group of price drivers. Since many cryptocurrencies lack conventional valuation anchors, investor expectations, narratives, media attention and online discussion may have a stronger influence on prices than in traditional financial markets. Kristoufek [13] shows that public interest is related to Bitcoin price dynamics. Sakariyahu et al. [18] find that economic and political uncertainty factors significantly affect crypto prices, while Bagh, Khan and Iftikhar [19] find that investor sentiment is a significant predictor of cryptocurrency returns and volatility. These findings support the inclusion of sentiment and news-based variables in empirical cryptocurrency modelling.

The literature also suggests that cryptocurrency markets move through different regimes rather than following one stable process. Such regimes may represent bull, bear, sideways or panic-like conditions, each associated with different return, volatility and risk characteristics. Koki, Leonardos and Piliouras [21] show that Hidden Markov Models can distinguish cryptocurrency regimes with different profit and risk characteristics, while Ma et al. [22] show that regime-switching models are useful when volatility dynamics differ across market states. This is important because regime identification connects market dynamics with model evaluation and economic interpretation.

Regime awareness is particularly relevant for machine learning-based prediction. If the relationship between explanatory variables and market direction changes across regimes, then aggregate performance metrics may hide important differences in model behaviour. Momentum variables may be more informative during trending markets, while volatility, sentiment or uncertainty indicators may become more important during panic periods [16,18,21,22]. Therefore, regime-based evaluation helps interpret not only whether a model predicts market direction, but also under which market conditions its decisions are more meaningful.

Overall, previous research shows that cryptocurrency price dynamics are shaped by internal market indicators, volatility, liquidity, sentiment, macro-financial conditions, regulatory information and technological events. The influence of these drivers is not stable across all periods or assets, but may change depending on market maturity, liquidity, investor attention and regime conditions. This supports the use of a modelling framework that combines market-based variables, news-related indicators, macro-financial information and regime features when analysing cryptocurrency price dynamics.

1.2.1. Technical Indicators and Derived Market Variables

Technical indicators and derived market variables are widely used in financial modelling because they transform raw price and volume observations into features that describe return behaviour, trend direction, volatility, momentum and market activity. In cryptocurrency research, these variables are especially important because many crypto-assets lack traditional valuation fundamentals. As a result, empirical models often rely on open, high, low, close and volume data to approximate short-term market dynamics and investor behaviour [11,13,23,24,25].

Returns and log returns are among the most important variables in financial time-series modelling because they allow price changes to be compared across assets and periods. Log returns are commonly used because they are additive over time and useful for volatility and risk analysis. However, log returns and simple returns should not be treated as identical. Hudson and Gregoriou [23] note that differences between logarithmic and simple returns become more important when volatility is high. This distinction is relevant in cryptocurrency markets, where large price changes are common.

Volatility measures approximate uncertainty and risk. For cryptocurrencies, rolling volatility indicators are useful because price changes are often large, clustered and regime-dependent. Earlier evidence on Bitcoin volatility shows that volatility is persistent and substantially higher than in major fiat currencies [14,20]. Therefore, short-term and medium-term realised volatility variables may help models distinguish calm periods from high-risk or panic periods.

Moving averages and momentum-related variables represent trend behaviour. Moving averages smooth short-term price fluctuations and may indicate the direction of the prevailing trend. Zhu and Zhou [24] argue that moving-average rules can serve as learning tools when investors face uncertainty about market fundamentals, while Han, Zhou and Zhu [25] show that moving averages may contain useful information for asset allocation when used as trend-following signals. These findings are relevant because cryptocurrency markets often experience strong bull and bear phases.

Momentum indicators capture the tendency of recent price movements to continue over short horizons, while volume-based indicators measure market activity, liquidity and investor participation. In cryptocurrency markets, momentum may be intensified by speculation, investor attention and online narratives [13]. Trading volume may indicate stronger demand, increased attention or market disagreement, but its interpretation depends on price direction and market context. Wei's [7] findings on liquidity and market efficiency are relevant here because more liquid cryptocurrencies tend to show lower return predictability.

Overall, technical indicators and derived market variables provide a structured representation of cryptocurrency market behaviour. Returns describe price changes, moving averages and momentum variables capture trend dynamics, volatility measures approximate risk, and volume-based indicators describe activity and liquidity. For this thesis, these variables are important not only as predictors, but also as interpretable signals that can later be evaluated through feature importance, SHAP values and other explainability methods.

1.3. News Sentiment and Alternative Data in Cryptocurrency Analysis

Investor sentiment and news-based information are important in cryptocurrency analysis because many crypto-assets do not have traditional valuation anchors such as dividends, earnings or cash-flow expectations. In this context, prices may react strongly to expectations, narratives, uncertainty, speculation and online attention. Therefore, cryptocurrency market dynamics cannot be fully understood by analysing price and volume data alone [13,18]. Sentiment variables are relevant because they may capture not only emotional reactions, but also demand pressure, uncertainty and changing investor expectations.

The literature shows that speculative sentiment can influence asset prices even when it is not directly linked to fundamentals. Davies [5] conceptualises speculative sentiment as aggregate uninformed and gambling-like demand. This idea is especially relevant for cryptocurrencies, where short-term trading, attention-driven demand and speculative expectations often play an important role. Crypto-specific studies support this argument: Sakariyahu et al. [18] show that global uncertainty and sentiment factors significantly affect cryptocurrency prices, while Bagh, Khan and Iftikhar [19] find that investor sentiment can predict cryptocurrency returns and volatility. Brauneis and Sahiner [26] also show that sentiment information extracted from crypto-market news can improve volatility forecasts in a meaningful share of cases. However, Freitas et al. [27] emphasise that sentiment-based predictability should be interpreted cautiously, especially in out-of-sample settings.

News-based sentiment analysis transforms unstructured textual information into structured variables that can be included in financial models. These variables may include article counts, average tone, sentiment scores, negative-news shares, topic-specific counts and sentiment dispersion. Janková [28] shows that textual data have become an important complement to traditional market variables in financial prediction. In cryptocurrency research, Lamon, Nielsen and Redondo [29] provide an early example of using news and social media sentiment to predict price fluctuations in Bitcoin, Ethereum and Litecoin. Shen et al. [30] also demonstrate how large-scale news databases can be transformed into structured media sentiment indicators such as tone, attention and emotional polarity.

However, news sentiment should not be treated as a single universal predictor. Article frequency may reflect attention and information intensity, but a larger number of articles can occur during both positive and negative market periods. Therefore, article counts are more informative when combined with sentiment tone, negative-news shares or topic-specific indicators. Topic information is also important because regulatory announcements, exchange hacks, fraud cases, institutional adoption, ETF-related news, technological upgrades and macroeconomic shocks may affect cryptocurrency markets through different mechanisms. Lyócsa et al. [16] show that Bitcoin volatility reacts to regulation-related news, regulatory sentiment and hacking attacks on cryptocurrency exchanges.

For this reason, cryptocurrency news indicators can be divided into crypto-specific and macroeconomic news variables. Crypto-specific news is directly related to the perceived safety, legality and adoption potential of digital assets. Regulatory news may affect expectations about legal status, taxation, supervision or institutional acceptance, while security-related news such as hacks, fraud and custody failures can reduce trust in the cryptocurrency ecosystem [16,31,32]. Macroeconomic news represents a broader information channel, including central bank policy, inflation, financial crises, geopolitical events, pandemic-related uncertainty and global risk sentiment. Sakariyahu et al. [18] show that global uncertainty and sentiment factors influence cryptocurrency returns, while Selmi [33] finds that the relationship between Bitcoin returns and economic policy uncertainty changed across major crisis periods, suggesting that macroeconomic effects may depend on market conditions.

The literature also emphasises several limitations of news-based sentiment indicators. News sentiment may be noisy, delayed, duplicated across sources or sensitive to the method used for sentiment measurement. In addition, high news frequency may reflect attention after a price movement has already occurred rather than information that predicts future movement [27,28]. For this reason, sentiment and news variables should be interpreted as explanatory signals rather than direct causal drivers of price changes.

Overall, previous research supports the inclusion of news and sentiment-related variables in cryptocurrency analysis, but also shows that their effect is unstable and context-dependent. Sentiment may capture investor expectations, uncertainty, speculative pressure and risk perception, while topic-specific indicators help distinguish between crypto-specific and macroeconomic information channels. For this thesis, these variables are relevant not only as predictors of market direction, but also as interpretable features that can later be evaluated through model explainability methods.

1.4. Machine Learning Methods for Cryptocurrency Market Prediction

Before machine learning methods became widely used in cryptocurrency research, market prediction was mainly based on traditional econometric and statistical time-series models, including regression, ARIMA-type and GARCH-type models. Their main advantage is interpretability, because coefficients, assumptions and statistical relationships can be examined more directly than in many machine learning models. For this reason, they remain useful as benchmarks in financial prediction studies.

Regression models can be used to analyse relationships between cryptocurrency returns and explanatory variables such as trading volume, liquidity, macroeconomic uncertainty or sentiment. However, standard linear regression assumes relatively stable relationships between variables, which is problematic in cryptocurrency markets where price dynamics may change across bull, bear, sideways and panic regimes. Earlier studies show that cryptocurrency behaviour depends on liquidity, sentiment, uncertainty and broader market conditions, making one stable linear relationship difficult to assume [7,18].

ARIMA-type models are useful for capturing autocorrelation and linear dependence in time-series data. In cryptocurrency research, they can serve as transparent baseline models for testing whether historical prices or returns contain predictable patterns. However, they are limited when market behaviour is nonlinear, affected by sudden shocks or structurally unstable. GARCH-type models are especially important for volatility analysis. Katsiampa [20] shows that Bitcoin volatility is better represented when both short-run and long-run volatility components are considered, while Baur and Dimpfl [14] demonstrate that Bitcoin volatility is substantially higher than that of major fiat currencies. However, standard GARCH models may still be limited when volatility changes because of structural breaks or regime shifts. Ma et al. [22] therefore show that allowing volatility dynamics to differ across regimes can improve cryptocurrency volatility forecasting.

Traditional econometric and statistical approaches therefore provide an important foundation for cryptocurrency analysis, but they are often too restrictive for modelling nonlinear, sentiment-sensitive and regime-dependent markets. Machine learning methods are increasingly used because they can handle many explanatory variables, model nonlinear relationships and capture interactions between market, sentiment, macroeconomic and regime-related signals [34]. This is relevant in cryptocurrency modelling, where the same prediction task may include historical returns, volatility, liquidity, technical indicators, news sentiment, macroeconomic uncertainty and market regime features.

Cryptocurrency-specific studies also show why machine learning has become common in this field. John, Binnewies and Stantic [35] show that cryptocurrency price prediction research has moved from earlier statistical approaches towards more advanced machine learning and deep learning models. Bouteska et al. [36] compare ensemble learning and deep learning models for several cryptocurrencies, including Bitcoin and Ethereum, and show that some machine learning models can outperform simpler benchmarks in forecasting and trading applications. However, these results should be interpreted cautiously, because model performance depends on the selected asset, prediction horizon, input variables, evaluation strategy and market period. Therefore, machine learning does not remove the need for careful validation and interpretability.

1.4.1. Classification versus Regression in Cryptocurrency Prediction

Cryptocurrency prediction studies usually follow either a regression or classification approach. Regression models aim to predict a numerical value, such as future price, return, log return or volatility. Classification models predict a discrete market state, most commonly whether the next-period movement will be upward or downward. This distinction is important because the modelling objective determines how the model is evaluated and how its outputs can be interpreted in practice.

Regression is useful when the research objective is to estimate the magnitude of future price or return changes. However, exact cryptocurrency price or return prediction is difficult because crypto markets are highly volatile, sentiment-sensitive and affected by sudden news, liquidity changes and regime shifts. Previous studies show that cryptocurrency volatility is persistent but unstable, while market efficiency may change across time and assets [7,20,22,38]. Therefore, numerical forecasts may be sensitive to outliers, extreme market movements and changing market conditions.

Classification provides a more decision-oriented perspective because it focuses on market direction rather than exact price magnitude. This is useful in financial decision-making, where investors often need to decide whether to buy, sell, hold or avoid a position. Bouteska et al. [36] show that ensemble and deep learning methods can be applied to cryptocurrency market direction forecasting. Omole and Enke [37] also connect Bitcoin price direction prediction with trading strategy evaluation, showing that directional models can be assessed not only through classification metrics, but also through the profitability of strategies based on their predictions.

However, classification also has limitations. A model may correctly predict direction but fail to capture whether the movement is economically meaningful. A small positive return and a strong positive return may both be labelled as “up”, although their investment implications differ. Classification results can also be affected by class imbalance, threshold selection, transaction costs and the frequency of generated trading signals. Therefore, accuracy alone is not sufficient. Metrics such as F1 score, class-specific performance and backtesting results are needed to evaluate whether directional predictions are useful in practice.

For this thesis, the UP/DOWN formulation is appropriate because it links model outputs with trading decisions, economic interpretation and explainability. It allows predictions to be transformed into buy, sell or risk-reduction signals, while SHAP analysis can later be used to identify which variables contributed to upward or downward predictions. Therefore, classification connects the predictive, interpretability and decision-making components of the thesis more directly than pure price-level regression.

1.4.2. Random Forest Models in Financial Classification Tasks

Random Forest is an ensemble learning method that combines many decision trees and aggregates their predictions. This structure reduces the risk of relying on a single unstable decision tree and usually produces more robust predictions. In classification tasks, Random Forest assigns the final class based on the collective decision of the trees, making it suitable for UP/DOWN market direction prediction. In cryptocurrency research, Basher and Sadorsky [39] apply tree-based machine learning classifiers, including Random Forest, to Bitcoin price direction forecasting and show that such models can be useful for directional prediction.

One reason why Random Forest is useful in financial classification is its ability to model nonlinear relationships and interactions. Financial markets are rarely driven by simple linear effects. The influence of volatility, trading volume or sentiment may depend on the current market environment, previous returns or liquidity conditions. Random Forest can capture such interactions because decision trees split the data into different regions and allow different relationships to be learned under different conditions [34,39]. This is especially relevant for cryptocurrency markets, where price behaviour may differ between calm, trending and high-volatility periods.

Random Forest is also practical when the dataset contains many explanatory variables. Cryptocurrency prediction models often include technical indicators, returns, volatility measures, volume-based variables, sentiment indicators, macroeconomic variables and regime-related features. Unlike linear models, Random Forest does not require the researcher to manually specify all interaction terms, making it suitable when relationships between variables are complex or uncertain [34,39]. However, this flexibility also requires careful validation, because the model can still learn unstable patterns that do not generalise to new market periods.

Another advantage of Random Forest is that it can provide feature importance measures. Basher and Sadorsky [39] find that tree-based methods can outperform logit models in Bitcoin price direction forecasting, while technical indicators appear among the most important predictors. This is relevant because Random Forest can be used not only for classification, but also as an initial tool for identifying which market indicators carry predictive information.

At the same time, Random Forest feature importance must be interpreted carefully. Mean decrease in impurity can show which variables were useful for splitting the data inside the model, but it does not prove causal influence. Variables with many possible split points or strong correlations with other predictors may appear more important. Therefore, feature importance is useful as an initial interpretability tool, but it should be supported by additional explainability methods such as SHAP [45]. This distinction is important because the thesis focuses not only on prediction, but also on the explanation of model decisions.

Overall, Random Forest is suitable for this thesis because it supports UP/DOWN classification, handles diverse explanatory variables, captures nonlinear relationships and provides a basis for model interpretation. However, its results must be evaluated using validation data, test-period performance, class-specific metrics, backtesting and explainability methods.

1.4.3. Hyperparameter Tuning and Model Validation

Model validation is especially important in cryptocurrency prediction because financial time series are ordered, non-stationary and sensitive to changing market conditions. Observations should not be randomly shuffled without caution, because future information may accidentally enter the training process and create data leakage. Therefore, cryptocurrency prediction models are usually evaluated using chronological train, validation and test splits, where the model is trained on earlier data, tuned on a later validation period and finally evaluated on an unseen test period.

This approach is consistent with the broader literature on time-series evaluation. Bergmeir and Benítez [40] note that standard cross-validation methods may create theoretical concerns in time-series problems because the data contain dependencies and time-evolving effects. Bergmeir, Hyndman and Koo [41] later show that cross-validation can be valid in some autoregressive settings, but only under specific assumptions, such as uncorrelated errors. In cryptocurrency modelling, where regimes, volatility and sentiment can change over time, chronological validation is usually more appropriate than random splitting.

Hyperparameter tuning is needed because model performance depends not only on the data, but also on configuration choices made before training. In Random Forest models, these include the number of trees, maximum tree depth, minimum samples per split, minimum samples per leaf, feature sampling and class-weighting logic. Poorly selected hyperparameters may lead to underfitting or overfitting. This is especially important in cryptocurrency prediction because a model can learn patterns specific to one market period and fail on new data.

Traditional tuning approaches include grid search and random search. Grid search evaluates a predefined set of hyperparameter combinations, while random search samples combinations more flexibly. Bergstra and Bengio [42] show that random search can be more efficient than grid search because not all hyperparameters are equally important. More recent approaches use Bayesian or sequential optimisation. Akiba et al. [43] describe Optuna as a hyperparameter optimisation framework with efficient search and pruning strategies. Such methods can make tuning more efficient, but they do not remove the need for proper validation.

The final test set must remain separate from both training and hyperparameter tuning. The validation period is used to select model settings, while the test period is used only to estimate out-of-sample performance. This distinction is especially important in cryptocurrency research, where market behaviour changes quickly and small improvements in classification metrics may not translate into economic value. Therefore, reliable model evaluation should combine chronological splitting, hyperparameter tuning, independent test-period assessment and, where relevant, economic backtesting. This helps determine whether the model captures useful market signals rather than historical noise.

1.5. Market Regime Identification in Cryptocurrency Markets

1.5.1. Market Regimes and Regime-Aware Evaluation

Financial markets do not behave uniformly across all periods. Instead, they may move through different states characterised by distinct return, volatility, liquidity and investor-behaviour patterns. In cryptocurrency markets, such regimes are especially relevant because price dynamics may shift rapidly between growth, decline, sideways movement and panic-like conditions. Bull regimes are usually associated with rising prices and stronger investor confidence, bear regimes with declining prices and weaker expectations, sideways regimes with limited directional movement, and panic regimes with sharp corrections, high volatility and increased uncertainty.

The idea that markets alternate between different latent states is well established in financial econometrics. Hamilton [44] introduced regime-switching models as a way to represent economic and financial time series that move between unobserved states with different statistical characteristics. This framework is relevant for cryptocurrency markets because previous studies show that crypto-asset volatility is persistent but unstable across periods [20,22]. Therefore, cryptocurrency markets should not be analysed as if one stable relationship between variables applied to the entire sample.

Regime-based analysis is also important for evaluating prediction models. A model that performs well during a bull market may not perform equally well during bear, sideways or panic conditions. During clear upward or downward trends, directional movements may be easier to identify, while sideways markets can be more difficult because price changes are weaker and more affected by noise. Panic regimes create a different challenge, since movements may be strong but driven by sudden shocks, liquidity pressure, fear and rapid sentiment changes [16,18,22]. Therefore, aggregate test-period accuracy may hide important differences in model behaviour across market states.

Previous cryptocurrency studies support the use of regime-aware analysis. Koki, Leonardos and Piliouras [21] show that cryptocurrency markets can be analysed through hidden regimes with different return and risk characteristics. Ma et al. [22] similarly demonstrate that volatility forecasting can benefit from allowing volatility dynamics to differ across regimes. These findings indicate that regimes are not only descriptive labels, but meaningful market conditions under which model behaviour, risk and predictive reliability may change.

Regime awareness is also relevant for explainability. If feature importance or SHAP values differ across regimes, then the same machine learning model may rely on different signals under different market conditions. Momentum variables may be more important during trending markets, while volatility, drawdown or sentiment-related variables may become more relevant during panic or uncertainty-driven periods. Therefore, regime-based evaluation allows model predictions to be interpreted in relation to the market environment rather than only at the full-sample level. SHAP-based explanations are useful in this context because they make it possible to examine how individual variables contribute to model predictions [45].

Overall, regime-aware modelling provides a more detailed view of cryptocurrency prediction models. Instead of evaluating only whether the model predicts market direction correctly on average, it allows the analysis to examine when the model works better or worse and under which market conditions its predictions are more reliable. This directly supports the empirical design of this thesis, where prediction performance and explainability results are interpreted in relation to HMM-identified market regimes.

1.5.2. Hidden Markov Models for Regime Detection

Hidden Markov Models are widely used for regime detection because they allow market states to be treated as latent processes. In financial markets, regimes such as bull, bear, sideways or panic periods are not directly observed. They must be inferred from observable variables such as returns, volatility, drawdown, trading volume or other market indicators. The logic of HMM is therefore suitable for market analysis: observed market variables are assumed to be generated by an underlying hidden state process, where each state has its own statistical characteristics.

The theoretical basis of this approach is closely related to Hamilton's [44] regime-switching framework, in which financial and economic time series can move between unobserved states. In contrast to static models, this allows the behaviour of observed variables to change depending on the hidden state. For cryptocurrency markets, this is important because returns and volatility are not stable across time, and periods of rapid growth, decline, low-activity movement and high volatility may have different statistical properties.

In HMM-based regime detection, variables such as returns, realised volatility, drawdown and trading volume can be used as observable signals for estimating latent regimes. Returns help distinguish positive and negative market phases, volatility captures uncertainty, drawdown reflects market stress, and volume may indicate investor participation or liquidity pressure. These variables do not define regimes mechanically, but together they provide information that helps the model separate market states with different risk-return profiles.

Empirical cryptocurrency studies support the use of HMM and related regime-switching approaches. Koki, Leonardos and Piliouras [21] apply Bayesian Hidden Markov Models to cryptocurrency markets and show that hidden states can be interpreted as regimes with different return and risk characteristics. Similarly, Ma et al. [22] show that regime-switching volatility models are useful when cryptocurrency volatility dynamics differ across market states. These findings support the view that cryptocurrency markets should not be analysed as one homogeneous period.

The value of HMM-based regime detection is not limited to identifying market phases. It can also strengthen the interpretation of predictive models. If model accuracy, feature importance or SHAP explanations differ across regimes, then model behaviour is not constant across market conditions. In this thesis, HMM regimes therefore provide an additional explanatory layer by connecting Random Forest predictions and XAI results with broader market states.

However, HMM-based regime detection also has limitations. The number of regimes must be selected by the researcher, and different choices may produce different interpretations.

In addition, labels such as “bull”, “bear” or “panic” are assigned after examining the statistical behaviour of each state; they are not automatically provided by the model. For this reason, HMM results should be interpreted using descriptive statistics, visual inspection and economic reasoning [21,22,44].

Overall, Hidden Markov Models are appropriate for cryptocurrency regime detection because they infer latent market states from variables that reflect return direction, volatility, market stress and participation. In this thesis, HMM-based regimes provide a basis for evaluating whether model performance and explainability results differ across changing cryptocurrency market environments.

1.6. The Need for Explainability in Machine Learning-Based Financial Market Analysis

Machine learning methods are increasingly used in financial market analysis because they can capture nonlinear relationships, interactions between variables and complex market signals. However, this flexibility also creates the black-box problem: a model may produce accurate predictions, while the reasoning behind these predictions remains difficult to understand. In financial applications, this is a serious limitation because model outputs can influence investment decisions, risk management, portfolio allocation and regulatory judgement. Therefore, predictive performance alone is not sufficient; users also need to understand why a model produced a particular result.

The explainable artificial intelligence literature treats the lack of transparency as both a practical and governance issue. Guidotti et al. [46] argue that many accurate decision-support systems operate as black boxes, where the reasoning behind the output is hidden from users. Rudin [47] takes a stricter position and argues that black-box models should be treated cautiously in high-stakes decision-making when interpretable alternatives are available. In market prediction, however, more flexible models may still be useful because financial market data often contain nonlinear relationships and interaction effects. In such cases, post-hoc explainability methods can reduce the interpretability gap, even if they do not make the model fully transparent.

This issue is especially relevant in cryptocurrency markets, where price movements are volatile, sentiment-driven and regime-dependent. A model that predicts an upward or downward movement may be technically useful, but without explanation it is difficult to judge whether the prediction is based on meaningful market signals, unstable noise or spurious relationships. Arsenault, Wang and Patenaude [48] note that the complexity of AI models can reduce trust and slow adoption in high-risk financial decision-making even when predictive performance is strong. Therefore, explainability supports not only transparency, but also model validation and economic interpretation.

For this thesis, explainability is important because the objective is not only to classify cryptocurrency market direction, but also to understand which market, sentiment, macroeconomic and regime-related factors influence the model’s predictions. This directly supports the object of the thesis: the explanation of machine learning model decisions in the analysis of cryptocurrency price dynamics. In this context, explainability methods are used to evaluate whether model decisions are based on economically interpretable signals and whether these decisions can be meaningfully analysed under changing market conditions.

1.6.1. Model Explainability Methods and Their Limitations

Explainability methods are commonly divided into global and local approaches. Global explainability focuses on the overall behaviour of the model and identifies which variables are generally most important for prediction. Local explainability explains one individual prediction and shows why the model produced a specific output for a specific observation. Both perspectives are useful in financial market analysis because analysts need to understand the model's general decision logic, but also the reasoning behind particular market signals [46,49,53]. In cryptocurrency prediction, this means examining both which variables generally influence UP/DOWN predictions and which variables explain a specific signal on a specific day.

Feature importance methods provide an initial way to interpret machine learning models by identifying which variables contribute most to prediction. In Random Forest models, one common measure is Mean Decrease in Impurity, which evaluates how much each variable contributes to reducing impurity across decision-tree splits. Basher and Sadorsky [39] show that variable importance can be used to assess which predictors are most influential within a Random Forest price-direction forecasting model. In cryptocurrency modelling, this can indicate whether the model is mainly driven by returns, volatility, volume, sentiment, macroeconomic variables or regime-related features. Such information helps evaluate whether the model's behaviour is economically plausible, for example when recent returns or volatility measures appear among the most important predictors, which is consistent with literature on momentum, volatility clustering and regime changes [13,20,21,22].

However, traditional feature importance has limitations. Mean Decrease in Impurity shows how useful variables are for splitting the data inside the model, but it does not prove that these variables have a causal effect on cryptocurrency prices. It can also be biased toward variables with many possible split points or toward correlated predictors. Strobl et al. [51] show that variable importance measures in Random Forests may be biased when predictor variables differ in scale or number of categories. Therefore, feature importance should be interpreted as model-based importance rather than direct economic causality.

SHAP values provide a more detailed post-hoc explanation of model behaviour. The method is based on Shapley values from cooperative game theory [52], which Lundberg and Lee [45] adapted into a unified additive feature-attribution framework for explaining machine learning predictions. SHAP assigns each feature a contribution value, showing how much it pushed a prediction away from the baseline model output. This is useful in cryptocurrency prediction because it shows not only which variables matter, but also whether their values push the model toward an UP or DOWN classification. Traditional feature importance can rank variables, but SHAP additionally provides the direction and magnitude of feature influence [45].

SHAP can be used for both global and local interpretation. At the global level, summary or beeswarm plots show which variables generally influence predictions and how their values affect the model output. At the local level, waterfall plots can explain one specific UP or

DOWN prediction by decomposing it into feature-level contributions. Dependence plots can also show whether the effect of a variable is linear, nonlinear or conditional on other variables [45]. These properties are relevant for cryptocurrency markets, where relationships between returns, volatility, sentiment, macroeconomic indicators and price direction may be nonlinear and regime-dependent. Černevičienė and Kabašinskas [54] also show that SHAP is among the commonly used explainability methods in AI applications in finance.

At the same time, SHAP and feature importance must be interpreted carefully. They explain how a trained model uses input variables, but they do not reveal true economic causality. This distinction is important because financial market variables are often correlated, unstable and affected by external shocks. Molnar [50] notes that interpretation methods can become less reliable when features are dependent, because the estimated contribution of one variable may partly reflect information contained in another variable. This is particularly relevant in cryptocurrency modelling, where many derived technical indicators are calculated from the same OHLCV data.

Post-hoc explanations may also create a false sense of certainty. Visual tools such as beeswarm plots, dependence plots or waterfall plots can make model behaviour appear intuitive, but they still represent the logic of a fitted model, not a verified financial theory. Guidotti et al. [46] emphasise that explanations must be evaluated in relation to their purpose and user context, because different explanation methods provide different types of information and may not fully capture model reliability. Therefore, explanation results should be interpreted together with validation performance, regime-based evaluation and economic backtesting.

Overall, feature importance and SHAP values are useful because they make machine learning models more transparent and help identify which variables influence predictions. However, in this thesis, these methods are interpreted as tools for explaining model behaviour rather than proving that specific variables cause cryptocurrency price movements. This cautious interpretation is necessary because the scientific problem addressed in the thesis is not only whether cryptocurrency market direction can be predicted, but whether machine learning model decisions can be explained and economically interpreted under changing market conditions.

1.7. Economic Evaluation of Cryptocurrency Prediction Models

Predictive performance metrics are necessary for evaluating whether a cryptocurrency classification model can distinguish between upward and downward market movements. In UP/DOWN prediction tasks, the model does not estimate the exact future return, but assigns each observation to a directional class. Therefore, evaluation must focus on classification quality rather than regression errors. Common metrics include accuracy, precision, recall, F1 score, macro F1 and class-specific performance indicators.

Accuracy measures the share of correctly classified observations among all predictions. It is simple to interpret and can be used as a first indicator of model performance. However, accuracy may be misleading when classes are imbalanced. For example, if upward movements occur more frequently than downward movements, a model may achieve acceptable accuracy by mostly predicting the dominant class. Saito and Rehmsmeier [55] emphasise that performance evaluation can be misleading under class imbalance, especially when relying on metrics that do not clearly reflect minority-class performance. This is relevant in cryptocurrency prediction because a model that performs well only for the more frequent class may have limited value for risk-management decisions.

Precision and recall provide a more detailed view of classification quality. Precision shows how many observations predicted as a given class were actually correct, while recall shows how many true observations of that class were successfully identified. In cryptocurrency prediction, this distinction has practical meaning. High precision for the UP class means that when the model predicts an upward movement, this signal is often correct. High recall for the UP class means that the model captures many actual upward movements. The same logic applies to the DOWN class, where recall may be especially important for risk reduction because it indicates how well the model identifies negative market movements [55].

The F1 score combines precision and recall into a single metric by calculating their harmonic mean. This is useful when both false positives and false negatives matter. In directional cryptocurrency prediction, a false UP signal may lead to entering or holding a position before a decline, while a false DOWN signal may cause the investor to miss a positive movement. Macro F1 is particularly useful because it calculates the F1 score separately for each class and then averages them without weighting by class frequency. This means that UP and DOWN performance contribute equally to the final metric. Therefore, macro F1 provides a more balanced assessment than accuracy when both positive market opportunities and negative risk periods are important [55].

Class-specific performance is also needed for economic interpretation. UP and DOWN predictions do not have the same practical meaning. UP predictions may be associated with buying or holding, while DOWN predictions may be associated with selling, avoiding exposure or reducing risk. Therefore, class-specific precision, recall and F1 scores allow the researcher to evaluate whether the model is better at identifying positive opportunities or negative market periods [55]. This is important because a model may achieve acceptable aggregate performance while still performing poorly for the class that is more relevant for risk reduction.

However, classification metrics alone do not show whether model predictions are economically useful. A model may correctly classify many small movements but fail during large price changes, or it may generate signals that are not useful once trading assumptions are considered. Therefore, model outputs should also be connected with practical decision-making through backtesting. In cryptocurrency prediction, an UP or DOWN class label can be translated into a simple trading rule: an UP prediction may be interpreted as a buy or hold signal, while a DOWN prediction may indicate selling, avoiding exposure or moving to cash.

The literature on directional cryptocurrency prediction often connects model outputs with trading strategy evaluation. Omole and Enke [37] compare Bitcoin price-direction models not only through predictive metrics, but also through trading strategies based on model outputs. This shows that a directional model should not be assessed only by whether it predicts class labels correctly, but also by whether its signals can support profitable or risk-reducing decisions. Backtesting provides the main tool for this type of economic evaluation because it shows how model-generated signals would have performed under defined trading assumptions.

A common benchmark for this evaluation is the buy-and-hold strategy. Buy-and-hold represents passive exposure to the cryptocurrency over the test period, while a model-based strategy adjusts exposure according to predicted signals. For example, the model-based strategy may hold the asset during UP signals and move to cash during DOWN signals. If the model-based strategy outperforms buy-and-hold, this suggests that the predictions may have practical value. If it underperforms, then good classification performance may not be sufficient for economic usefulness. However, backtesting results should be interpreted cautiously because strategy performance can be affected by the selected period, trading assumptions and possible overfitting to historical data [56].

Backtesting also reveals weaknesses that may not be visible in aggregate classification metrics. A model may have acceptable F1 score but perform poorly if false UP signals occur before large market declines. Similarly, a model may correctly predict many small downward movements but miss major crashes. This means that the timing and economic size of correct and incorrect predictions matter. Bouteska et al. [36] show that cryptocurrency forecasting models can be evaluated through both forecasting performance and trading outcomes, which supports the need to connect predictive results with economic performance.

Overall, economic evaluation provides a bridge between machine learning prediction and practical interpretation. Classification metrics show whether the model identifies market direction, while backtesting shows whether these predictions can be translated into economically meaningful trading behaviour. For this thesis, this step is necessary because Random Forest classification outputs are evaluated not only as statistical predictions, but also as market signals compared with a buy-and-hold benchmark. This supports the broader research problem of the thesis: prediction accuracy alone is not sufficient unless model decisions can also be interpreted and assessed in terms of economic relevance.

1.8. Summary of the Literature Review and Justification of the Thesis Topic

1.8.1. Main findings from previous research

The reviewed literature shows that cryptocurrency markets differ from traditional financial markets because they are decentralised, traded continuously, highly volatile and strongly affected by uncertainty, speculation and investor attention. Nakamoto's [1] Bitcoin concept introduced a decentralised peer-to-peer payment system, while the Bank for International Settlements [2] later emphasised that decentralisation creates both technological opportunities and practical limitations. These characteristics make cryptocurrency markets difficult to analyse using only traditional financial assumptions and justify the use of data-driven approaches that can account for complex and changing market behaviour.

Previous research also shows that cryptocurrency price formation is multidimensional. Ciaian, Rajcaniova and Kancs [12] show that Bitcoin prices are influenced by both supply-demand factors and cryptocurrency-specific attractiveness factors. Kristoufek [13] similarly finds that Bitcoin price dynamics are shaped by fundamental, speculative and technical drivers, and that these relationships vary across time horizons. This means that cryptocurrency prices cannot be explained by one factor alone. Market-based indicators, liquidity, technical variables, investor attention, macroeconomic conditions and external information may all contribute to price movements.

A consistent finding in the literature is that cryptocurrency markets are volatile and unstable across time. Katsiampa [20] shows that Bitcoin volatility has persistent components and can be modelled using GARCH-type approaches, while Baur and Dimpfl [14] demonstrate that Bitcoin volatility is substantially higher than that of major fiat currencies. Regime-switching studies further indicate that cryptocurrency markets may move through different states, such as bull, bear, calm or high-volatility regimes [21,22]. Therefore, cryptocurrency market dynamics should not be treated as constant across the whole sample period.

The literature also supports the relevance of investor sentiment and news-based information. Studies show that cryptocurrency prices and volatility can be affected by online attention, sentiment, global uncertainty, regulatory news and security-related events [16,18,19,29,30]. Lyócsa et al. [16] find that Bitcoin volatility reacts to regulation-related news and hacking incidents, while Sakariyahu et al. [18] show that global uncertainty and sentiment factors affect cryptocurrency returns. These findings justify the inclusion of news and sentiment indicators together with market-based variables.

At the same time, previous research shows that cryptocurrency predictability is limited and unstable. Wei [7] finds that return predictability decreases as cryptocurrency liquidity increases, while Tran and Leirvik [9] show that crypto-market efficiency changes over time. This suggests that predictive signals may exist, but their strength depends on asset maturity, liquidity, market regime, sentiment conditions and the analysed period. Therefore, models used for cryptocurrency prediction should be evaluated not only by aggregate predictive performance, but also by whether their decisions remain interpretable under changing market conditions.

Machine learning methods are useful in this context because they can model nonlinear relationships, interactions between variables and complex market signals [34,35,36,39]. However, the literature does not show that machine learning automatically solves the problem of cryptocurrency prediction. Model performance remains sensitive to the selected features, time period, validation design and market conditions. More importantly, complex models may produce predictions without clearly showing which variables influenced the decision. This creates the need for explainable artificial intelligence methods, which help evaluate whether model outputs are based on meaningful market signals rather than unstable or spurious relationships [46,47,48].

Finally, the reviewed literature shows that statistical predictive performance does not necessarily imply economic usefulness. A model may achieve acceptable accuracy or F1 score, but still fail when its predictions are translated into trading decisions. For this reason, model-based signals should be evaluated through backtesting and compared with simple benchmarks such as buy-and-hold [36,37]. However, economic evaluation should also be interpreted cautiously, because trading performance may depend on the selected period, threshold logic and market conditions [56].

Overall, previous research indicates that cryptocurrency markets are volatile, sentiment-sensitive, nonlinear and regime-dependent. Machine learning methods can help analyse such complex relationships, but prediction alone is not sufficient. Explainable artificial intelligence methods are needed to interpret model decisions, while regime-aware analysis and backtesting are necessary to evaluate whether these decisions are stable and economically meaningful. These findings provide the foundation for this thesis, which combines directional cryptocurrency prediction, HMM-based regime identification, SHAP-based explainability and trading-signal backtesting.

1.8.2. Identified Research Gap

The reviewed literature shows that cryptocurrency market dynamics have been analysed from several perspectives, including volatility modelling, sentiment analysis, market efficiency, machine learning prediction, regime identification and trading-signal evaluation [7,9,16,18,20,21,22,35,36,37]. However, these perspectives are often examined separately. Many studies focus on whether a model can improve forecasting or directional classification performance compared with statistical benchmarks, but prediction accuracy alone does not explain how the model forms its decisions or whether those decisions are meaningful in different market conditions.

This creates a scientific gap between machine learning-based prediction and the interpretation of model decisions in cryptocurrency market analysis. A model may classify UP and DOWN movements with acceptable accuracy, but this does not show which market, sentiment, macroeconomic or regime-related variables drive the prediction. It also does not show whether the model behaves differently during bull, bear, sideways or panic-like regimes, or whether the generated signals have practical economic meaning. Therefore, the problem is not only whether cryptocurrency market direction can be predicted, but whether machine learning model decisions can be explained and evaluated in relation to market dynamics.

Explainable artificial intelligence methods address part of this gap by making complex models more transparent [46,47,48]. However, in cryptocurrency prediction research, explainability is often treated as an additional post-model step rather than as a central part of market-dynamics analysis. Feature importance or SHAP values may be used to identify influential predictors, but less attention is given to connecting these explanations with market regimes and economic signal interpretation. As a result, it may remain unclear whether model predictions are based on economically meaningful market behaviour or on patterns that are unstable across time.

A further gap concerns the integration of different information sources. Previous studies show that technical variables, sentiment, uncertainty, macroeconomic conditions and regime shifts can influence cryptocurrency returns and volatility [13,16,18,19,21,22,26,30]. However, fewer studies combine these components in one empirical framework and then examine how they contribute to machine learning model decisions. This is important because the meaning of a variable may change depending on market conditions. For example, volatility, drawdown or negative news may have different implications during panic periods than during stable or trending markets.

Economic interpretation is also part of the research gap. Some studies connect directional predictions with trading strategies and backtesting, but backtesting is not always analysed together with explainability and regime-aware evaluation [36,37]. Therefore, it may remain unclear whether a model-based strategy performs well because it captures meaningful market dynamics or because it fits a specific historical period. For this reason, economic evaluation should be interpreted together with model explanations, market-state analysis and the risk of backtest overfitting [56].

This thesis addresses the identified gap by analysing machine learning model decisions in cryptocurrency price-dynamics prediction rather than focusing only on predictive accuracy. The empirical framework combines Random Forest UP/DOWN classification, HMM-based market regime identification, SHAP-based explainability and backtesting of model-generated trading signals. This allows the study to evaluate whether model decisions are interpretable, whether they differ across market regimes and whether they can be connected with economically meaningful trading-signal behaviour.

2. Research Methodology for Explainable Analysis of Cryptocurrency Market Dynamics

2.1. Research Design and Object of the Study

The object of this research is the market dynamics of three cryptocurrencies: Bitcoin, Ethereum and Solana. These assets were selected because they are actively traded but differ in market maturity, historical development and volatility. Bitcoin represents the most established cryptocurrency market, Ethereum reflects a major blockchain ecosystem, and Solana represents a newer and more volatile asset with a shorter available trading history. Therefore, the assets were analysed separately in order to preserve their individual market characteristics.

The empirical research was designed to evaluate cryptocurrency market dynamics not only through prediction accuracy, but also through model interpretability and economic meaning. For this reason, the study combines daily market, news, macro-financial and regime-related variables with machine learning classification and explainability methods [34,36,46,48]. The target variable was formulated as a binary UP/DOWN market direction indicator rather than an exact price or return forecast, because directional predictions can be more directly connected with trading-signal interpretation and risk-management decisions [37,39].

The research followed a chronological empirical design. First, cryptocurrency market data, news-based indicators and macro-financial variables were collected, cleaned and merged at daily frequency. Then, market regimes were identified using Hidden Markov Models, and separate Random Forest classification models were developed for Bitcoin, Ethereum and Solana [21,22,39]. Finally, the models were evaluated using classification metrics, interpreted using explainability methods, and assessed economically through backtesting [45,46,55,56].

A chronological train-validation-test split was used to reduce the risk of data leakage. The training set was used for model fitting, the validation set for model selection and hyperparameter tuning, and the test set for final out-of-sample evaluation. This structure was selected because cryptocurrency observations are time-ordered, and random splitting could lead to overly optimistic results by allowing future information to influence model development [40,41].

The main stages of the empirical research design are summarised in Table 1.

Table 1. Overview of the empirical research design

Component	Methodological choice	Reasoning
Research object	Bitcoin, Ethereum and Solana	Assets differ in maturity, volatility and market behaviour
Prediction task	Directional classification	More suitable for UP/DOWN trading-signal interpretation than exact price forecasting [37,39]
Model type	Random Forest classifier	Handles nonlinear relationships and interactions between variables [34,39]
Data frequency	Daily observations	Allows market, news, macroeconomic and regime variables to be aligned

Data split	Chronological train-validation-test split	Reduces data leakage risk in time-series modelling [40,41]
Regime analysis	Hidden Markov Models	Identifies latent market states with different risk-return characteristics [21,22,44]
Explainability	Feature importance and SHAP	Explains global and local model behaviour [45,46,51]
Economic evaluation	Backtesting against buy-and-hold	Tests whether predictions have practical economic meaning [37,56]

2.2. Cryptocurrency Market Data

Cryptocurrency market data were collected for Bitcoin, Ethereum and Solana. The market data were collected from the Binance exchange using the `ccxt` Python library. The analysed trading pairs were BTC/USDT, ETH/USDT and SOL/USDT. The same USDT quote currency was used for all three assets to keep the price series expressed in a comparable dollar-denominated form.

The collected data consisted of daily OHLCV candlesticks, including opening price, highest daily price, lowest daily price, closing price and trading volume. Daily frequency was selected because the broader modelling dataset also included daily cryptocurrency news indicators, macroeconomic news variables and macro-financial indicators. This allowed all data sources to be aligned by date and merged into one modelling dataset. OHLCV-based variables are commonly used in cryptocurrency modelling because they capture market-generated information such as returns, volatility, trend behaviour, liquidity and trading activity [11,13,23,24,25].

Bitcoin and Ethereum observations were available from the beginning of the selected collection period, while Solana observations started later because Solana became public later than monitoring start. This difference was retained during data preparation instead of being artificially filled, because creating pre-launch market data would introduce artificial information. As a result, Solana has a shorter available modelling history than Bitcoin and Ethereum. The asset coverage and available observation periods are summarised in Table 2

Table 2. Cryptocurrency market data collected for the study

Asset	Trading pair	First available observation	Last available observation
Bitcoin	BTC/USDT	2019-01-01	2026-03-31
Ethereum	ETH/USDT	2019-01-01	2026-03-31
Solana	SOL/USDT	2020-08-11	2026-03-31

The raw OHLCV dataset contained 2647 observations and 16 columns, covering the period from 1 January 2019 to 31 March 2026. The expected OHLCV columns for Bitcoin, Ethereum and Solana were checked before feature creation, and no required raw OHLCV columns were missing. This confirmed that the collected market dataset was structurally suitable for the next stage of feature engineering.

2.2.1. Construction of Derived OHLCV-Based Market Variables

After collecting the raw OHLCV data, additional market variables were constructed separately for Bitcoin, Ethereum and Solana. This step was necessary because raw price levels alone do not fully describe cryptocurrency market behaviour and are less suitable for comparing assets with different price scales. Therefore, the feature engineering stage focused on relative and transformed variables describing price change, volatility, trend, momentum, drawdown and trading activity [11,13,23].

For each asset, ten derived variables were created: simple return, log return, 7-day and 30-day rolling volatility, 7-day and 30-day moving average, 30-day momentum, drawdown, volume change and close-to-30-day-moving-average ratio. Since the same variables were created for BTC, ETH and SOL, the feature engineering stage produced 30 new variables in total.

Simple return was calculated as the relative change in closing price between two consecutive days [23]:

$$\boldsymbol{R_t} = \frac{\boldsymbol{P_t} - \boldsymbol{P_{t-1}}}{\boldsymbol{P_{t-1}}} \qquad \textbf{(2.1)}$$

where P_t is the closing price at day t, and P_{t-1} is the closing price on the previous day. This variable measures the daily percentage price change and directly represents market direction and movement size.

Logarithmic return was calculated as [23]:

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \qquad (2.2)$$

where $r_t$is the log return at day t. Log returns were included because they are commonly used in financial time-series analysis and represent proportional price changes in a form suitable for volatility and risk-related calculations [23].

Rolling volatility was calculated using the standard deviation of log returns over a selected rolling window [14,20]:

$$\sigma_{t,n} = SD(r_{t-n+1}, \dots, r_t) \qquad (2.3)$$

where n represents the rolling window length. In this study, $n = 7$ and $n = 30$ were used. The 7-day volatility captures short-term uncertainty, while the 30-day volatility reflects broader monthly market risk conditions. These variables were included because cryptocurrency markets often experience strong volatility clustering and rapidly changing risk levels [14,20].

Moving averages were calculated as the average closing price over a rolling window [24,25]:

$$MA_{t,n} = \frac{1}{n}\sum_{i=0}^{n-1} P_{t-i} \qquad (2.4)$$

where $MA_{t,n}$ is the moving average at day t, calculated over n days. In this study, 7-day and 30-day moving averages were created. The shorter moving average captures recent price behaviour, while the longer moving average represents a smoother medium-term trend.

To express the current price relative to its recent trend, the close-to-moving-average ratio was calculated [24,25]:

$$Ratio_t = \frac{P_t}{MA_{t,30}} \tag{2.5}$$

This variable shows whether the current closing price is above or below its 30-day average. It is useful because it represents the asset's position relative to its recent trend rather than relying on the absolute price level.

A 30-day momentum variable was calculated as [13,24,25]:

$$Momentum_{t,30} = \frac{P_t - P_{t-30}}{P_{t-30}} \tag{2.6}$$

This variable measures the percentage change in price over the previous 30 days. It was included to capture medium-term trend continuation or reversal behaviour, which may be relevant during bull and bear market periods.

Drawdown was calculated as the decline from the previous running maximum closing price [14,20]:

$$Drawdown_t = \frac{P_t - \max(P_1, \ldots, P_t)}{\max(P_1, \ldots, P_t)} \tag{2.7}$$

This variable measures how far the asset price has fallen from its previous peak. It captures downside market stress and provides information that daily returns alone may not show, because an asset can remain in a deep drawdown even when daily movements are moderate.

Volume change was calculated as the relative change in trading volume [7,11]:

$$VolumeChange_t = \frac{V_t - V_{t-1}}{V_{t-1}} \tag{2.8}$$

where $V_t$is trading volume at day t. This variable was included because changes in trading volume can indicate shifts in market activity, liquidity and investor attention, especially during rallies, sell-offs or news-driven events [7,11].

The feature engineering process naturally created missing values at the beginning of the time series. Return variables require a previous observation, while rolling volatility, moving averages and momentum variables require a complete historical window. These missing values were treated as a natural consequence of the calculations and data availability, not as data errors.

2.3. Macroeconomic and Financial Market Data Collection

Macroeconomic and financial market variables were included to represent the broader external environment in which cryptocurrency markets operate. This data layer was added because cryptocurrency prices may be influenced not only by internal market behaviour, but also by inflation, liquidity conditions, monetary policy, bond yields, the strength of the United States dollar, stock market performance, financial uncertainty and safe-haven asset dynamics. Since Bitcoin, Ethereum and Solana are traded as speculative financial assets, broader macro-financial conditions may affect investor risk appetite and therefore the probability of upward or downward market movement [15,18].

The macroeconomic and financial market data were collected from FRED and Yahoo Finance APIs and later merged into a daily dataset. Variables collected, their meaning and reason for inclusion is available in table 3.

Table 3. Macroeconomic and Financial Market Variables Used in the Study

Variable	Source Type	Meaning	Reason for inclusion	Supporting literature
M2_money_stock	FRED	Broad money supply measure	Represents liquidity conditions and monetary expansion	[15,18,33]
CDI_all_items	FRED	Consumer Price Index	Captures inflationary environment	[18,33]
federal_funds_rate	FRED	Federal funds rate variable	Represents short-term monetary policy stance	[15,18,33]
us10y_yeald_fred	FRED	US 10-year Treasury yield	Captures long-term interest-rate expectations	[15,18,33]
broad_usd_index	FRED	Broad US dollar index	Represents general US dollar strength	[3,15,33]
vix_fred	FRED	Volatility/risk sentiment indicator	Captures financial-market uncertainty	[3,15,18]
sp500_fred	FRED	S&P 500 market indicator	Represents broader risky-asset market performance	[3,15,18]
dxy_yf	Yahoo Finance	US Dollar Index	Captures dollar-market movements	[3,15,33]
gold_fut	Yahoo Finance	Gold futures price	Represents safe-haven and store-of-value dynamics	[3,14,39]
sp500_yf	Yahoo Finance	S&P 500 index	Captures equity-market performance	[3,15,18]
us10y_yield	Yahoo Finance	US 10-year yield series	Captures bond-market and rate expectations	[15,18,33]
vix_df	Yahoo Finance	VIX volatility index	Represents market uncertainty and risk aversion	[3,15,18]

2.4. Cryptocurrency News Data Collection and Aggregation

Cryptocurrency-specific news data were collected from the Global Database of Events, Language and Tone using Google BigQuery. This data source was included to capture the external information environment surrounding Bitcoin, Ethereum, Solana and the broader cryptocurrency market. The query covered the period from 1 January 2019 to 31 December 2025 and used cryptocurrency-related keywords and organisation names, including terms related to Bitcoin, Ethereum, Solana, the general crypto market, exchanges and institutional actors. The raw extraction contained 4,056,039 article-level observations with GDELT metadata fields such as publication date, document identifier, themes, organisations and tone.

During preprocessing, publication dates were converted to daily format, tone scores were extracted, theme and organisation fields were cleaned, and duplicate records were removed. Articles were then assigned to BTC, ETH, SOL or GENERIC_CRYPTO depending on whether they mentioned a specific asset or the broader cryptocurrency market. The generic category was retained because many articles refer to the cryptocurrency market as a whole rather than to one individual asset.

Event-category indicators were created using keyword patterns for five crypto-related news groups: regulation, hacks and fraud, ETF or institutional adoption, market crashes, and technology or network developments. These categories were selected because previous studies show that cryptocurrency markets may react to regulation-related news, hacking incidents, security events, institutional developments and market-wide sentiment changes [16,31,32]. A negative-tone flag was also created when the GDELT tone score was below zero. The processed article-level data were then aggregated by date and asset category. The final daily variables included article count, mean tone, negative-news share and topic-specific article counts for each event category.

A complete date-asset grid was created for all dates and asset categories, so that days without relevant news were retained in the final dataset. Missing article counts and topic counts were filled with zero, and missing tone-based values were also set to zero. The main limitation of this approach is that the filtering and event classification rely on keyword rules and automated GDELT tone scores, which may leave some irrelevant articles in the dataset or miss relevant articles without the selected terms.

2.5. General Macroeconomic News Data Collection and Aggregation

General macroeconomic news data were collected from GDELT to represent the broader information environment surrounding cryptocurrency markets. This data layer was included because cryptocurrency prices may react not only to crypto-specific events, but also to global uncertainty, monetary policy, geopolitical shocks, financial crises, energy-market developments and pandemic-related news. Previous research shows that global uncertainty, macro-financial conditions and crisis-related events can affect cryptocurrency returns, volatility and hedging behaviour [15,18,33].

The macroeconomic news query used selected keywords and themes related to global economic and geopolitical uncertainty. The selected topic groups included pandemic-related news, war, financial crises, monetary policy, energy crises and supply-chain disruptions. Organisation-based indicators were also created for institutions and groups such as the Federal Reserve, the European Central Bank, the World Health Organization, OPEC and central banks. These variables were included to capture the intensity of macroeconomic and geopolitical news coverage, since news-based uncertainty and sentiment may influence investor risk appetite and market reactions [18,30,33].

After extraction, the article-level GDELT fields were cleaned and transformed. Publication dates were converted to daily format, tone scores were extracted, duplicate records were removed, and only variables required for aggregation were retained. The cleaned records were aggregated by date into daily total article count, mean tone, tone standard deviation, topic-specific article counts and organisation-based counts. These variables were designed to measure the intensity and tone of macroeconomic and geopolitical news coverage rather than actual macroeconomic outcomes.

The final macroeconomic news dataset covered the period from 1 January 2019 to 31 December 2025 and was later merged with cryptocurrency market data, crypto-specific news indicators, macro-financial variables and market-regime features. As with cryptocurrency news data, the main limitation is that the approach relies on keyword filtering and automated GDELT metadata, so the resulting indicators may contain some noise or miss articles without the selected terms [28,30].

2.6. Dataset Standardisation, Exploratory Data Analysis and Data Quality Assessment

Before modelling, all collected datasets were standardised and examined through exploratory data analysis. This step was necessary because the study combines several data sources with different structures: cryptocurrency OHLCV data, crypto-specific news indicators, macroeconomic news indicators and macro-financial variables. Since the final modelling dataset required one observation per day, all date columns were converted into a common `date` format and the datasets were aligned at daily frequency. This was important because time-series modelling requires observations to be ordered and aligned consistently over time [40,41].

The standardisation stage included four main input datasets: cryptocurrency news data, OHLCV market data, macro-financial market data and general macroeconomic news data. The cryptocurrency news dataset was originally stored in long format, where each date had separate rows for BTC, ETH, SOL and `GENERIC_CRYPTO`. For modelling purposes, it was pivoted into wide format, so that each date had one row and asset-specific news variables appeared as separate columns, such as `BTC_article_count`, `ETH_article_count` and `GENERIC_CRYPTO_mean_tone`.

After standardisation, the datasets were merged by date. The merged exploratory dataset contained 2,647 observations and 115 variables, covering the period from 1 January 2019 to 31 March 2026. This broader dataset was used to evaluate data completeness, duplicate records, variable distributions and correlation structure before creating the cleaned after-EDA datasets for modelling.

The duplicate-date check showed that the main datasets did not contain duplicated date rows after standardisation. In the original crypto-news long-format dataset, repeated dates occurred because each date was represented by several asset categories, but duplicate date-asset pairs were not present. After pivoting to wide format, the dataset also had one row per date. This confirmed that the data structure was suitable for merging and later time-series modelling.

Missing-value analysis was performed to identify whether gaps were caused by errors or by expected data-availability differences. The largest missing-value group was related to Solana variables, because SOL/USDT data started later than BTC and ETH.

Macroeconomic and financial market variables also contained some missing values. This was expected because traditional financial market and macroeconomic indicators do not always follow the same calendar as continuously traded cryptocurrency markets. Monthly macroeconomic variables, such as money supply, CPI and interest-rate indicators, also required alignment with daily data. These gaps were documented before the final modelling dataset was prepared.

Correlation analysis was used to identify strongly related or duplicated variables. This was important because highly correlated features can reduce the clarity of feature importance and SHAP explanations. The analysis showed strong relationships between some market-derived variables, especially across major cryptocurrencies. For example, BTC and ETH return-based variables were highly correlated, and several volatility and momentum variables also showed cross-asset correlation. These findings were expected because major cryptocurrencies often move together during broad market trends. Correlation analysis was therefore used as a diagnostic tool during feature selection, not as a purely automatic deletion rule [50,51].

General variable distributions were also reviewed, especially for news, macroeconomic and market indicators. This helped identify extreme values, low-variation variables and differences in scale between feature groups.

Overall, the standardisation and EDA stage ensured that the modelling data were date-aligned, structurally consistent and suitable for supervised machine learning. The main issues identified were expected and explainable: Solana had a shorter available history, macro-financial variables had calendar-related gaps, and some macro-news indicators were constant or zero-only.

2.7. Market Regime Identification Using Hidden Markov Models

Market regimes were identified using Gaussian Hidden Markov Models. The purpose of this stage was to capture latent statistical states in cryptocurrency market behaviour and to use these states as additional explanatory variables in the modelling dataset. The regimes were estimated separately for Bitcoin, Ethereum and Solana, because each asset has different volatility, price history and market maturity. The use of regime-switching logic is consistent with financial time-series literature, where market behaviour may alternate between unobserved states with different statistical characteristics [44]. In cryptocurrency research, HMM and regime-switching approaches have also been used to identify return and volatility regimes [21,22].

For each cryptocurrency, the HMM was fitted using selected OHLCV-derived variables that describe short-term market behaviour. The input variables included logarithmic return, 7-day rolling volatility and drawdown. These variables were selected because they represent price direction, recent uncertainty and downside pressure. Before fitting the models, the input variables were standardised to prevent variables with larger numerical scales from dominating the estimation. The use of return and volatility-related inputs is consistent with previous cryptocurrency regime and volatility studies [20,21,22].

Candidate models with different numbers of hidden states were compared using AIC and BIC values. The final regime structure was selected by considering both model-selection criteria and the need to avoid excessive fragmentation of the time series. A four-state structure was retained for each cryptocurrency and later used consistently in the empirical analysis.

The HMM output was not treated as a direct classification of real-world market states. The numerical regime labels are arbitrary and do not have an automatic economic meaning. Therefore, the regimes were interpreted only through their descriptive statistics, such as average return, volatility and drawdown, and through their location in the price series. In the modelling dataset, the inferred regimes were one-hot encoded and merged by date with the other explanatory variables. Latent regimes require economic labelling after estimation rather than being automatically defined by the model [21,22,44].

2.8. Feature Selection and Correlation Handling

Feature selection was applied to reduce the broad initial set of candidate variables to a more stable and interpretable modelling set. The dataset contained several feature groups: OHLCV-derived market variables, macro-financial indicators, macroeconomic news variables, crypto-specific news indicators and HMM-based regime variables. This structure was used because each group represented a different information channel: internal market behaviour, broader financial conditions, macroeconomic uncertainty, crypto-related information flow and latent market states.

The main reason for feature selection was to reduce noise, redundancy and interpretation problems. Feature selection is commonly used in machine learning to improve model efficiency, reduce irrelevant inputs and support a clearer understanding of the analysed process [34]. This was particularly important in this thesis because the final objective was not only predictive accuracy, but also explainability through feature importance and SHAP analysis [45,46].

A key principle was to avoid unnecessary reliance on absolute price levels, raw volumes and raw index levels where possible. Cryptocurrency prices can move into very different ranges over time, so raw levels may increase the risk of distribution shift between training and testing periods. Therefore, transformed variables were preferred when possible, including returns, log returns, percentage changes, volatility measures, moving-average ratios, drawdown indicators and lagged return variables. These variables describe relative market behaviour more clearly than raw levels [23].

Correlation analysis was used to identify variables carrying highly similar information. The Pearson correlation coefficient was used:

$$r_{XY} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}} \quad (2.9)$$

where r_{XY} measures the linear relationship between variables (X) and (Y). Values close to (1) or (-1) indicate strong association, while values close to (0) indicate weak linear association. In this study, correlation analysis was used as a diagnostic tool rather than an automatic deletion rule [50].

Correlation handling was especially important because many derived OHLCV variables are calculated from the same underlying price series. Moving averages, momentum, volatility measures and price ratios may partially overlap. Cross-asset variables may also be correlated because major cryptocurrencies often move together during broad market cycles. Strongly correlated variables were therefore reviewed carefully, as they can make feature importance and SHAP explanations less stable by distributing explanatory contribution across similar predictors. This issue is also relevant for Random Forest models, since variable importance measures can be biased when predictors are correlated or differ in scale [50,51].

The final feature lists were selected from the broader candidate set by considering economic meaning, correlation structure, distribution-shift risk and interpretability. Variables were more likely to be retained if they represented a distinct information channel, such as market direction, volatility, drawdown, news sentiment, macroeconomic conditions or market regime. Variables were more likely to be excluded if they duplicated other predictors, had weak variation, relied too heavily on raw levels or reduced interpretability.

2.9. Random Forest Classification Model

The predictive modelling stage was based on the Random Forest classifier. Random Forest is an ensemble learning method that combines multiple decision trees and makes the final classification decision by aggregating their outputs. For a classification task, each tree $h_b(x)$predicts a class label, and the Random Forest prediction is obtained by majority voting [39]:

$$\hat{y} = \mathrm{mode}\,\{h_1(x), h_2(x), \dots, h_B(x)\}, \quad (2.10)$$

where B is the number of trees, x is the input feature vector, and $\hat{y}$ is the predicted class. This approach reduces the instability of individual decision trees and allows the model to capture nonlinear relationships and interactions between variables [34,39].

Random Forest was selected because cryptocurrency market direction may depend on complex interactions between market-based indicators, news sentiment, macroeconomic variables and regime features. The model is suitable for this setting because it does not require a predefined linear relationship between predictors and the target variable. It can also handle mixed feature groups and provides feature-importance measures, which supports the explainability objective of the thesis [39].

Separate Random Forest models were trained for Bitcoin, Ethereum and Solana. This asset-specific approach was used because each cryptocurrency has different market maturity, volatility, liquidity and reaction to external information. Training one common model for all assets could hide these differences and make interpretation less clear.

The prediction task was formulated as binary UP/DOWN classification. Instead of estimating the exact future price or return, the model classified whether the future market movement was positive or negative. This formulation was chosen because the empirical objective was to analyse directional market signals and connect predictions with later backtesting logic [37,39].

The modelling procedure used a chronological train-validation-test split. The training set was used to fit the model, the validation set was used for feature selection and hyperparameter tuning, and the test set was reserved for final out-of-sample evaluation. Chronological splitting was necessary because cryptocurrency observations are time ordered; random splitting could introduce data leakage by allowing future information to influence model development [40,41].

Feature scaling was applied before modelling to keep the pipeline consistent across feature groups and later explainability steps. Although Random Forest does not strictly require scaling, because tree-based splits are not distance-based, standardisation helped maintain a uniform preprocessing structure for market, news, macroeconomic and regime variables.

Model performance was evaluated using accuracy, macro F1, class-specific F1 for UP and DOWN, classification report and confusion matrix. Accuracy measured the overall share of correct predictions, while macro F1 and class-specific F1 scores were used to evaluate whether the model performed consistently across both directional classes. Final model performance results are reported in the results chapter.

2.9.1. Hyperparameter Tuning and Model Selection

After defining the Random Forest modelling framework, hyperparameter tuning was used to select model settings and feature combinations based on validation performance. This step was necessary because Random Forest models can behave differently depending on parameters such as tree depth, number of trees, minimum leaf size and feature sampling. If these values are chosen arbitrarily, the model may either underfit by failing to capture relevant market patterns or overfit by learning noise from the training period [42].

Optuna was used for automated hyperparameter optimisation. This tool was selected because it allows a more efficient search through the parameter space than manual trial-and-error or a fixed grid search. Instead of testing only a predefined set of combinations, Optuna can explore alternative configurations and focus the search on more promising areas. This was useful because the modelling pipeline had to be repeated separately for BTC, ETH and SOL, and each asset could require different parameter settings due to different volatility, liquidity and market-history characteristics [43]

The optimisation objective was validation macro F1 score. This metric was chosen because the task was directional UP/DOWN classification, where both classes are important. Macro F1 was calculated as [55]:

$$F1_{macro} = \frac{1}{K}\sum_{k=1}^{K} F\,1_k \tag{2.11}$$

where K is the number of classes and $F1_k$ is the F1 score for class k. In this study, $K = 2$, corresponding to the UP and DOWN classes. Optimising only accuracy could favour the more frequent class, while macro F1 gives equal weight to UP and DOWN performance. This was especially relevant because the model outputs were later used for economic interpretation and backtesting, where both missed upward opportunities and missed downward risk periods matter.

Feature selection was also treated as part of model development. The initial feature pool included several groups of variables: OHLCV-derived indicators, macro-financial variables, macro-news indicators, crypto-news indicators and HMM regime variables. During development, different feature combinations were tested using group comparison, ablation logic and random subset search. This was done to identify whether adding a feature group improved validation performance or only added noise and reduced interpretability. The search process was therefore used not simply to maximise metrics, but to select a model configuration that remained understandable and economically meaningful.

The search stage was not treated as a final result. It was used only to define the final model configuration. After the best feature set and hyperparameter combination were selected on the validation period, this configuration was fixed and then evaluated on the test period. This separation was important because using the test set during tuning would create an overly optimistic estimate of model performance. Therefore, the validation set supported model selection, while the test set was reserved for final out-of-sample assessment [40,41].

2.10. Explainability Analysis

Explainability analysis was performed after training the final Random Forest classification models. This step was necessary because the aim of the thesis was not only to predict cryptocurrency market direction, but also to understand which variables influenced the model’s decisions. Three interpretation approaches were used: impurity-based Random Forest feature importance, SHAP analysis and PDP/ICE analysis [45,46].

First, Random Forest feature importance was calculated using the impurity-based importance measure. This method evaluates how much each variable contributes to reducing classification impurity across the decision trees in the forest. A variable receives higher importance if it is frequently used in splits that improve class separation. This provided an initial global view of which features contributed most to the UP/DOWN classification task. However, impurity-based importance was treated only as a preliminary interpretation tool because it can be affected by correlated variables and does not show the direction of feature influence [51].

To obtain a more detailed explanation, SHAP values were calculated using the Tree Explainer approach, which is designed for tree-based models such as Random Forest. SHAP assigns each feature a contribution value for a model prediction. In simplified form, the model output can be decomposed as [45,52]:

$$f(x) = \phi_0 + \sum_{j=1}^{M} \phi_j \tag{2.12}$$

where $f(x)$is the model prediction for observation x, $\phi_0$is the baseline model output, $M$is the number of features, and $\phi_j$is the SHAP contribution of feature j. A positive SHAP value means that the feature pushed the prediction toward the analysed class, while a negative SHAP value means that the feature pushed the prediction away from that class.

SHAP values were calculated on the test set, because the purpose was to explain out-of-sample model behaviour rather than only the fitted training data. Since the models were binary classifiers, the interpretation focused on class 1, corresponding to the UP class. This means that SHAP values were interpreted as contributions pushing the model toward or away from an upward market prediction [45].

Several SHAP visualisations were used. SHAP beeswarm plots were used for global interpretation because they show both the importance of variables and the direction of their influence across test-set observations. SHAP dependence plots were used to examine how the value of the most important feature was related to its SHAP contribution. Local SHAP waterfall plots were used to explain individual predictions by showing how the model moved from the baseline probability to the final predicted probability for one selected observation [50].

PDP and ICE plots were also used as complementary interpretability tools. Partial Dependence Plots show how the average predicted probability of the UP class changes when one selected feature varies, while other features are averaged over. Individual Conditional Expectation plots show the same relationship separately for individual observations. Therefore, PDP provides the average model response, while ICE shows whether this response is consistent across different observations [50]. In this thesis, PDP/ICE plots were used mainly for the most important return-based variables in order to verify whether the model response supported the SHAP interpretation.

These explanation methods were used to connect model predictions with market interpretation. They allowed the analysis to evaluate whether the models relied more on return-based variables, volatility indicators, news sentiment, macroeconomic indicators or market-regime variables. However, all explainability results were interpreted cautiously. Feature importance, SHAP and PDP/ICE explain how the trained model used the available variables, but they do not prove that these variables caused cryptocurrency price movements. Therefore, the explainability analysis was treated as model-behaviour interpretation rather than causal inference.

2.11. Backtesting Procedure

Backtesting was used to evaluate whether Random Forest classification outputs could be translated into economically meaningful trading signals. This step was necessary because classification metrics such as accuracy or F1 score show whether the model predicts UP and DOWN movements correctly, but they do not show whether these predictions would improve investment performance. Therefore, the model-based strategy was compared with a passive buy-and-hold benchmark. This follows the logic that directional prediction models should be evaluated not only statistically, but also through their practical trading implications [36,37].

The model output used for trading-signal construction was the predicted probability of the UP class. A long signal was generated only when this probability exceeded a predefined confidence threshold [37]:

$$Signal_t = \begin{cases} 1, & \text{if } P(UP_t) \geq \tau \\ 0, & \text{if } P(UP_t) < \tau \end{cases} \qquad (2.13)$$

where $P(UP_t)$is the model-predicted probability of an upward movement at time t, and $\tau$is the confidence threshold. The threshold was included to avoid acting on weak model predictions. A higher threshold can reduce noisy trades, but it can also reduce the number of trading opportunities.

Additional filtering rules were applied to reduce unstable signal switching. A persistence filter required the same signal to appear for a selected number of consecutive observations before a position change was accepted. A minimum holding-period rule was also used to prevent excessive trading. These rules were included because cryptocurrency markets are volatile, and daily predictions may fluctuate during sideways or uncertain periods. The exact threshold, persistence and minimum holding-period parameters differed across BTC, ETH and SOL, reflecting asset-specific validation behaviour and market characteristics.

To avoid look-ahead bias, signals were shifted forward before being applied to returns. A signal generated using information available at day (t) was executed at the opening price of day (t+1). This means that the model was not allowed to trade using information from the same day on which the signal was produced. In simplified form, the model-based strategy return can be written as [36,37]:

$$R_{t+1}^{strategy} = Position_t \cdot R_{t+1}^{asset} \qquad (2.14)$$

where $Position_t$is the position determined by the signal at day t, and $R_{t+1}^{asset}$is the next-period asset return. If $Position_t = 1$, the strategy is exposed to the cryptocurrency; if $Position_t = 0$, the strategy remains out of the market.

The model-based strategy was compared with a buy-and-hold benchmark. The buy-and-hold return was calculated by assuming continuous exposure to the asset throughout the same test period [36]:

$$R_{t+1}^{BH} = R_{t+1}^{asset} \qquad (2.15)$$

This benchmark was used because it represents the simplest passive investment alternative. If the model-based strategy outperforms buy-and-hold, it suggests that the model signals may add economic value. If it underperforms, then good classification performance may not be sufficient for practical usefulness.

Strategy performance was evaluated using total return, annualised Sharpe ratio, maximum drawdown, long exposure percentage and number of trades. Total return measured cumulative strategy performance over the test period [37]:

$$TotalReturn = \prod_{t=1}^{T}(1 + R_t^{strategy}) - 1 \tag{2.16}$$

The annualised Sharpe ratio was used to evaluate return relative to volatility [36]:

$$Sharpe = \frac{\bar{R}^{strategy}}{\sigma(R^{strategy})} \times \sqrt{365} \tag{2.17}$$

where $\bar{R}^{strategy}$is the average daily strategy return, $\sigma(R^{strategy})$is the standard deviation of daily strategy returns, and 365 was used because cryptocurrency markets trade continuously. Maximum drawdown measured the largest peak-to-trough decline in cumulative strategy value [56]:

$$MaxDrawdown = \min_t \left(\frac{V_t - \max(V_1, \ldots, V_t)}{\max(V_1, \ldots, V_t)} \right) \tag{2.18}$$

where $V_t$is the cumulative portfolio value at time t. Long exposure percentage measured the share of days when the strategy held the cryptocurrency [37]:

$$Long\% = \frac{\sum_{t=1}^{T} Positio\,n_t}{T} \times 100 \tag{2.19}$$

The number of trades measured how often the strategy changed position and was included to evaluate turnover.

The backtesting results were interpreted cautiously. Transaction costs, bid-ask spreads, slippage and exchange-specific liquidity constraints were not fully incorporated into the backtest. Therefore, the results represent gross strategy performance rather than a complete real-world trading simulation. This limitation is important because frequent trading can reduce or eliminate apparent profitability after costs. In addition, Bailey et al. [56] warn that backtest results may become overfitted when many strategy parameter combinations are tested and the best-performing configuration is selected. For this reason, the backtest in this thesis is interpreted as an economic evaluation of model signals under defined assumptions, not as proof of a universally profitable trading strategy.

3. Empirical Analysis of Bitcoin, Ethereum and Solana Market Dynamics Using Explainable Artificial Intelligence

3.1. Market Regime Identification Results

Following the methodology described in Section 2.7, Gaussian Hidden Markov Models were estimated separately for Bitcoin, Ethereum and Solana. The purpose of this stage was to identify latent market states that describe changes in return direction, short-term volatility and downside risk. The regimes were not assigned manually; they were inferred from the observed market variables and then interpreted according to their statistical characteristics and position in the price history.

The results of this stage are presented in three steps. First, the number of regimes is evaluated using model-selection criteria for candidate HMMs with two to five hidden states. Second, the selected four-regime structure is interpreted for each cryptocurrency using regime-specific return, volatility and drawdown characteristics. Third, the regime sequence is visualised over time to assess whether the inferred states correspond to economically meaningful market phases, such as upward-trending, declining, sideways or high-volatility periods.

It is important to note that HMM regime labels are arbitrary. For example, Regime 0 does not automatically represent a “better” or “worse” state than Regime 1. The economic meaning of each regime is assigned only after analysing its statistical behaviour and its position in the market timeline. Therefore, the following section interprets the regimes based on their observed return, volatility and drawdown patterns rather than on the numerical label itself.

3.1.1. Selection of the Number of Regimes

Table 4. HMM Model Selection Results

Asset	Number of regimes	AIC	BIC
BTC	2	18511.71	18599.37
BTC	3	15770.20	15922.14
BTC	4	14365.77	14593.68
BTC	5	15516.92	15832.49
ETH	2	18379.62	18467.28
ETH	3	16131.24	16283.18
ETH	4	15179.55	15407.46
ETH	5	13292.32	13607.89
SOL	2	13859.58	13943.30
SOL	3	13873.35	14018.47
SOL	4	9708.06	9925.75
SOL	5	9384.44	9685.86

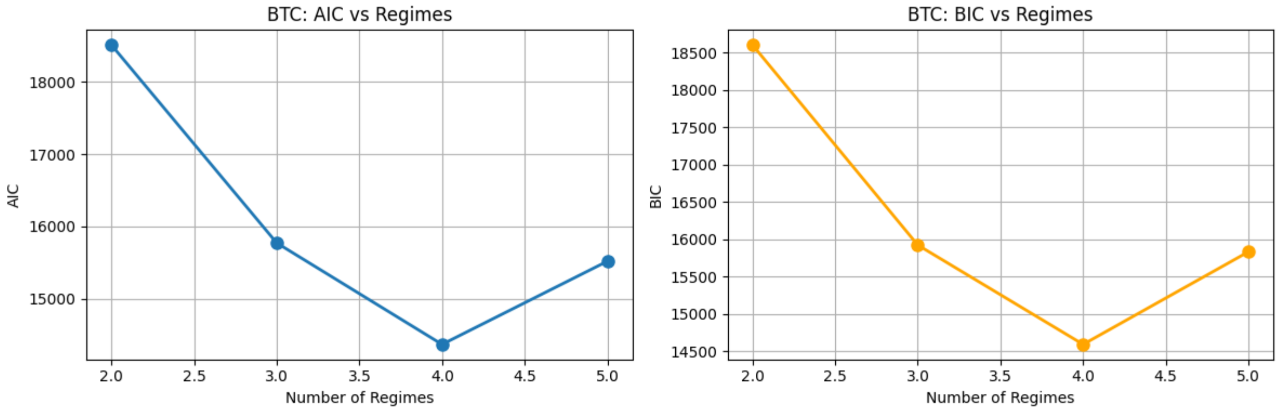

1 Fig. AIC and BIC Values for Candidate Bitcoin HMM Regime Models

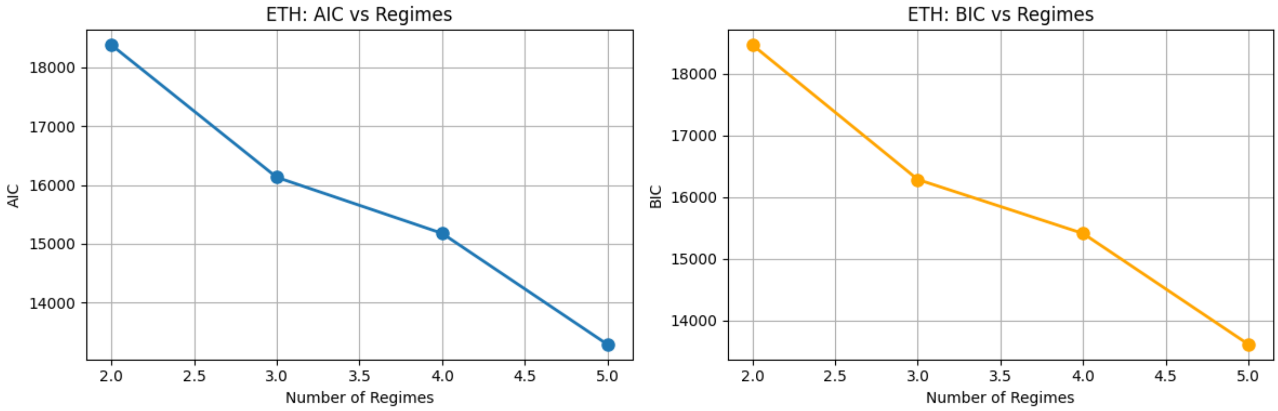

2 Fig. AIC and BIC Values for Candidate Ethereum HMM Regime Models

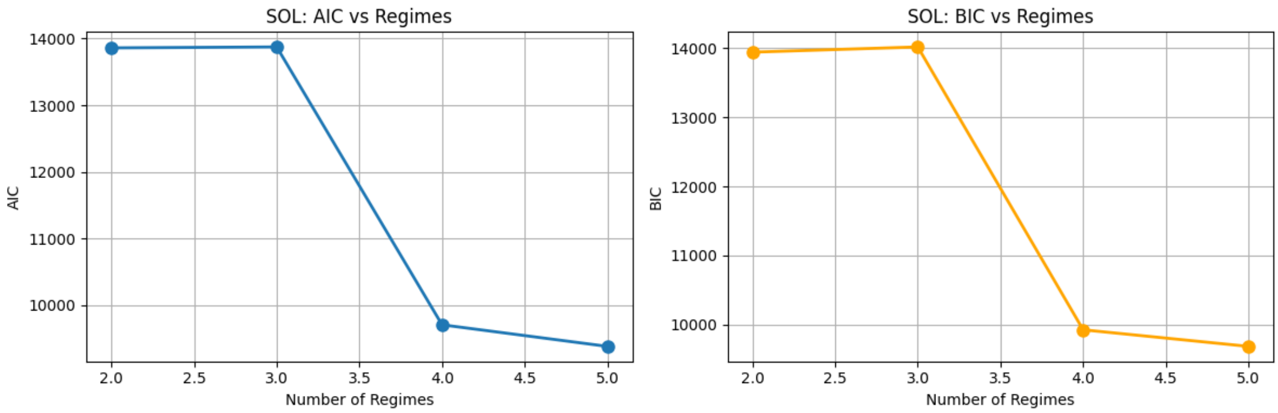

3 Fig. AIC and BIC Values for Candidate Solana HMM Regime Models

Candidate HMMs with two, three, four and five hidden states were compared for each cryptocurrency. The comparison was based on AIC and BIC values together with economic interpretability of the resulting states. Although information criteria provide useful guidance, the final number of regimes also had to remain interpretable, because the regimes were later used as explanatory variables and as a basis for model-performance analysis. A four-regime structure was selected because it provided a balance between statistical fit and meaningful market-state separation.

Although the five-regime model improved statistical fit in some cases, the four-regime model was retained because it produced a more interpretable structure and avoided excessive fragmentation of market states.

3.1.2. Interpretation of Identified Regimes

Table 5. Descriptive Statistics of HMM-Identified Market Regimes

Asset	Regime	Count	Share	Mean return	Mean volatility	Mean drawdown
BTC	0	897	35.18%	0.00150	0.01883	-0.12044
BTC	1	467	18.31%	0.00044	0.02036	-0.64461
BTC	2	542	21.25%	0.00424	0.04026	-0.12068
BTC	3	644	25.25%	-0.00117	0.03788	-0.41508
ETH	0	289	11.33%	-0.00213	0.08022	-0.27637
ETH	1	736	28.86%	0.00129	0.02248	-0.58362
ETH	2	1069	41.92%	0.00344	0.03297	-0.20765
ETH	3	456	17.88%	-0.00224	0.04676	-0.54708
SOL	0	413	21.05%	0.01279	0.06130	-0.13586
SOL	1	659	33.59%	-0.00057	0.03673	-0.42459
SOL	2	381	19.42%	-0.00401	0.08870	-0.59661
SOL	3	509	25.94%	0.00062	0.04200	-0.90104

After selecting the four-regime HMM structure, the identified regimes were interpreted using regime-specific descriptive statistics. The interpretation was based on three variables: mean return, mean volatility and mean drawdown. Mean return indicates the average direction of price movement within the regime, mean volatility reflects the level of short-term uncertainty, and mean drawdown shows the average decline from previous price peaks. Since HMM regime numbers are arbitrary, the regime labels were not interpreted as ordinal values. Instead, each regime was interpreted according to its statistical characteristics.

For Bitcoin, Regime 0 was the most frequent state, covering 35.18% of observations. It was characterised by a small positive mean return, relatively low volatility and a moderate drawdown. This regime can be interpreted as a relatively stable positive or recovery-type market state. Regime 1 covered 18.31% of observations and had a small positive mean return, moderate volatility and the deepest average drawdown among Bitcoin regimes. This suggests a recovery or depressed-market regime, where prices were still far below previous peaks despite slightly positive average returns. Regime 2 had the highest positive mean return and the highest volatility, indicating a high-volatility growth regime. Regime 3 showed a negative mean return, elevated volatility and a deep drawdown, which makes it the closest Bitcoin equivalent of a bearish or high-stress regime.

For Ethereum, Regime 2 was the dominant state, accounting for 41.92% of observations. It had the highest positive mean return, moderate volatility and a relatively smaller drawdown compared with the most stressed regimes. This can be interpreted as Ethereum's main growth regime. Regime 1 also showed a positive mean return, but with lower volatility and a deeper drawdown, suggesting a recovery or moderate-growth state after previous market declines. In contrast, Regime 0 and Regime 3 had negative mean returns. Regime 0 had the highest volatility but a less severe drawdown than Regime 3, suggesting a high-volatility correction regime. Regime 3 combined negative mean return, elevated volatility and deep drawdown, and can therefore be interpreted as Ethereum's bearish or stressed drawdown regime.

For Solana, the regime structure reflects the asset's more volatile and shorter historical profile. Regime 0 had the highest positive mean return and relatively high volatility, while drawdown remained comparatively limited. This regime can be interpreted as a strong growth regime. Regime 1 was the most frequent state, covering 33.59% of observations, with a slightly negative mean return, moderate volatility and a notable drawdown. This suggests a weak or declining market state. Regime 2 had the strongest negative mean return, the highest volatility and a deep drawdown, making it the clearest Solana panic or high-volatility stress regime. Regime 3 had a small positive mean return, moderate volatility, but the deepest average drawdown. This suggests a depressed recovery regime, where prices may stabilise or slightly improve while remaining far below previous peaks.

Overall, the regime statistics show that the HMM separated market conditions into economically meaningful states. Across the three assets, some regimes represent relatively stable or growth-oriented periods, while others capture high-volatility stress, bearish movement or recovery from deep drawdowns. The results also show that the same regime number does not have the same interpretation across assets. For example, Regime 2 represents a high-return, high-volatility state for Bitcoin, a dominant growth state for Ethereum, and a panic-like stress state for Solana. Therefore, the regimes should be interpreted asset by asset rather than treated as directly comparable numerical categories.

These regime classifications are useful for the later modelling analysis because they allow model performance to be evaluated under different market conditions. Instead of assessing the Random Forest classifiers only on the full test period, the regime variables make it possible to examine whether prediction quality differs during growth, decline, recovery or high-stress regimes. This is especially important in cryptocurrency markets, where model reliability may vary depending on volatility, drawdown and market direction.

3.1.3. Regime Dynamics Over Time

The regime plots show how the HMM-classified states are distributed across the cryptocurrency price histories. These visualisations are important because they allow the statistical regime labels to be compared with observable market phases. Since HMM regimes are inferred from return, volatility and drawdown variables, they should not appear as random labels across the timeline. Instead, economically meaningful regimes should cluster around periods with similar market behaviour, such as upward trends, prolonged drawdowns, sideways movement or high-volatility corrections.

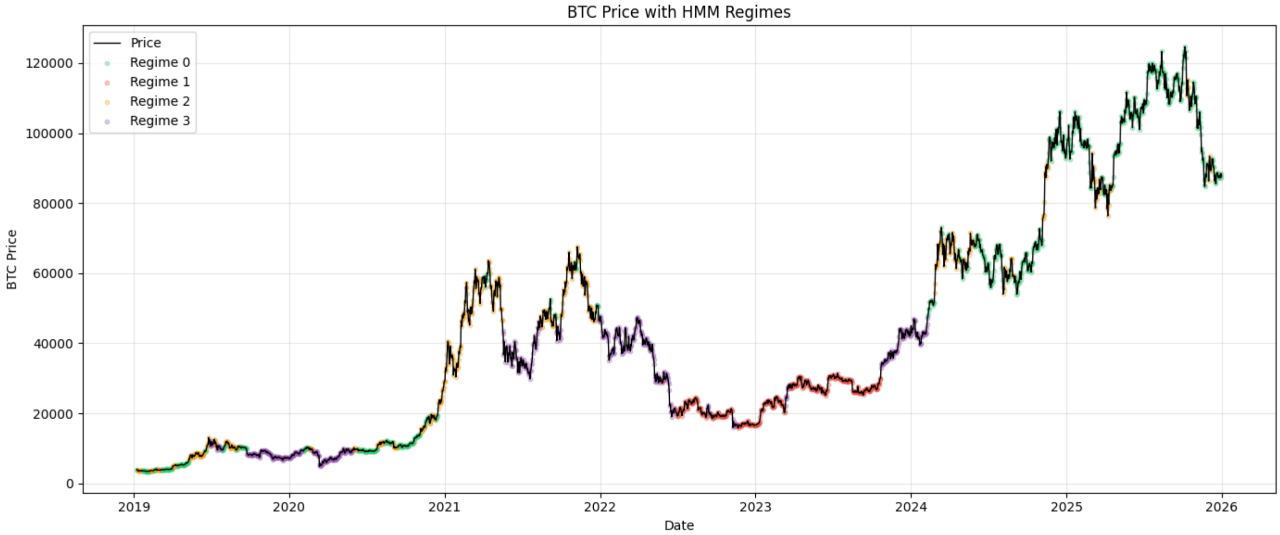

4 Fig. Bitcoin Price with HMM-Identified Market Regimes

For Bitcoin, the regime sequence follows several visible market phases. Regime 2 appears frequently during strong upward price movements, especially around the 2021 rally and later high-volatility growth periods. This is consistent with the descriptive statistics, where BTC Regime 2 had the highest mean return and the highest volatility. Regime 3 appears more frequently during declining or stressed periods, including the market correction after the 2021 peak and later drawdown phases. Regime 1 is concentrated around depressed market conditions, especially during the lower-price period in 2022-2023, which is consistent with its deep average drawdown. Regime 0 appears during more stable positive or recovery periods and becomes more common during the later upward market phase.

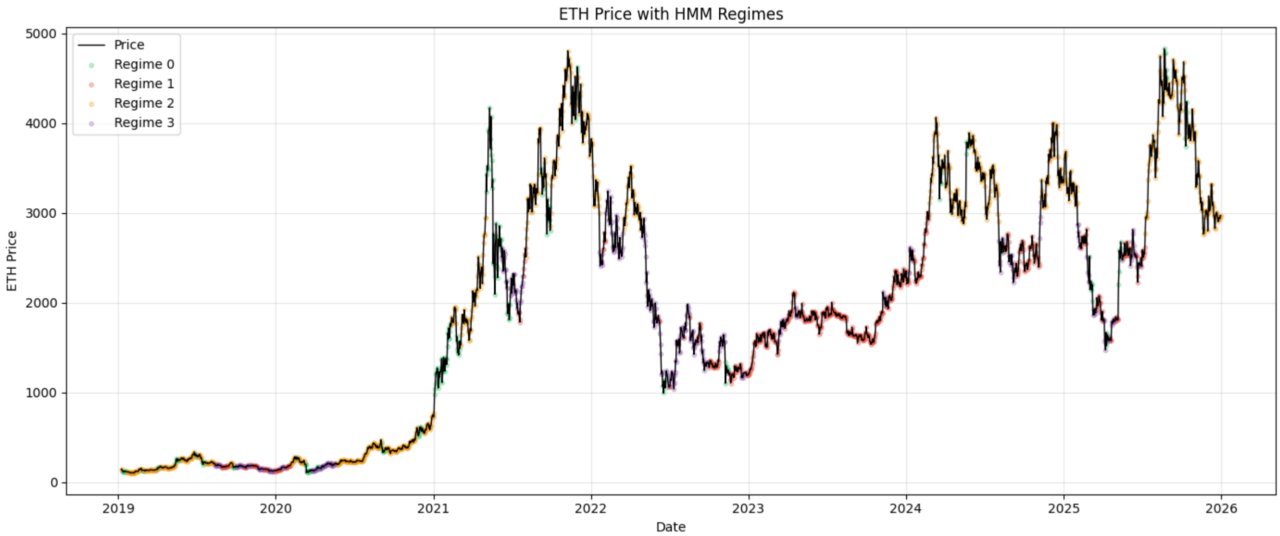

5 Fig. Ethereum Price with HMM-Identified Market Regimes

For Ethereum, the regime dynamics are similar in broad structure, but not identical. Regime 2 appears during major growth phases, including the strong 2021-2022 expansion and parts of the later recovery. This matches the regime statistics, where ETH Regime 2 had the highest mean return and was the most frequent state. Regime 3 is more visible during deeper drawdown periods, especially after the 2021-2022 peak, and can be interpreted as a stressed or bearish drawdown regime. Regime 0 appears around sharper volatile corrections, while Regime 1 is more associated with moderate recovery or less extreme market conditions. This confirms that the Ethereum HMM did not simply reproduce Bitcoin's regime labels, but identified asset-specific market-state dynamics.

6 Fig. Solana Price with HMM-Identified Market Regimes

For Solana, the regime sequence reflects a more volatile and compressed market history. Regime 0 appears during strong upward movements, especially during the rapid 2021 increase and later high-growth episodes. Regime 2 is associated with high-volatility stress and sharp negative movements, consistent with its negative mean return and highest volatility. Regime 3 appears during the prolonged depressed period after the 2021-2022 decline, when Solana remained far below previous highs despite some stabilisation. Regime 1 is more common in weaker or moderately declining periods. Compared with Bitcoin and Ethereum, Solana's regimes show stronger alternation between rapid growth, deep drawdown and high-volatility stress, reflecting the asset's shorter and more volatile market history.

Overall, the regime plots support the use of HMM regimes as market-state indicators. The regimes are not evenly or randomly distributed across time; instead, they cluster around visually distinct market phases. This is important for the later modelling analysis because it suggests that the regime variables capture meaningful changes in market conditions. In the next stages of the results, these regimes are used to evaluate whether classification performance and model explanations differ between growth, recovery, drawdown and high-volatility states.

3.1.4. External Shock Periods and Market Dynamics

The inferred HMM states should not be interpreted as predefined real-world market regimes, because the model separates observations according to the statistical behaviour of selected variables rather than according to named economic events. However, the timing of some state changes and larger price movements can still be interpreted in relation to broader external market conditions. This is relevant because previous research shows that cryptocurrency markets may be affected by global uncertainty, investor sentiment, monetary-policy conditions, geopolitical events and broader financial-market turbulence [15,18,33].

In the analysed period, several external and market-wide events coincided with visible changes in cryptocurrency price dynamics. The COVID-19 period was associated with increased global uncertainty and financial-market stress, while the later 2021 expansion, 2022 decline, 2023 recovery and 2024-2025 market cycle were connected with changing liquidity conditions, inflation and interest-rate expectations, regulatory developments, institutional adoption and broader crypto-market stress. For Bitcoin, the 2024 upward movement coincided with the introduction of spot Bitcoin ETFs, while the later 2025 correction coincided with weaker risk sentiment and ETF outflows. For Solana, the price dynamics were more volatile and reflected stronger asset-specific sensitivity during both the 2021 rise and the later correction periods.

This interpretation remains contextual rather than causal. The HMM states were derived only from return, volatility and drawdown-related variables, so they cannot be directly labelled as COVID, crisis, bull or bear regimes. Instead, external events provide additional economic context for understanding why some periods show different return, volatility and drawdown behaviour.

3.2. Random Forest Classification Results

The Random Forest classifiers were trained separately for Bitcoin, Ethereum and Solana in order to preserve asset-specific market behaviour. Table 6 summarises the test-period classification results. The models were evaluated using accuracy, precision, recall, macro F1 score and class-specific F1 scores for the UP and DOWN classes. Macro F1 was included because it gives equal weight to both directional classes and therefore provides a more balanced view of model performance than accuracy alone.

Table 6. Test-Set Performance of Random Forest Classification Models

Asset	Test Accuracy	Macro Precision	Macro Recall	Macro F1	F1 UP	F1 DOWN	Top Feature
BTC	~63%	0.620	0.620	~0.63	~0.64	~0.62	BTC_log_ret (~61%)
ETH	~75%	0.757	0.755	~0.75	~0.72	~0.76	ETH_ret (~56%)
SOL	~76%	0.763	0,763	~0.76	0.76	~0.76	SOL_ret (~85%)

The results show that all three models achieved above-random directional classification performance during the test period. The Bitcoin model produced the weakest result among the three assets, with accuracy and macro F1 of approximately 0.63. This indicates that the model was able to identify some directional structure in Bitcoin price movements, but the signal was relatively moderate. The Ethereum and Solana models achieved stronger results, with accuracy and macro F1 values of approximately 0.75-0.76. This suggests that, during the analysed test period, ETH and SOL direction was more effectively captured by the selected feature sets than BTC direction.

The difference between Bitcoin and the two other assets may reflect differences in market maturity and predictability. Bitcoin is the most established and liquid cryptocurrency in the sample, and its short-term direction may therefore be harder to classify using daily market and external variables. Ethereum and Solana, especially Solana, may contain stronger short-term directional patterns because their market behaviour is more volatile and less mature. However, this interpretation should be treated cautiously, because the results also depend on the selected time period, feature set and class distribution.

The obtained results can also be positioned in relation to previous cryptocurrency direction-prediction research. Basher and Sadorsky [39] applied tree-based machine learning methods to Bitcoin price-direction forecasting and reported that bagging and Random Forest models achieved approximately 75-80% accuracy for a five-day prediction horizon, while 10-20-day forecasts exceeded 85%. However, these results are not directly comparable with the present study because the prediction horizon and evaluation setup differ. This thesis focuses on daily UP/DOWN classification and uses a chronological train-validation-test structure, which provides a stricter out-of-sample setting for short-horizon prediction. Therefore, the lower Bitcoin result in this thesis should be interpreted as a more conservative daily-direction result rather than as direct underperformance.

The comparison is still useful because it shows that tree-based models can capture directional patterns in cryptocurrency markets, but that performance depends strongly on the prediction horizon and modelling design. Longer horizons may produce stronger accuracy because short-term daily movements are noisier and more affected by volatility. In addition, this thesis evaluates macro F1 together with accuracy, which penalises uneven classification performance across UP and DOWN classes. Overall, the results suggest that the proposed Random Forest framework captured meaningful short-term directional information, especially for Ethereum and Solana, while Bitcoin remained more difficult to predict at the daily horizon.

Detailed interpretation of the variables driving these predictions is presented in the explainability section, where feature importance, SHAP values and PDP/ICE plots are used to examine how the models reached their decisions.

3.3. Model Signals and Backtesting results

After evaluating the classification metrics, the next step was to examine how the Random Forest probability outputs were converted into trading signals during the test period. This analysis is important because a model can achieve acceptable classification performance but still produce unstable or economically weak signals. Therefore, the predicted probability of the UP class, the resulting buy/sell signals, and the cumulative performance of the strategy were analysed for Bitcoin, Ethereum and Solana.

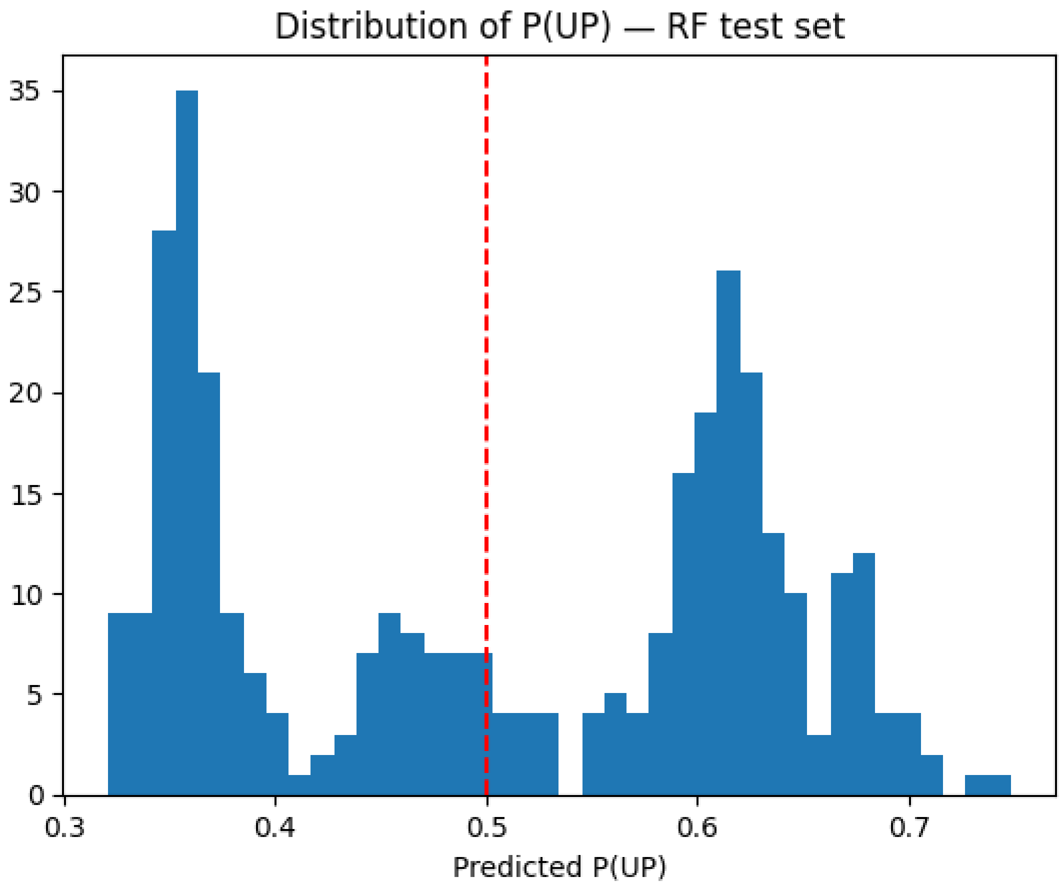

7 Fig. Distribution of Predicted UP Probabilities for the Bitcoin Random Forest Model

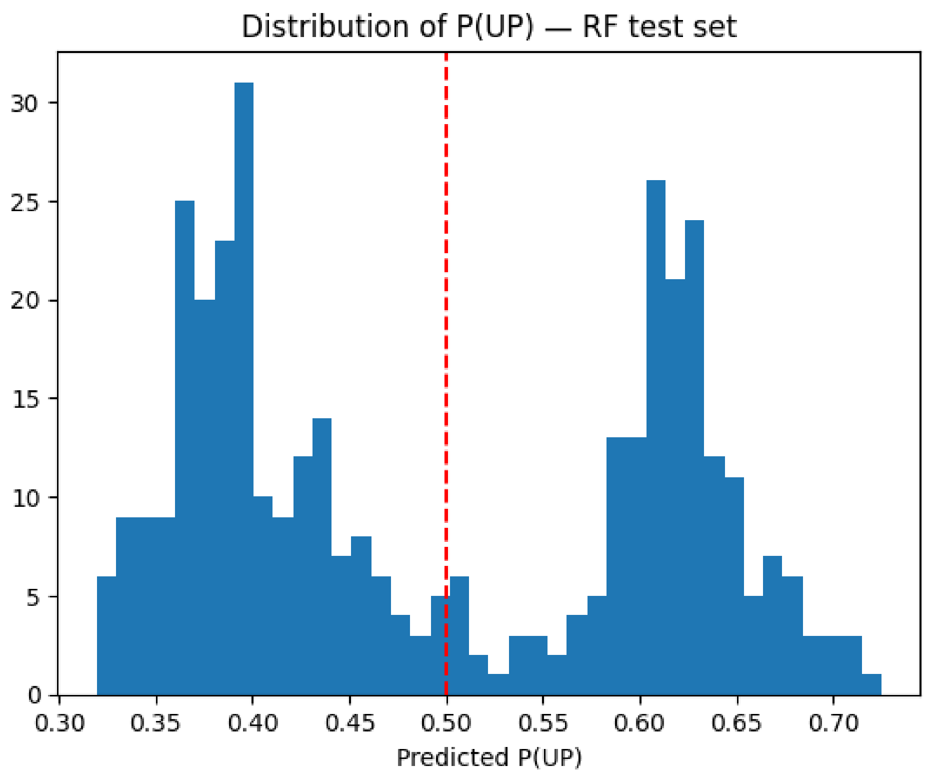

8 Fig. Distribution of Predicted UP Probabilities for the Ethereum Random Forest Model

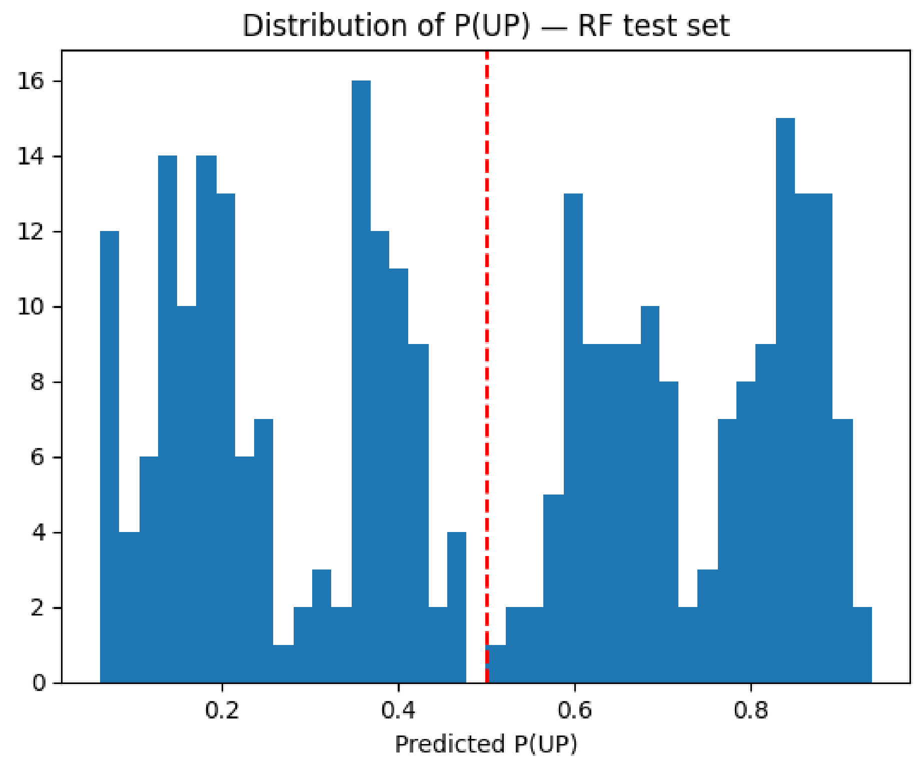

9 Fig. Distribution of Predicted UP Probabilities for the Solana Random Forest Model

The probability distributions show that the three models produced different levels of directional confidence. For Bitcoin, the predicted UP probabilities were mostly concentrated in two areas: one group around approximately 0.35-0.40 and another around approximately 0.58-0.65. This indicates that the BTC model often separated observations into relatively clear DOWN-oriented and UP-oriented probability regions, although some observations still remained close to the 0.5 decision threshold. For Ethereum, the distribution was also split into two broad regions, with a lower-probability concentration around approximately 0.35-0.45 and a higher-probability concentration around approximately 0.58-0.65. This suggests that the ETH model produced a similar two-sided probability structure, but with a meaningful number of observations near the threshold. Solana showed the widest probability spread, with values ranging from very low probabilities below 0.2 to high probabilities above 0.8. Compared with BTC and ETH, the SOL model therefore produced more extreme probability estimates, suggesting stronger confidence separation between predicted UP and DOWN observations.

These probability distributions justify the use of probability-based trading rules rather than simple class labels alone. If the predicted probability is close to the threshold, the model signal is less certain and may not be strong enough to justify a position change. For this reason, the backtesting procedure used confidence thresholds and persistence rules to filter weaker or unstable signals. The strategy parameters differed across assets because BTC, ETH and SOL had different probability distributions and market behaviour. In this setup, the Random Forest models produced UP probabilities, while the backtesting layer defined how strong and persistent the signal had to be before it was converted into a long or cash position.

For Bitcoin, the backtesting configuration used a confidence threshold of 0.5 and a two-day persistence rule. The BTC strategy generated relatively frequent position changes, with 18 buy and 18 sell signals during the test period. This means that the Bitcoin strategy was more active than the Ethereum strategy and reacted more often to changes in predicted direction.

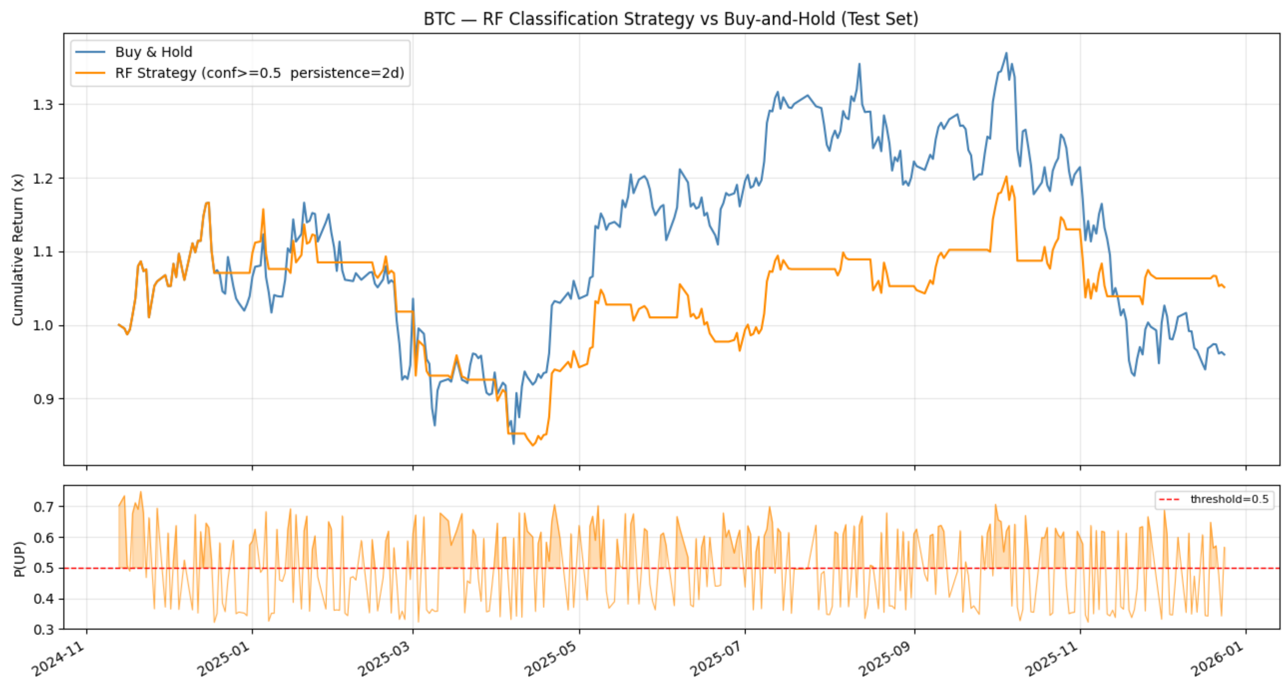

10 Fig. Bitcoin Random Forest Strategy versus Buy-and-Hold During the Test Period

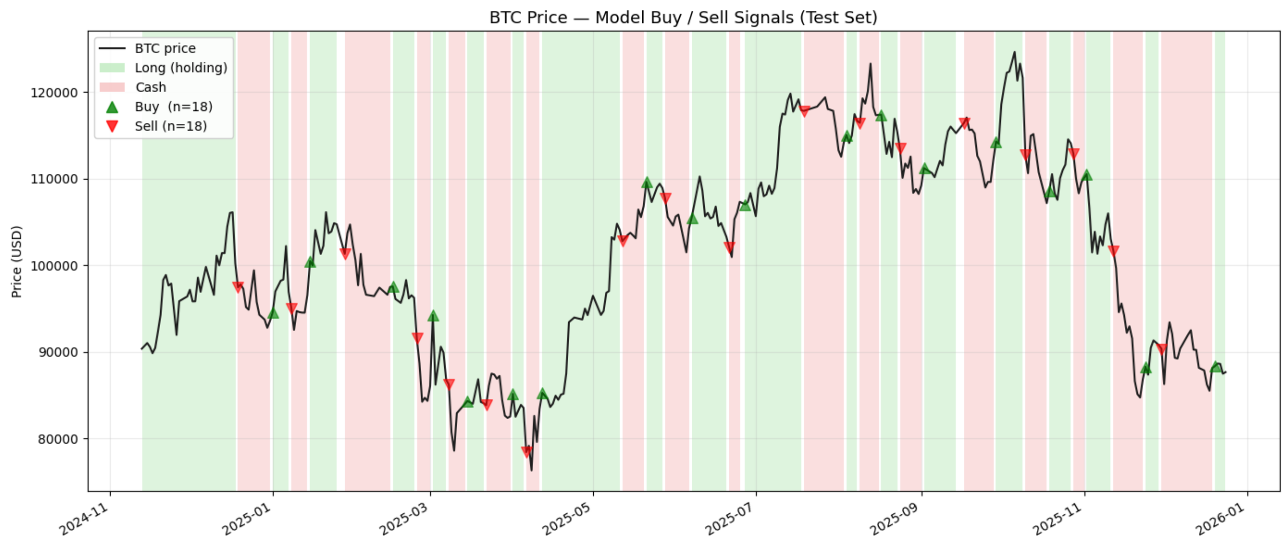

11 Fig. Bitcoin Price with Random Forest Buy and Sell Signals During the Test Period

The BTC cumulative return chart shows that buy-and-hold captured a stronger upward move during the middle of the test period, but later experienced a substantial decline. The Random Forest-based strategy captured less of the strongest rally, but avoided part of the later drawdown and finished above the buy-and-hold benchmark. This suggests that, for Bitcoin, the strategy was more useful as an exposure-management tool than as a pure return-maximisation strategy. The signal chart supports this interpretation, as the strategy alternated between long and cash periods rather than remaining continuously invested.

For Ethereum, the backtesting configuration used a lower confidence threshold of 0.44 and a four-day persistence rule. This made the ETH strategy less reactive and more selective, producing only 3 buy and 4 sell signals during the test period. The longer persistence requirement reduced short-term switching and created longer holding and cash periods.

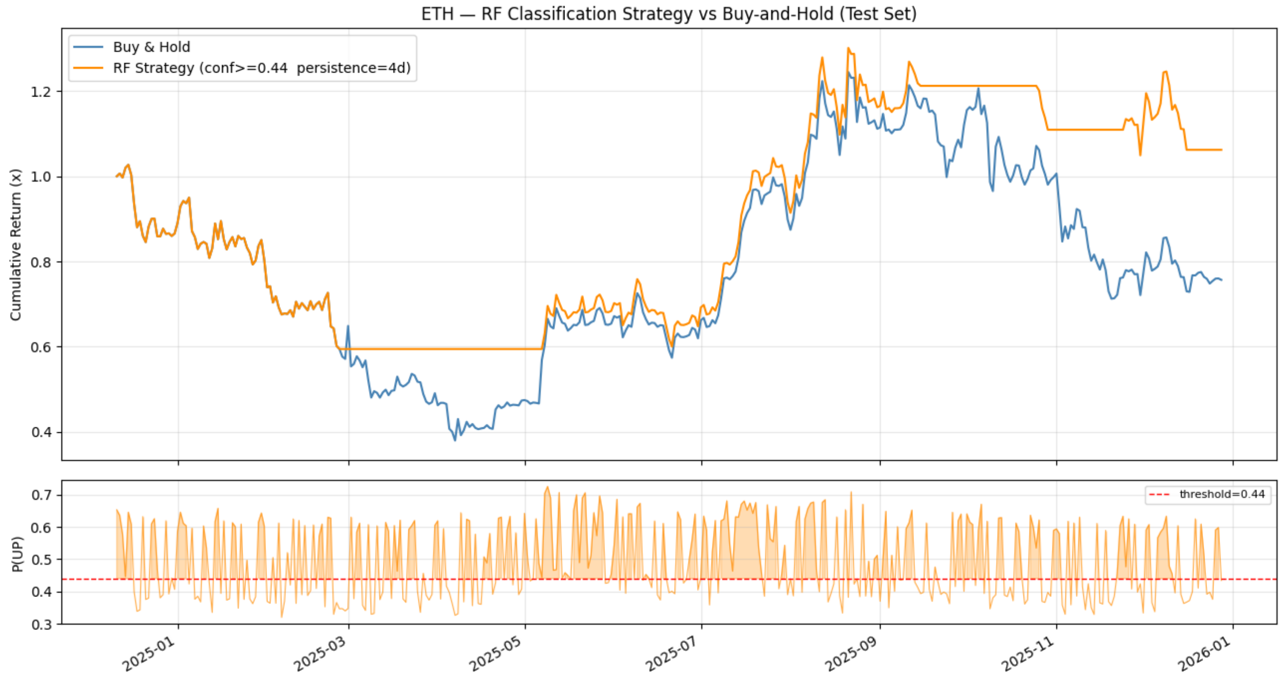

12 Fig. Ethereum Random Forest Strategy versus Buy-and-Hold During the Test Period

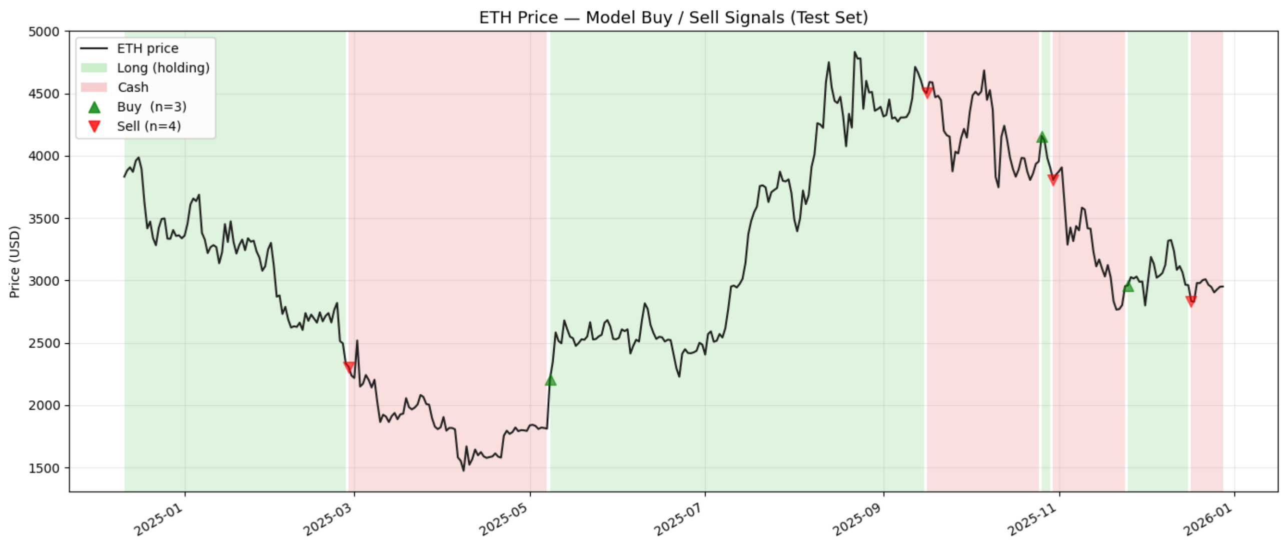

13 Fig. Ethereum Price with Random Forest Buy and Sell Signals During the Test Period

The ETH strategy avoided a large part of the early drawdown visible in the buy-and-hold benchmark and later participated in the recovery phase. Although it did not remain fully exposed throughout the whole upward movement, it finished above buy-and-hold in the displayed test period. This indicates that the ETH strategy benefited mainly from avoiding prolonged weak periods and re-entering the market during stronger conditions. Compared with BTC, the ETH strategy was less active but more stable in its position changes.

For Solana, the backtesting configuration used a 0.5 confidence threshold and a three-day persistence rule. The SOL strategy generated 11 buy and 10 sell signals, making it more active than ETH but less active than BTC. This was consistent with the wider probability distribution of the SOL model, where probabilities were more strongly separated and therefore more often crossed the trading threshold with clearer confidence.

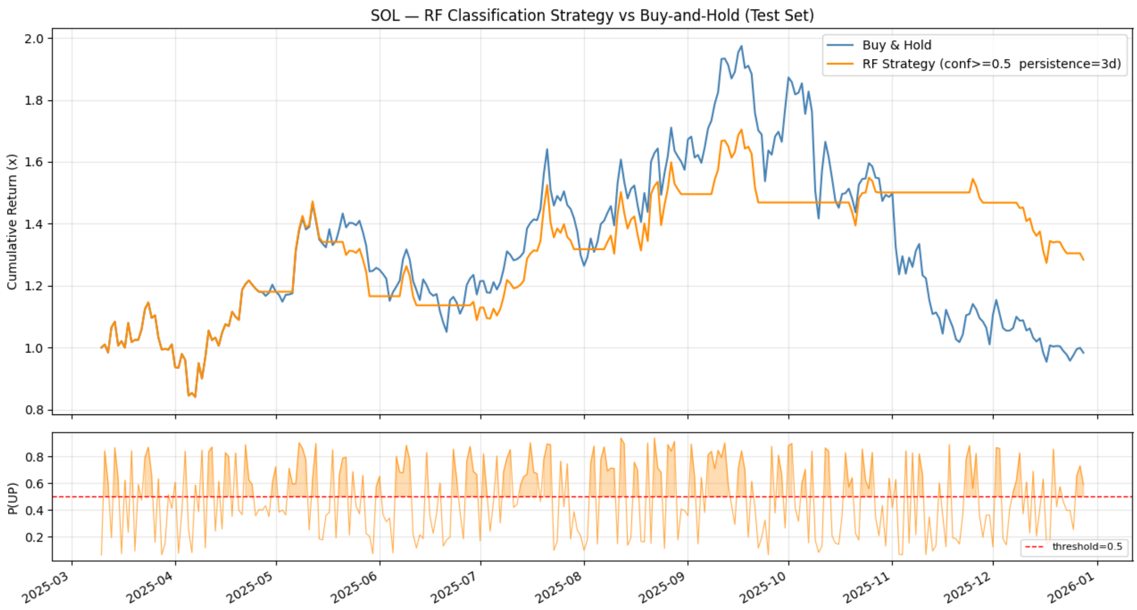

14 Fig. Solana Random Forest Strategy versus Buy-and-Hold During the Test Period

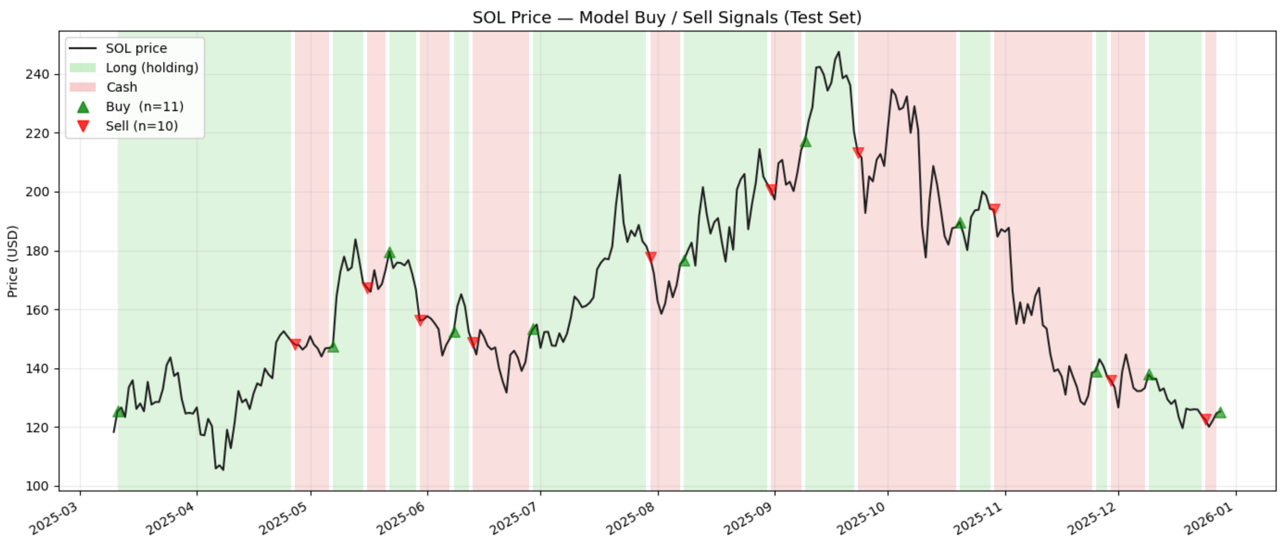

15 Fig. Solana Price with Random Forest Buy and Sell Signals During the Test Period

The SOL strategy showed the strongest visual separation from the buy-and-hold benchmark. It captured part of the upward movement during the middle of the test period, but more importantly avoided a substantial part of the later decline. While buy-and-hold fell sharply after the peak, the strategy reduced exposure and preserved more of its accumulated return. This suggests that, for Solana, the model-based signal was especially useful for managing downside exposure during unstable market conditions.

Overall, the trading signal analysis shows that the final strategy behaviour differed across cryptocurrencies. These differences were not caused only by the Random Forest probability outputs, but also by the asset-specific backtesting rules applied after prediction. The models produced UP probabilities, while the strategy converted those probabilities into long or cash positions using selected thresholds, persistence filters and holding-period assumptions. Because BTC, ETH and SOL had different probability distributions and market behaviour, the final strategy parameters were not identical across assets. This is a standard approach in strategy design, where prediction outputs are translated into asset-specific trading rules rather than forcing the same rule on markets with different volatility and signal characteristics.

At the same time, these results should be interpreted cautiously. The backtesting rules influence the final trading path, and the selected parameters may not remain optimal in future market conditions. In addition, transaction costs, slippage and exchange-specific liquidity constraints were not fully incorporated. Therefore, the backtest should be interpreted as an economic evaluation of model signals under defined assumptions, not as evidence of a complete real-world trading system.

3.4. Explainability Analysis of Random Forest Predictions

3.4.1. Global Feature Importance

After examining classification performance and trading signal behaviour, the next step was to analyse which variables contributed most to the Random Forest models. For this purpose, impurity-based feature importance was used as an initial global explainability measure. In Random Forest models, feature importance reflects how much each variable contributes to reducing classification impurity across the ensemble of decision trees. Therefore, this analysis provides an overall view of which predictors were most influential in the UP/DOWN classification task.

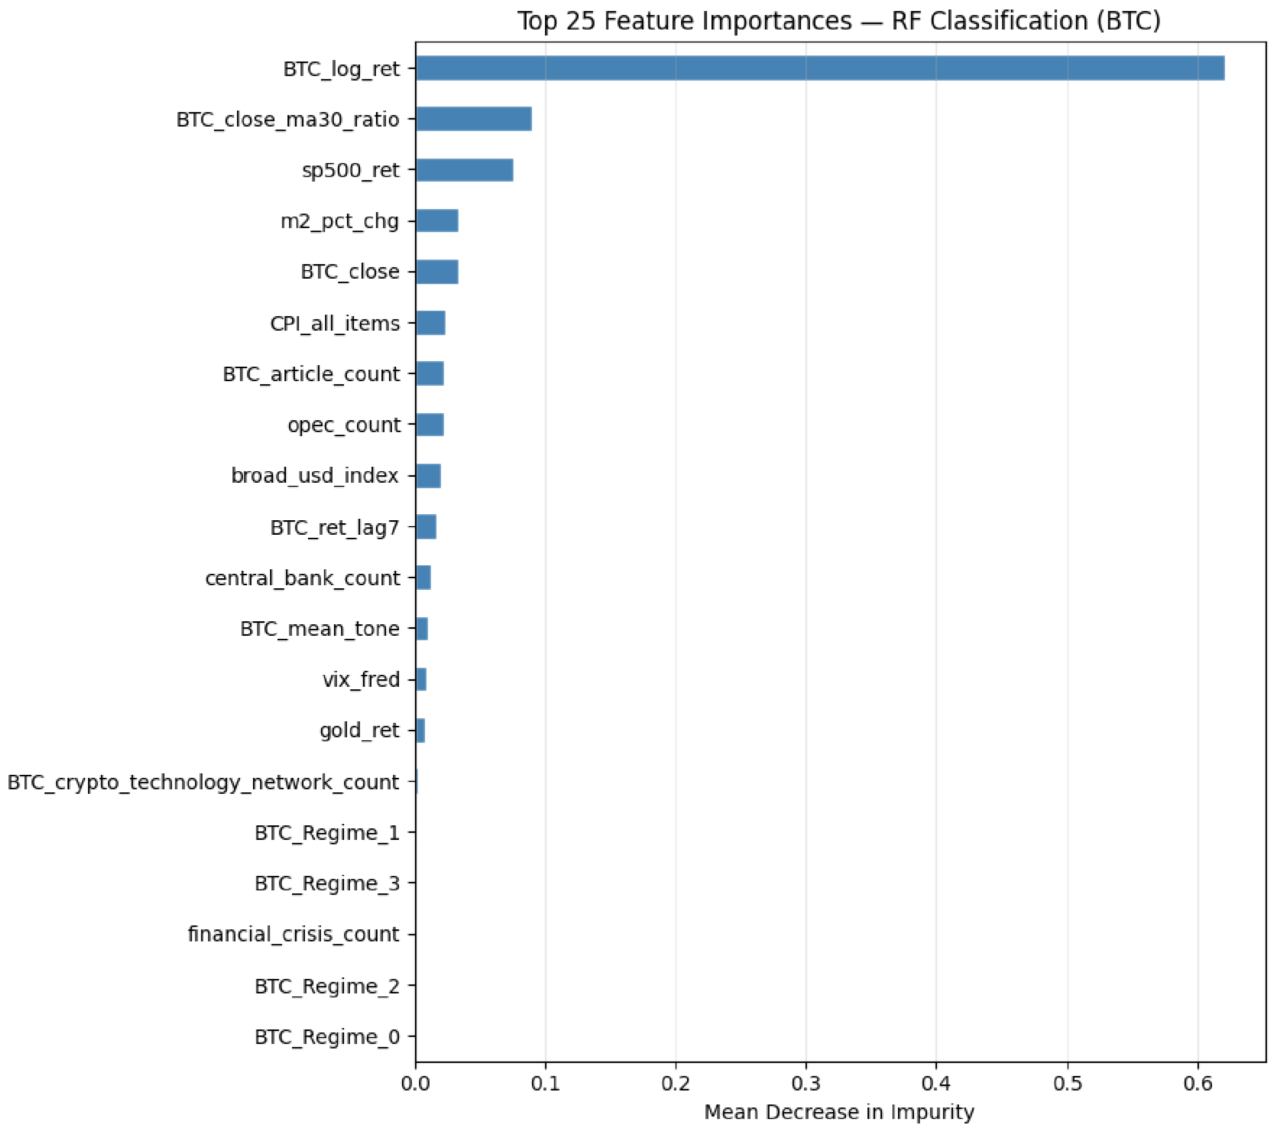

16 Fig. Feature Importance of the Bitcoin Random Forest Classification Model

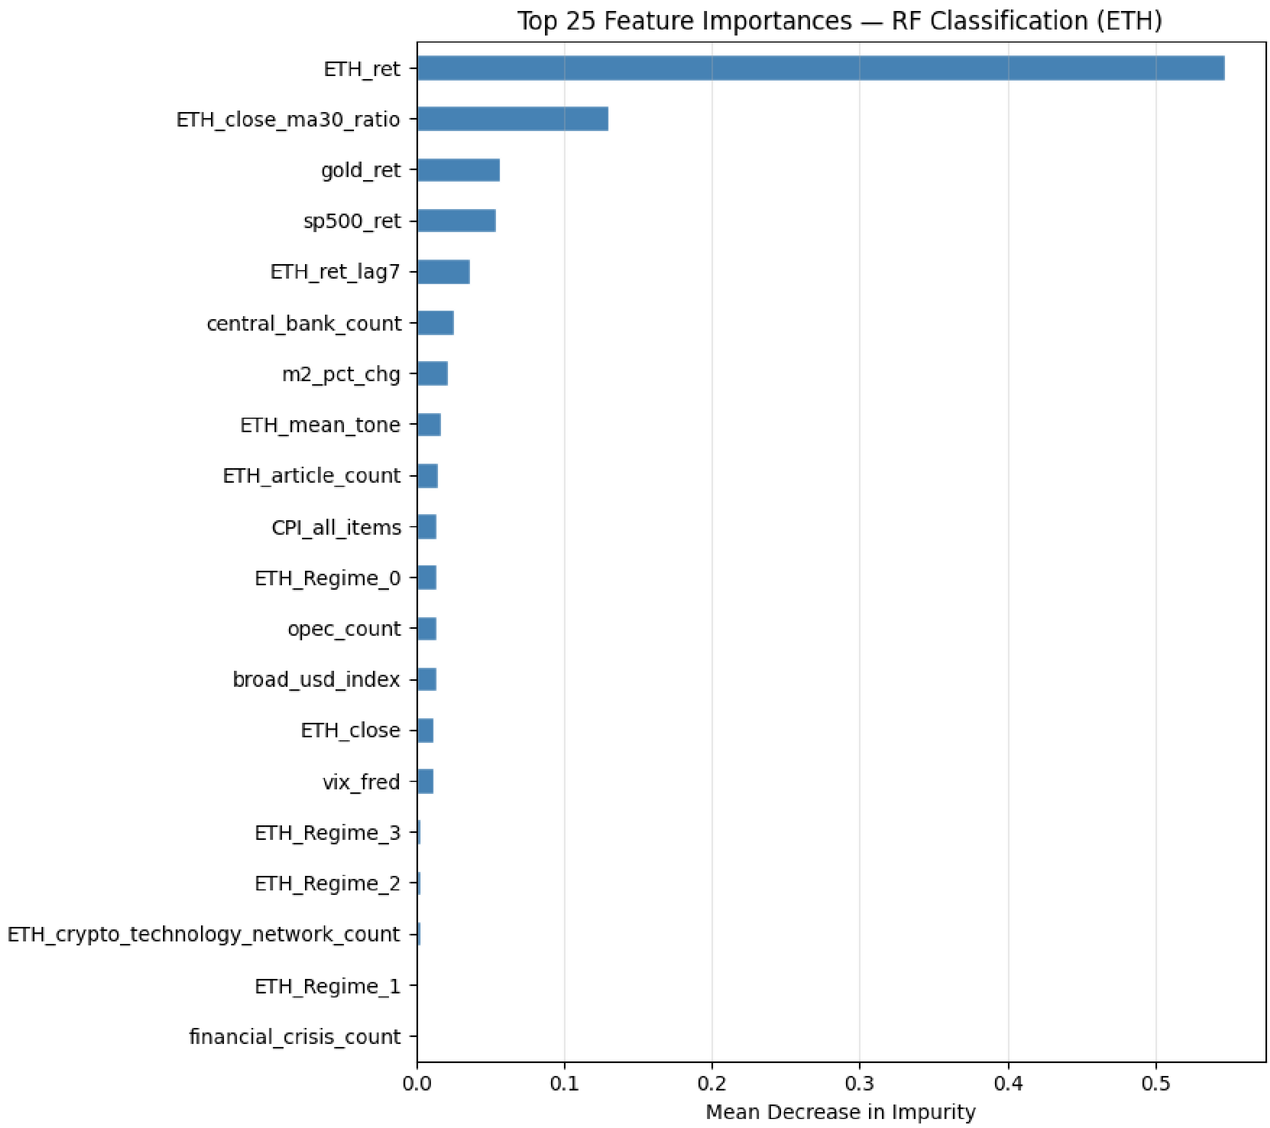

17 Fig. Feature Importance of the Ethereum Random Forest Classification Model

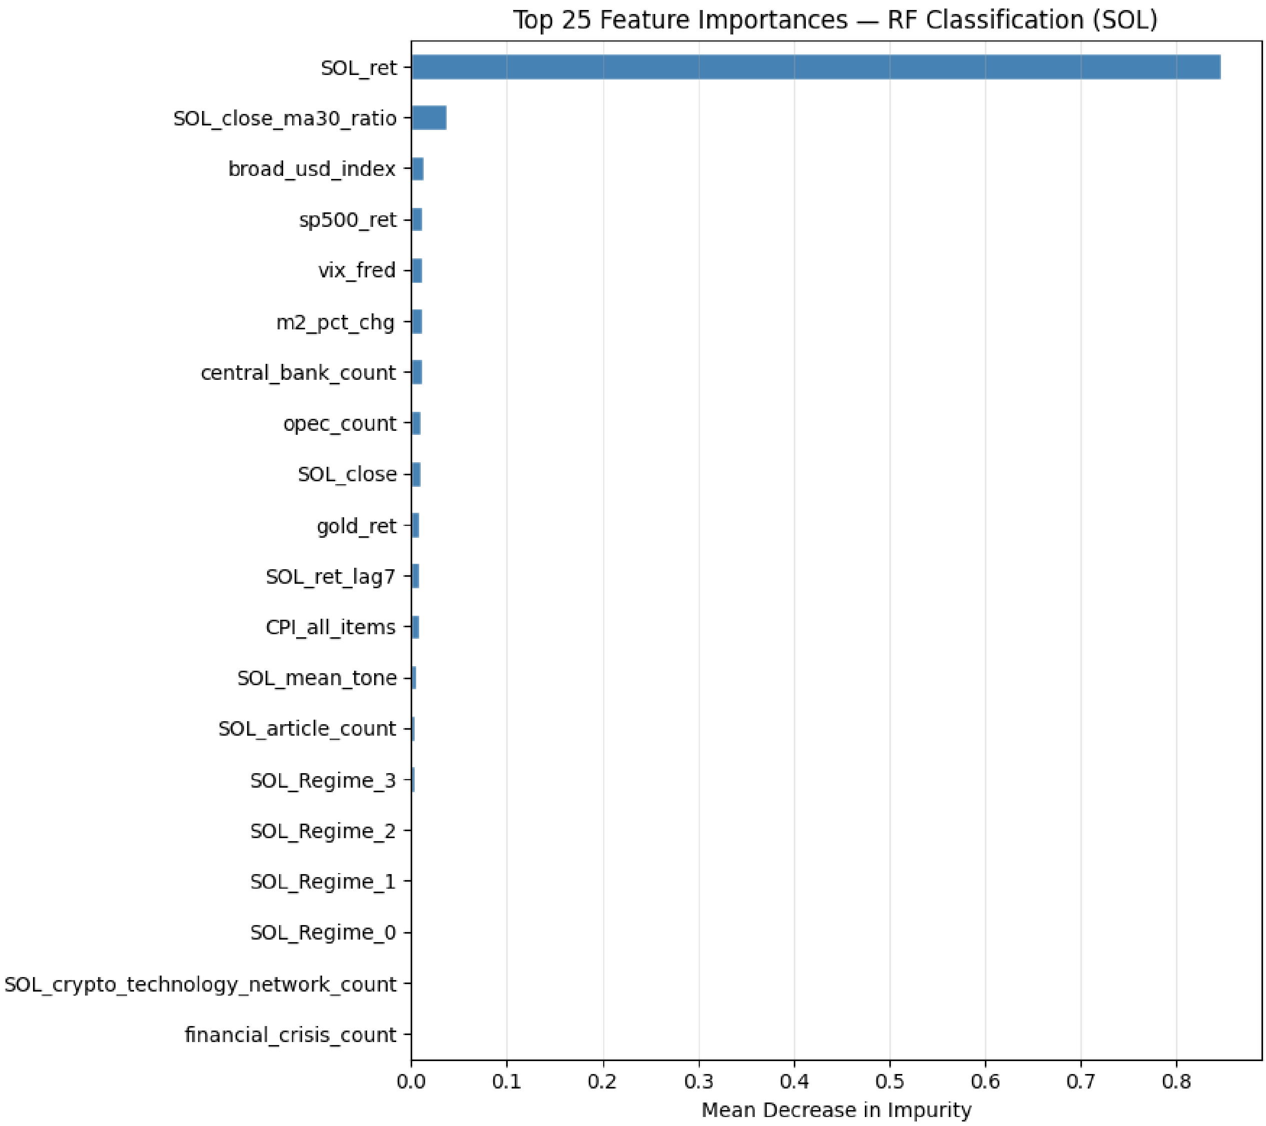

18 Fig. Feature Importance of the Solana Random Forest Classification Model

The feature-importance results show a clear common pattern across all three cryptocurrencies: the most influential variable in each model was the asset's own recent return indicator. For Bitcoin, the dominant feature was Bitcoin log returns (BTC_log_ret), with substantially higher importance than all other predictors. For Ethereum, the most important variable was Ethereum returns (ETH_ret), while for Solana the strongest feature was Solana returns (SOL_ret). This indicates that the models relied primarily on recent market movement when classifying future direction.

A second common pattern is that the ratio between the current closing price and the 30-day moving average also played an important role in all three models. For Bitcoin, BTC_close_ma30_ratio was the second most important variable; the same was true for ETH_close_ma30_ratio and SOL_close_ma30_ratio in the corresponding models. This suggests that, in addition to recent return behaviour, the models also used the asset's position relative to its recent trend as an important signal. In economic terms, this variable may reflect whether the asset is trading above or below its medium-term trend level.

Despite this common structure, the degree of concentration differed across assets. In the Bitcoin model, BTC_log_ret clearly dominated the feature ranking, but the remaining importance was still distributed across several additional variables, including sp500_ret, m2_pct_chg, BTC_close, CPI_all_items, BTC_article_count, opec_count and broad_usd_index. This suggests that the BTC model relied mainly on recent Bitcoin return, but also incorporated some information from broader macroeconomic, financial-market and news-related variables.

The Ethereum model showed a similar but somewhat more balanced structure. Although ETH_ret was by far the most important feature, the second-ranked variable, ETH_close_ma30_ratio, also had meaningful importance, followed by gold_ret, sp500_ret, ETH_ret_lag7, central_bank_count and m2_pct_chg. Compared with Bitcoin, the Ethereum model appears to have placed relatively more weight on broader financial-market and macroeconomic variables, even though the recent return signal still remained dominant.

For Solana, the concentration of importance was the strongest. The feature SOL_ret accounted for the overwhelming majority of impurity-based importance, while all other variables had only minor contributions. The second-ranked variable, SOL_close_ma30_ratio, was still much smaller, and the remaining features such as broad_usd_index, sp500_ret, vix_fred, m2_pct_chg, central_bank_count and opec_count had very limited importance. This suggests that the Solana model relied much more heavily on recent asset-specific return information than on macroeconomic, sentiment or regime-related variables.

Another notable result is that several news-based variables, such as article counts, mean tone and crypto-technology indicators, had low importance relative to return and trend-based features. This suggests that within the tree-splitting logic of the Random Forest, they were less frequently selected as the main variables for reducing impurity than recent return and trend position measures.

From an analytical perspective, these results show that the global predictive structure of the models was strongly market-driven. The models relied primarily on recent return behaviour and the relationship between price and recent trend. Macroeconomic and financial variables had secondary importance, while regime and news-based features played a more limited role in the global impurity-based ranking. This pattern was strongest for Solana, somewhat less concentrated for Bitcoin, and relatively more diversified for Ethereum.

At the same time, impurity-based feature importance should be interpreted cautiously. This method indicates which variables were globally useful for the Random Forest model, but it does not show whether high or low values of a feature push the prediction toward the UP or DOWN class. It is also sensitive to correlation structure between predictors. Therefore, the feature-importance ranking is treated here as an initial global overview, while the next section uses SHAP analysis to provide a more detailed interpretation of the direction and magnitude of feature effects.

3.4.2. Global SHAP Analysis

While impurity-based feature importance showed which variables were most useful to the Random Forest models overall, it did not indicate how these variables influenced predictions toward the UP or DOWN class. To address this, SHAP analysis was applied to the test set. In this study, SHAP values were interpreted for the UP class, meaning that a positive SHAP value increased the predicted probability of an upward movement, while a negative SHAP value decreased it.

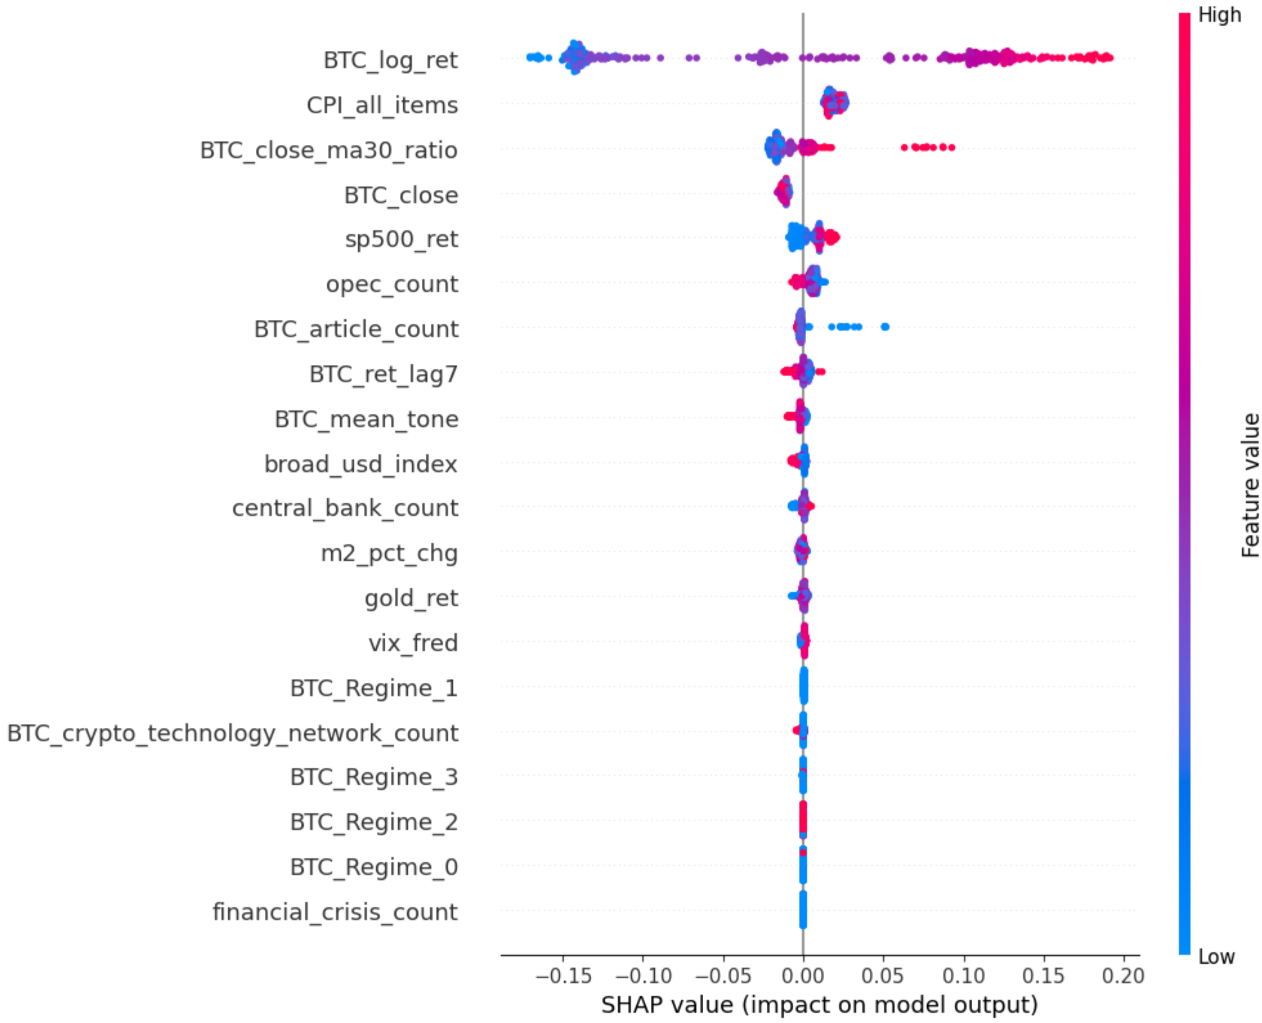

19 Fig. SHAP Beeswarm Plot for the Bitcoin Random Forest Classification Model

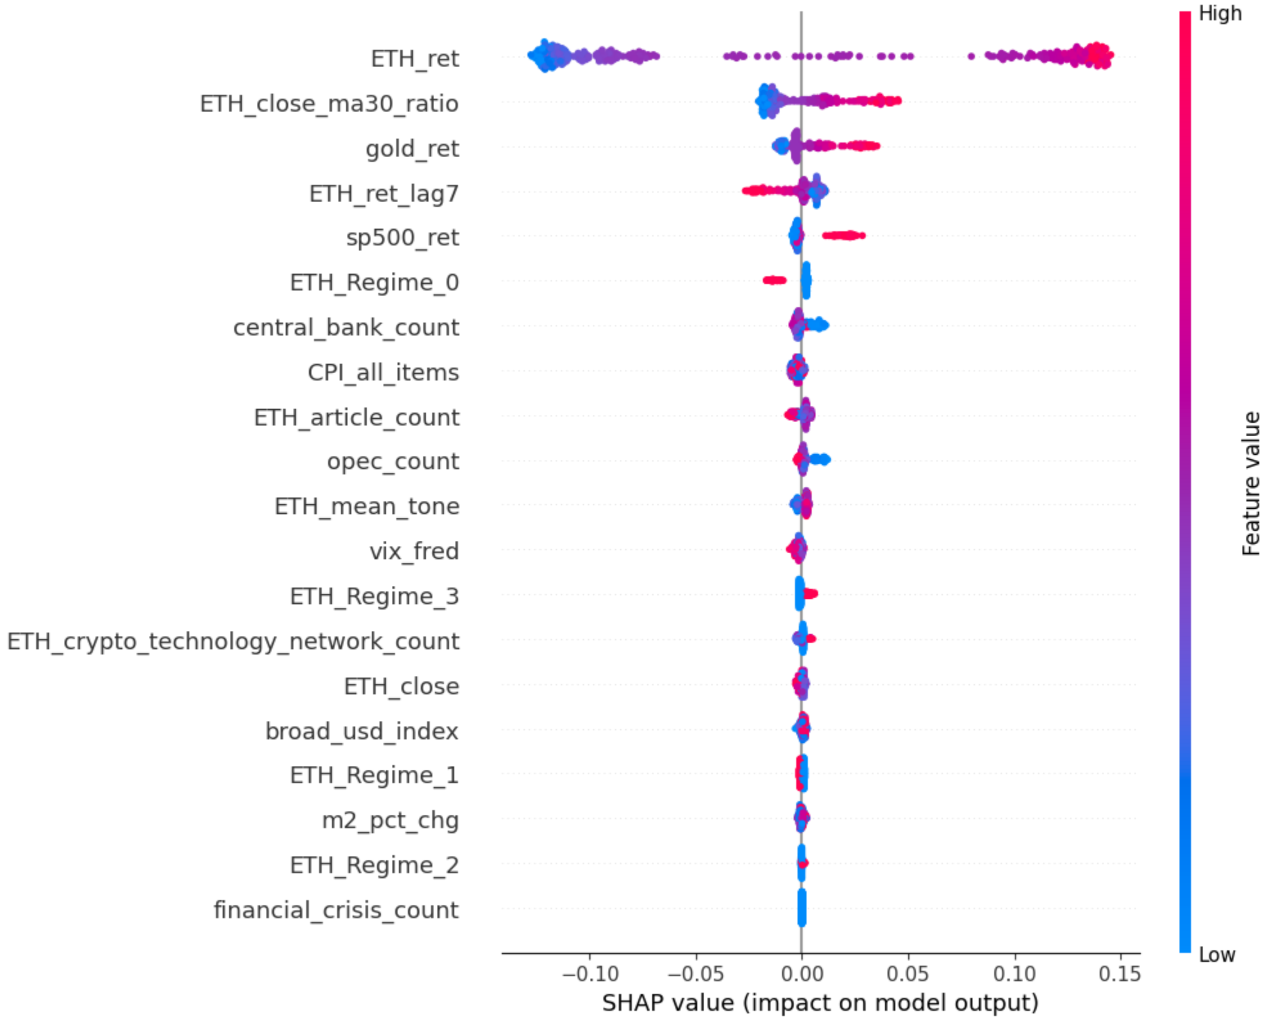

20 Fig. SHAP Beeswarm Plot for the Ethereum Random Forest Classification Model

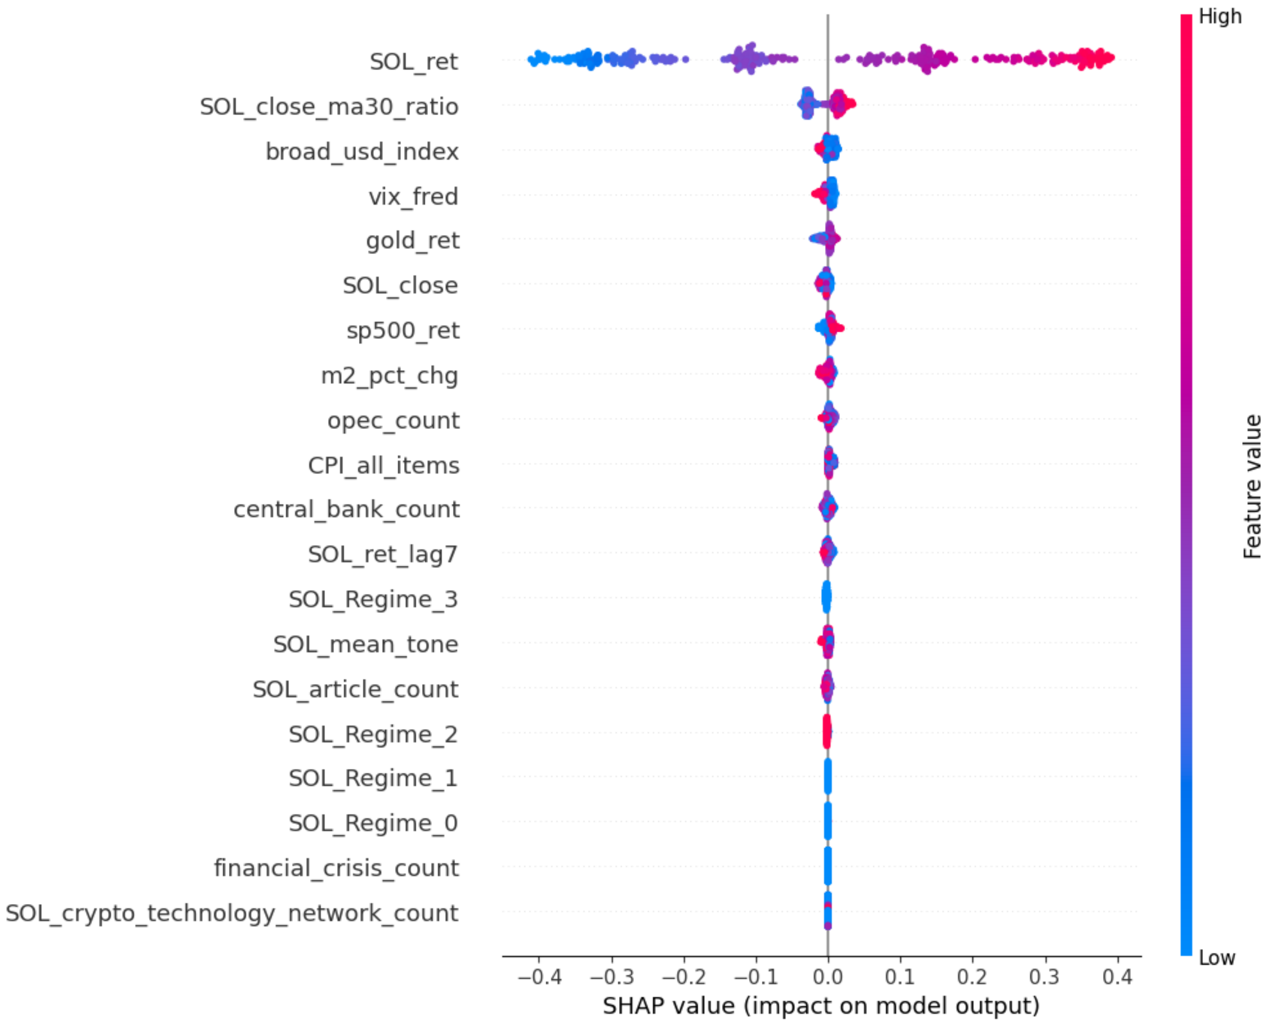

21 Fig. SHAP Beeswarm Plot for the Solana Random Forest Classification Model

The SHAP beeswarm plots confirm the main conclusion already suggested by the impurity-based feature importance analysis: in all three models, the dominant explanatory variable was the asset's own recent return measure. For Bitcoin, BTC_log_ret was clearly the most influential feature; for Ethereum, the strongest feature was ETH_ret; and for Solana, the dominant variable was SOL_ret. In all three cases, higher values of the return feature were associated with positive SHAP values, meaning that stronger recent returns pushed the model toward predicting an UP movement. Conversely, lower or more negative return values were associated with negative SHAP values, pushing the model away from the UP class.

For Bitcoin, the SHAP plot shows that BTC_log_ret had by far the widest spread of SHAP values, indicating that it was the main driver of variation in model output across observations. High values of BTC_log_ret strongly increased the probability of the UP class, while low values had a clear negative effect. The next most influential features were CPI_all_items and BTC_close_ma30_ratio. The latter showed a pattern in which higher values tended to push predictions toward the UP class, which is consistent with the interpretation that a price positioned above its recent moving-average trend is associated with stronger market conditions. Other variables, such as sp500_ret, BTC_article_count, BTC_ret_lag7 and BTC_mean_tone, had noticeably smaller effects.

For Ethereum, the overall structure was similar but slightly more balanced. ETH_ret again dominated the model, with high return values strongly associated with positive SHAP values and low return values with negative SHAP values. ETH_close_ma30_ratio was the second most important variable and showed a broadly similar directional pattern to Bitcoin, where higher values were associated with higher UP-class probability. Additional features such as gold_ret, ETH_ret_lag7 and sp500_ret had secondary importance, but their SHAP spreads were still much smaller than that of the main return feature.

For Solana, the concentration of explanatory power was even more pronounced. SOL_ret dominated the SHAP ranking and had the broadest SHAP range among all three assets. High values of SOL_ret strongly increased the probability of an UP prediction, while low values strongly reduced it. This confirms that the Solana model relied very heavily on recent market movement. SOL_close_ma30_ratio was the second most important feature, but its contribution was much smaller than that of SOL_ret. The remaining variables, including broad_usd_index, vix_fred, gold_ret, sp500_ret and macroeconomic indicators, had only modest SHAP influence, while regime and news-related variables again remained close to zero for most observations.

A comparison across the three assets suggests that the predictive logic of the models was broadly similar, but not equally concentrated. All three models relied most strongly on recent return information, and in all cases higher return values supported UP predictions. However, the dominance of the return feature was strongest for Solana, somewhat less extreme for Bitcoin, and relatively more diversified for Ethereum. This indicates that the Solana model was driven most heavily by immediate asset-specific market signals, while the Ethereum model incorporated slightly more influence from trend-position and broader market variables.

The SHAP beeswarm plots also provide an important methodological insight. Although macroeconomic, sentiment and regime variables were included in the modelling framework, their influence on the final test-set predictions was generally much smaller than that of recent return and trend-based variables. This does not imply that such variables are universally unimportant in cryptocurrency analysis. Rather, it suggests that within the final Random Forest models estimated in this study, the strongest predictive structure came primarily from internal market information rather than from external news or latent regime indicators.

At the same time, these results should be interpreted carefully. SHAP explains how the trained model used the available features, but it does not prove causal relationships. For example, the finding that high recent returns push predictions toward the UP class does not mean that recent return itself causes future growth in a structural economic sense. Instead, it means that the model learned a decision pattern in which recent return information was a strong predictor of future direction within the given sample.

Overall, the global SHAP analysis confirms and extends the feature-importance results. The models were driven primarily by recent return and trend-position variables, while macroeconomic, sentiment and regime-related features played a secondary role. The next subsection examines these relationships in greater detail using SHAP dependence plots, which allow the shape of the feature-prediction relationship to be analysed more directly.

3.4.3. SHAP Dependence Analysis

SHAP dependence plots were used to examine the relationship between the most influential feature and its contribution to the UP-class prediction. For all three cryptocurrencies, the most important feature was the asset's own recent return measure. Therefore, the dependence analysis focused on *BTC_log_ret* for Bitcoin, *ETH_ret* for Ethereum and *SOL_ret* for Solana. These plots show how changes in the return variable affected the SHAP value of that feature, meaning how strongly it pushed the model toward or away from an UP prediction.

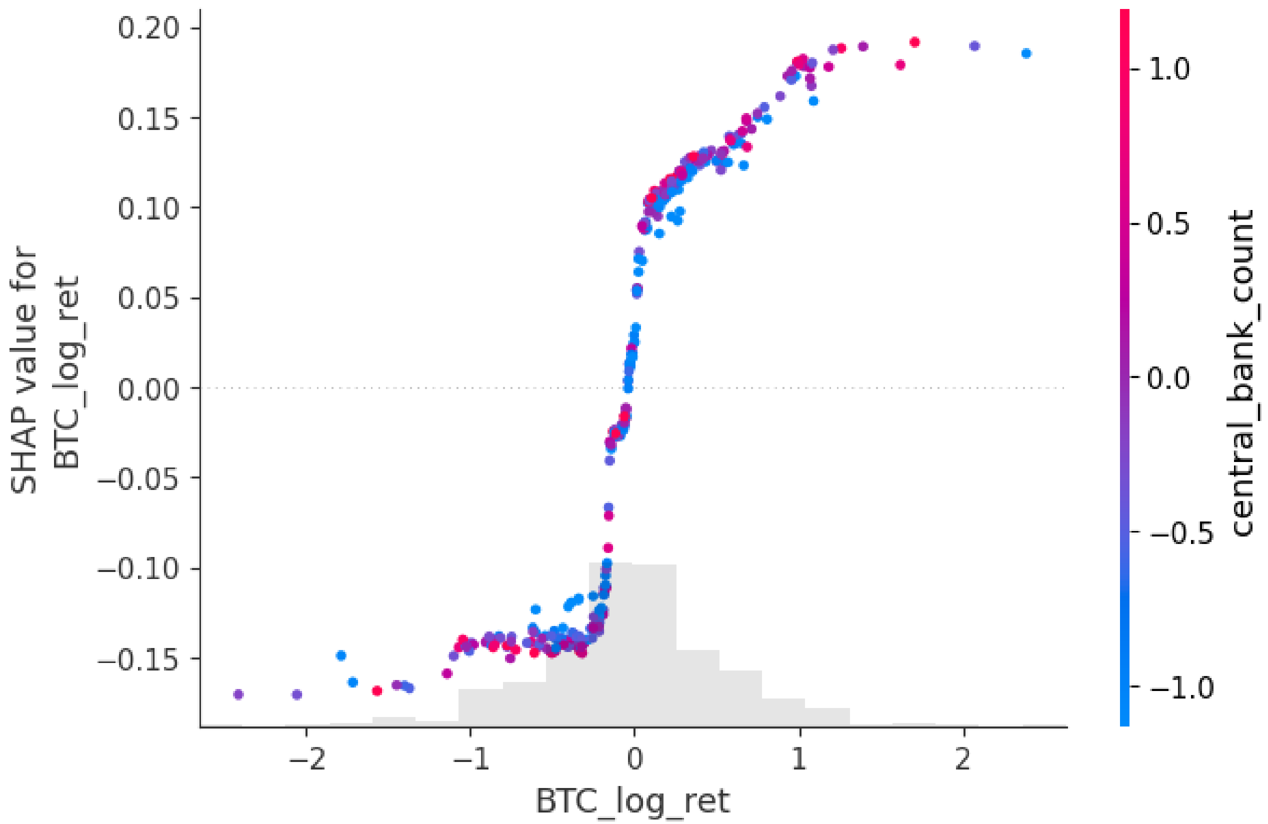

22 Fig. SHAP Dependence Plot for BTC log returns

or Bitcoin, the dependence plot shows a clear nonlinear relationship between Bitcoin logarithmic returns and its SHAP contribution. Negative values of log returns generally had negative SHAP values, meaning that they reduced the predicted probability of an UP movement. As the value approached zero, the SHAP contribution changed sharply, and positive values of log returns increasingly pushed the model toward the UP class. This suggests that the BTC model treated recent return as a threshold-like signal: negative or weak recent returns reduced the probability of an upward prediction, while stronger positive returns increased it.

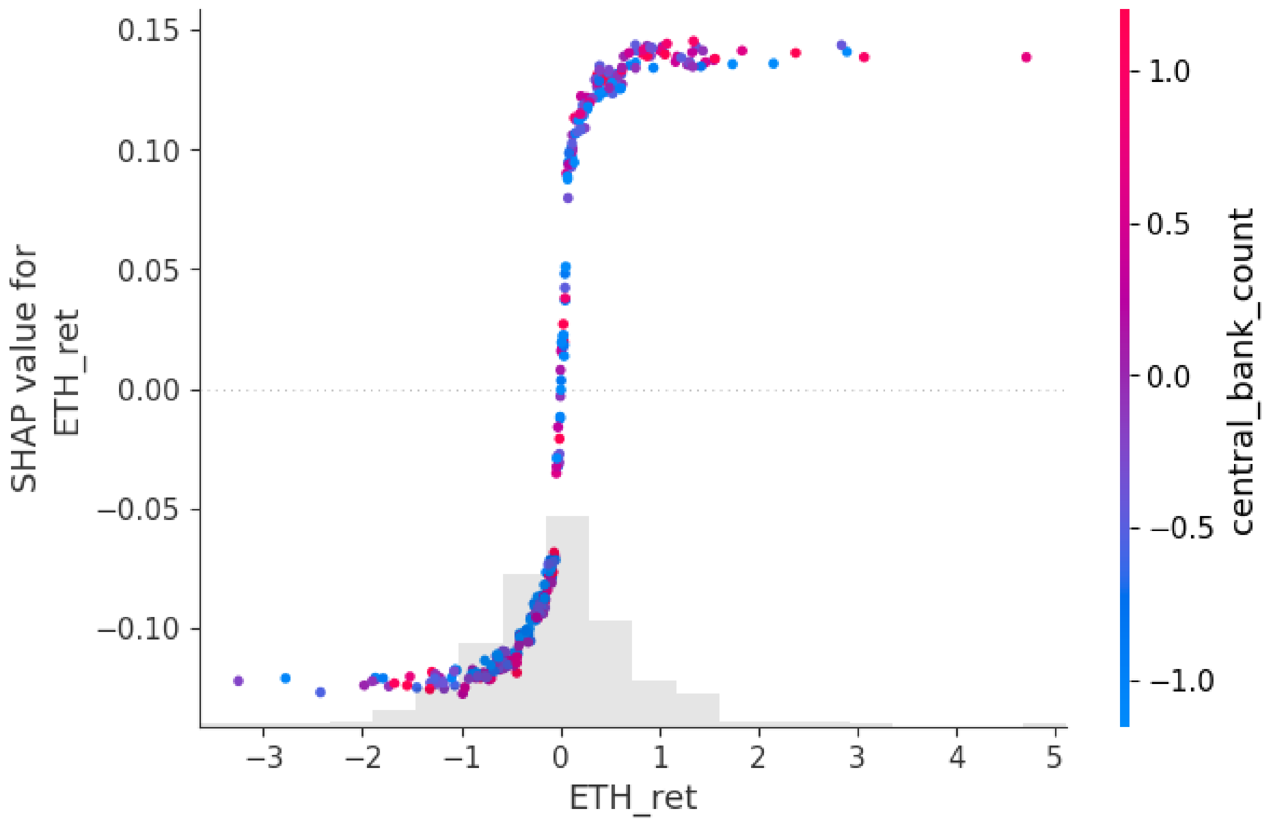

23 Fig. SHAP Dependence Plot for ETH log returns

The Ethereum dependence plot shows a similar pattern. Negative *ETH_ret* values mostly produced negative SHAP contributions, while positive values were associated with positive SHAP values. The transition around zero was especially sharp, indicating that the ETH model strongly distinguished between negative and positive recent return conditions. After the return became positive, the SHAP contribution quickly stabilised at a higher level, suggesting that positive recent ETH returns consistently increased the probability of an UP prediction.

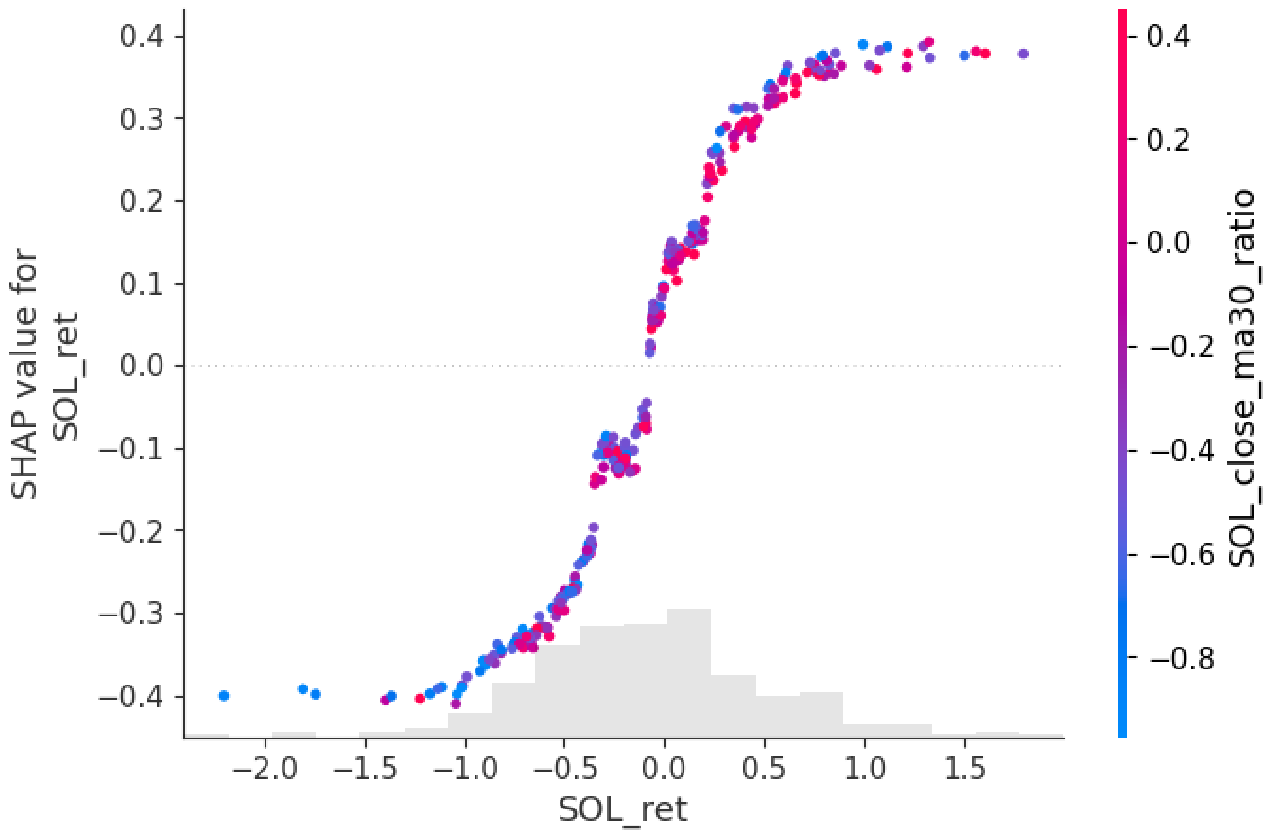

24 Fig. SHAP Dependence Plot for SOL log returns

For Solana, the relationship was even stronger and wider in magnitude. Negative *SOL_ret* values produced large negative SHAP contributions, while positive values produced large positive contributions. Compared with BTC and ETH, the SOL plot shows a broader SHAP range, which is consistent with the earlier feature-importance and beeswarm results showing that *SOL_ret* dominated the Solana model. The pattern indicates that the SOL classifier relied heavily on recent return behaviour when assigning UP or DOWN probabilities.

Across all three assets, the dependence plots show that the Random Forest models learned a nonlinear return-based decision structure. The relationship was not simply gradual across the full range of values; instead, the strongest shift occurred around the transition from negative to positive recent returns. This is consistent with a momentum-like model logic, where recent positive movement increases the probability of a future UP classification, while recent negative movement decreases it. However, this should be interpreted as a model-based relationship rather than causal evidence. The SHAP dependence plots explain how the trained Random Forest models used return variables, but they do not prove that recent returns cause future cryptocurrency price increases.

The dependence analysis also confirms that the models were mainly driven by internal market information. Although news, macroeconomic and regime variables were included in the modelling framework, the strongest and clearest marginal effect came from the asset's own recent return variable. This finding is important for the interpretation of the overall results: the models achieved directional classification performance mainly by exploiting recent market movement and trend-related information, while external variables played a more secondary role.

3.4.4. Partial dependence and individual conditional expectations of Return-Based Effects

To complement the SHAP analysis, Partial Dependence Plots (PDP) and Individual Conditional Expectation (ICE) plots were used for the most important return-based variable in each model. While SHAP values explain how features contributed to individual predictions, PDP and ICE plots show how the predicted probability of the UP class changes when a selected feature varies. This provides an additional view of the model's learned response function.

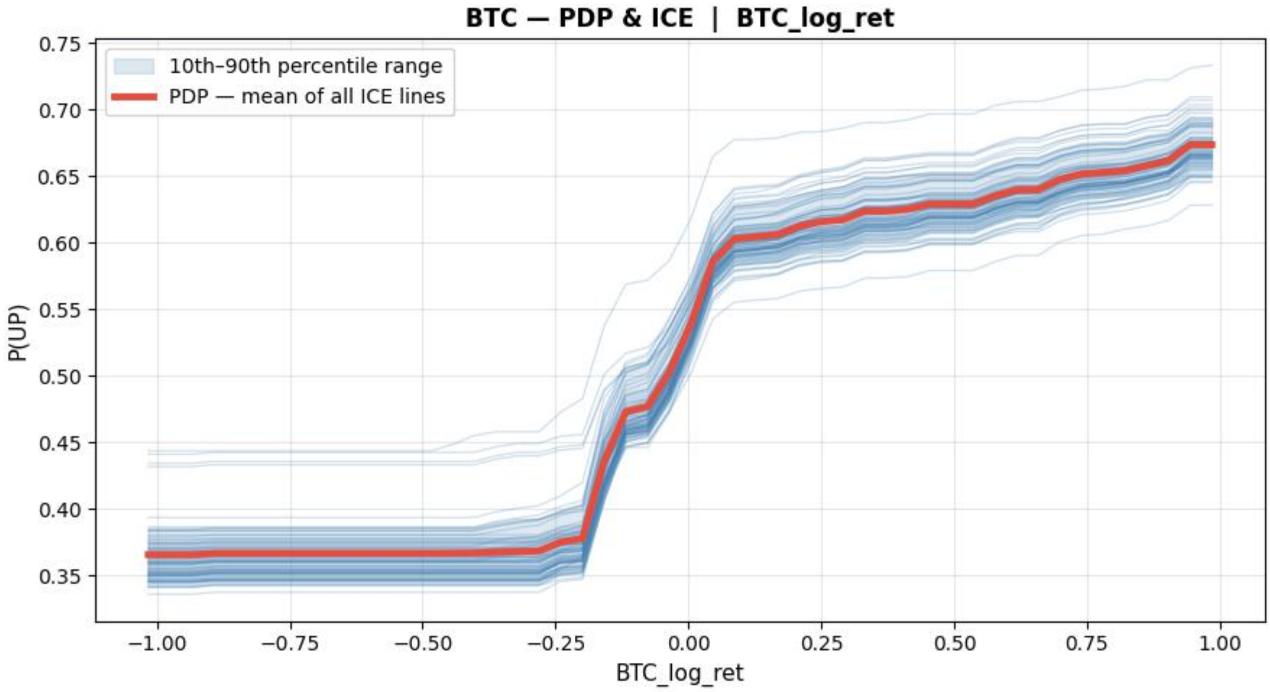

25 Fig. Partial Dependence and Individual Conditional Expectation of Bitcoin UP Probability

For Bitcoin, the PDP and ICE plot shows that the predicted probability of the UP class increases as *BTC_log_ret* rises. When *BTC_log_ret* is clearly negative, the average predicted probability remains low, around 0.35-0.40. The strongest increase occurs around the transition from negative to slightly positive log returns. After this point, the predicted probability rises above 0.60 and continues to increase gradually for higher positive return values. This pattern supports the SHAP dependence result: the BTC model treated recent log return as a nonlinear directional signal, where the shift from negative to positive return values had the strongest effect on the UP prediction.

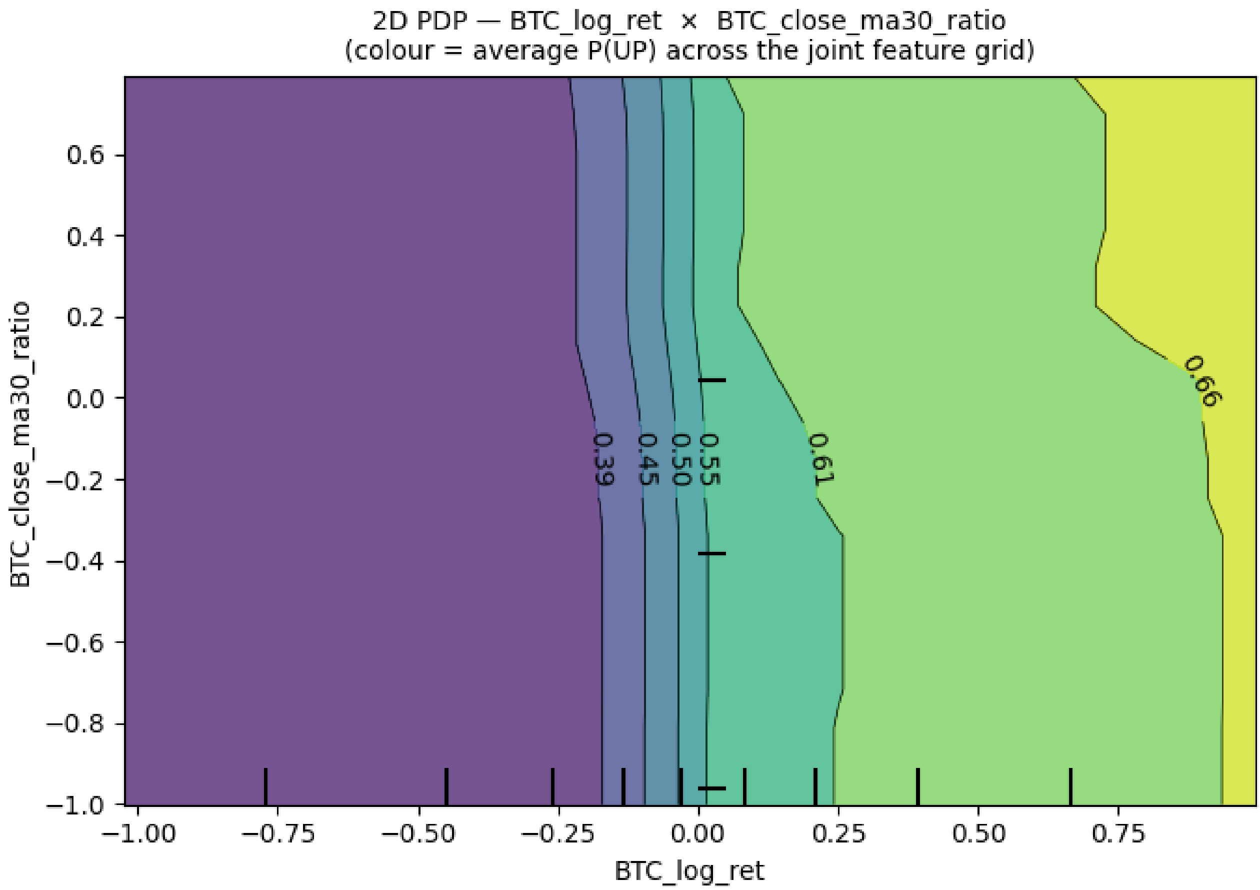

26 Fig. Two-Dimensional Partial Dependence Plot for Bitcoin Log Return and Bitcoin Close-to-30-Day Moving Average Ratio

The two-dimensional PDP for Bitcoin shows the joint effect of *BTC_log_ret* and *BTC_close_ma30_ratio*. The probability of an UP prediction is lowest when recent log return is negative, regardless of the trend-ratio value. When *BTC_log_ret* becomes positive, the predicted probability increases substantially. The moving-average ratio adds additional information, but the horizontal movement across the return axis is the dominant effect. This confirms that recent return was the main driver, while the price position relative to the 30-day moving average acted as a secondary trend-related signal.

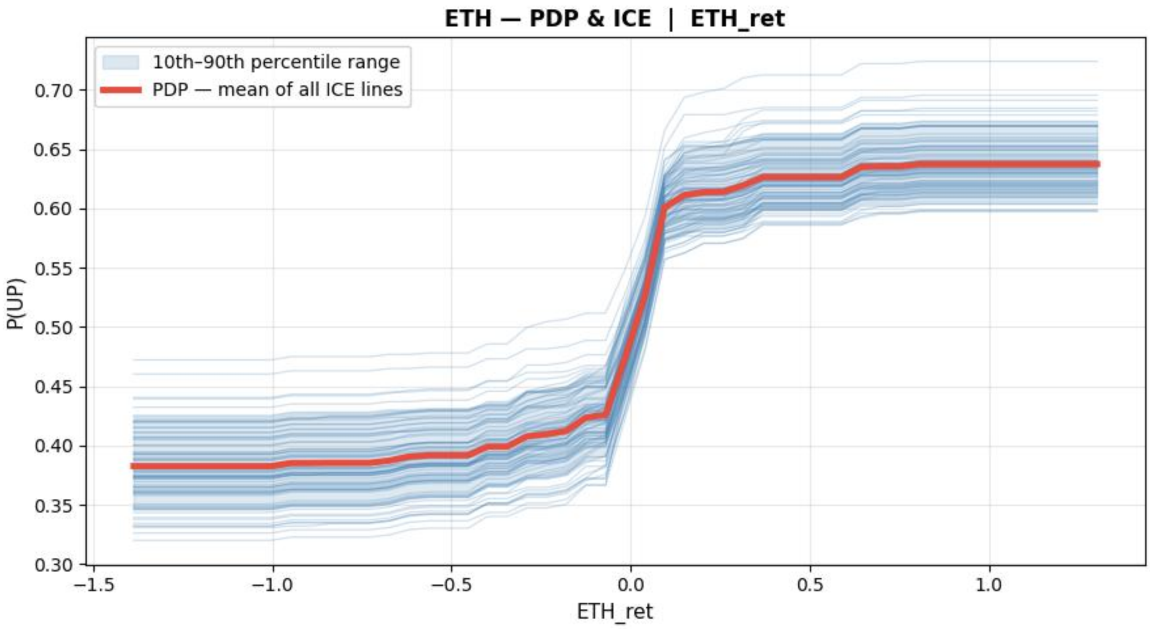

27 Fig. Partial Dependence and Individual Conditional Expectation of Ethereum UP Probability

For Ethereum, the PDP and ICE plot shows a similar threshold-like pattern. When *ETH_ret* is negative, the average predicted UP probability remains below 0.45. Around zero, the probability increases sharply, rising to approximately 0.60. After the return becomes positive, the probability stabilises at a higher level and increases more slowly. The ICE curves follow a relatively similar direction, suggesting that the return effect was consistent across many test-set observations. This indicates that the ETH model also relied on recent return as a central signal for distinguishing between UP and DOWN conditions.

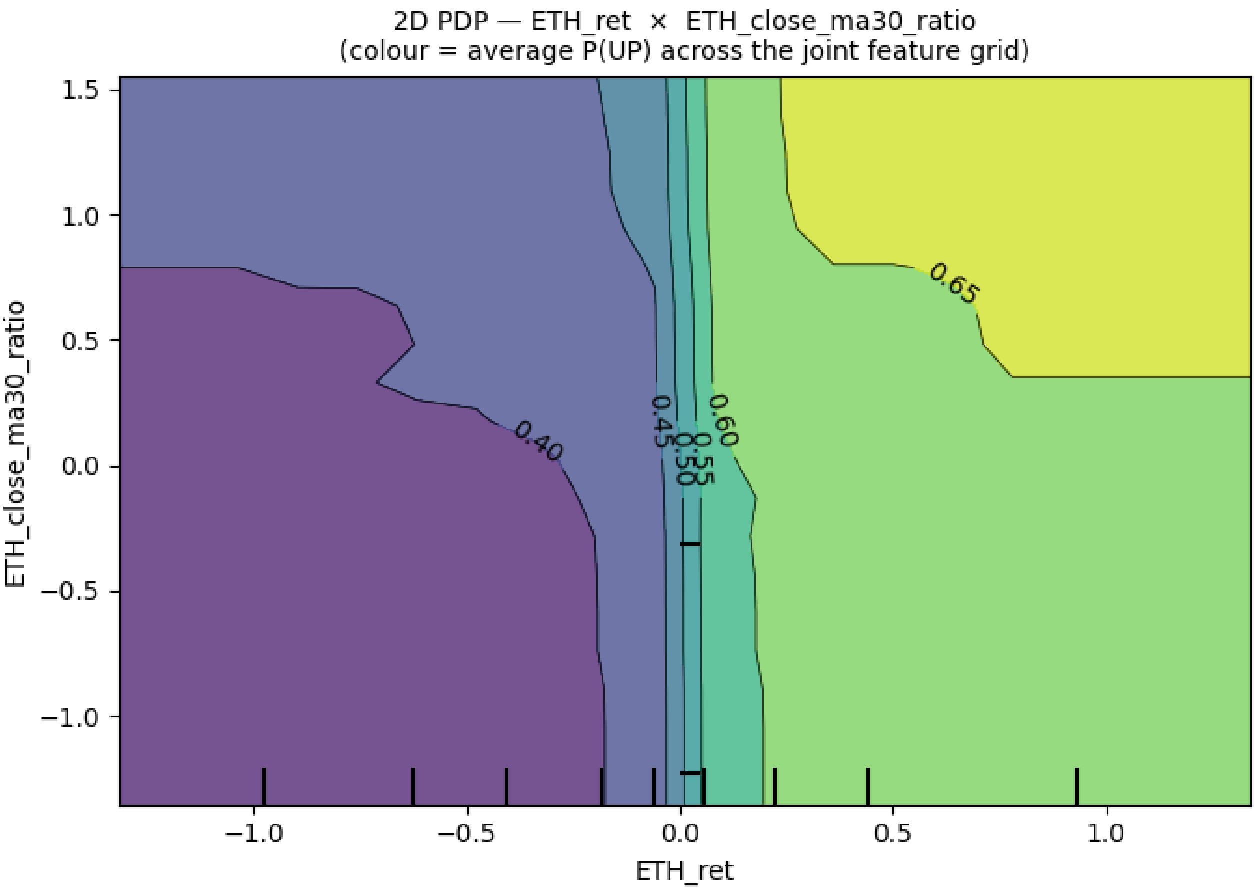

28 Fig. Two-Dimensional Partial Dependence Plot for Ethereum Log Return and Ethereum Close-to-30-Day Moving Average Ratio

The Ethereum two-dimensional PDP again shows that the return variable has the strongest effect on the predicted UP probability. Negative *ETH_ret* values are associated with lower probabilities, while positive values are associated with higher probabilities. The *ETH_close_ma30_ratio* provides additional variation, especially when the return is close to zero or moderately positive, but it does not dominate the effect of *ETH_ret*. Economically, this suggests that the model mainly used recent price movement, while the asset's position relative to its medium-term trend helped refine the prediction.

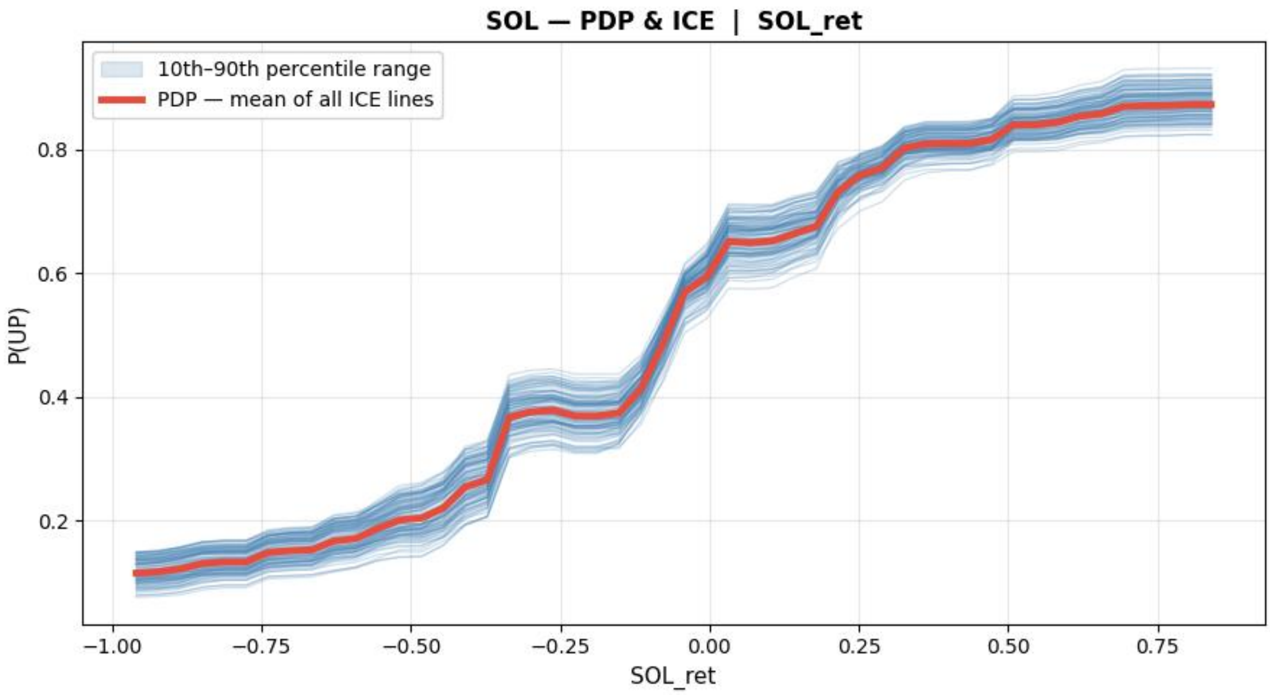

29 Fig. Partial Dependence and Individual Conditional Expectation of Solana UP Probability

or Solana, the PDP and ICE plot shows the strongest and widest probability response among the three assets. When *SOL_ret* is strongly negative, the predicted probability of the UP class is very low. As *SOL_ret* increases, the predicted probability rises steadily and reaches high levels when the return becomes clearly positive. Compared with BTC and ETH, the Solana plot shows a larger change in predicted probability across the range of the return variable. This is consistent with the earlier feature-importance and SHAP results, where *SOL_ret* was the most dominant predictor.

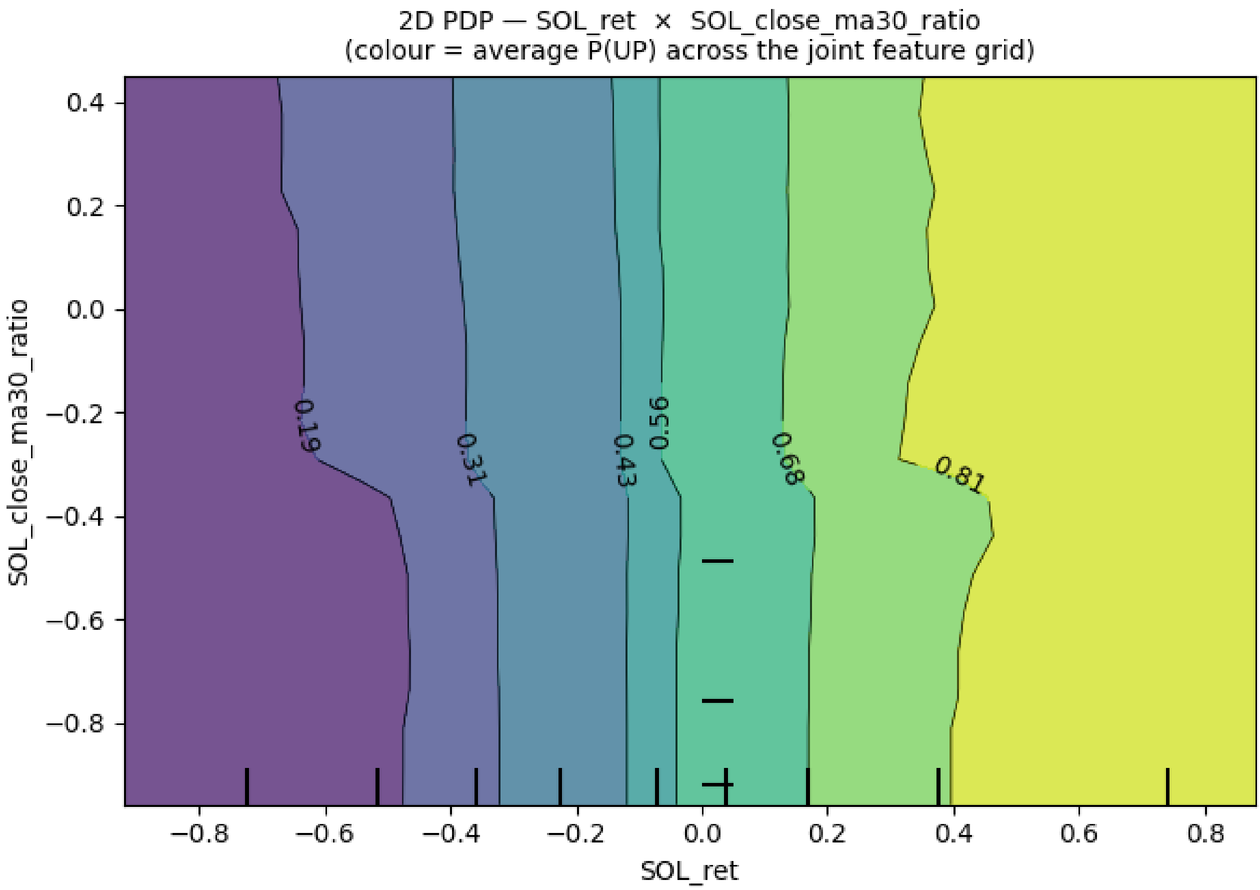

30 Fig. Two-Dimensional Partial Dependence Plot for Solana Log Return and Solana Close-to-30-Day Moving Average Ratio

The two-dimensional PDP for Solana confirms that *SOL_ret* strongly controls the model's UP probability. The probability increases from low values when return is negative to high values when return is positive. The *SOL_close_ma30_ratio* also contributes to the model response, but its effect is secondary compared with the return variable. The highest UP probabilities are observed when the return variable is positive, while low or negative return values are associated with low predicted probabilities even when the trend-ratio variable changes.

Overall, the PDP and ICE results support the conclusions from feature importance and SHAP analysis. Across all three assets, recent return was the dominant driver of predicted UP probability. The relationship was nonlinear and threshold-like: negative or weak recent returns were associated with lower UP probabilities, while positive recent returns strongly increased the probability of an upward prediction. The effect was visible for all assets, but it was strongest for Solana, which is consistent with the model's high reliance on *SOL_ret*.

The two-dimensional PDPs also show that the close-to-30-day-moving-average ratio added trend-position information, but it did not replace the role of recent return. In all three models, the return variable remained the main axis along which predicted UP probability changed. Therefore, the PDP/ICE analysis reinforces the broader economic interpretation of the models: the Random Forest classifiers primarily captured short-term continuation or momentum-like market behaviour, while trend-position variables provided secondary support. As with SHAP results, these patterns should be interpreted as model behaviour rather than causal evidence about cryptocurrency price formation.

3.5. Economic Interpretation of Predictive Signals and Market Dynamics

After the technical evaluation of classification metrics, trading signals and explainability outputs, the results can be interpreted from an economic perspective. This is important because the objective of the thesis is not only to build predictive models, but also to analyse what the models reveal about cryptocurrency market dynamics and whether their outputs can support practical decision-making.

The first important finding is that the models were mainly driven by recent asset-specific market movement. Feature importance, SHAP analysis and PDP/ICE plots all showed that *BTC_log_ret, ETH_ret* and *SOL_ret* were the dominant predictors in their respective models. In all three cases, negative recent returns reduced the predicted probability of an UP movement, while positive recent returns increased it. Economically, this suggests that the models mostly captured short-term continuation or momentum-like behaviour rather than deeper fundamental valuation effects. Therefore, the models should be interpreted as short-horizon directional signal models.

A second finding is that the importance of external variables was weaker than initially expected. The modelling framework included macroeconomic indicators, crypto-specific news variables, macro-news indicators and HMM regime variables, but these variables generally had smaller influence than return- and trend-based features. This does not mean that macroeconomic conditions, sentiment or regimes are irrelevant for cryptocurrency markets. Rather, it suggests that, within the tested daily classification framework, the strongest direct signal came from internal market behaviour. External variables may still provide useful context, but their contribution to short-horizon directional classification was secondary.

One notable finding is that sentiment-based news variables did not have a strong influence on model decisions. Variables such as mean news tone and negative-news share were included to capture whether positive or negative information flow affected cryptocurrency direction, but they did not appear among the dominant predictors in the final models. This suggests that, in the tested daily classification framework, aggregated news sentiment was not as informative as recent market movement. A possible explanation is that broad daily sentiment measures may be too noisy or too general to capture the exact timing of market reactions. Cryptocurrency prices may react quickly to news, meaning that part of the sentiment effect may already be reflected in returns by the time the daily feature is observed.

Among the alternative and external information variables, macro-news attention related to central banks appeared more relevant than general crypto-news tone. This is economically interesting because it suggests that the models responded more to broader monetary-policy and financial-market attention than to average sentiment in cryptocurrency news. Central bank news can proxy expectations about liquidity, interest rates and risk appetite, which are important for speculative assets such as cryptocurrencies. However, this result should be interpreted cautiously. The finding does not mean that central bank news directly caused cryptocurrency movements; it only shows that the trained models found this variable more useful than most other news-based indicators within the selected feature set and test period.

In short-term decision-making, recent price behaviour may contain more immediately actionable information than broad daily macro-news or sentiment aggregates. News and macroeconomic variables may influence the market indirectly, during specific events, or over longer horizons, but their effect may be more difficult to capture in a daily classification model. Therefore, the weaker role of external variables should be interpreted as a limitation of their direct predictive power in this framework, not as evidence that they have no economic relevance.

In the test periods shown, the strategies did not always capture the full upside of buy-and-hold during strong rallies, but they reduced exposure during several weaker or declining phases. This was especially visible for Ethereum and Solana, where the model-based strategies avoided a significant part of the buy-and-hold drawdown. From a business or investment perspective, this distinction is important. In highly volatile markets such as cryptocurrencies, reducing downside exposure can be valuable even if the strategy does not fully participate in every upward movement. A model that helps identify when exposure should be reduced may support risk management, capital preservation and more disciplined decision-making.

The results also show that explainable artificial intelligence adds value beyond predictive metrics. Without SHAP and PDP/ICE analysis, the models would only show whether UP or DOWN predictions were correct. Explainability methods showed that the models relied mainly on recent return and trend-position variables, while external variables played a smaller role. This makes the model behaviour more transparent and allows the results to be interpreted economically. For example, the dominance of return-based variables suggests that the models were not primarily driven by news sentiment or macroeconomic narratives, but by short-term internal market dynamics.

Overall, the empirical results suggest that explainable artificial intelligence can support cryptocurrency market analysis in two main ways. First, it can identify which variables drive directional predictions and whether these drivers are economically plausible. Second, it can connect model outputs with practical exposure decisions through backtesting and regime-aware evaluation. In this study, the main economic insight is that short-term cryptocurrency direction was driven primarily by recent asset-specific market movement, while external information variables had a secondary role. The practical value of the models therefore lies mainly in supporting conditional exposure management under changing market conditions rather than in providing a universally profitable trading strategy.

4. Conclusions

1. The research problem analysis showed that cryptocurrency market dynamics are difficult to evaluate because these markets are highly volatile, continuously traded, sensitive to news and affected by changing market regimes. This confirms the relevance of explainable artificial intelligence, as cryptocurrency prediction should not only indicate future market direction, but also explain which variables influence model decisions and whether the results have economic meaning.

2. The theoretical review showed that cryptocurrency market behaviour is influenced by internal market variables, technical indicators, investor sentiment, macro-financial conditions and regime changes. It also showed that machine learning models can be useful for financial prediction, but their value depends on whether predictions can be interpreted through explainability methods, regime-aware evaluation and economic testing.

3. The research methodology was developed by combining daily market data, derived technical indicators, news-based variables, macro-financial indicators and HMM-based market regimes into one modelling framework. Random Forest classifiers were used for UP/DOWN prediction, while feature importance, SHAP, PDP/ICE analysis and backtesting were applied to evaluate model interpretability and economic relevance.

4. The empirical evaluation showed that Random Forest models achieved above-random prediction results for Bitcoin, Ethereum and Solana. Bitcoin reached accuracy and macro F1 of about 63%, while Ethereum and Solana reached about 75–76%. HMM-based state detection also captured statistically distinct market periods, including sharp changes around the pandemic-related shock period. Explainability analysis showed that recent asset-specific return variables were the main prediction drivers, which are accounting for about 60-85% of feature importance, while news, macroeconomic and regime variables had weaker direct influence. Backtesting confirmed that the signals were economically interpretable, mainly by helping reduce exposure in some weaker market periods rather than consistently outperforming buy-and-hold strategies.

This study has several limitations. The news variables were constructed using GDELT data, keyword-based filtering and automated tone indicators, so they may contain noise, duplicated information or incomplete representation of relevant market narratives. The HMM states should also be interpreted cautiously, because they represent statistical states inferred from selected market variables rather than directly observable economic regimes.

As return-based variables dominated the model explanations, future research could test whether stronger news-processing methods, on-chain indicators or regime-specific models improve the contribution of external information. Since the backtesting results showed economic interpretability but not consistent outperformance over buy-and-hold, future research should also evaluate alternative probability thresholds and position-sizing rules. This would help determine whether explainable model signals can be used not only for interpretation, but also for more robust risk-management decisions.

List of references

1. NAKAMOTO, Satoshi. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008.
2. BANK FOR INTERNATIONAL SETTLEMENTS. Cryptocurrencies: looking beyond the hype. In: BIS Annual Economic Report. 2018. Chapter 5.
3. LEVANTESI, Susanna; PISCOPO, Gabriella; ROVIELLO, Alba. Cryptocurrency in global dynamics: Analyzing the Crypto Volatility Index and financial markets with machine learning. Physica A: Statistical Mechanics and its Applications. 2025. Vol. 674. DOI 10.1016/j.physa.2025.130770.
4. LO, Andrew W. Efficient Markets Hypothesis. In: The New Palgrave Dictionary of Economics. 3rd ed. 2018.
5. DAVIES, Shaun William. Speculation Sentiment. Journal of Financial and Quantitative Analysis. 2022. Vol. 57, no. 7. DOI 10.1017/S0022109022000291.
6. MANAO, Alwinda; WASPADA, Ikaputera; SARI, Maya. Efficient Capital Market Theory and its Implications for Investment Decision Making in Financial Markets. Dinasti International Journal of Economics, Finance & Accounting. 2025. Vol. 5, no. 6. DOI 10.38035/dijefa.v5i6.3682.
7. WEI, Wang Chun. Liquidity and market efficiency in cryptocurrencies. Economics Letters. 2018. Vol. 168. DOI 10.1016/j.econlet.2018.04.003.
8. JAGRIČ, Timotej; LUETIĆ, Davor; MUMEL, Damijan; HERMAN, Aljaž. Detecting Potential Investors in Crypto Assets: Insights from Machine Learning Models and Explainable AI. Information. 2025. Vol. 16, no. 4. DOI 10.3390/info16040269.
9. TRAN, Vu Le; LEIRVIK, Thomas. Efficiency in the markets of crypto-currencies. Finance Research Letters. 2020. Vol. 35. DOI 10.1016/j.frl.2019.101382.
10. CHU, Jeffrey; ZHANG, Yuanyuan; CHAN, Stephen. The adaptive market hypothesis in the high frequency cryptocurrency market. International Review of Financial Analysis. 2019. Vol. 64. DOI 10.1016/j.irfa.2019.05.008.
11. EASLEY, David; O'HARA, Maureen; YANG, Songshan; ZHANG, Zhibai. Microstructure and Market Dynamics in Crypto Markets. SSRN Electronic Journal. 2024. DOI 10.2139/ssrn.4814346.
12. CIAIAN, Pavel; RAJCANIOVA, Miroslava; KANCS, d'Artis. The economics of BitCoin price formation. Applied Economics. 2016. Vol. 48, no. 19. DOI 10.1080/00036846.2015.1109038.
13. KRISTOUFEK, Ladislav. What are the main drivers of the bitcoin price? Evidence from wavelet coherence analysis. PLoS ONE. 2015. Vol. 10, no. 4. DOI 10.1371/journal.pone.0123923.
14. BAUR, Dirk G.; DIMPFL, Thomas. The volatility of Bitcoin and its role as a medium of exchange and a store of value. Empirical Economics. 2021. Vol. 61, no. 5. DOI 10.1007/s00181-020-01990-5.
15. IYER, Roshan. New Evidence on Spillovers Between Crypto Assets and Financial Markets. IMF Working Papers. 2023. Vol. 2023, no. 213. DOI 10.5089/9798400256622.001.

16. LYÓCSA, Štefan; MOLNÁR, Peter; PLÍHAL, Tomáš; ŠIRAŇOVÁ, Mária. Impact of macroeconomic news, regulation and hacking exchange markets on the volatility of bitcoin. Journal of Economic Dynamics and Control. 2020. Vol. 119. DOI 10.1016/j.jedc.2020.103980.
17. ZHANG, Pengcheng; XU, Kunpeng; QI, Jiayin. The impact of regulation on cryptocurrency market volatility in the context of the COVID-19 pandemic — evidence from China. Economic Analysis and Policy. 2023. Vol. 80. DOI 10.1016/j.eap.2023.08.015.
18. SAKARIYAHU, Rilwan; LAWAL, Rodiat; ADIGUN, Rasheed; PATERSON, Audrey; JOHAN, Sofia. One crash, too many: Global uncertainty, sentiment factors and cryptocurrency market. Journal of International Financial Markets, Institutions and Money. 2024. Vol. 94. DOI 10.1016/j.intfin.2024.102028.
19. BAGH, Tanveer; KHAN, Muhammad Asif; IFTIKHAR, Kainat. Investor sentiment and cryptocurrency market dynamics: A Vector Autoregressive and Granger causality approach. Social Sciences and Humanities Open. 2025. Vol. 12. DOI 10.1016/j.ssaho.2025.102128.
20. KATSIAMPA, Paraskevi. Volatility estimation for Bitcoin: A comparison of GARCH models. Economics Letters. 2017. Vol. 158. DOI 10.1016/j.econlet.2017.06.023.
21. KOKI, Constandina; LEONARDOS, Stefanos; PILIOURAS, Georgios. Exploring the predictability of cryptocurrencies via Bayesian hidden Markov models. Research in International Business and Finance. 2022. Vol. 59. DOI 10.1016/j.ribaf.2021.101554.
22. MA, Feng; LIANG, Chao; MA, Yuanhui; WAHAB, M. I. M. Cryptocurrency volatility forecasting: A Markov regime-switching MIDAS approach. Journal of Forecasting. 2020. Vol. 39, no. 8. DOI 10.1002/for.2691.
23. HUDSON, Robert S.; GREGORIOU, Andros. Calculating and comparing security returns is harder than you think: A comparison between logarithmic and simple returns. International Review of Financial Analysis. 2015. Vol. 38. DOI 10.1016/j.irfa.2014.10.008.
24. ZHU, Yingzi; ZHOU, Guofu. Technical analysis: An asset allocation perspective on the use of moving averages. Journal of Financial Economics. 2009. Vol. 92, no. 3. DOI 10.1016/j.jfineco.2008.07.002.
25. HAN, Yufeng; ZHOU, Guofu; ZHU, Yingzi. A trend factor: Any economic gains from using information over investment horizons? Journal of Financial Economics. 2016. Vol. 122, no. 2. DOI 10.1016/j.jfineco.2016.01.029.
26. BRAUNEIS, Alexander; SAHINER, Mehmet. Crypto Volatility Forecasting: Mounting a HAR, Sentiment, and Machine Learning Horserace. Asia-Pacific Financial Markets. 2024. DOI 10.1007/s10690-024-09510-6.
27. FREITAS, Guilherme V. A. D.; LIMA, Matias A.; YAMACHI, Gustavo Y. O.; MACIEL, Leandro S. Market sentiment and the predictability of cryptocurrency risk premium using technical indicators. Finance Research Open. 2025. Vol. 1, no. 4. DOI 10.1016/j.finr.2025.100067.
28. JANKOVÁ, Zuzana. Critical Review of Text Mining and Sentiment Analysis for Stock Market Prediction. Journal of Business Economics and Management. 2023. Vol. 24, no. 1. DOI 10.3846/jbem.2023.18805.

29. LAMON, Connor; NIELSEN, Eric; REDONDO, Eric. Cryptocurrency Price Prediction Using News and Social Media Sentiment. SMU Data Science Review. 2017. Vol. 1, no. 3.
30. SHEN, Shulin; XIA, Le; SHUAI, Yulin; GAO, Da. Measuring news media sentiment using big data for Chinese stock markets. Pacific-Basin Finance Journal. 2022. Vol. 74. DOI 10.1016/j.pacfin.2022.101810.
31. ROGNONE, Lavinia; HYDE, Stuart; ZHANG, S. Sarah. News sentiment in the cryptocurrency market: An empirical comparison with Forex. International Review of Financial Analysis. 2020. Vol. 69. DOI 10.1016/j.irfa.2020.101462.
32. CHEN, Yu Lun; CHANG, Yung Ting; YANG, J. Jimmy. Cryptocurrency hacking incidents and the price dynamics of Bitcoin spot and futures. Finance Research Letters. 2023. Vol. 55. DOI 10.1016/j.frl.2023.103955.
33. SELMI, Refk. A war in a pandemic: The recent spike in economic uncertainty and the hedging abilities of Bitcoin. Economics Bulletin. 2022. Vol. 42, no. 3.
34. HOANG, Daniel; WIEGRATZ, Kevin. Machine learning methods in finance: Recent applications and prospects. European Financial Management. 2023. Vol. 29, no. 5. DOI 10.1111/eufm.12408.
35. JOHN, David L.; BINNEWIES, Sebastian; STANTIC, Bela. Cryptocurrency Price Prediction Algorithms: A Survey and Future Directions. 2024.
36. BOUTESKA, Ahmed; ABEDIN, Mohammad Zoynul; HAJEK, Petr; YUAN, Kunpeng. Cryptocurrency price forecasting – A comparative analysis of ensemble learning and deep learning methods. International Review of Financial Analysis. 2024. Vol. 92. DOI 10.1016/j.irfa.2023.103055.
37. OMOLE, Oluwadamilare; ENKE, David. Deep learning for Bitcoin price direction prediction: models and trading strategies empirically compared. Financial Innovation. 2024. Vol. 10, no. 1. DOI 10.1186/s40854-024-00643-1.
38. KOK, S. H.; ABDULLAH, Azween; JHANJHI, N. Z.; SUPRAMANIAM, Mahadevan. Prevention of crypto-ransomware using a pre-encryption detection algorithm. Computers. 2019. Vol. 8, no. 4. DOI 10.3390/computers8040079.
39. BASHER, Syed Abul; SADORSKY, Perry. Forecasting Bitcoin price direction with random forests: How important are interest rates, inflation, and market volatility? Machine Learning with Applications. 2022. Vol. 9. DOI 10.1016/j.mlwa.2022.100355.
40. BERGMEIR, Christoph; BENÍTEZ, José M. On the use of cross-validation for time series predictor evaluation. Information Sciences. 2012. Vol. 191. DOI 10.1016/j.ins.2011.12.028.
41. BERGMEIR, Christoph; HYNDMAN, Rob J.; KOO, Bonsoo. A note on the validity of cross-validation for evaluating autoregressive time series prediction. Computational Statistics and Data Analysis. 2018. Vol. 120. DOI 10.1016/j.csda.2017.11.003.
42. BERGSTRA, James; BENGIO, Yoshua. Random search for hyper-parameter optimization. Journal of Machine Learning Research. 2012. Vol. 13.
43. AKIBA, Takuya; SANO, Shotaro; YANASE, Toshihiko; OHTA, Takeru; KOYAMA, Masanori. Optuna: A Next-generation Hyperparameter Optimization Framework. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2019. DOI 10.1145/3292500.3330701.

44. HAMILTON, James D. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. Econometrica. 1989. Vol. 57, no. 2. DOI 10.2307/1912559.
45. LUNDBERG, Scott M.; LEE, Su-In. A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems. 2017.
46. GUIDOTTI, Riccardo; MONREALE, Anna; RUGGIERI, Salvatore; TURINI, Franco; GIANNOTTI, Fosca; PEDRESCHI, Dino. A survey of methods for explaining black box models. ACM Computing Surveys. 2019. Vol. 51, no. 5. DOI 10.1145/3236009.
47. RUDIN, Cynthia. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. 2019.
48. ARSENAULT, Pierre Daniel; WANG, Shengrui; PATENAUDE, Jean Marc. A Survey of Explainable Artificial Intelligence (XAI) in Financial Time Series Forecasting. ACM Computing Surveys. 2025. Vol. 57, no. 10. DOI 10.1145/3729531.
49. RIBEIRO, Marco Tulio; SINGH, Sameer; GUESTRIN, Carlos. "Why should I trust you?" Explaining the predictions of any classifier. In: NAACL-HLT 2016: Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Demonstrations Session. 2016. DOI 10.18653/v1/n16-3020.
50. MOLNAR, Christoph. Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. 2022. Available from: https://christophm.github.io/interpretable-ml-book/.
51. STROBL, Carolin; BOULESTEIX, Anne-Laure; ZEILEIS, Achim; HOTHORN, Torsten. Bias in random forest variable importance measures: Illustrations, sources and a solution. BMC Bioinformatics. 2007. Vol. 8. DOI 10.1186/1471-2105-8-25.
52. SHAPLEY, L. S. A Value for n-Person Games. In: Contributions to the Theory of Games (AM-28), Volume II. 2016.
53. LI, Meng; SUN, Hengyang; HUANG, Yanjun; CHEN, Hong. Shapley value: from cooperative game to explainable artificial intelligence. 2024.
54. ČERNEVIČIENĖ, Jurgita; KABAŠINSKAS, Audrius. Explainable artificial intelligence (XAI) in finance: a systematic literature review. Artificial Intelligence Review. 2024. Vol. 57, no. 8. DOI 10.1007/s10462-024-10854-8.
55. SAITO, Takaya; REHMSMEIER, Marc. The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. PLoS ONE. 2015. Vol. 10, no. 3. DOI 10.1371/journal.pone.0118432.
56. BAILEY, David H.; BORWEIN, Jonathan M.; LÓPEZ DE PRADO, Marcos; ZHU, Qiji Jim. The probability of backtest overfitting. Journal of Computational Finance. 2017. Vol. 20, no. 4. DOI 10.21314/JCF.2016.322.

1 Appendix

Comparison of cryptocurrency prices over time

2 Appendix

Target distribution for each cryptocurrency

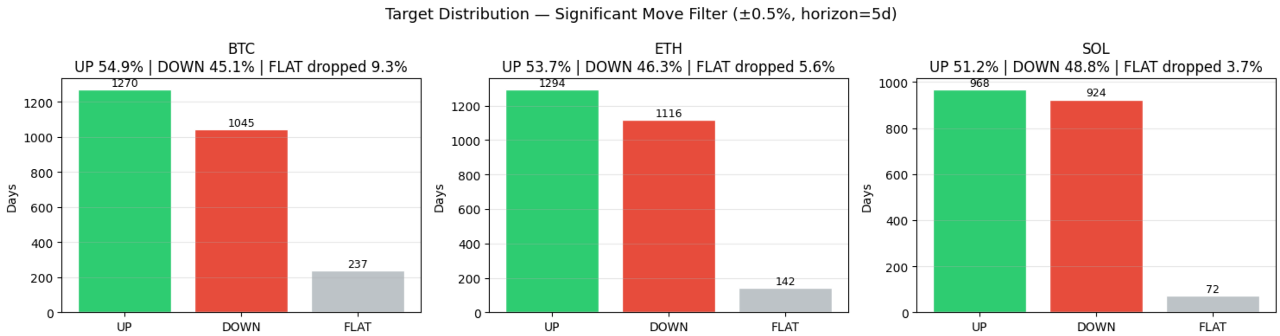

3 Appendix

Data splitting

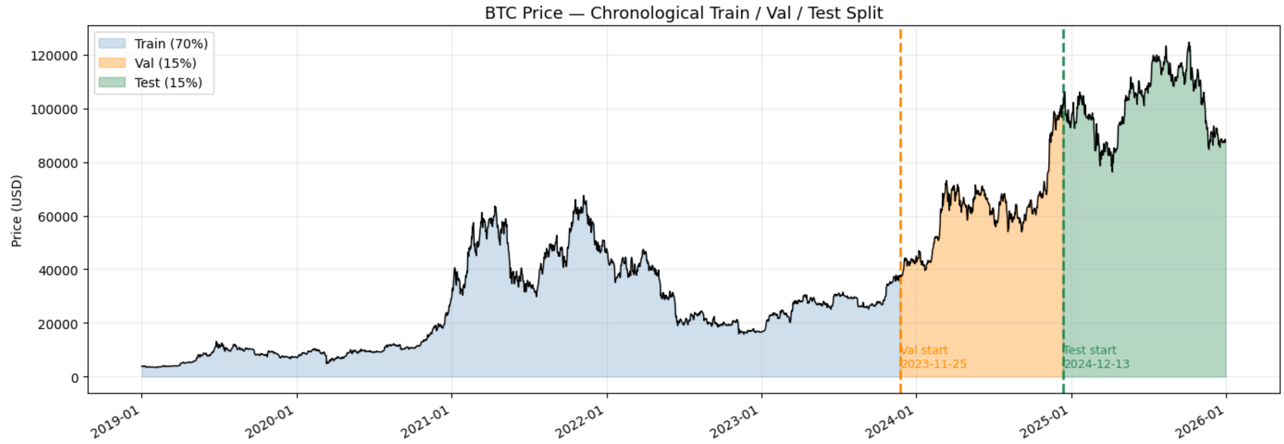

Crypto articles count and mean tone

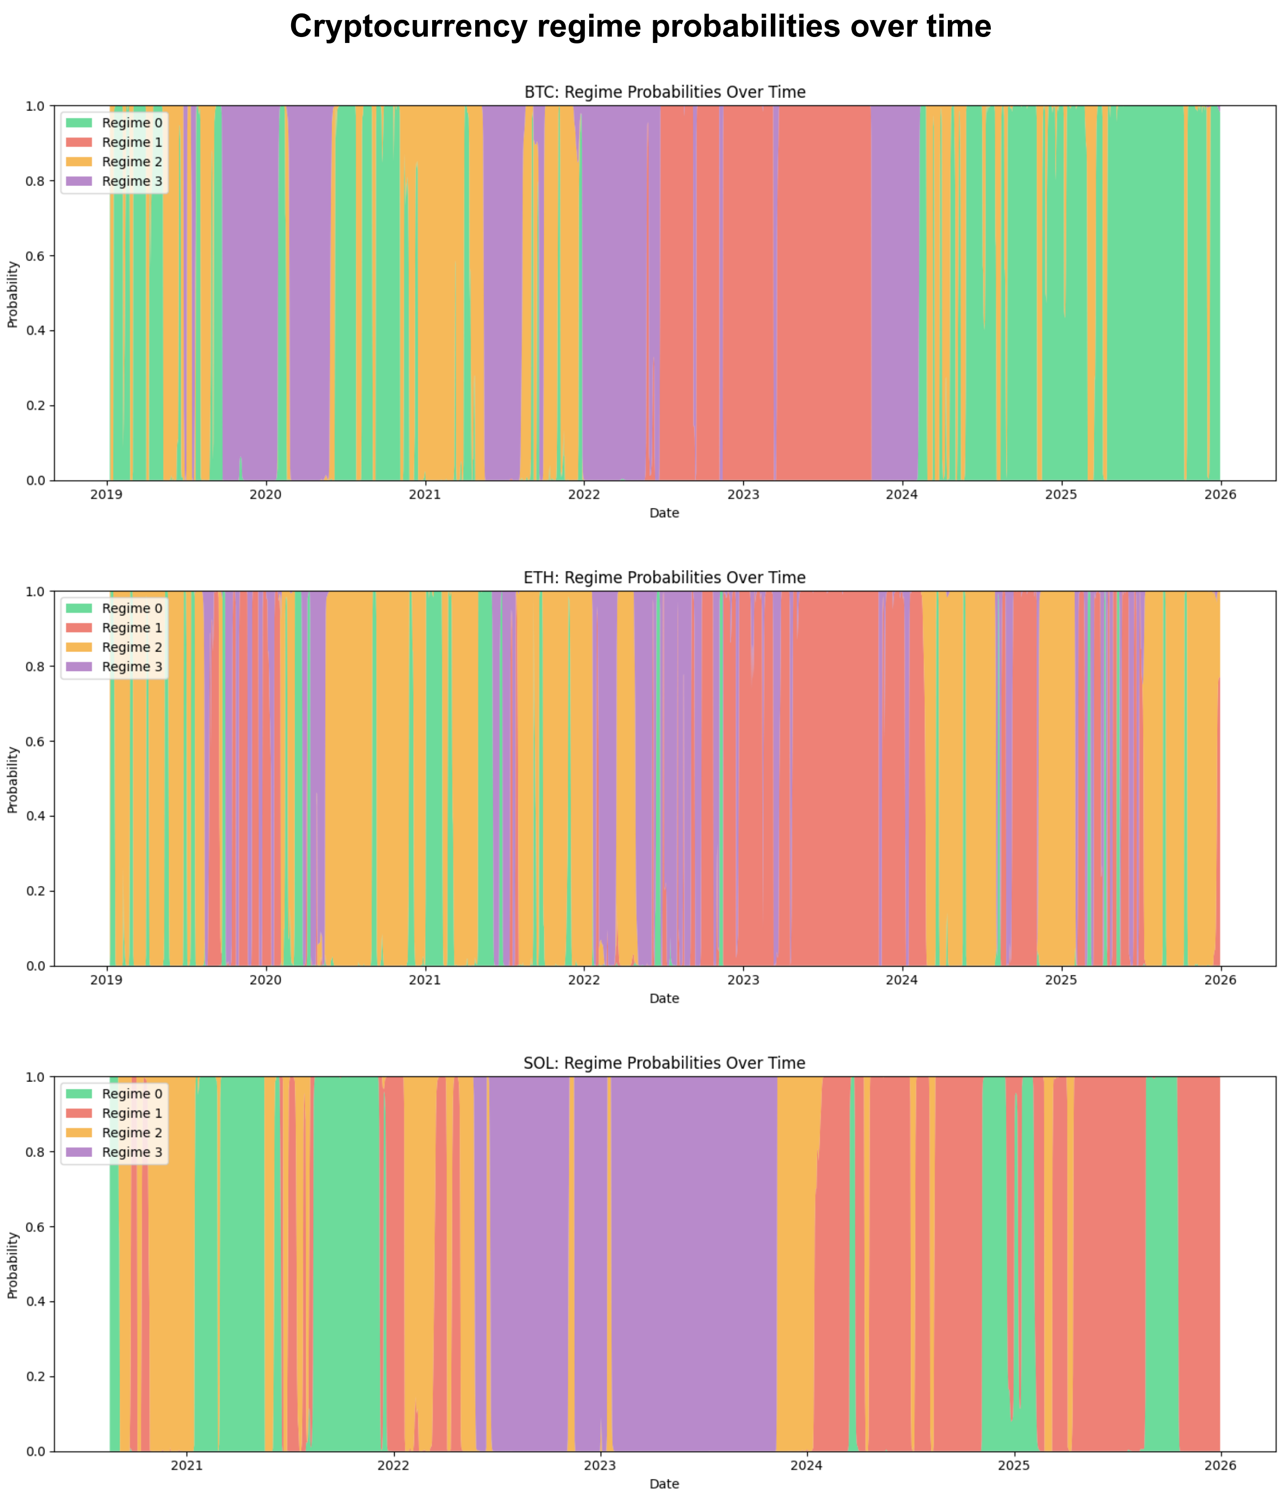
Cryptocurrency regime probabilities over time
BTC: Regime Probabilities Over Time
Regime 0
Regime 1
Regime 2
Regime 3
Probability
1.0
0.8
0.6
0.4
0.2
0.0
2019
2020
2021
2022
2023
2024
2025
2026
Date
ETH: Regime Probabilities Over Time
Regime 0
Regime 1
Regime 2
Regime 3
Probability
1.0
0.8
0.6
0.4
0.2
0.0
2019
2020
2021
2022
2023
2024
2025
2026
Date
SOL: Regime Probabilities Over Time
Regime 0
Regime 1
Regime 2
Regime 3
Probability
1.0
0.8
0.6
0.4
0.2
0.0
2021
2022
2023
2024
2025
2026
Date

6 Appendix

Empirical regime transition matrixes

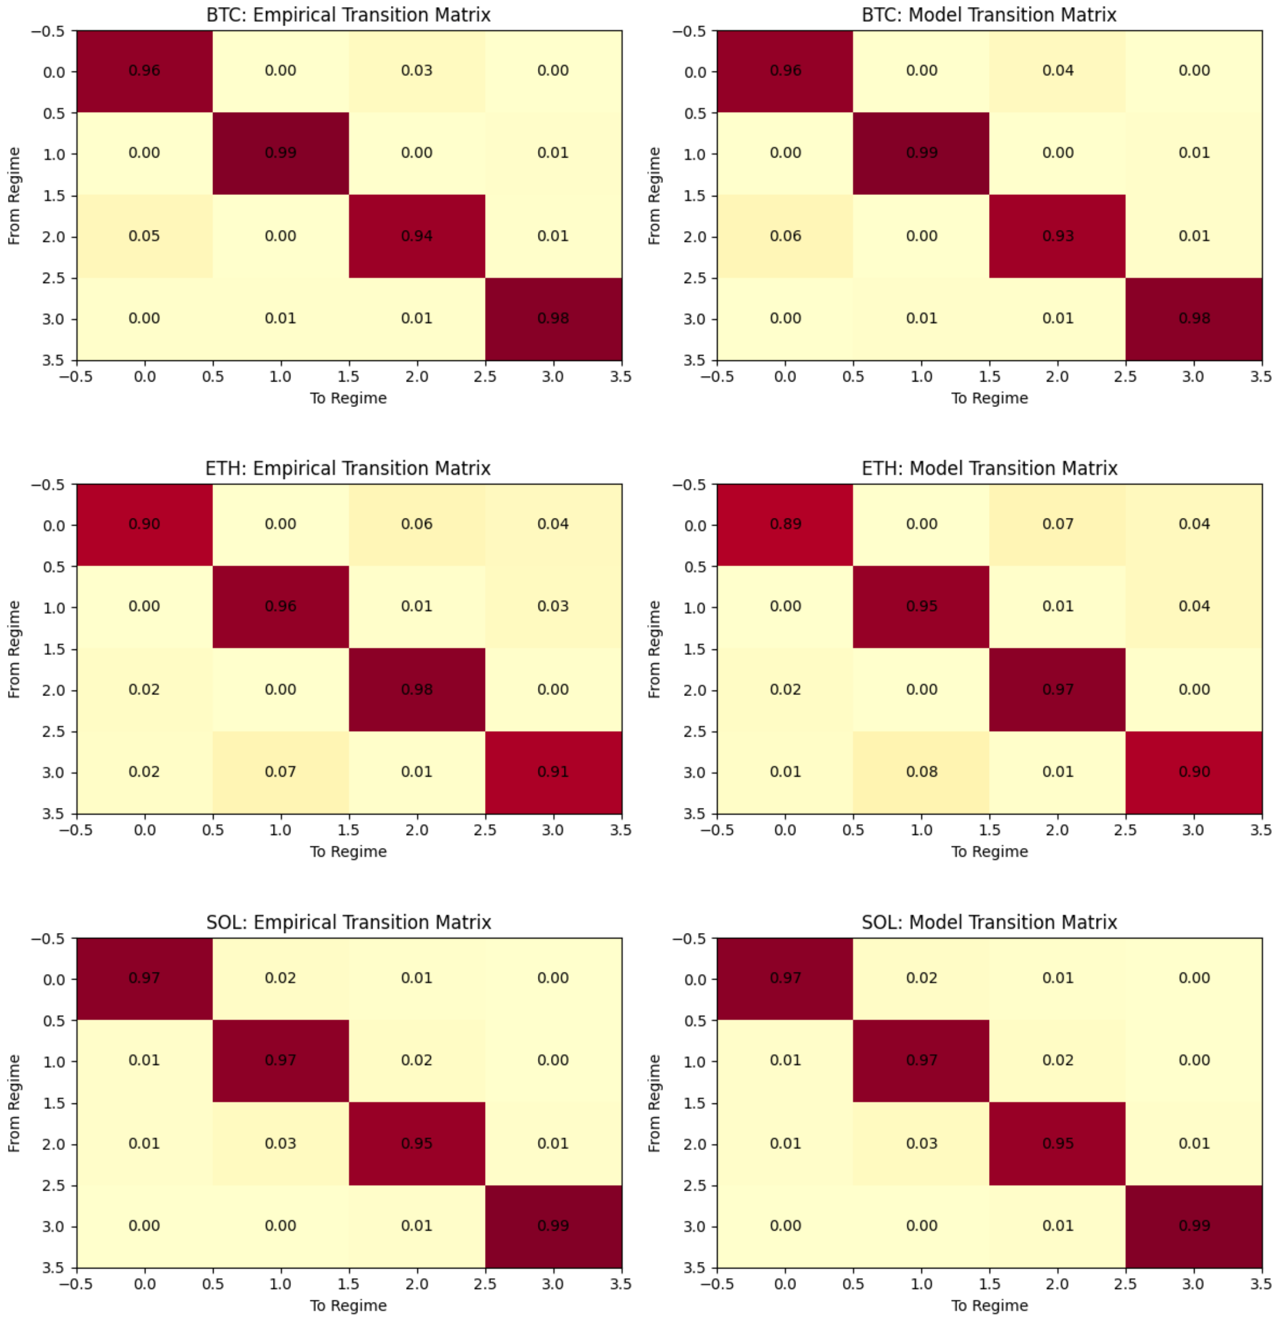

7 Appendix

Optuna results

	n_estimators	max_depth	Min_sample_leaf	Max_features	Max_samples	Best val score
BTC	150	3	14	0.4	~0.8	0.6844
ETH	100	3	50	sqrt	~0.84	0.7415
SOL	200	6	44	0.6	~0.8	0.7929

Example day waterfall of which variable influence the probability of UP for specific day

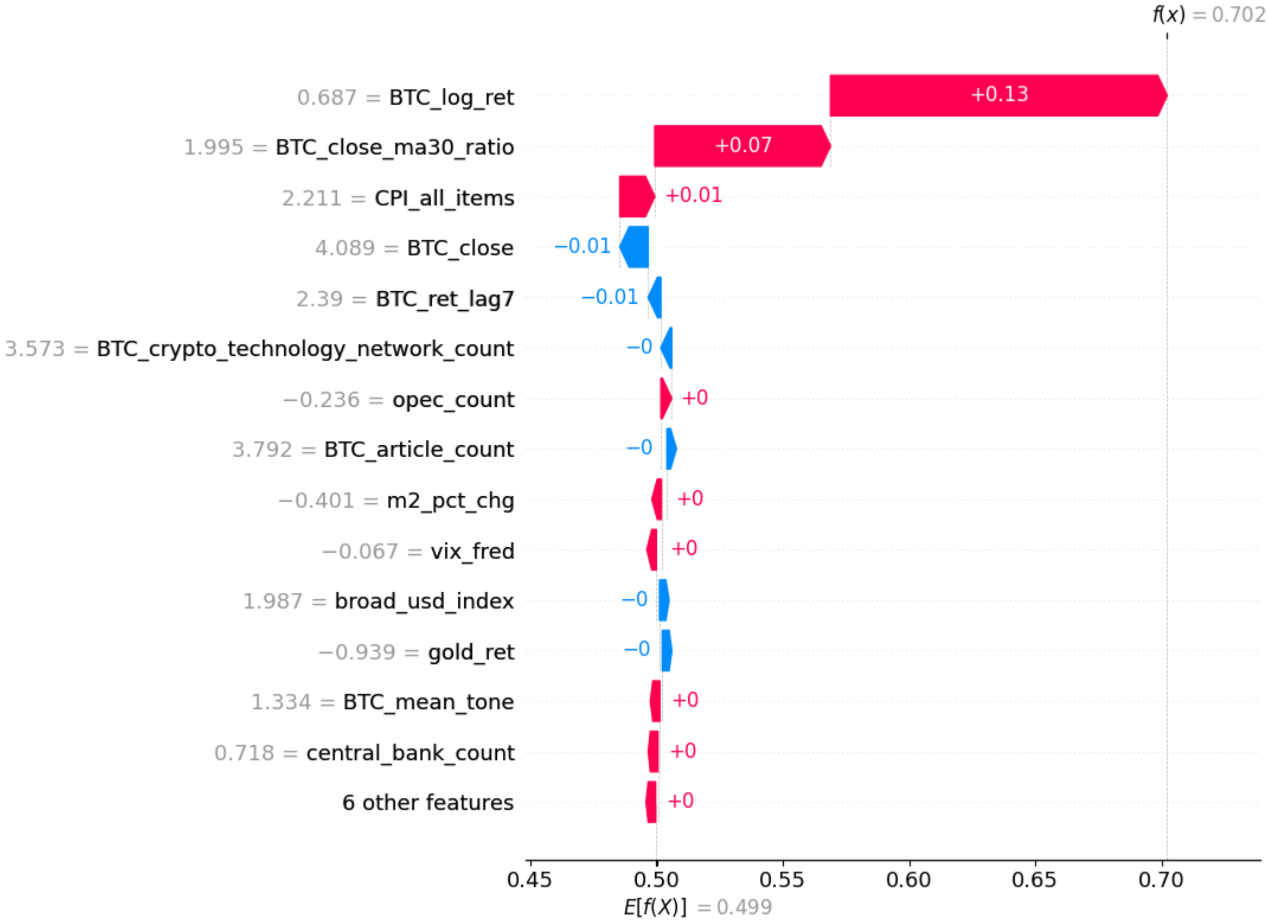

9 Appendix

Model performance per HMM regime on test set